NASA
High-End Computing Program User Needs Assessment 2020
Appendix D: User Inputs

HEC Program Manager: Tsengdar Lee

December 11, 2020

Coordinator: Michael M. Little

Executive Committee
Tsengdar Lee, NASA Headquarters
Aaron Pina, NASA Headquarters
Piyush Mehrotra, NASA Ames Research Center
William Thigpen, NASA Ames Research Center
Daniel Duffy, NASA Goddard Spaceflight Center
Robert Ferraro, Jet Propulsion Laboratory

Appendix D
User Needs Inputs

This Appendix contains all the input received from users and other data obtained from HEC Monthly Reports. It is organized around the sessions at the User Needs Assessment Workshop including the slides for each presentation and the notes taken. Slide decks have been edited to unused pages.

Directory of User Inputs

Appendix D0 Plenary Session Slides

ID	Speaker	Title
D0.1	Dr. Tsengdar Lee	HEC Needs Workshop
D0.2	Mike Little	How it works and guidance to presenters
D0.3	Dr. Karen Willcox	Computing Needs for Digital Twins
D0.4	Dr. Jack Kaye	HEC Needs Assessment
D0.5	Dr. Ellen Gertsen	SMD Strategy for Data Management and Computing for Groundbreaking Science 2019-2014

Appendix D1 Workshop Session 6-1 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-1.1	SMD	ESD	Da Silva	OSSE for Earth Science Space Missions
6-1.2	SMD	ESD	Nearing	Deep Learning for Hydrology Modeling
6-1.3	SMD	Helio	Guhathakurta	Heliophysics Overview
6-1.4	ARMD		Doebler	X-59 Sonic Doom Propagation Simulations
6-1.5	SMD	ESD	Larour	ISSM: Ice-sheet and Sea-level System Model
6.1-6	SMD	Helio	Kitiashvili	Realistic Modeling of the Sun

Appendix D2 Workshop Session 6-2 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-2.1	STMD		Korzun	Retropropulsion for Atmospheric Environments
6-2.2	ARMD, HEOMD, STMD		Kleb	Viscous Adaptive-Mesh CFD
6-2.3	STMD, ARMD		Nielsen	Exascale Code Development for Emerging Architectures High Performance Computing's Impact on Aerospace Prediction
6-2.4	STMD	GCD	Edquist	Supersonic Retropropulsion

6-2.5	ARMD	CST	Carter	Low Doom Flight Demonstration Signature Validation
6-2.6	NESC		Streett	LS Duffet / Aeroacoustics
6-2.7	ARMD		Roozeboom	Red Rover: Connecting Experimental and Computational Facilities and Communities

Appendix D3 Workshop Session 6-3 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-3.1	SMD	ESD	Giles	Aeronet NRT and Historical Inversion Processing
6.3-2	SMD	ESD	Partain, Cronk	GeoCarb Data Operations Center
6-3.3	SMD	ESD	Dungan	NASA Earth Exchange
6-3.4	SMD	ESD	Tenenbaum, Jenkins	TESS Science Processing Operations Center
6-3.5	SMD		von Allmen	First-Principles Modeling of Thermoelectric Materials
6-3.6	SMD	ESD	Vandal	Deep Learning for Geosationary Data
6-3.7	SMD	HECC	Ciotti	NASA Advanced Supercomputing

Appendix D4 Workshop Session 6-4 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-4.1	ARMD	TACP/ TTT	Wagner	Femera (Finite Element Analysis for Elasticity)
6-4.2	STMD	HPCI	Warner	Uncertainty Quantification (UQ) for Expensive, High Fidelity Models
6-4.3	STMD, SMD, HEOMD		Hill	Aerothermal Analysis of Planetary Entry Systems
6-4.4	ARMD	IASP, AAVP	Khorrami	System-level Airframe Noise Simulations for Civl Transports

Appendix D5 Workshop Session 6-8 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-8.1	ARMD	AAVP	Allan	RVLT High Fidelity Modeling – Lift+Cruise UAM Concept Vehicle
6-8.2	ARMD, HEOMD		Kenway	High Efficiency Database Generation with LAVA
6-8.3	ARMD, HEOMD		Cadieux	Scale-resolving Simulations with LAVA Cartesian AMR
6-8.4	ARMD	RVLT	Wang	Multidisciplinary Analysis and Optimization for Rotorcraft Aeromechanics

6-8.5	HEOMD	SLS	Gomez	HEO/Launch and Entry Aerosciences
6-8.6	ARMD	TACP	Malik	TTT-RCA Technical Challenge of Accurate CL_{max} Prediction
6.8-7	STMD, SMD		Darnhardt	Modeling Parachute Dynamics for Entry Systems
6-8.8	ARMD	TAC	Rogers	Transformational Tools and Technologies (T^3) Project HEC Usage
6-8.9	ARMD	TTT	Lawson	Computational Materials & Molecular Simulations

Appendix D6 Workshop Session 6-9 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6-9.1	SMD	ESD	Peters-Lidard, Kumar	Land Information System (LIS)
6-9.2	SMD	ESD	Le Moigne	AIST New Observing Strategies (NOS)
6-9.3	ARMD	AAVP	Jacobson	Computational Aeroelasticity
6-9.4	ARMD	AAVP, TACP	Choudhari	Laminar-Turbulent and Transition Physics & Modeling AAVP (HTP) and TACP (TTT)
6-9.5	SMD	ESD	Ramachandran	Data Production (Forward processing and Reprocessing)
6-9.6	SMD	ESD	Ramachandran for Maskey	Scaling pixel level smoke detection for high temporal geostationary satellite data
6.9-7	SMD	ESD	Owen	Large-scale InSAR analysis
6.9-8		HECC	Ranjan	AI/ML at Scale

Appendix D7 Workshop Session 6-19 Inputs

Session	Mission Directorate	Division/ Program	Presenter	HEC Project Title
6.19-1	SMD	ESD	Putman	Global Simulations of increasing resolution and complexity
6-19.2	SMD	ESD	Privé	Meteorological OSSEs
6-19.3	SMD	ESD	Todling	Forefront Numerical Weather Analysis & Prediction
6-19.4	SMD	ESD	Gelaro	Coupled Analysis/Reanalysis: MERRA-3 and Deyond
6-19.5	SMD	ESD	Molod	S2S2D Prediction and Predictability: Retrospective Forecast Suite
6-19.6	SMD	ESD	Pawson	GEOS-CF
6-19.7	SMD	ESD	Colarco	Chemistry-Climate Modeling
6-19.8	SMD	ESD	da Silva	OSSE for Earth Science Missions (presented on 6-1)

6-19.9	SMD	ESD	Gelaro	GMAO support to JCSDA

Appendix D8 Non-workshop Inputs

ID	Mission Directorate	Division/ Program	Presenter	HEC Project Title
D8.1	SMD	Astro	Kenyon	Simulation of Planetary Formation
D8.2	SMD	Astro	Powell	TESS Full Frame Light Curves
D8.3	SMD	Astro	Jiang	A Global 3-D Radiation Magneto-hydrodynamic simulation of super-Eddington Accretion Disks
D8.4	SMD	Astro	Brooks	Simulation Galaxy Formation and Evolution
D8.5	SMD	Astro	Cen	Galaxy Formation and Evolution
D8.6	SMD	ESD	Su	JPL OSSE Support
D8.7	SMD	ESD	Chu	OCO-2 and OCO-3 Science Data Processing
D8.8	SMD	Helio	DeVore	Solar Corona and Space Weather Models
D8.9	SMD	Helio	Chen	MMS/RBSP/THEMIS/Psyche
D8.10	ARMD		Malik	Need for Modern HPC Hardware to Solve NASA Challenge Problems (Draft)
D8.12	ARMD		Capps	Industry Inputs into Future NASA Project Planning (9/2/19)
D8.13		DT	McLarney	DT Artificial Intelligence and Machine Learning APMC 6/11/20
D8.14	ARMD		Malik	NASA CFD Vision 2030 Study – A path to Revolutionary Computational Aerosciences (9/30/14)
8-15	ARMD	AOSP	Oza	Machine Learning for Aviation Safety
8-16	ARMD	TTT, AATT	Gregory/ Schneck	Supporting Safe, Robust Aerospace Structures Through HPC Enhanced Computational NDE
D8.17	SMD	HECC	Thigpen	HECC Overview, March 2019
D8.18	ARMD		Capps	Industry Inputs into Future NASA Project Planning, March 6, 2019
D8.19	OCIO	DT	McLarney	Digital Transformation: AI/ML Briefing to NASA Administrators Program Management Council
D8.20	ARMD		Malik	NASA CFD 2030 Study September, 2014
D8.21	OCE	HPCI	Gregory	Supporting Safe, Robust Aerospace Structures Through HPC Enhanced Computational Non-destructive Testing (CNDE)
D8.22	SMD	ESD	Fenty	ECCO-ACCESS: a Cloud-Native Storage and Data Analysis System for Ocean Climate Research

D8.23	SMD	NAS	Mehrotra	NASA's Drive Towards Exascale Computing, July 19, 2020
D8.24	SMD	ESD	Donnellan	QuakeSim: Multi-Source Synergistic Data Intensive Computing for Earth Science

Appendix D9 HEC Monthly Status Report Project Highlights

Note these documents are in the form of a 1-page highlight

ID	Mission Directorate	Division/ Program	HEC Project Title
D9.1	HEOMD	Commercial Collaboration	Simulating Dream Chaser ® Spaceplane Aerodynamics
D9.2	SMD	ESD	Simulations Probe the Sun's Effect on Climate
D9.3	SMD	ESD	Upgrade Markedly Improves Skill of NASA Seasonal Prediction System
D9.4	ARMD		Predicting High-Altitude Relight in Aircraft Engines
D9.5	SMD	ESD	NCCS Hosted Models Look at Fires on Two Continents
D9.6	ARMD		New CFD Methods for Predicting Quadcopter Drone Noise
D9.7	SMD	ESD	Realistic Simulations of the Coupled Atmosphere-Ocean-Ice System
D9.8	HEOMD		Building Booster Separation Aerodynamic Databases for Artemis II
D9.9	SMD	ESD	GMAO Visualizes Smoke Transport from Australian Fires
D9.10	SMD	ESD	Applications Performance & Productivity Team Helps Improve Ice Sheet Model Performance
D9.11	ARMD		Minimizing Sonic Boom Through Simulation-Based Design: The X-59 Airplane
D9.12	SMD	ESD	NCCS Accelerates Simulation of Atmospheric Chemistry Using Machine Learning
D9.13	ARMD		Simulating a Full-Scale, Large Airliner Landing
D9.14	SMD	ESD	NCCS Accelerates Fine-Scale Forest Modeling Across the North American Boreal Zone
D9.15 (2 pg)	SMD	Helio	Researchers Model Superflare from Sun-Like Star at NCCS
D9.16	STMD		Using Retrorockets for Human Exploration of Mars
D9.17	SMD	ESD	OCO-2 and GEOS Team Up to Produce a New View of Carbon Dioxide
D9.18	SMD	ESD	Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability
D9.19	SMD	ESD	Enabling NASA Earth Exchange (NEX) Petabyte Data Production Systems
D9.20	SMD	ESD	GMAO and NCCS Begin ORACLES Support
D9.21	SMD	ESD	Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability

D9.22	SMD	ESD	Using Katrina and Sandy Data to Improve Hurricane Prediction Tools
D9.23	SMD	ESD	Cloud Resolving Models – Cloud Library
D9.24	SMD	ESD	Projecting Sea Level Rise by Modeling the Evolution of Ice Sheets
D9.25	SMD	ESD	Using GEOS-5 Nature Run Results to Characterize Impact of Sampling Frequency on DSCOVR-Derived Energy Budget of the Earth
D9.26	SMD	ESD	KORUS-AQ Field Campaign Support
D9.27	SMD	ESD	Estimating Trees and Shrubs in the Sub-Sahara With Satellite Imagery
D9.28	SMD	ESD	GISS Simulations Show Farms Are a Major Source of Air Pollution
D9.29	SMD	ESD	High Resolution Particulate Matter Studies Are Fueled by 1-km MODIS AOD, Generated from MAIAC Algorithm on Discover
D9.30	SMD	ESD	Virtual Earth Observing via GMAO's GEOS-5 Global 1.5-km Simulation
D9.31	SMD	ESD	30-Year Global Fire Weather Analysis Based on MERRA and Ground Station Observations Sets Stage for Contributions from GPM, TRMM, and GPCP
D9.32	SMD	ESD	Estimating Woody Biomass on South Side of the Sahara at the 40-50 cm Scale Using Amazon Web Services (AWS): Progress for Intel Challenge "Head in the Clouds" Project
D9.33	SMD	ESD	March 2016 Field Campaign Support (ENRR, AfriSAR)
D9.34	SMD	ESD	Discover Supports Twice-Daily 7-Day GEOS-5 Experimental 6-km Global Forecasts
D9.35	SMD	ESD	Pleiades-Enabled Reprocessing by OCO-2 Provides a Consistent Ten-Month Data Record
D9.36	SMD	ESD	GEOS-5 Chemistry Climate Model Simulation of Stratospheric Ozone Intrusion Featured at SIGGRAPH 2015 Conference
D9.37	SMD	ESD	Modeling the Carbon Balance in Arctic Ecosystems
D9.38	SMD	ESD	Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring
D9.39	SMD	ESD	Pleiades' computational power has enabled improved understanding of methane (CH_4) fluxes
D9.40	SMD	ESD	Discover Expansion Powers Progress in Dynamical Downscaling Studies with the NASA Unified-WRF Model
D9.41	SMD	ESD	NASA Releases Detailed Global Climate Change Projections
D9.42	SMD	ESD	NASA Releases Detailed Global Climate Change Projections
D9.43	SMD	ESD	Climate Model Data Services (CDS) Support NASA's Response to Nepal Earthquake

D9.44	SMD	ESD	Real-Time Assimilation of ISS-RapidScat Observations
D9.45	SMD	ESD	Pleiades-Run Simulations Provide A Virtual Telescope on the World's Oceans and Sea Ice
D9.46	SMD	ESD	CISTO Climate Data Services Staff Cited in Federal Lab Consortium's 2015 Interagency Partnership Award for UV-CDAT
D9.47	SMD	ESD	HECC is Used to Clear 5 Years of UAVSAR Data Processing Backlog and Reduce Latency
D9.48	SMD	ESD	1.5km Global Simulation with GEOS-5 on SCU10 at NCCS
D9.49	SMD	ESD	GEOS-5 Aerosol and Chemistry Forecasts Support CalWater 2/ACAPEX Field Campaigns
D9.50	SMD	ESD	Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring
D9.51	SMD	Astrophysics	FOGGIE: Simulating the Cosmic Fog Around Galaxies

Appendix D0 Plenary Session Slides
Guidance to Presenters

Why HEC Needs Workshop? Why Now?

1. The last HEC Needs workshop was done in 2013. SMD Data Management and Computing Strategy (2020) calls for routine assessments.
2. We're not prepared for the end of Moore's Law – David Rotman, MIT Technology Review (Feb, 2020)
   - Moore's law has ended. We cannot rely on the automatic performance gain anymore.
   - Finding successors to today's silicon chips will take years of research. If you're worried about what will replace Moore's Law, it's time to panic.
3. Big data analytics, including ML/AI, has disrupted the general processor development and sped up the specialization in computing processors.
4. New HEC acquisition strategy may be needed in the next 10-20 years.

HEC Needs Workshop 2

Current HEC Landscape

- The Decline of Computers as a General Purpose Technology: Why Deep Learning and the End of Moore's Law are Fragmenting Computing (Neil Thompson & Svenja Spanuth, 2018)

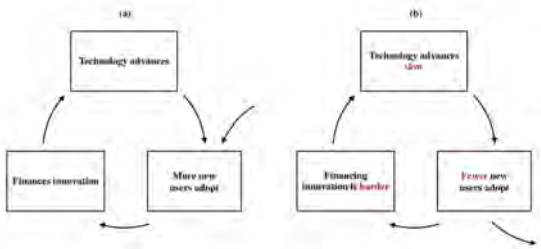

- Left Behinds:
  - Applications whose current algorithms are ill-suited to specialization
  - Applications with insufficient demand
  - Fragmented application users that fail to coordinate

Post Moore's Law Computational Modeling Strategy (Thompson et. al.)

- Performance gain will not come from semiconductor physics and silicon-fabrication technology
  - Past software development paradigm will not work
  - Legacy software will have a hard time realizing new specialized computing architecture
- Performance gains come from the entire computing stack (software, algorithm, hardware)
  - Software blot will need to be excised
  - Algorithms will need to be rethought
  - Hardware will become less general-purpose

What's Included in the "Needs?"

- Computing architecture
- Computing cycles (be specific on CPU or GPU), storage, networking
- Data analytics and/or data sharing infrastructure
- Code modernization with specific performance target
- Algorithm development

Thank You!

Introduction for Dr. Karen Willcox

Director of the Oden Institute for Computational Engineering and Sciences.

Dr. Karen Willcox is currently the director of the Oden Institute for Computational Engineering and Sciences and a professor of Aerospace Engineering and Engineering mechanics at UT/Austin. She was born in Auckland, New Zealand and attended the University of Auckland where she earned a Bachelor's of Science in 1994. She then attended MIT receiving a Master's Degree in Aeronautics and Astronautics in 1996 and a PhD in 2000. In 2008 she helped MIT found the MIT Center for Computational Engineering and had visiting appointments in international institutions.

Dr. Willcox's research work started out as Reduced-Order Aerodynamic Models for Aeroelastic Control of Turbomachines and has evolved in several different ways which you'll hear about today.

Dr. Willcox's honors include a Member of the New Zealand Order of Merit and a Fellow of the Society for Industrial and Applied Math. She also made the shortlist of NASA Astronaut candidates in 2009 and 2013.

We're very glad to have her here today to keynote our discussions for the next 10 days of High-end Computing Needs. Dr. Willcox's talk will introduce the exciting new area of Digital Twins and computational needs to support them.

The Team

Michael Kapteyn
MIT

Stefanie Salinger
UT Austin

Dr. David Knezevic
CTO, Akselos

Luwen Huang
MIT Mapping Lab

Cory Kays
Aurora Flight Sciences

Funding sources: US Air Force **Computational Math** Program (F. Fahroo);
US Department of Energy **AEOLUS MMICC** (S. Lee, W. Spotz);
US Air Force **Dynamic Data Driven Application Systems** Program (E. Blasch);
The **Boeing** Company;
SUTD-MIT **International Design Centre**

Outline

1 **A Role for AI in Science & Engineering**

from forward simulation to predictive data science

2 **Towards Predictive Digital Twins**

via component-based reduced-order models and interpretable machine learning

3 **Conclusions & Outlook**

Beyond forward simulation

Outer-Loop Applications

"Computational applications that form outer loops around a forward simulation."

Examples: optimization, design space exploration, uncertainty quantification, inverse problems, control, data assimilation, decision support

Scientific AI / Scientific ML could clearly play a large future role in our outer-loop applications

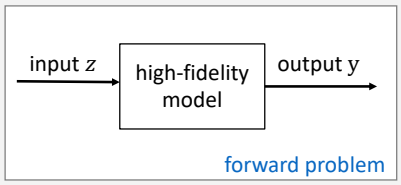

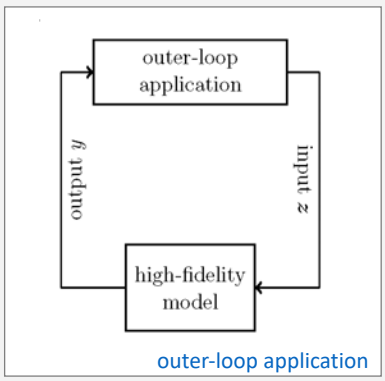

4

Scientific Machine Learning

"Scientific machine learning (SciML) is a core component of artificial intelligence (AI) and a computational technology that can be trained, with scientific data, to augment or automate human skills.

Across the Department of Energy (DOE), SciML has the potential to transform science and energy research. Breakthroughs and major progress will be enabled by harnessing DOE investments in massive data from scientific user facilities, software for predictive models and algorithms, high-performance computing platforms, and the national workforce."

5

AI For Science

"we use the term 'AI for Science' to broadly represent the next generation of methods and scientific opportunities in computing, including the development and application of AI methods (e.g., machine learning, deep learning, statistical methods, data analytics, automated control, and related areas) to build models from data and to use these models alone or in conjunction with simulation and scalable computing to advance scientific research"

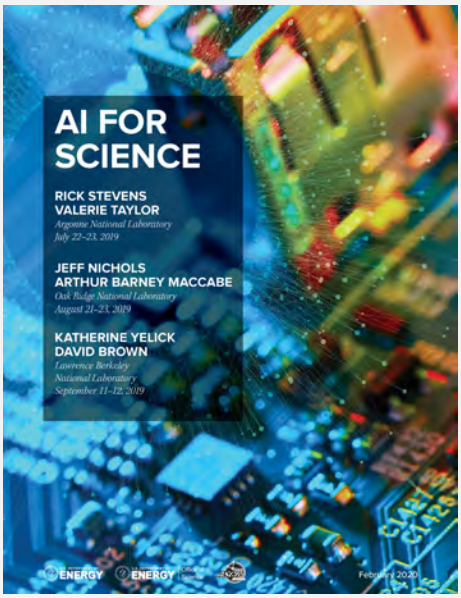

6

Scientific Machine Learning

What are the opportunities and challenges of machine learning in complex applications across science, engineering, and medicine?

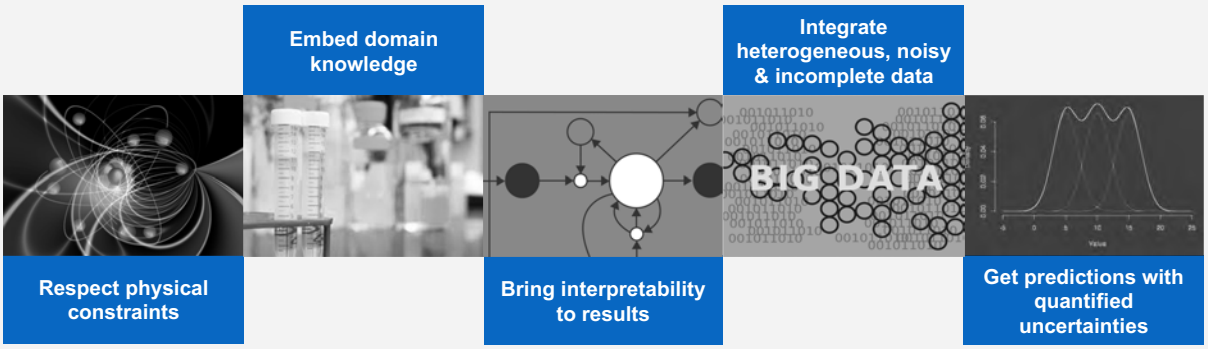

7

ODEN INSTITUTE
FOR COMPUTATIONAL ENGINEERING & SCIENCES

How do we harness the explosion of data to extract knowledge, insight and decisions?

BIG DECISIONS

need more than just big data…

They need **BIG MODELS** too.

Inspired by
Coveney, Dougherty, Highfield "Big data need big theory too"

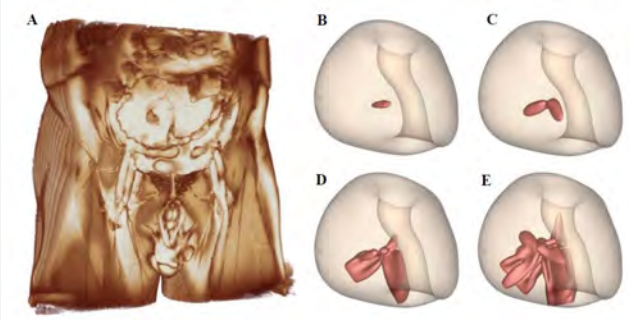

Patient-specific prostate tumor modeling (T. Hughes)

Predictive Digital Twin (K. Willcox)

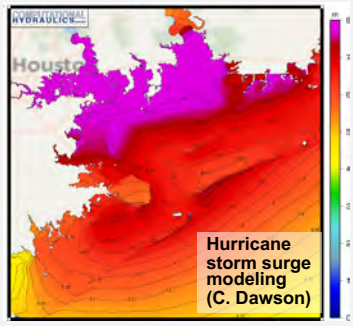
Hurricane storm surge modeling (C. Dawson)

BIG DECISIONS

need more than just big data…

1. **Complex multiscale multiphysics phenomena**
   driving the dynamics of high-consequence applications
2. **High dimensional parameters**
   underlying the characterization of scientific and engineering systems
3. **Data are sparse, intrusive and expensive to acquire**
   especially in the most critical regimes
4. **Rare events**
   design decisions certified against small probabilities of failure (10^{-6} – 10^{-9})
5. **Uncertainty quantification**
   in model inference and certified predictions in regimes beyond training data

BIG DECISIONS

need **BIG MODELS** too.

"Computational Science
or Computational Science & Engineering (CSE)

is an interdisciplinary field that uses mathematical modeling and advanced computing to understand and solve complex problems. At its core CSE involves developing models and simulations to understand physical/natural systems.

BIG DECISIONS must incorporate the **predictive power**, **interpretability**, and **domain knowledge** of physics-based models.

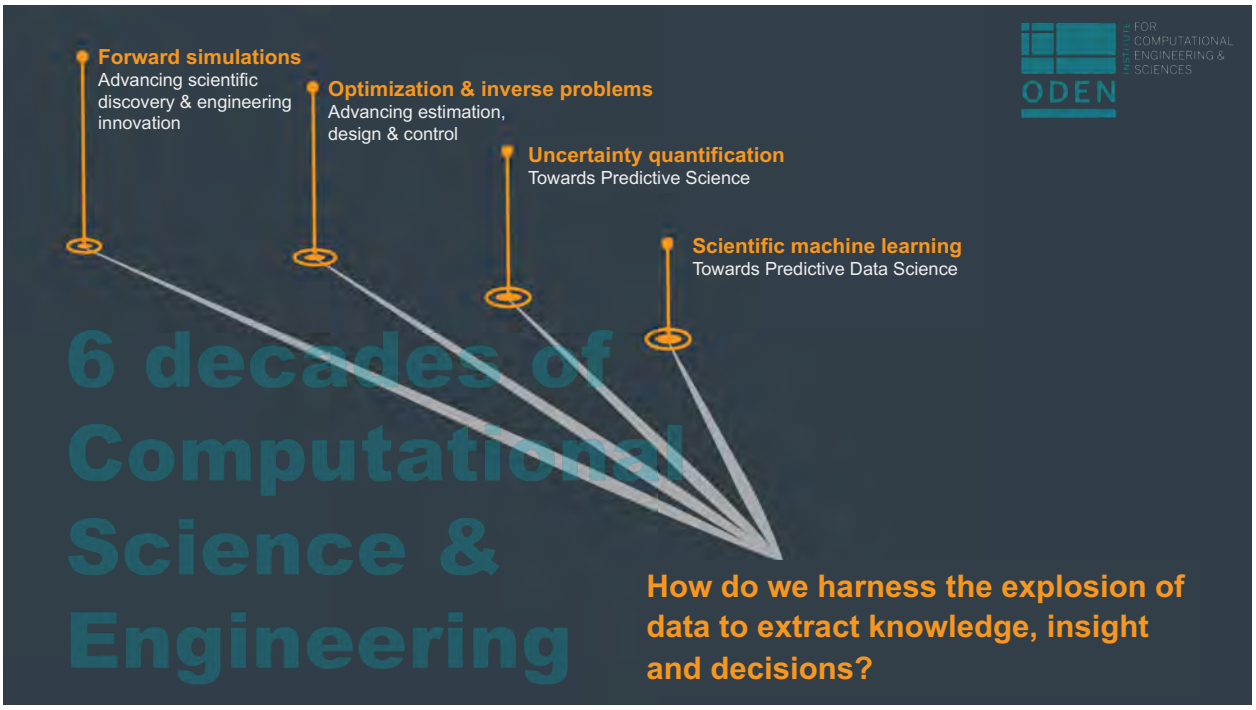

1 AI in Science & Engineering

2 Predictive Digital Twin

3 Conclusions & Outlook

Towards Predictive Digital Twins

via component-based reduced-order models and interpretable machine learning

Digital twins enable data-driven decisions

High-consequence decisions require digital twins that are
predictive • reliable • explainable

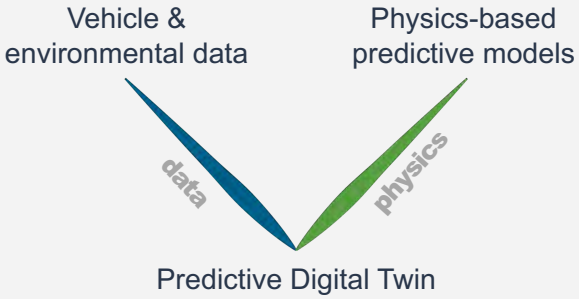

self-aware aircraft

Our **digital twin adapts** to the **evolving UAV structural health…**

providing **near real-time capability estimates** that enable **dynamic decision-making**

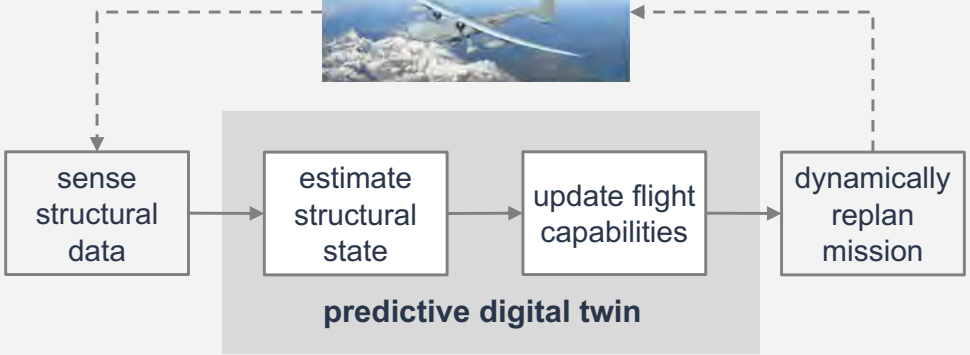

self-aware aircraft

Physics-based models

- physical laws define low-dimensional manifolds
- simulate new previously unseen scenarios
- have quantifiable uncertainty
- parameters represent real-world quantities

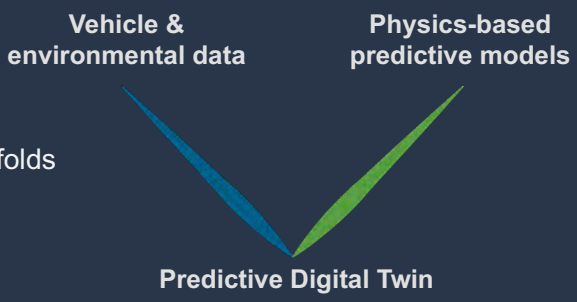

But physics-based models are too complex and too expensive for use in near real-time onboard decision-making…

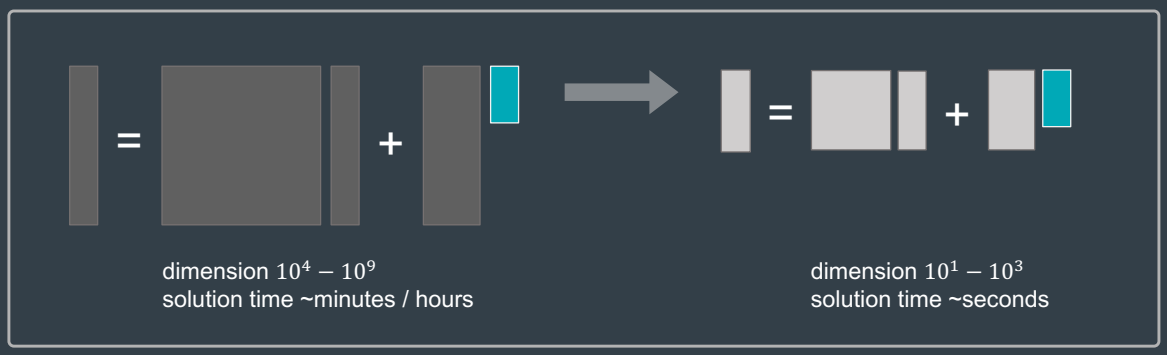

Projection-based model reduction

1 **Train**: Solve PDEs to generate training data
2 **Identify structure**: Compute a low-dimensional basis
3 **Reduce**: Project PDE model onto the low-dimensional subspace

Machine learning

"The scientific study of algorithms & statistical models that computer systems use to perform a specific task without using explicit instructions, relying on patterns & inference instead." [Wikipedia]

Reduced-order modeling

"Model order reduction (MOR) is a technique for reducing the computational complexity of mathematical models in numerical simulations." [Wikipedia]

What is the connection between reduced-order modeling and machine learning?

Model reduction methods have grown from CSE, with a focus on ***reducing* high-dimensional models** that arise from physics-based modeling, whereas machine learning has grown from CS, with a focus on ***creating* low-dimensional models** from black-box data streams.

[Swischuk et al., *Computers & Fluids*, 2018]

Machine learning

Perspectives stem from Computer Science (CS), with focus on ***creating* efficient models** from black-box data streams, relying on patterns & inference rather than explicit instructions

Reduced-order modeling

Perspectives stem from Computational Science & Engineering (CSE), with focus on ***reducing* high-dimensional models** that arise from physics-based modeling

Machine learning & physics-based modeling
Can we get the best of both worlds?

Discover hidden structure	Embed governing equations
Non-intrusive implementation	Structure-preserving
Black-box & flexible	Predictive (error estimators)
Accessible & available	Stability-preserving

18

Reduced-order modeling leads to low-cost physics-based models that enable predictive digital twins

Challenges & limitations

- **training** is expensive
- scaling to **high-dimensional parameters**
- dealing with **discontinuous parameter** dependence

Approach

- Static-Condensation Reduced-Basis-Element (SCRBE) method [Huynh 2013]

"Divide and conquer"

AKSELOS

Example component: section of a wing

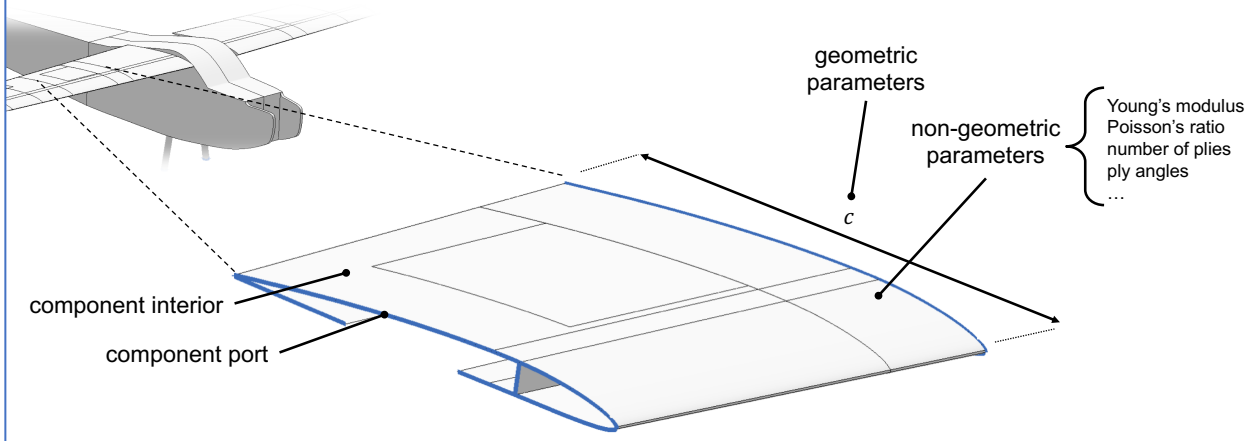

Example component: section of a wing

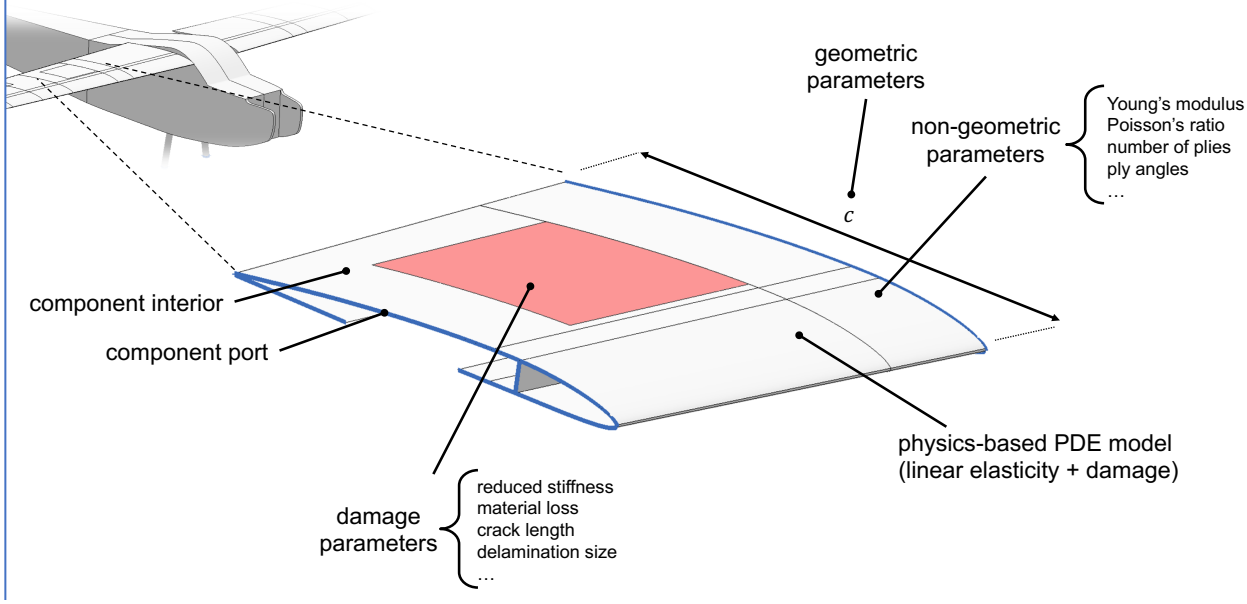

Example component: section of a wing

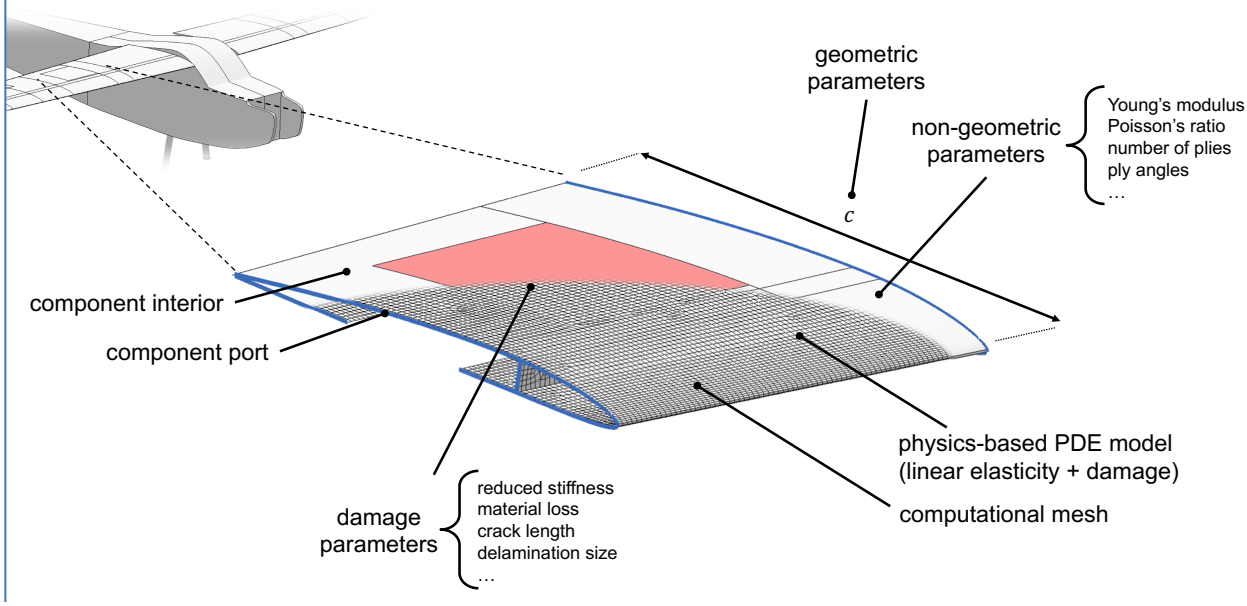

local effects
(component parameters)

+

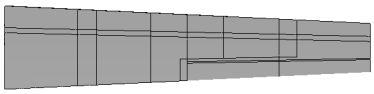

interactions
(assembly parameters)

+

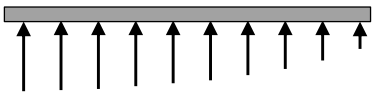

context
(loads parameters)

A complex nonlinear system is more than just the sum of its pieces

Predictive Digital Twin

via component-based ROMs and interpretable machine learning

ROMs embed predictive modeling and reduce the cost of training

Offline:

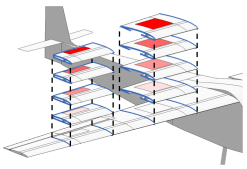

Construct library of ROMs representing different asset states

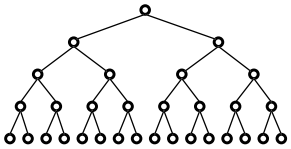

Use model library to train a classifier that predicts asset state based on sensor data

Online:

sensor data

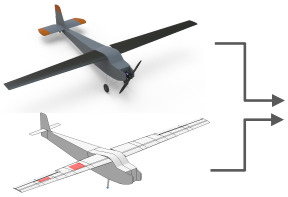

current Digital Twin

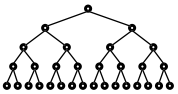

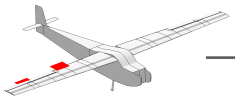

updated Digital Twin

Analysis
Prediction
Optimization

[Kapteyn, Knezevic, W. AIAA Scitech 2020]

Customized 12ft Telemaster aircraft

- Complex structure with multiple materials
- Custom wing sets: pristine & damaged
- Custom sensor suite

Temperature, pressure and humidity sensors

3 axis accelerometer
3 axis gyro

Dual high-frequency dynamic strain and vibration sensors

Internal structure

*Willcox has a family member who is co-founder of Jessara. Purchase of the sensors for use in the research was reviewed and approved in compliance with all applicable MIT policies and procedures.

Physics-based Digital Twin

Finite element model: multiple material types (carbon fiber, carbon rod, plywood, foam) & multiple element types (solid, shell, beam)

Reduced model: 0.03 seconds per structural analysis (cf. 55 seconds for the finite element model)

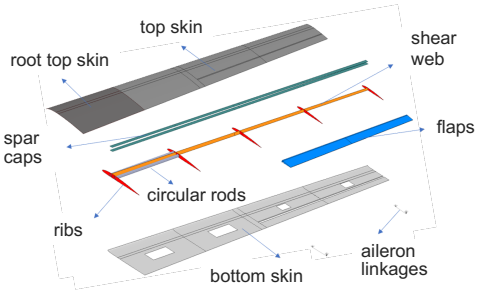

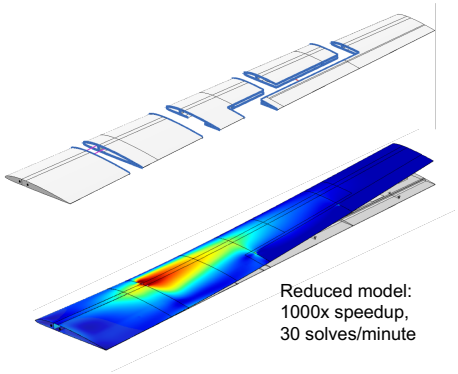

From component-based model to digital twin: **Physics-based library**

Offline: Construct a library of damage states for each component

- Create multiple copies of each component
- Train components for parameter ranges of interest (local + interactions)
- Compute associated aircraft structural load constraints (context)

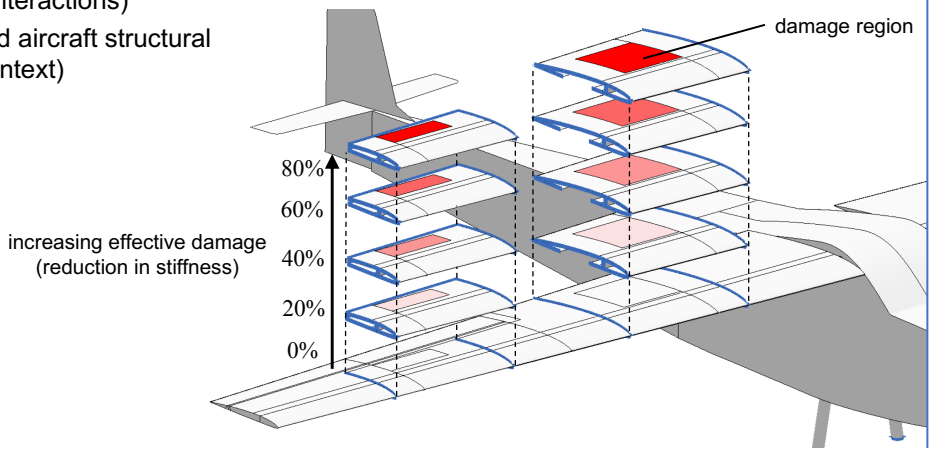

Onboard sensors inform which model is used in the digital twin

Data-driven digital twin:

Onboard sensors are used to select a reduced-order model from the library

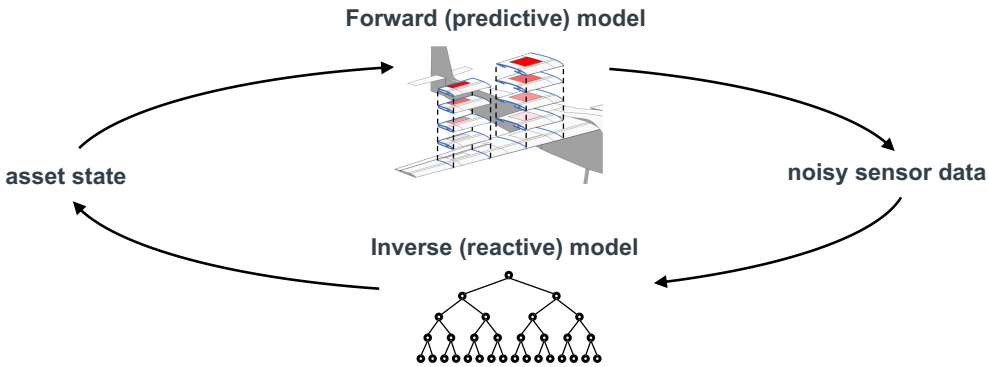

- Use **predictive models** to **generate training data**
- Use **machine learning** to train an **interpretable, explainable reactive model**

13

From component-based model to digital twin: **Interpretable machine learning**

Offline: Train a classifier using simulation data

- Optimal Classification Trees [Bertsimas & Dunn, 2017]
- Highly interpretable
- Natural framework for sensor selection
- Online classification is rapid

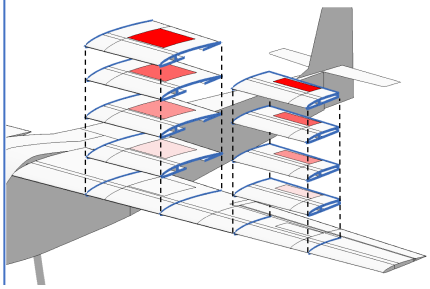

Component 1

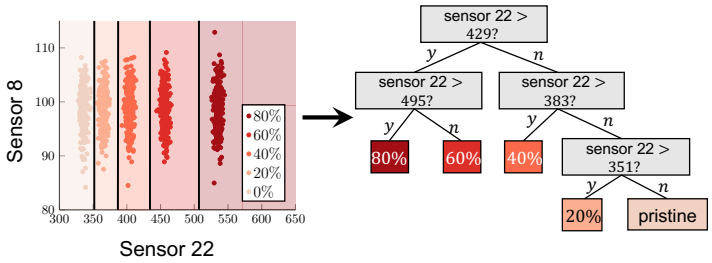

Component 2

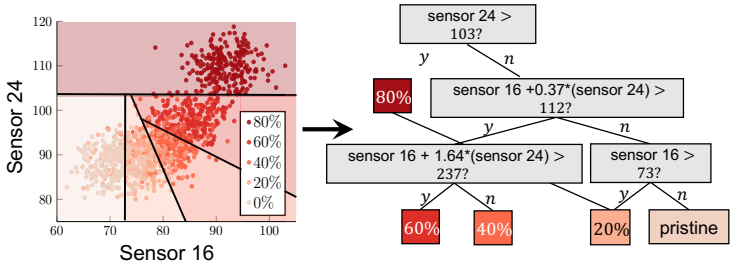

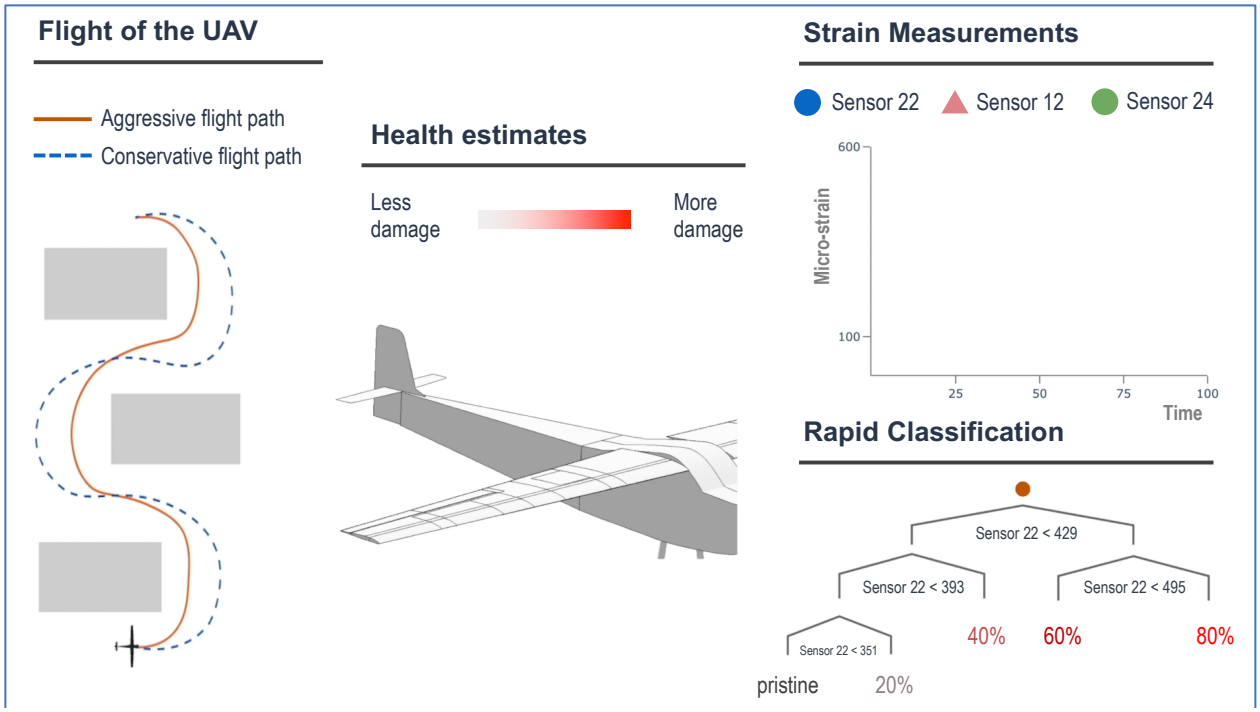

1 AI in Science & Engineering

2 Predictive Digital Twin

3 Conclusions & Outlook

Conclusions & Outlook

From forward simulations to predictive data science

Scientific Machine Learning

Learning from data through the lens of models is a way to exploit structure in an otherwise intractable problem

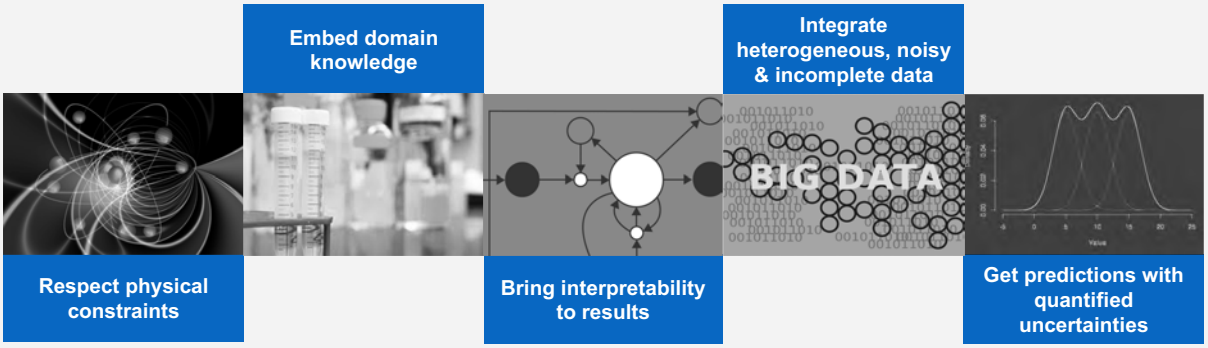

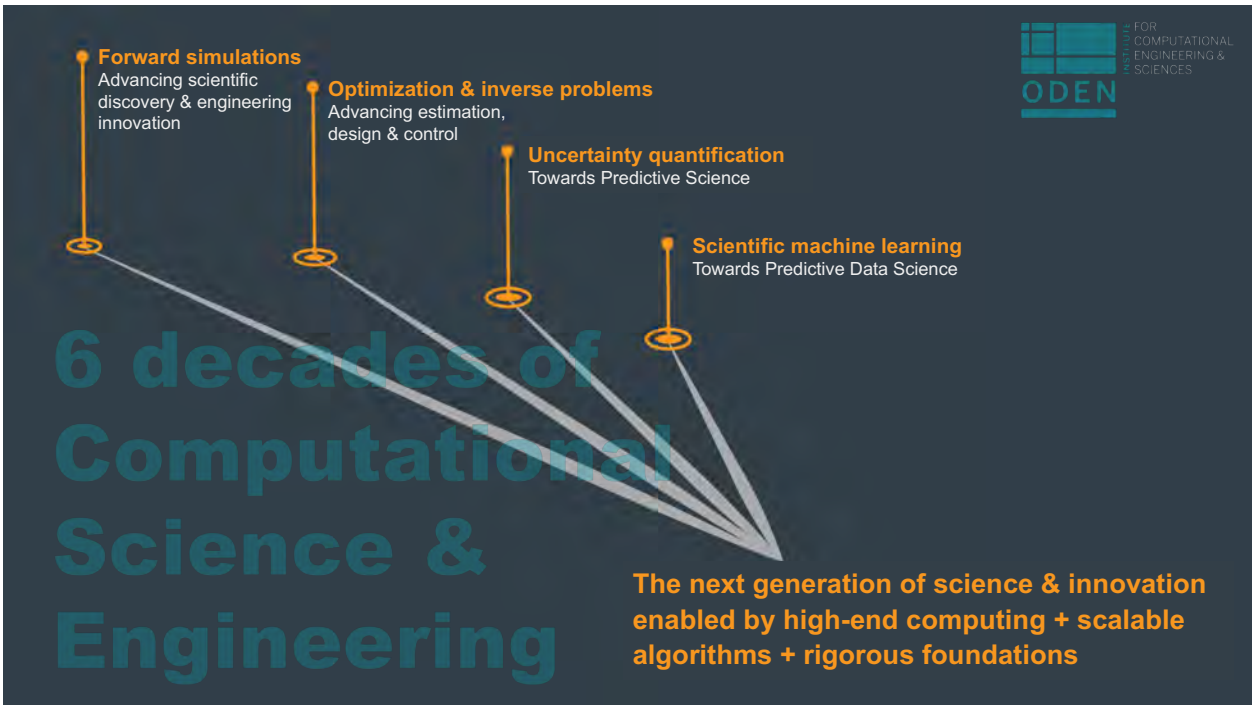

Interdisciplinary research & education in computational engineering & sciences

developing high-performance computing solutions to society's big problems

ODEN.UTEXAS.EDU

ODEN INSTITUTE
FOR COMPUTATIONAL ENGINEERING & SCIENCES

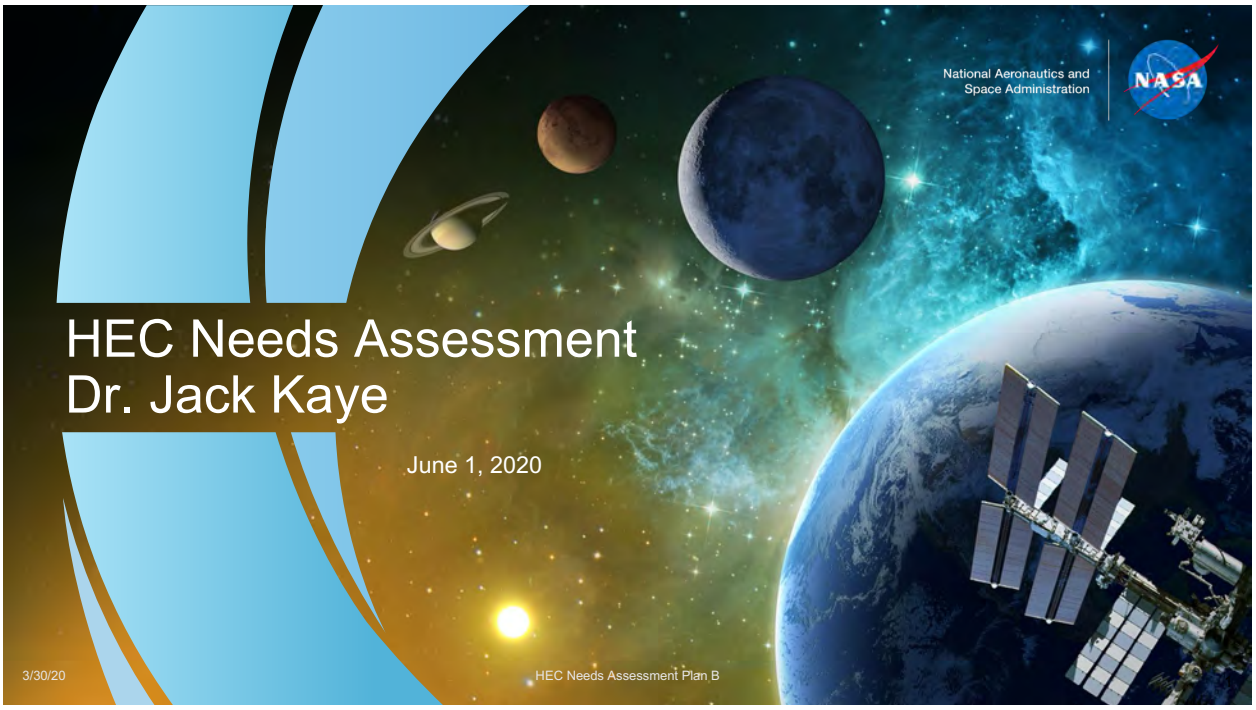

Background

In January 2015, the Big Data Task Force (BDTF) was chartered through the NASA Advisory Council (NAC) to study and identify best practices in big data. The final report of the BDTF was delivered to the NAC Science Committee in November 2017.

Strategic data management and science computing across SMD was identified as a priority for assessment and action during the SMD Senior Leadership retreat in May 2017.

In February 2018, SMD chartered a working group with representatives from each division to develop a new SMD-wide strategy to enable greater scientific discovery over the next five years by leveraging advances in information technology to improve SMD science computing and data archives.

Today's presentation will cover the working group's findings and recommendations, responses to related products, and proposal for a new directorate-wide Strategic Data and Computing Initiative.

2

Approach to Developing Strategy

As defined by the working group, the SMD strategy was guided by four principles:

- Improve discovery and access for all SMD data to immediately benefit science data users and improve the overall user experience
- Leverage current technology for the discovery, access, and effectiveness of NASA's data, as well as enable new technology and analysis techniques for scientific discovery
- Identify large-scale and cross-disciplinary/division science users and use cases to inform future science data system capabilities
- Champion robust theory programs that are firmly based on NASA's observations

Given the breadth and potential impacts across the SMD community as a result of this strategy, the team used several mechanisms to collect stakeholder feedback and to promote data sharing and information gathering:

- Archives Processing and Data Exploitation Meeting, August 9-11, 2018
- NAC Science Committee draft analysis and findings in response to the report of the BDTF, August 28, 2018
- NASEM *Open Source Software Policy Options for NASA Earth and Space Sciences*, September 25, 2018
- Workshop on Maximizing the Scientific Return of NASA Data, October 30-31, 2018
- Request for Information (RFI): Strategic Plan for Scientific Data and Computing, September 18-November 1, 2018

3

Core Tenets/Guiding Principles

- The status quo will not work. The rate of change in this area exceeds our current capacity and our current systems are not set up to allow us to be aspirational in the next five years without significant investment
- Create a strong foundation that enables SMD to be responsive as the field changes
- Centralized constraint model that sets policies for all of SMD, deals with exogenous risks/opportunities, and shares best practices across the entire community – Managed by an SMD Data Officer
- Do not want to go to a fully centralized approach – Consistent with the recommendation of the NAC Science Committee, divisions must be responsible for the specialized components and implementation of policies to meet the needs of their communities
- May want to reassess this model over time – Periodic evaluation, considering the needs and best practices in the future, is appropriate. Divisions should also be encouraged to conduct similar reviews over time to evolve their specialized capabilities.
- Management of equipment and hardware can be centralized.
- Want people who know the data to manage the data where it sits

Vision: To enable transformational open science through continuous evolution of science data and computing systems for NASA's Science Mission Directorate.		
Mission: Lead an innovative and sustainable program supporting NASA's unique science missions with academic, international and commercial partners to enable groundbreaking discoveries with open science data. Continually evolve systems to ensure they are usable and support the latest analysis techniques while protecting scientific integrity.		
Goal 1: Develop and Implement Capabilities to Enable Open Science	**Goal 2**: Continuous Evolution of Data and Computing Systems	**Goal 3:** Harness the Community and Strategic Partnerships for Innovation
Strategy 1.1: Develop and implement a consistent open data and software policy tailored for SMD	**Strategy 2.1**: Establish standardized approaches for all new missions and sponsored research that encourage the adoption of advanced techniques	**Strategy 3.1:** Develop common metadata standards for all NASA science data
Strategy 1.2: Upgrade capabilities at existing archives to support machine readable data access using open formats and data services	**Strategy 2.2**: Integrate investment decisions in High-End Computing with the strategic needs of the research communities using this capability	**Strategy 3.2**: Utilize the full capacity of advances in High-End Computing to achieve SMD's research goal
Strategy 1.3: Develop and implement a SMD data catalog to support discovery and access to complex scientific data across divisions	**Strategy 2.3**: Invest in capabilities to use commercial cloud environments for open science	**Strategy 3.3:** Promote opportunities for continuous learning as the field evolves through collaboration
Strategy 1.4: Increase transparency into how science data are being used through a free and open unified journal server	**Strategy 2.4**: Invest in the tools and training necessary to enable breakthrough science through application of AI/ML	
	Strategy 2.5: Cultivate a strong community of practice across SMD, the science archives, and the broader research community	

Findings and Recommendations

- Open Data/Open Software
  - 1a. Standard open data policy for all new missions
  - 1b. Standard open data requirements in ROSES
  - 1c. Lifecycle approach for long-term data curation after KDP-F
  - 2a. Open source software requirement for new software development
  - 2b. Open source software requirement for ROSES
  - 2c. Open source software requirement for new missions
- High-End Computing (HEC) Program
  - 3. HEC assessments no less than every five years
- Archives Modernization
  - 4a. All digital data required to go to a NASA archive for long-term curation and public availability
  - 4b. Evaluation criteria for future ROSES solicitations to assess the adequacy of data management plans
  - 5. Collaboration and cooperation of data professionals to enable cross-cutting scientific discovery
  - 6. Work with the NASA Centers to increase their recruitment of data science professionals
  - 7. Data and software usability, discoverability, and accessibility to be evaluated as part of future seniors reviews
  - 8. Facilitate greater discoverability of like data holdings in various archives
  - 9. Create a free and open, unified journal server to make science papers more accessible to the public
- Advanced Capabilities
  - 10. Explore AI/ML and other novel computational techniques through various avenues
  - 11. make investments to incentivize and educate the community on how to use AI/ML
- Management
  - 12. Appointment of an SMD-level Data Officer

BACK UP

7

Strategic Data Management Working Group

Name	Affiliation
Ellen Gertsen, Co-Chair	SMD Front Office
Kevin Murphy, Co-Chair	ESD
Hashima Hasan	APD
Jeffrey Hayes	HPD
Pat Knezek	APD
Bill Knopf	PSD
Janet Kozyra	HPD
Tsengdar Lee	ESD
Jared Leisner	HPD
Rebecca McCauley Rench	PSD
Viet Nguyen	JASD
Mariel Borowitz*	SIMD

*IPA ended in Fall 2018

8

NAC SC Recommendations

Overall, NAC SC agrees with the BDTF that SMD data archive programs and projects are performing well and are properly taking steps to modernize. However, the volume, variety and velocity of NASA science data is taxing established methods and technologies. The SC finds that SMD should:

1) make investments in hardware, software, training and education to accelerate modeling workflows,
2) participate in the Department of Energy's (DOE) exascale computing program,
3) implement server-side analytics (SSA) capabilities (with caution),
4) forge a joint program with the National Science Foundation's (NSF) Big Data Innovation Regional Hubs and Spokes program, and,
5) incorporate data science and computing advisory positions in the SMD advisory committees.

In all efforts, the SC underscores that it is important that data science and computing experts work closely and collegially with domain scientists to implement effective solutions that are based on an understanding of the domain.

As to the future, the SC commends that an SMD Strategic Data Working Group has been set up that will bring forward these ideas, without interfering with how each division manages data.

For more info, reference "NAC SC Big Data Product" that contains SC feedback on each of the SC's Ad Hoc Task Force on Big Data findings/recommendations: https://science.nasa.gov/science-committee/meetings

9

NAS Open Code Recommendations

Recommendation: NASA Science Mission Directorate should develop internal policies and external legal language conducive to the swift release of open source scientific software, and the full participation of NASA employees in internal and external open-source projects, without jeopardizing national security or incurring legal liability.
Recommendation: NASA Science Mission Directorate should encourage the use of standard open source licenses, but not mandate a particular license. Non-standard licenses should be justified in the software management plan.
Recommendation: Any open source software policy that NASA Science Mission Directorate develops should not impose an undue burden on researchers; therefore, any policy should be as simple as possible and fully fund any mandates
Recommendation: NASA Science Mission Directorate should consider a variety of policy options depending on discipline and software type and transition to greater openness over time.
Recommendation: NASA Science Mission Directorate should support the infrastructure, governance, and maintenance of a healthy open source community, taking advantage of existing community resources to the greatest extent possible.
Recommendation: NASA Science Mission Directorate should support open source community-developed libraries that advance NASA science.
Recommendation: NASA Science Mission Directorate should foster career credit for scientific software development by encouraging publications, citations, and other recognition of software created as part of NASA funded research.
Recommendation: NASA Science Mission Directorate should initiate and sponsor programs to educate and train researchers in open source best practices. Topics could include, but are not limited to export controls, licensing and intellectual property, workflows, and software development. These resources could be made available to the community via in-person trainings as well as webpages, screencasts, and webinars.
Recommendation: NASA Science Mission Directorate should explicitly recognize the scientific value of open source software and incentivize its development and support, with the goal that open source science software becomes routine scientific practice.

10

Stakeholder Participation by Organization

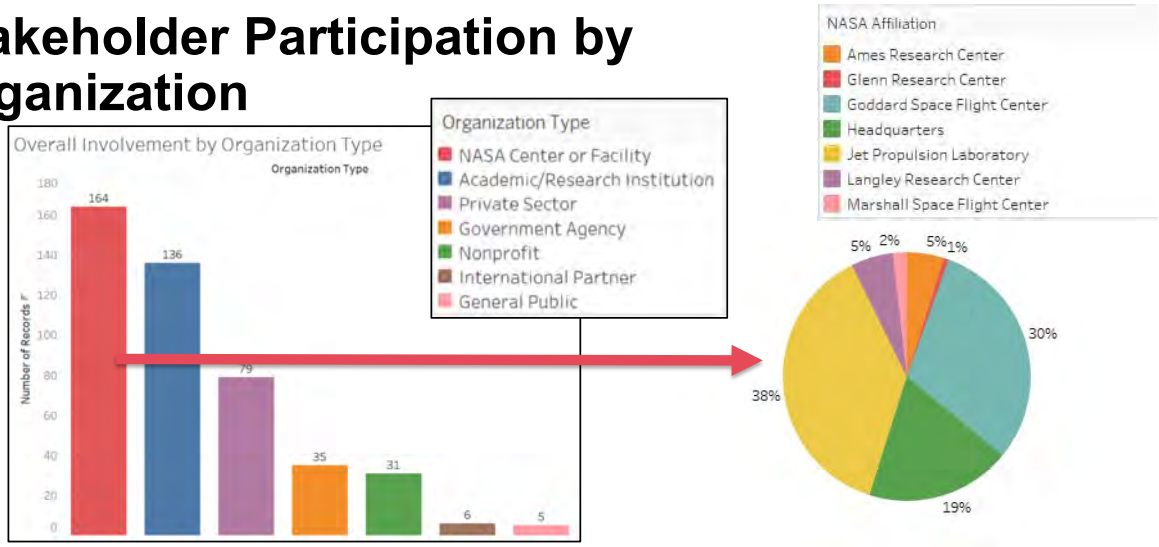

Over 450 people from academia, other Federal government agencies, foreign space agencies, commercial providers, professional societies, and the general public shared their ideas with NASA to inform the new SMD-wide Science Plan.

Stakeholder Participation by Discipline

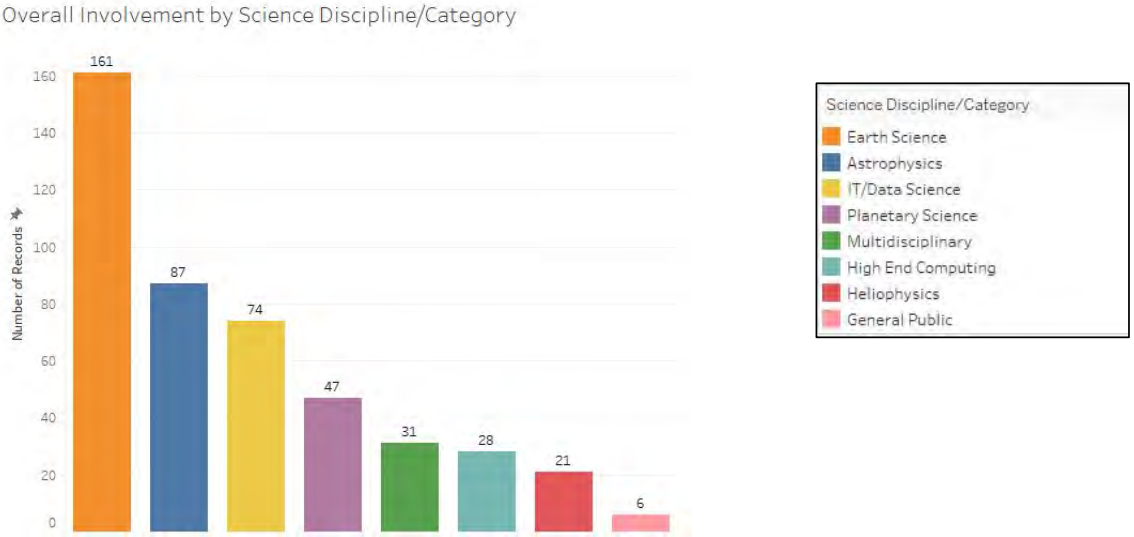

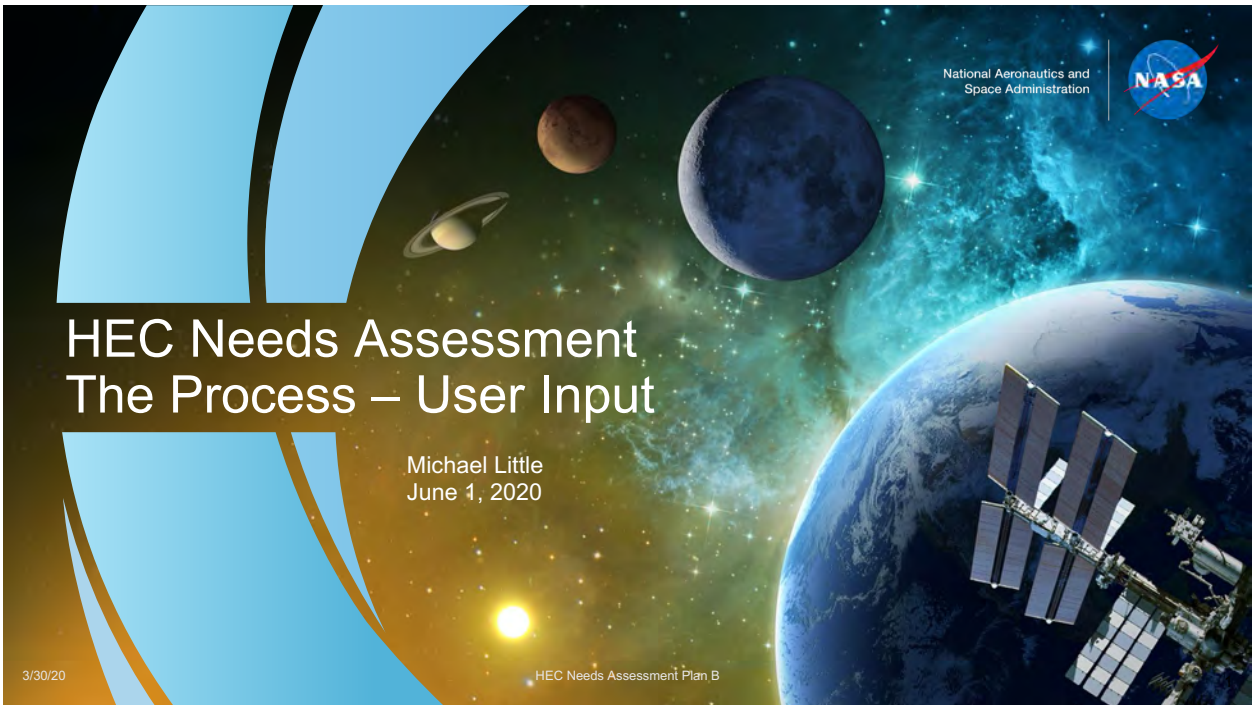

HEC Needs Assessment – User Input Feeds the Report

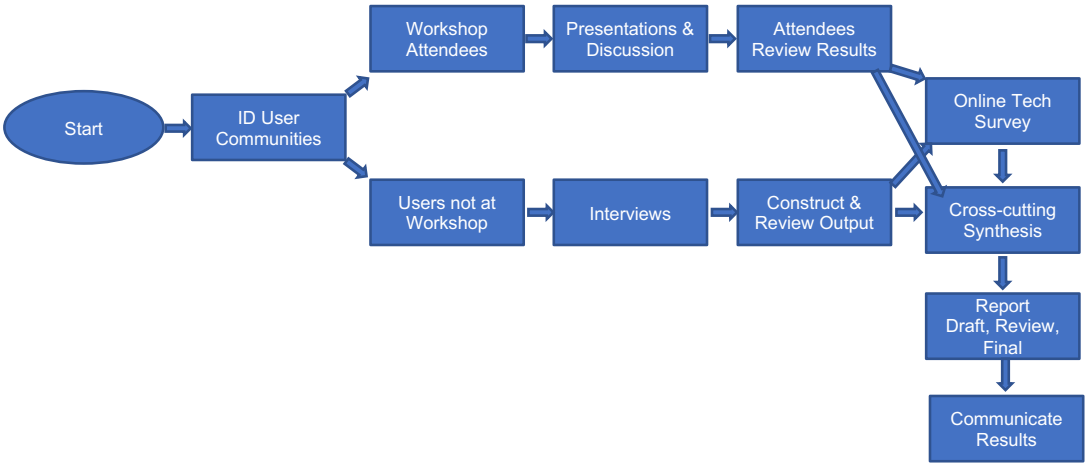

Guidance for Preparing Use Case

Instructions

1. PURPOSE: to document key Use Cases which use high performance computing (NASA or other) and how they support NASA's mission, research and development objectives
   a) The user input for the HEC Needs Assessment will consist of two parts: (a) Use Case presentation to be discussed during the Affinity Group and (b) a survey of Use Case technical characteristics which will not be presented but will be used to forecast future demands. Instructions for the survey will be provided later.
   b) The sessions are planned to be 3 hours with 10 participants Each use case presentation will be limited to 10 minutes with 5 minutes for clarifying questions. Unfortunately, the remote nature of this workshop limits the extent of discussion of each case to only collecting the data needed by the Needs Assessment.
   c) Please try to summarize the Use Case at a high enough level to stay within the 10 minute schedule.
2. The following 3 slides are the template for the presentation. Please try to summarize so it fits in these 3 slides: a) Mission Objective, b) Current application characteristics and c) Key Questions for Future
3. Headings in black are the major areas that should be addressed in the presentation. Items in red or in [ ] brackets are suggested items and should be removed and replaced by your bullets
4. Each participant will have the opportunity to review and clarify the notes from the workshop for a week after their workshop is held.
5. If you have questions, email m.m.little@nasa.gov

3/30/20 HEC Needs Assessment Plan B 2

Use Case: Mission Need

- Mission: [Briefly describe Agency-level need for this application or capability]
  - How can you, your program and NASA measure the impact of the HEC on this mission need?
- Criticality: [How essential is NASA's HEC to accomplishing this work]
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? (i.e., buy a system, turn over to partner, rent time on a commercial facility, Commercial Cloud Computing?)
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - [What is impact of inadvertent release of data? (i.e., public confusion, professional embarrassment, National Security, Privacy, etc?)]
  - [If HEC availability is interrupted regularly, how does that affect your ability to achieve objectives?]
  - [What is mission need for external collaborators (without NASA credentials) to access system or to run jobs?]
- HEC Function: [Briefly describe how HEC supports performance of this work]
  - [How often do you need to run the application to support mission objectives]
  - [How is the output used to achieve mission objective?]
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - [If NASA supercomputing was no longer available to meet your needs, how would you perform this work? (i.e., buy a system, rent time on a commercial facility, use Commercial Cloud Computing?)]
- Mission Concept: [Provide a graphic if possible, depicting this application or use case]

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - What triggers this workflow?
  - Is there a need for interaction by user, operator or other system in the computing workflow?
  - Does this use case need quick reaction or near-real time computing? Explain briefly.
- Data:
  - In: [Summarize Data used to feed into the code/application in terms of volume, latency, source]
  - Out: [Summarize data produced by computation and what happens to it]
    - [Does the output data require retention? How is that selected?]
    - [How is the data analyzed to provide mission results and where? (statistical or AI analysis, visualization?)]
  - Transfer:
    - [Where does the output go after it is generated? (i.e., local to HEC storage system or to another location)]
- Computation:
  - Type: [Summarize the type of computing used/needed and how HEC supports this today]
    - [Is this capacity computing or are special HEC capabilities needed for this computation?]
- Evaluation: [Summarize how you evaluate the effectiveness of HEC in supporting this work]

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.]

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?]
  - Impact: [How does this affect fulfilling the mission need?]
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?]

- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?]
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?]
  - Re-use: [Will the number of people running this code/application grow?]
  - Re-hosting: [What is needed to migrate code to future computing architectures?]

Information about the Survey

- The following slides describe the kind of data that the Needs Assessment will collect via a survey
- This material is needed to supplement the description of the Use Case presented to the Affinity Group but will not be presented.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones
    - Programming languages used (If multiple languages used, provide estimate percentage of the code in each languages.)
    - Number of lines of code
    - Workflow mechanism used to control code execution (perl, jupyter, shell script, etc.)
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage is covered below
    - Run in batch, interactive, or near-real time

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other)
  - How much input and output are generated per run
  - Summary of the total amount of input and output data per year
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena?
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used.
  - How much data needs to be curated (archived)?
  - How much data needs to be distributed and, generally, to whom?
- What are the key computing obstacles to effective execution of the current application?
- Are there any analytic functions performed on the output of the application?
  - Timing, computing load, workflow
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.)
    - Increase in resolution and/or timesteps
    - Rewrite to improve performance
    - Refactor for emerging platforms
    - Rehost onto cloud
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements
    - Run in batch, interactive, or near-real time

Appendix D1 Workshop Session 6-1 Inputs

Arlindo da Silva	GSFC/610.1	Observing System Simulation for Earth Science Space Missions
Grey Stephen Nearing	Upstream Tech	Deep Learning for Hydrology Modeling
Irina N. Kitiashvili	NASA ARC	3D Realistic Modeling of the Sun
Madhulika Guhathakurta	NASA HQ/Helio	Heliophysics Overview
Will Doebler	LARC/D321	X-59 Sonic Boom Propagation Simulations
Eric Larour	JPL	ISSM: Ice-Sheet and Sea-Level System Model

National Aeronautics and Space Administration

Workshop Session 6-1

1 June 2020

National Aeronautics and Space Administration

Use Case: Observing System Simulation for Earth Science Space Missions

Arlindo da Silva
1 June 2020
GSFC Global Modeling and Assimilation Office, Code 610.1
Science Mission Directorate

Observing System Simulations: Mission Need (1 of 2)

- **Mission Objective**:
  - Detailed simulation of global observing systems is critical for pre-launch algorithm development and assessment, and for conducting quantitative science trade studies during mission formulation.
  - Measuring impact of HEC: The maturity and readiness of algorithms at launch, prediction of performance after launch

- **Use Case Concept**

Nature Run: Earth System Model

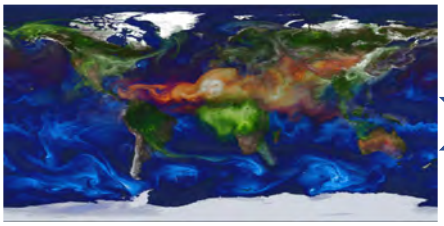

Sampled at Instrument Footprint

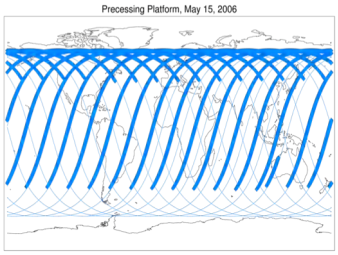

Detailed Instrument Simulator

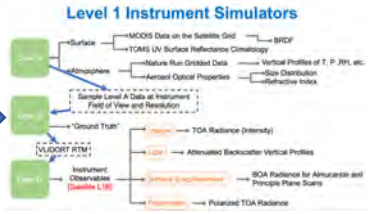

Observing System Simulations: Mission Need (1 of 2)

- **Criticality**: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - HEC resources are essential for realistic observing system simulations. Without these capabilities one would need to resort to a limited number of cases and conditions, or to rely on low-resolution, low-fidelity simulations, all affecting the effectiveness of the activity.
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - While algorithm development and science trade studies are generally the domain of the science team, the inadvertent release of science data does not pose any major risk. However, some level of confidentially is warranted while a system is still in development.
  - Continuous availability of HEC resources is only marginally important as this is not a RT applications. However, capacity will be needed to perform high resolution simulations on demand
  - Output files will need to be accessed by external collaborators (assuming a diverse Science Team including US Universities and international participants). Because of massive size of datasets, there will be a need for partners to execute algorithms on the same nodes where data resides.
- **HEC Function**:
  - A global, very resolution Nature Run will be performed for a period of at least 1 year, with checkpoints at discrete intervals(very little I/O otherwise) – earth system model is the data compression algorithm.
  - For specific missions and applications (e.g., clouds/aerosols for ACCP), model will restart at key periods and produce carrying out in-line, as the model reruns, mission specific calculations
- **Impact of HEC**: See "Criticality" above.

Current/Future HEC Application Characteristics

- Current Work flow:
  - A very high-resolution Nature Run with detailed microphysics is performed for a period of at least one year, saving to disk full output for all anticipated uses – compromises are made, such as output frequency and variables
  - For a given Space Mission application, files are sampled at instrument footprint, saved to disk, and detailed instrument simulators are performed off-line
  - Nature Run Generation is a highly parallel application meant for batch processing
  - Instrument simulators are embarrassingly parallel applications, often I/O bound.
- Future Work flow:
  - Phase I: A very high-resolution Nature Run with detailed microphysics is performed for a period of at least one year, with frequent checkpoints, and limited output (mostly for browsing/case selection).
  - Phase II: For a given Space Mission application, "forecasts" are run from saved restarts, sampling model parameters at instrument footprint, performing detailed instrument simulation on-line, in parallel.
  - User provides instrument simulator plug-ins, prepares configuration files, submit jobs.
  - Both phases are highly parallel applications meant for batch processing.

Current/Future HEC Application Characteristics

	Current	Future
Data	7 km horizontal resolution, 72 level model **In:** • Nature Run: minimum input requirements O(Gb) • Observing System Simulations: ~ 500 Tb/year for instrument/orbit **Out:** • Nature Run: ~ 800 Tb/year • Observing System Simulations: e.g. PACE: 5 Tb/year for OCI L1 • Statistical analysis usually performed where the data resides. **Transfer:** • Output stays at NCCS data portal where it can be accessed via https:// or OPeNDAP	3 km horizontal resolution, 132 level model **In:** • Nature Run: minimum input requirements O(Gb) • Observing System Simulations: ~ 13 Tb/year **Out:** • Nature Run: ~ 13 Tb/year (weekly restarts) • Observing System Simulations: Depends on Observing System • Statistical analysis usually performed where the data resides. **Transfer:** • Output stays at NCCS data portal where it can be accessed via https:// or OPeNDAP
Computation	• GEOS-5 Nature Run (G5NR) was performed at NCCS discover • Several Observing System Simulations (e.g., PACE/ACCP) are being performed at NCCS because of "data locality"	• Simulations can be performed at any NASA HEC Center or in the pub;lic Cloud, provided some form of "data portal" is provided. • Both NR and Observing System Simulation need to be performed at a HEC center, but not necessary NASA's (input is small).
Eval	Nature Run must have resolution and microphysics complexity to simulate the Observing System at hand. HEC capability must be able to support simulations long enough to fully document impact of the new observing platform.	

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - GEOS development cycles: reanalysis & NR production, development
  - AGU, AMS and other highly attended meetings if resources are not dedicated
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Under the new paradigm, observing system simulations require regular access to large number of cores for production
  - Impact: Given current queuing structure it may have a large impact on time-to-solution. A public cloud solution may be applicable.
  - External Sources: N/A
- How will this Use Case evolve over the next decade?
  - Maturity: NR capability relatively mature, while in a state of continuously development. New paradigm for in-line instrument simulator needs definition of protocols and user training (plug-ins, etc.)
  - Growth: increase in resolution will lead to smaller time steps; with sampling being done on-line there is dependence on output frequency.
  - Re-use: this capability should extensible and highly re-usable, hopefully engaging new space missions
  - Re-hosting: This will be address within GMAO's general strategy for evolving the GEOS modeling system.

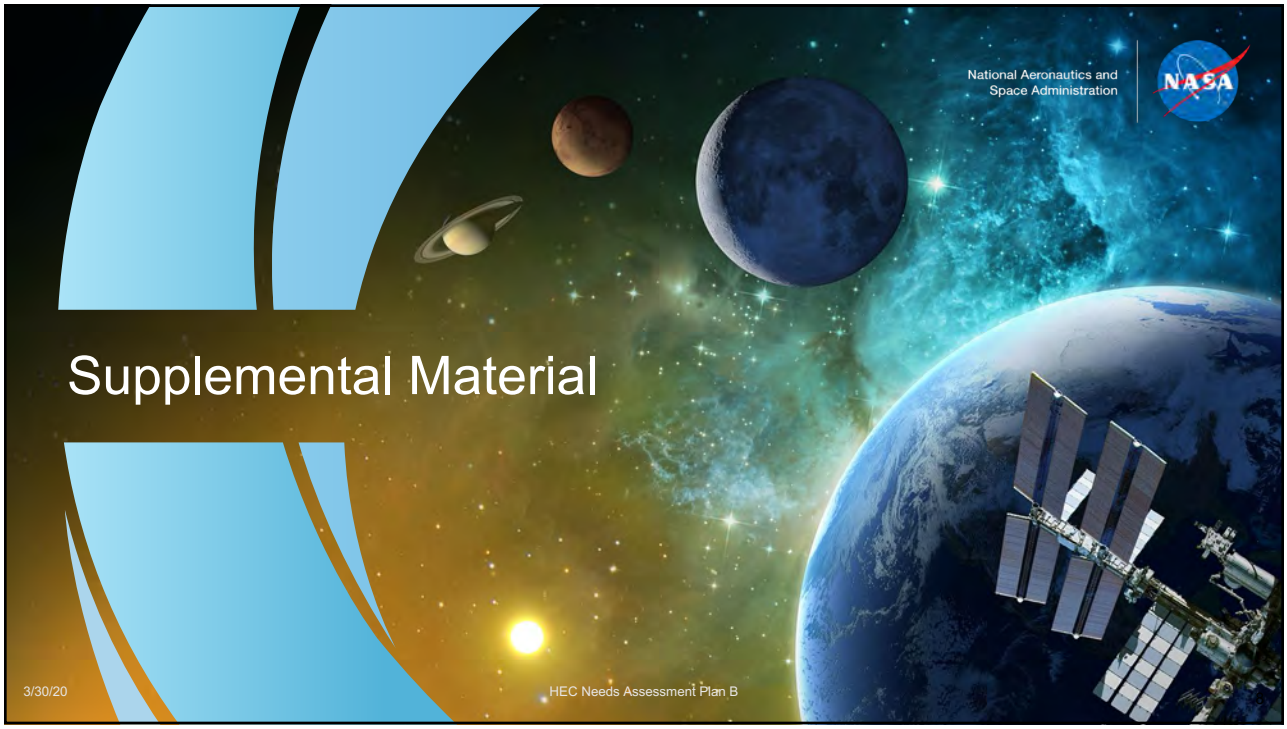

O.S.S.E

- **Observing System**
- **Simulation**
- **Experiment**

Model-based OSSE

A framework for numerical experimentation in which *observables* are simulated from fields generated by an earth system model, including a *parameterized* description of the *observational error* characteristics.

Simulations are performed in support of an experimental goal.

The "E" in OSSE
A Spectrum of OSSEs

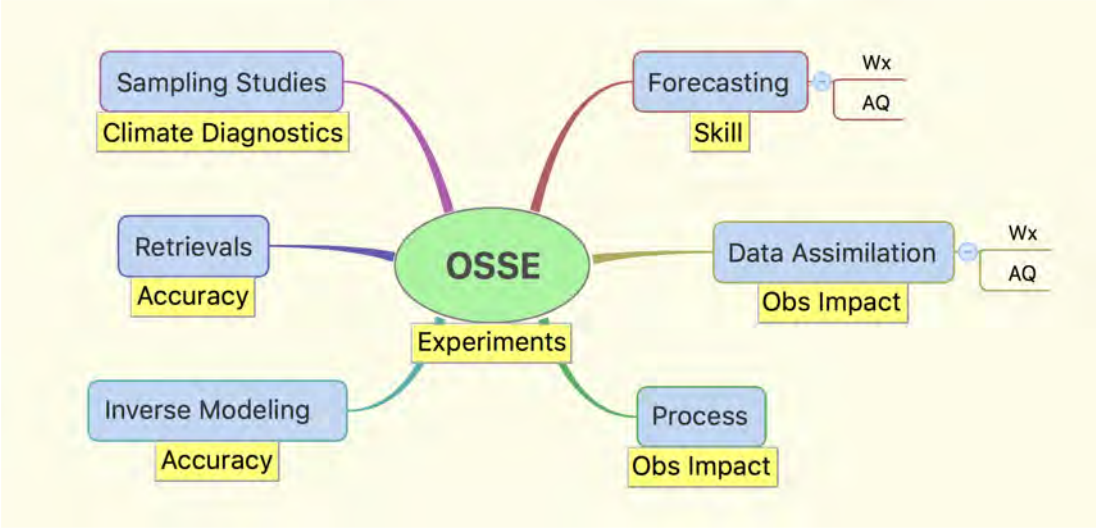

Impact of Model Resolution

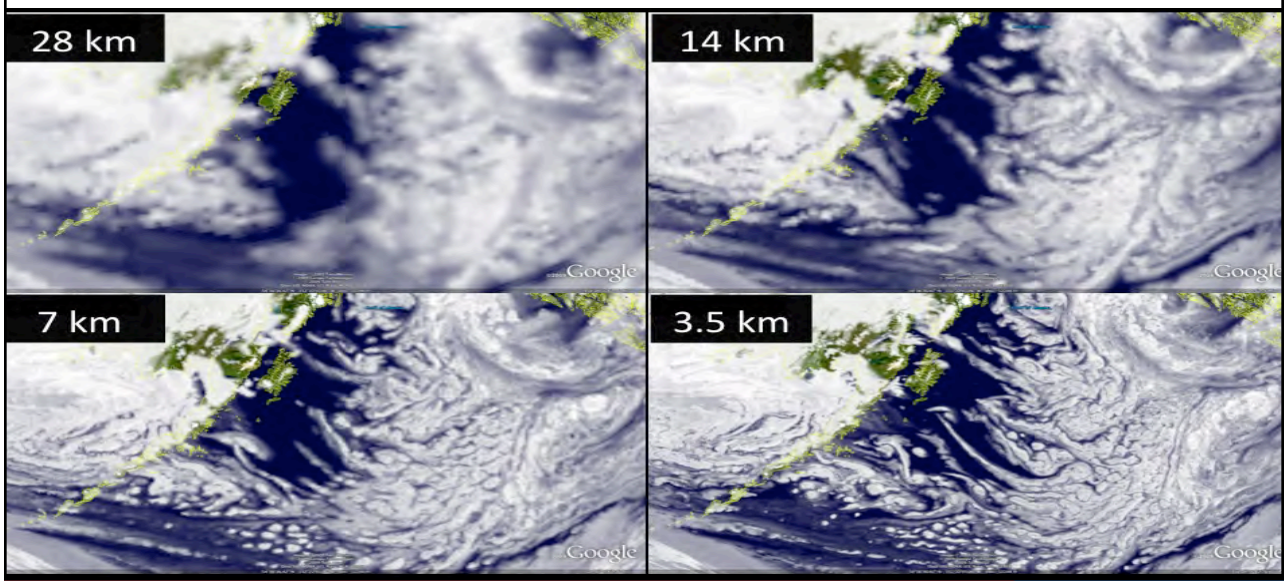

Instrument Simulators

- Level 1
  - Detailed radiative transfer calculation in the presence of clouds, aerosols, ice, etc.
  - Instrument characteristics
  - Observables: polarized radiances, backscatter
- Level 2
  - Retrieved quantities at observation location
  - Averaging kernels, error characteristics
- Level 3
  - Hourly to seasonal mean statistics sampled at the instrument footprint

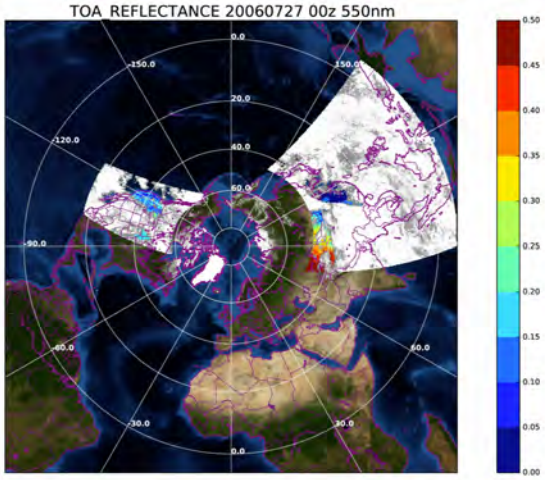

Level 1 Instrument Simulators

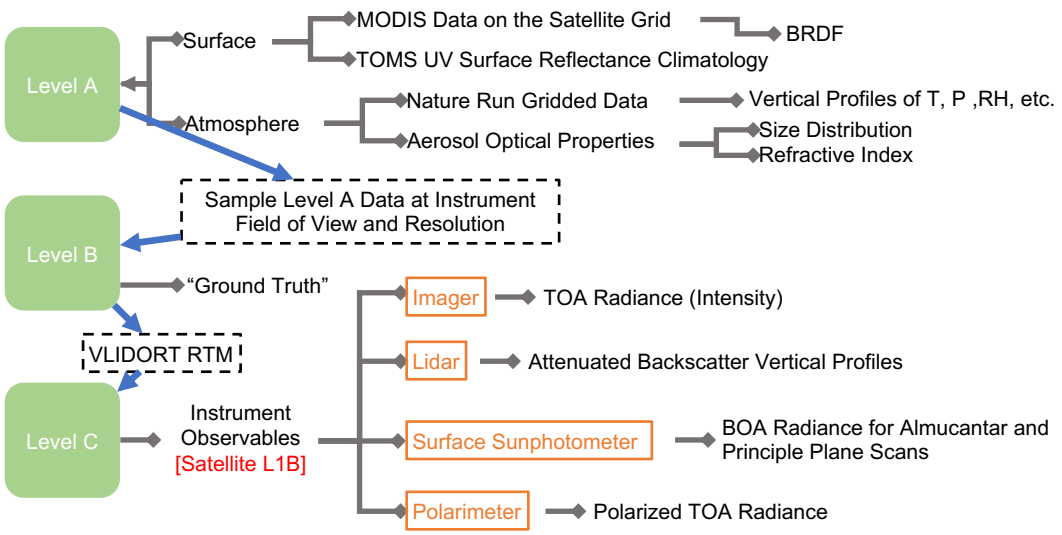

Level 1 Instrument Simulators

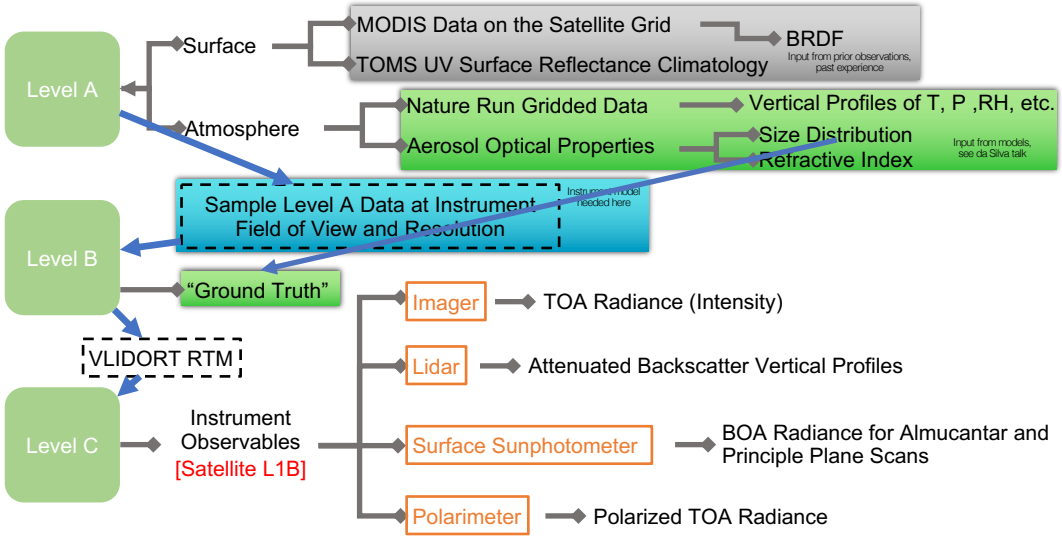

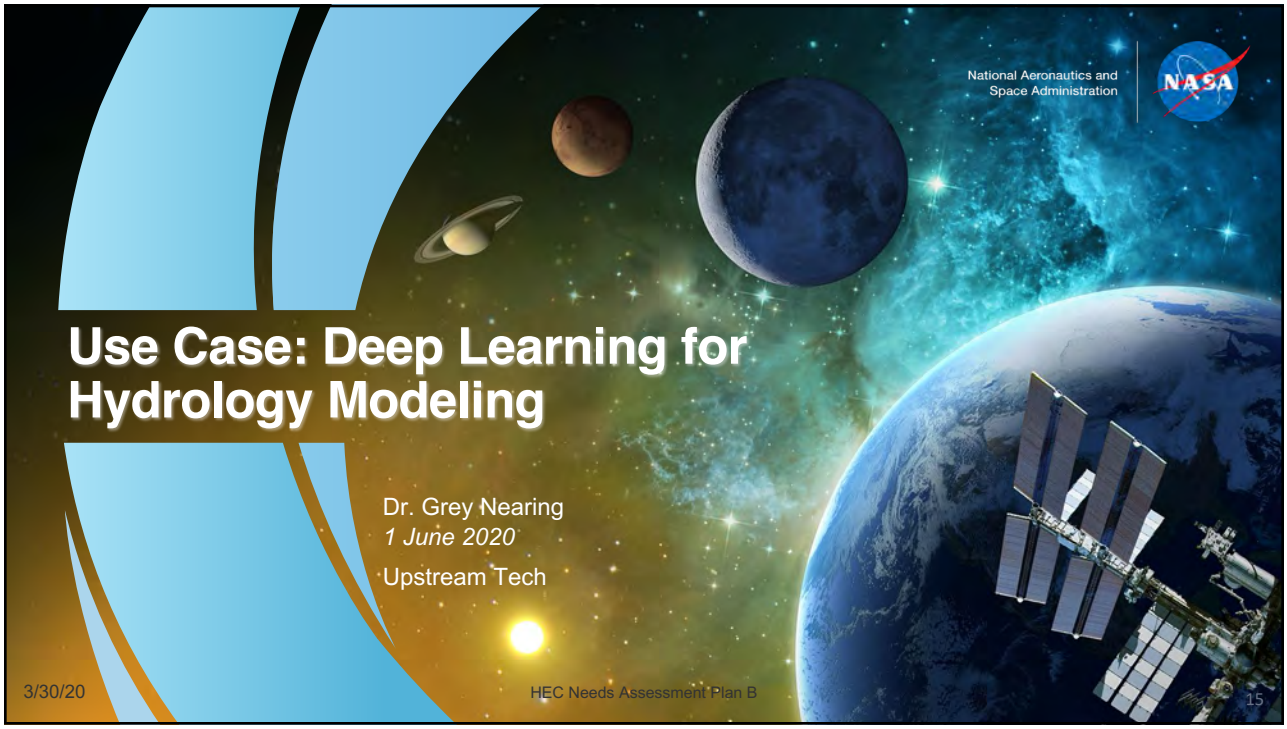

Current Use Case (Survey)

- **Achieving the mission objective**
  - Application: **Training & Testing Deep Learning Models for Hydrology Simulation**
    - Antecedents: **None**
    - Programming languages used: **Python (PyTorch, Tensorflow)**
    - Number of lines of code: **O(10^4)**
    - Workflow mechanism used to control code execution: **Shell Script**
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): **1 node, O(10^1) cores**
    - Accelerator usage (type and memory requirement): **GPU (more GPU memory than exists)**
    - Amount of memory required (overall or per process): **TBs would be useful**
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other: **Shared Memory, NVLink**
    - Run in batch, interactive, or near-real time: **Batch + Interactive GPU access for development & prototyping. We also run an operational (near-real-time) version of this project in the private sector using GCP.**

Current Use Case (Survey)

- **Describe the data requirements for the application**

  Basic Situation: Hydrology is 3-D, Atmosphere is 4-D. We usually have much lower data requirements.

  - Data Flow: **raw files in different formats -> database in standard format -> dataloader (pytorch, tensorflow) -> node memory -> cpu workers -> gpu memory**
  - How are the data loaded in the application: **Storage -> Memory: usually serial; Memory -> GPU: parallel**
  - How much input and output are generated per run: **GBs**
  - Summary of the total amount of input and output data per year: **O(10^1) TB**
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? **N/A**
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other: **Cloud storage available to users from many institutions. High speed I/O (e.g., PCI/E storage) is nice, but not necessary.**
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used. **Gridded -> Watersheds, Time (dis)aggregation**
  - How much data needs to be curated (archived)? **O(10^1) TB**

Future of the Current Use Case (Survey)

- **Describe the computing needs of the application as it evolves over the next 5-10 years**
  - What are the key computing obstacles that could affect these goals?

    As we scale up to globally-connected simulations, distributed memory (GPUs) will become an issue.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in resolution and/or timesteps: **Larger storage and memory requirements when moving to finer spatial and temporal resolutions**
    - Rewrite to improve performance: **Refactoring for**
    - Rehost onto cloud: Likely **Prohibitively Expensive**
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores): **Dozens of GPUs**
    - Accelerator usage (type and memory requirement): **GPUs (others? TPUs, Neuromorphic?)**
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements: **Likely not a limiting factor**
    - Run in batch, interactive, or near-real time: **All of the above**

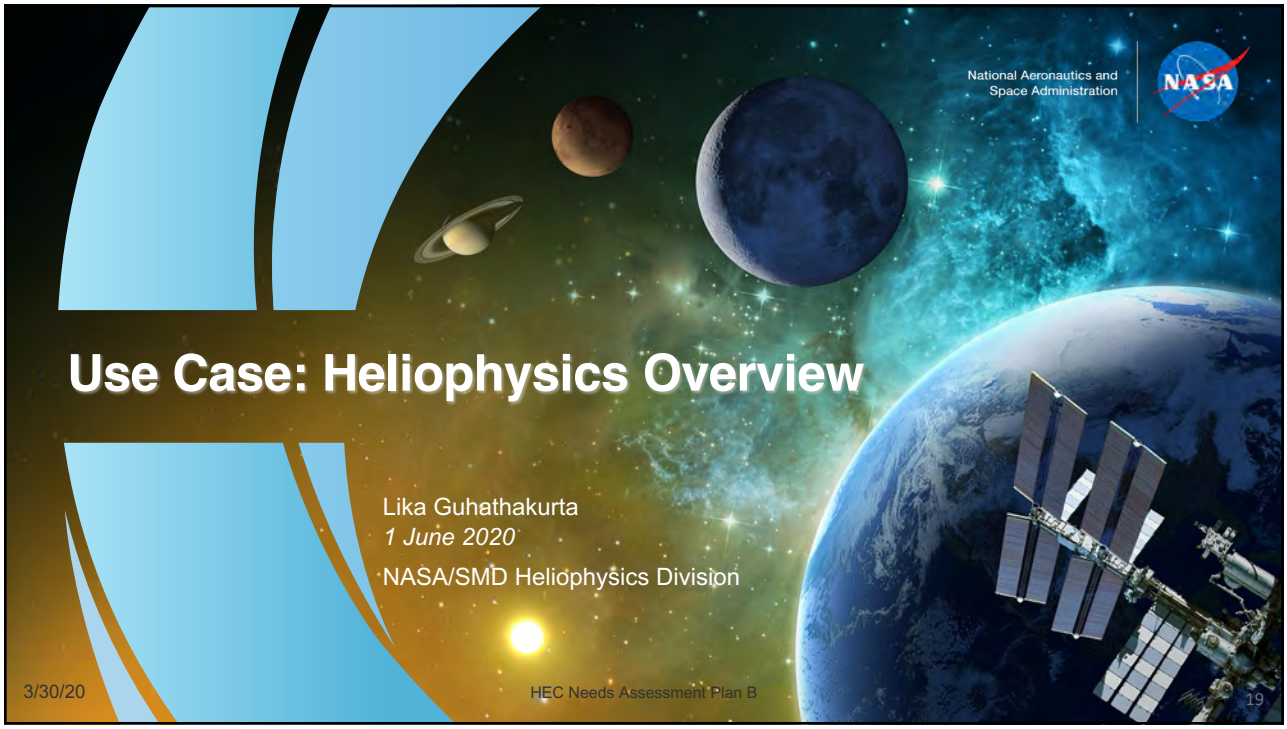

Heliophysics Past

- During the last decade, the heliophysics community has gone through a radical transformation with vast increases in data rates, archival volumes, and the use of sophisticated instruments. *Example*: SDO data rate ~2 TB/day (cf. 10^2 times TRACE, 10^3 times SOHO-EIT), producing images/data of unprecedented quality and complexity.

- Progress in our understanding relies critically on combining data analysis with computer simulations, and data assimilation.

- Computer simulations and AI tools are the *third* leg of discovery. They help us explore where we cannot sense, can give us a comprehensive view of sparsely sampled environments, and place the heliophysics domain into a controllable laboratory setting---all of this at a cost that is a fraction of the cost of a mission (ACCEHS Report, 2010). https://lwstrt.gsfc.nasa.gov/images/pdf/ACCEHS.pdf

3/30/20 20

Computational Grand Challenge

- Broad interdisciplinary challenge straddling Heliophysics, Planetary, and Exoplanetary Science (with implications for habitability: “Are we alone in this Universe?”)
- Traditional global models based on MHD have largely run their course: reconnection, turbulence, dissipation and energetics not accurately modeled. *Ad hoc* treatments (like anomalous resistivity and viscosity) are not accurate recipes and will not be reliably predictive.
- The science involves Intricate interplay between small and large scales: small scale dynamics have major global implications, global dynamics drive local processes.

Multi-Scale Grand Challenge

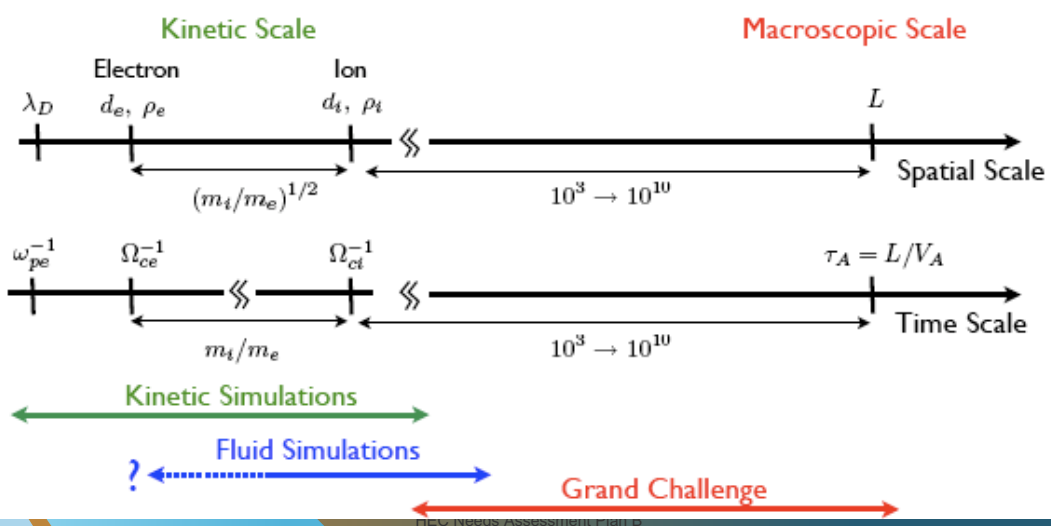

Desired Progress Needed

We can build on the NASA-NSF Partnership on Strategic Capabilities in Space Weather since 2005 that emphasizes computer code development, data assimilation and mining, and validation against observations.

- To study processes occurring at ion and electron scales, with scale separation between ions and electrons,
- To include new equations of state for electrons and ions, with physics-based closure
- To handle multiple ion species (such as hydrogen *and* oxygen), thus enabling the coupled treatment of composition, wave and instability dynamics in the magnetosphere, heliosphere, solar environment and its implications for space weather events,
- Efficient and flexible computer simulation codes and frameworks that use state-of-the-art algorithms and scale to up to thousands of processors on modern computational architectures, and sophiscated AI derived model products/input.
- Validation and verification plan to model solar, helio & geospace that can be tested with mission data,
- Open source and accessible to community (e.g., github)

Data Science, Computing & AI

- Organize “Data Science, Computing and AI” in NASA SMD in a way that will integrate software *and* hardware management spanning all SMD divisions, and in partnership with each division setting its science priorities.

- What we learned from George Siscoe one of the authors of the Heliophysics Text Books: Do not wait for the perfect simulation code, which is impossible without continuous validation and continuous guidance from data. Prediction imperative should guide physics imperative, not the other way around (as we have learned from meterology). And this is now also applicable to implementation of AI tools.

National Aeronautics and Space Administration

Use Case: X-59 Sonic Boom Propagation Simulations

Will Doebler
June 1, 2020
LaRC D321 Structural Acoustics Branch
Commercial Supersonic Technology Project

Use Case: Mission Need

- **Mission: X-59 Quiet Super Sonic Technology (QueSST)**
  - HEC/Pleiades enabled completion of exit criterion for a Level 1 Milestone:
    - "Atmospheric Effects on Low-Noise Sonic Booms Quantified"
  - Other X-59 Mission HEC Uses:
    - Turbulence effects on sonic booms, aircraft CFD, uncertainty quantification
- **Criticality: Medium/High**
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - Rent time elsewhere, use Commercial Cloud Computing, and/or reduce resolution and run on LaRC's smaller K-Cluster
- **Mission Security Needs: (Confidentiality, Integrity, Availability)**
  - Impact of inadvertent release of data: public confusion, professional embarrassment, loss of intellectual property
  - If HEC availability is interrupted regularly, may need to consider other platforms to meet mission objectives
  - "Nice to have" external collaborator access to NASA HEC
- **HEC Function: Use of 100s of cores on Pleiades, use of storage space on Lou**
  - This specific case was a "one and done," but it may be exercised in the future
  - Output informs decisions for future X-59 flight test planning

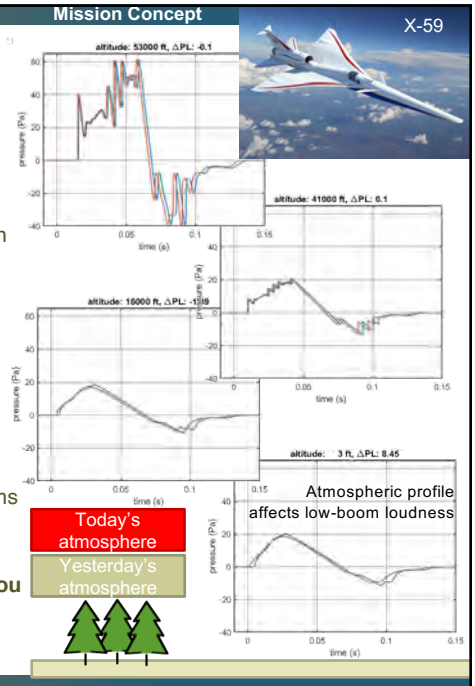

Use Case: Current HEC Application Characteristics

- **Work flow:**
  1. Create working folder for each core and copy aircraft info and executable into each working folder
  2. Write input file containing the assigned atmospheric profile information for each core
  3. Run PCBoom executable to propagate sonic boom waveforms through the current atmosphere
  4. Copy output files from working folder to data storage folder then delete input/output files from working folder
  5. Repeat steps 2 through 5 for all atmospheres assigned to the core
  6. Tar the data storage folder and send to Lou using shiftc
  7. Copy data to LaRC K-Cluster for post-processing in MATLAB using LaRC license
  8. Visualize on personal computer
  - *What triggers this workflow?* This specific case was triggered by Project need. Future triggers could be new/updated CFD or data at different global locations
  - *Is there a need for interaction by user, operator or other system in the computing workflow?* Yes, for post-processing and visualization
  - *Does this use case need quick reaction or near-real time computing?* No
- **Data:**
  - In: atmospheric profile information (datetime, location), aircraft information (CFD cylinder, heading, trajectory)
    - Stored on lustre – 5 yr. USA atmospheric data ~50 GB, aircraft info ~10 MB, executable 2 KB
  - Out: Text files containing sonic boom waveforms, boom arrival position and time. Text files are moved to data collection folder, tarred, and put onto Lou ~5 TB. Then sent to K-Cluster for post-processing in MATLAB. Final results are visualized on personal computer.
    - Project requested all output data be retained.
    - Analyzed results were reported during Technical Quality Review and will be disseminated at various scientific conferences and future publications.
  - Transfer:
    - Lou for storage and LaRC's K-Cluster for analysis using a code written in MATLAB
- **Computation:**
  - Capacity computing, allowing batch executions to run in parallel
- **Evaluation: Successful**
  - Would appreciate more HEC training opportunities

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - Flight test campaigns, new/updated code inputs (CFD or atmospheric model), publication cycles, contract deadlines, unforeseen project needs
- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - Constraints on Use Case: would appreciate reduced queue time and increased CPU allotment. Would appreciate expanded MATLAB availability. Would be nice to have more lustre storage space
  - Impact: reduced turnaround time would improve flight test planning and decision making
  - External Sources: None
- **How will this Use Case evolve over the next decade?**
  - Maturity: Likely will have on-going development for the next few years. X-59's flight test campaign ends within this decade
  - Growth: Move from USA-only to global resolution, increase resolution of inputs
  - Re-use: Possibly increase users, as others may want determine results for their aircraft designs. Possibly re-run the current use case as improved input data become available
  - Re-hosting: Need to maintain Python versions or support translation to other codes, may need to recompile executable for new architecture

National Aeronautics and Space Administration

Use Case: ISSM: Ice-sheet and Sea-Level System Model

Eric Larour, Erik Ivins, Surendra Adhikari, Lambert Caron, Nicole Schlegel, Helene Seroussi, Mathieu Morlighem, Josh Cuzzone
June 1, 2020
JPL

- **ISSM: Ice-Sheet and Sea-Level System Model**
- **HEC Needs Assessment**

- **Team at JPL:**
- **Sea-level: Eric Larour, Erik Ivins, Surendra Adhikari, Erik Ivins, Lambert caron**
- **Ice-sheets: Eric Larour, Nicole Schlegel, Helene Seroussi, Mathieu Morlighem (UCI), Josh Cuzzone (UCLA)**
- **Funding: MAP, N-SLCT, Cryosphere, IceSat-2, IDS, JPL R&TD.**

- **Capabilities: FEM based, fully unstructured meshes (no grids)**
- **Sea-level solver: unstructured meshes, dense connectivity.**
- **GIA computations: spherical harmonics spectral decomposition based solvers**
- **Ice-sheet dynamics solver: FEM , sparse connectivity.**
- **UQ runs: repeat above solvers many times. Using Parallel Sandia National Labs Dakota samplers.**
- **run 5 days, 1000 cpus, 10 samples for 100 cpus each.**
- **Adjoint-based methods -> out of core adjoint computations for transient models.**

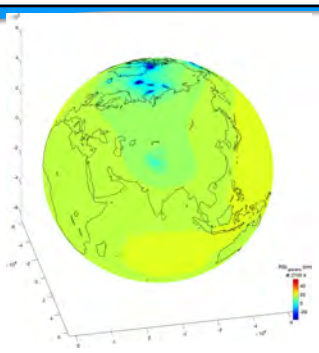

Sea-level glacier fingerprint computed using ISSM sea-level solver.

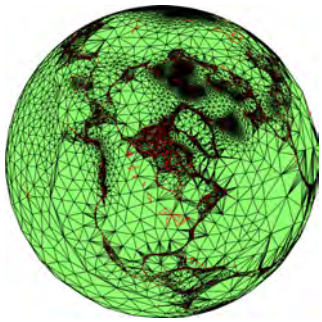

ISSM unstructured earth mesh.

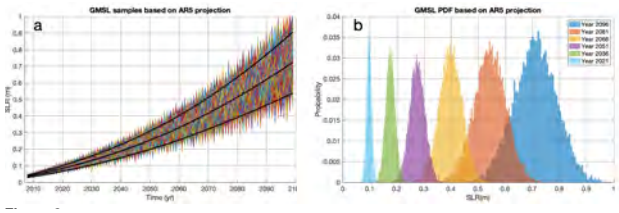

ISSM probabilistic projection of sea-level using ice-sheets, glaciers and ocean thermal contributions.

Requirements

- Unstructured meshes: very hard to use GPU.
- Very fast infiniband because of FEM approach.
- Challenge of using sparse direct solvers and spare iterative solvers.
- Challenge of full connectivity dense operations from each cpu core to all other cpu cores.
- Larse memory requirements on each node.
- Adjoint based methods: out of core needs, 50 Tb per run (at a minimum). Very fast LFS requirements.

National Aeronautics and Space Administration

Use Case: 3D Realistic Modeling of the Sun

Irina N. Kitiashvili
06/01/2020
NASA ARC/TNC
SMD

Mission Objective

- **Mission: Obtain high-fidelity 3D radiative models to reproduce the multiscale solar dynamics from the interiors to the corona to enhance predictive capabilities of activity manifestations, impacts on the space weather conditions and Earth's environment.**
  - To accomplish the mission goal, significant computational and storage capabilities are required
- **Criticality: The mission goals cannot be accomplished without the HEC support.**
  - *If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?*
    1) investment in a large computer cluster and storage system;
    2) investigate possibilities to get access to different supercomputing systems (NCAR, DOE, XSEDE, Blue Waters)
- **Mission Security Needs: (Confidentiality, Integrity, Availability)**
  - HEC is critically important for this mission because of the new instrument developments, new science requirements to the model performance and needs to produce realistic synthetic observations for data calibration and interpretation. In a case of interruption of HEC availability, no realistic-type models can be obtained, the ability to perform physics-based analysis of available data will be very limited.
  - Have the fast and secure access to identified data sets for external collaborators (without NASA credentials)
  - Development of a data request system
- **HEC Function:**
  - To support the mission objectives it is requires to run application on a non-stop basis
  - The mission output used for calibration and interpretation of observations from NASA's space missions, understanding the observed phenomena, development capabilities to predict space weather conditions, verification and validation new data analysis methods, support new instrumentation development, design to identify requirements and specifications

Work flow

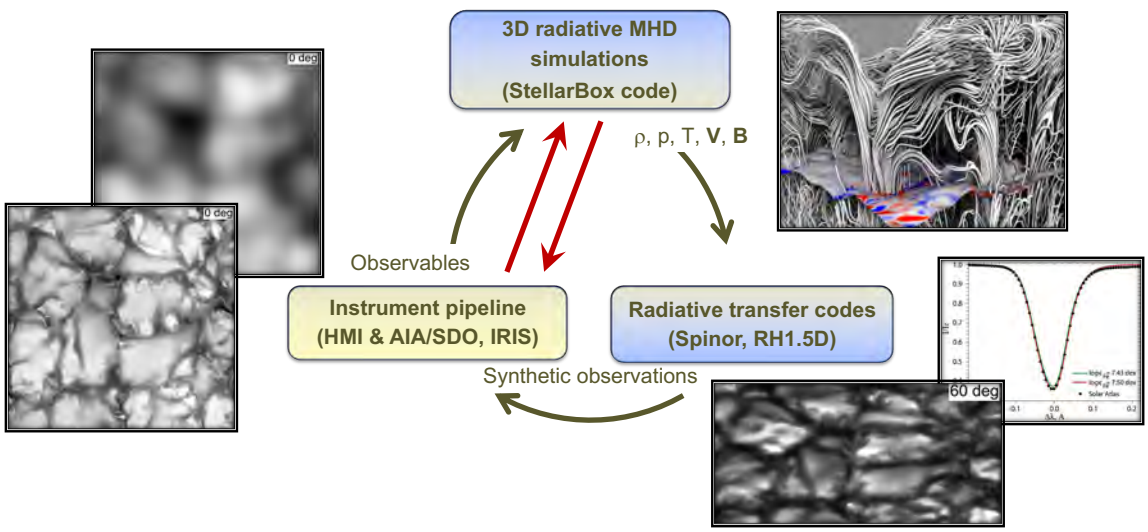

Current HEC Application Characteristics

- **Data:**
  - In: StellarBox code (Wray et al., 2018) takes into account compressibility of the highly stratified plasma, 3D multi-group radiative energy transfer between the fluid elements, a real-gas equation of state, magnetic effects, subgrid turbulent transport.
  - Out: The output data used for understanding the observed solar phenomena (e.g., eruptive activity, self-organization processes) and link the observations to physical processes which cannot be observed directly, validation of data analysis techniques (e.g., helioseismology, spectro-polarimetry) and discovered phenomena
    - The output data require retention and additional processing to convert the data into the NASA's missions synthetic observables
    - The data are analyzed using different approaches such as feature identification and tracking, statistical analysis, ML techniques, and 3D visualization.
  - Transfer: From local to the HEC storage system (Lou system, nobackup disk), local disk, data portal
- **Computation: Type: Cluster computing.**
- **Evaluation:**
  - Performance of computational efficiency, available disk storage, time in queue, work of the support team
  - Tools for the data analysis, visualization and expertise support

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect theload on computing assets and cause it to vary?**
  - Large-scale, non-linear models and post processing of the data, application testing and revision require substantial multi-core computational resources (4,000-10,000 CPU). This requirement causes extremely long queue wait time.
- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - Constraints on Use Case: big data analysis (including ML), disk space, computational allocation, queue, network, I/O
  - Impact: insufficient performance to obtain the required models and data analysis.
  - External Sources: no. (previously used: XSEDE system, Stampede)
- **How will this Use Case evolve over the next decade?**
  - Growth 1: Increase the computational domain for investigation of large-scale solar eruptions and extreme events
  - Growth 2: Application of the 3D radiative MHD modeling to the full spherical Sun to model activity cycles from first principles

HEC Needs Assessment Workshop

Affinity Group Sessions

Monday 6/1/20 1pm – 4pm (AG6-1) (nominal affinity: Earth Modeling and Helio) (10am – 1pm Pacific)

Facilitator: Peter Williams <peter.williams6850@gmail.com>
Scribes: Jarret Cohen <jarrett.s.cohen@nasa.gov> and Sean Keefe <rosemary.e.keefe@nasa.gov>

NOTE: This set of notes contains all clarifications, changes, and comments received from Use Case presenters as of June 22, 2020. All changes are accepted/incorporated and all comments retained.

KEY POINTS

Mission Use Cases covered during HEC Need Assessment

Observing system simulation, Deep Learning models for hydrology simulation, 3D modeling of solar dynamics, heliophysics, commercial supersonic technology, Ice sheet and sea level system modeling

Impacts of HEC on Mission Needs: Range of Impacts and Significance

- Dependent: HEC resources are essential for realistic observing system simulations to allow robust simulation cases and conditions with high-resolution and high-fidelity. Model the observed system then the observing instrument to improve understandings of both.
- Dependent: System models demand HEC, but the Machine Learning and Deep Learning about those systems and the models, needed to improve understanding of both, depend on HEC, especially GPU-based systems
- Efficiency and Effectiveness: Cost and risk of a mission far more than those of simulations.
- Potential call-out: "Computer simulations and AI tools are the third leg of discovery. They help us explore where we cannot sense, can give us a comprehensive view of sparsely sampled environments, and place the heliophysics domain into a controllable laboratory setting---all of this at a cost that is a fraction of the cost of a mission." Lika Guhathakurta
- *Fundamental* research <u>and</u> *applied* research: Often simultaneously
- *"Reliability and scalability of the cloud is good for real-time, customer-driven operations [but] development cycles can't be done on the cloud because of the need for full-control of environment, large number of runs, and tens of thousands of GPU hours."* Grey Nearing
- HEC is critical to understanding high-risk, low-occurrence/low-probability events including rare, yet catastrophic events.

Mission Security Needs relevant to HEC: Range and Significance

Security needs discussed confirm the five areas mentioned in the Use Case outline (bullets) and adds the importance of <u>stability</u> (system instability caused by security threats compromises mission integrity) and <u>protection of intellectual property</u> (of particular concern in public-private partnerships)

- Confidentiality
- Integrity
- Availability: Varies, with some missions dependent and others less so
- Public Confusion/Controversy
- National Security

Criticality of HEC to Mission Needs

- Financial consequences
- Mission Security

Workflow Issues and Needs

- This section and the subsequent one about "key factors affecting load" might be integrated

Data Input, Output, and Transfer: Issues and Needs

- Possibly keep this section focused on processing speed (FLOPS) and data volume (BYTES) as a specific look at these two elements of workflow issues or needs
- *Input, output, and transfer* seem generally associated with--if not dependent upon--computation and storage. Could combine this section with the following section or explain why these are distinctly important given the purposes of this needs assessment (The idea is that either lumping or splitting are fine as long as understandably explained based on meaningful considerations)
- *Computation and Storage:* A principle aspect of this challenge is that output from some work becomes input for other work, thus storage between runs becomes central to workflow. Storage on or as part of HEC system has costs, as does moving off/on (transfer) because the volume of data is significant.

Evaluation: HEC Effectiveness and Value Measures

- Might roll this category under the earlier "Impacts" section [or cross-walk impacts and evaluation to ensure each reinforces the other – evaluate based on critical impacts]
- Performance of computational efficiency, available disk storage, <u>time in queue</u>, work of the support team
- Tools for the data analysis, visualization and expertise support

Key Factors Affecting Current Load on Computing Assets

- Publication cycles and University calendars
-

Anticipated Evolution of Use Case and Mission Need

- Maturity
- Growth
- Re-Use
- Re-Hosting
- Other(s)

Functional Gaps: Existing and Anticipated

- Constraints: big data analysis (including ML), disk space, computational allocation, queue, network, I/O
- Impact of Constraint or Gap: insufficient performance to obtain the required models and data analysis, mission failure,
- External Sources

Attendees

1. Mike Little, Convener, NASA Goddard, ESTO and HEC Program
2. Peter Williams, Facilitator
3. Presenter #1: Arlindo da Silva, GMAO, NASA Goddard
4. Presenter #2: Grey Nearing, University of Alabama and Upstream Tech
5. Presenter #3: Irina Kitiashvilli, NASA Ames
6. Presenter #4: Lika Guhathakurta, NASA HQ, Heliophysics Division
7. Presenter #5: William Doebler, NASA Langley, Structural Acoustics Branch
8. Presenter #6: Eric Larour, JPL
9. Aaron Pima, NASA HQ, Weather Strategy Plan
10. Dan Duffy, NCCS
11. Daniel Hedin, NASA HQ, Agency Institutional Readiness Project
12. David Considine, NASA HQ, Modeling and Analysis and Prediction Program
13. Blaise Hartmann, NAS
14. Elizabeth Hook, NASA HQ, HEC Program
15. Ellen Salmon, NCCS
16. Henry Jin, NAS
17. Jarrett Cohen, NASA Goddard, HEC Program
18. Johnny Chang, NASA Ames
19. Jon Jenkins, NASA Ames, Kepler/TESS/Pathfinder
20. Karen Petraska, NASA HQ
21. Laura Carriere, NCCS
22. Leigh Ann Tanner, NAS
23. Piyush Mehrotra, NAS
24. Robert Ferraro, JPL, Discipline Program Manager of Earth Science
25. Robert Hood, NAS
26. Sean Keefe, NCCS
27. Sharad Gavali, NASA Ames
28. Subhash Saini, NAS
29. Ted Manning, NAS
30. Tsengdar Lee, NASA HQ, HEC Program
31. Vidya Bobbiligama, NAS
32. William Thigpen, NAS, HECC Manager
33. Yan-Tyng Chang, CSRA

1. Use Case: Observing System Simulation for Earth Science Space Missions

Arlindo da Silva
arlindo.m.dasilva@nasa.gov
GSFC Global Modeling and Assimilation Office, Code 610.1
SMD

Main points:

P1: Simulating future NASA space-based Observing Systems based high-fidelity earth system model and instrument simulators. Earth system model simulations must span a full seasonal cycle at global scale. Criticality: this must be done an a HEC system for this type of model.

P2: Earth system model simulations can be shared with the public; early observing system simulations and trade studies may not be appropriate for public release, but there are no critical security issues.

P3: Future flow: frequent checkpoints with limited output. I/O speed becomes a bottleneck with higher resolutions if one continues to produce a full set of output. ML can be used to speedup model parameterizations and instrument simulator. Minimize I/O, use model as a compression algorithm: simulations can be restarted from any point during the year, with calculations performed online with full access to model state.

P4: Future place: HEC Center or public Cloud. Trading disk space to processing ability.

P5: Requires regular access to large number of cores. Instrument simulator plug-in protocol and user training required.

Raw notes:
Slide 2: Observing System Simulations: Mission Need (1 of 2)
Specific topic chosen for formulation of Earth science missions.

Mission concept: want to simulate that observing system with a high degree of fidelity. The flow is high-resolution, high-fidelity global physics model –> sample output of model –> apply instrument simulator, create a data stream similar to the instrument.

The goal is global coverage and fidelity.

Slide 3: Observing System Simulations: Mission Need (1 of 2)
It must be done on a HEC system. There is no way to do it on a desktop.

They are sharing data among the science team, not so much the public.

Slide 4: Current/Future HEC Application Characteristics
With increasing resolution and complexity, how does the workflow need to evolve?
It becomes a game where you have to read data from your files. I/O becomes more and more of a bottleneck.

You have to make the high-resolution run sequential. You do so with relatively limited I/O.

The user would provide instrument plug-ins that you would dynamically load libraries into as you go along.

Use model as the data compression algorithm. The idea is that instead of writing huge amount of 3D data every few minutes one would dump the "model state" much less frequently (say, weekly), and restart the model from anywhere along this trajectory. Once one restarts the model and integrates it forward in time you have full access to the model state, every time step, and by means of plug-ins can then deploy instrument simulators on-line.
This calls for running the earth system model, sequentially for a year, with weekly state dumps (checkpoints). A mission designer may want to simulate their observing system for select days/weeks during the year. So, she develops a plug-in for their instruments and restarts the model for the periods of interest, running it for a relatively short amount of time.

Slide 6
Massively parallel, but I/O gets in the way.

Q&A:
Q from Daniel Duffy (Int) to Everyone: 1:38 PM
Can the Nature Run simulation be broken into multiple streams, or does it require a single stream to process the entire time period?
A: Ideally we would like to do that in a single stream, because in reanalyses we have the observational data to bring it back to same point. With a freely running model, you will not have that; it will be [asynchronous, starting at different points?].

Q from William Thigpen (Int) to Everyone: 1:40 PM
Do you have a number of cores you require?
A: See slide 14 with Table 2.1. It is 40,000 cores for 6 km, perhaps less with optimization.

2. Use Case: Training and Testing Deep Learning Models for Hydrology Simulation

Grey Stephen Nearing
grey@upstream.tech
University of Alabama and Upstream Tech

Main points:

P1: Provide hydrology info. to various groups (hydropower, flood forecasting, water management), building and developing a large-scale hydrology model, development in Deep Learning, one set of code for R&D, another set for customers. Moving to ADAPT soon.

P2: The majority of computational requirements is in the training and R&D process. Real-time operations is comparatively cheap in terms of computational cost.

P3: Reliability & scalability of the cloud is good for real-time customer-driven operations. Development cycles can't be done on cloud, need full control of environment, large #s of runs, expensive to do in cloud, tens of thousands of GPU hours. Using only Python now for Deep Learning. DL codes are smaller in magnitude (lines of code, and computational requirements) than other hydro models like NASA Land Information System

P4: Biggest individual training run can be done on a laptop in a few hours, but scale up to multiple GPUs, which makes them fast enough to train many more models. Hydro models take less CPU and storage than GEOS. [Grey Nearing: Hydro models take *less* CPU and storage than atmospheric models because we only have two spatial dimensions, where the atmosphere has three. My point here was that although our problems are smaller, we are still pushing limits on GPU availability. This limits the number of models we can train, not the training of any individual model, whereas the situation *might* be different in atmospheric modeling where every individual model will be exponentially larger.]

P5: We need many people from many institutions to have access to this data (about a terabyte of data for small runs, but scaling up to tens of TB for higher spatiotemporal resolutions.) [Grey Nearing: I did not say this during the talk, but it is important. We are currently increasing spatiotemporal resolution, and this requires sharing of larger datasets that are starting to become a little slower to transfer. Global models at hourly resolution will require on the order of 100 TB of pre-processed data.]

P6: Limiting factor: bandwidth between GPUs. We could scale up to hundreds or (likely) thousands of GPUs. Just starting to use NCCS GPUs. [Grey Nearing: Just to add to this, we are very GPU limited right now. GPUs are expensive and we need more of them than we could possibly get. The reason we need a lot of GPUs is to train more models to help us do things like sensitivity analyses, test model configurations, test ensemble configurations, and do hyperparameter searches. These are standard tasks in any ML workflow, and we are a long way away from having enough GPU access to reach a limit of what we could do if we weren't computationally constrained.

But … this is only in the R&D phase. The operational phase is cheap by comparison. This is notably different than some more traditional modeling tasks, which are still expensive to run operationally. For ML, it is the training that takes a lot of time, then production runs are much cheaper.]

P7: Traditional ML or DL studies: all work done in development cycle. Hyperparameter searches will take as many GPUs as you can throw at it.[Grey Nearing: Hyperparameter searches (and most of the R&D workflow) are heavily parallelizable. There was a question about number of GPUS per node (connected by NVLINK). Right now we could use up to about 1TB of total GPU memory per node and see linear or super-linear speedup (less data shuffling between CPU and GPU). But as we increase resolution and move to global models we

will see potential benefit from 10s of TB of connected GPU memory (i.e., with NVLINK). If/when we start processing GEOS output directly in the deep learning architecture (planned for future), we would see benefit for 100s or TB of connected GPU memory.

The bottom line that I want to stress is that as more of the NASA projects start adopting hard-core deep learning work, GPUs (and NVLINK-connected GPUs) will be the limiting factor. This has been my experience working in ML labs over the last few years.]

Raw notes:
Provide information to regulators. We are building a large-scale hydrology model based in deep learning.

Two codes: One for public; one for selling by company (Upstream Tech).

Migrating to NCCS ADAPT cluster this week. Mostly work on university clusters. HPC needs are in the training portions. Real-time phase is a tiny fraction of the computation; computing costs are much lower.

We run our real-time system on the cloud for reliability and scalability. We can't do the development cycles in the clouds; doing large numbers of runs is prohibitively expensive when you need 10s of thousands of hours. We use GPUs.

Hydrology has significantly less data requirements than atmospheric modeling; hydrology is 3D, and atmosphere is 4D (e.g., LIS vs. GEOS). We need more GPUs than anyone could ever buy. We have about 6 years of GPU work in the queue. The limiting factor is the bandwidth between the GPUs. We are starting to use the ADAPT GPUs this week.

Q&A:
Q from Robert Ciotti (Int) to Everyone: 1:51 PM
Nodes can be built with from 1 to many GPUs (8 or even 16), Is there a particular size you prefer
A: It is always too small. We could easily go up to 1,600 nodes on a single shared memory GPU node. More GPUs on the same NVLink node.

Q from Robert Ciotti (Int) to Everyone: 1:52 PM
You mentioned some bandwidth limitations in connectivity?
A: Storage to memory and memory to GPU. With PCIe storage, we can get from storage to memory pretty fast. More memory is more important than more CPUs. The CPUs are used for data shuffling to the GPU memory. Having more GPUs allows putting more of your dataset on the GPU so that you have to do less shuffling from GPU memory to CPU. Within the node, NVLink is much faster than the network.

Q from Robert Ciotti (Int) to Everyone: 1:53 PM
Bandwidth requirements for I/O?
A: I don't have numbers.

Q from Robert Ciotti (Int) to Everyone: 1:53 PM
Data set sizes?
A: Smaller than 1 terabyte. We are not data limited. We get most of our data from Earth Engine. We do a lot of processing on that data. We have a Google Bucket. We would love to use the cloud as much as possible to avoid shuffling data around. Right now, cloud GPUs are too expensive. We get data from the gridded scale down (especially from global models) to the watershed scale; that takes a lot of time. We need a lot of people from a lot of institutions to have access to this data. That is where working with NCCS comes in.

Datasets include GFS data from NOAA and remote sensing from different instruments. Long-term outputs from global climate models are too big for us to store on the cloud.

3. Use Case: 3D Realistic Modeling of the Sun

Irina N. Kitiashvili
irina.n.kitiashvili@nasa.gov
NASA ARC/TNC
SMD

Main points:

P1: The HEC support is critical for achieving the mission goals because large-scale realistic modeling of the solar dynamics and activity require significant computational resources and data storage. Because of the large size of resulting data sets, the model analysis and visualization cannot be performed on local workstations.

P2: Security needs: It is important to provide fast and secure access to simulation data.

P3: Because of the large size of data sets that cannot be analyzed (or even opened) on local workstations, it is crucial to develop data request system, where a user can specify the data needs required for a particular application. It is challenging to make the data publicly available because in the current system the users who make their datasets publicly available have to host them in their personal /nobackup directory, thus reducing their data space for research. As more data are released to the community, less storage space is available for ongoing projects. Therefore, collecting the public data sets in a dedicated storage space (with backup) will be very beneficial for the community. The data sets with restricted access for collaboration can be stored in a dedicated /nobackup disk space.

P4: It is important to run applications on a non-stop basis. The simulation output is used for the calibration and interpretation of mission observations, development of requirements for new instrument designs, validation and verification of new data analysis approaches, as well as for the understanding of the physical processes that drive the solar activity and its impacts.
It is critical have a simplified procedure to get new allocation.

P5: The output data requires retention and additional processing as new observations become available.

P6: Requirements: More disk space, sufficient computational resources (4-10K CPUs), disk space, optimization queue, stable network, fast I/O for massive computation, ability to modify properties of a job during waiting time in queue for depended jobs.
It is critical to have a simplified and fast procedure to get a new allocation. In the present time, getting allocation for a new project takes several months to obtain after a project start date. The awarded allocation is always significantly smaller than requested in a proposal.
The unused allocation should not disappear in the next allocation period.
The allocation decision should be based not only on actual system usage but also the system stability, queue. Sometimes the allocation time cannot be efficiently used because of a few weeks of waiting time in queue.

Raw notes:

Slide 2: Mission Objective
Mission: Why do we need this?
This is needed because of the strong impact of solar activity on the space environment. The project objectives are provide theoretical support for the interpretation of observations and development of predictive capabilities for the radiation environment and extreme space weather events.

The use case produce high-fidelity 3D radiative models to reproduce the multiscale solar dynamics from the interiors to the corona to enhance the predictive capabilities of activity manifestations, impacts on the space weather conditions and Earth's environment.

Criticality: The project needs significant computational time, storage and tools to analyze massive data sets. Even a single 3D snapshot or a restart file of radiative MHD data is too big to open on a local machine. Therefore, the HEC resources are critical for obtaining the high-fidelity 3D radiative models of the solar MHD processes, and their analysis.

Mission Security Needs: The resulting data can be applied to advance a wide range of problems. Therefore, it is essential to improve the data access for researchers without NASA credentials. The dedicated disk space is required to avoid interference with the ongoing projects.
Because of the continuous development of new knowledge and instrumentation, we need to run the application on a non-stop basis by adding new regimes of solar activity and conditions and (we need to) reanalyze the existing models. Long waiting time in the queue can slow down the mission deliverables and jeopardize the research process. Thus, the optimization of the scheduling system is critical.

For the mission efficiency, it will be very beneficial to have access to the basic information about the current and coming computer resource reservations, such as a reservation schedule, the number and type of nodes which will be taken out of the system pool.

Slide 3: Workflow
The main computational effort in this project is focused on the generation and processing of high-fidelity realistic 3D radiative MHD models, which simulate different conditions on the Sun. This type of models is computationally expensive (millions of the CPU-hours) and months of running time. Each case requires significant data storage (over 100TB). For example, a minimal requirement for a hydrodynamic model for a helioseismic analysis in support of the SDO mission requires 52TB data for a low-resolution model. One MHD case requires a minimum of 70TB of data.
The models require long-term storage because of a substantial amount of physical processes taken into account, which can be used for solving a broad range of different scientific problems. An important part of this effort is the generation of synthetic observations which reproduce the instrumental characteristics and observing conditions and directly support the NASA's space missions. The generation of synthetic observations usually represents running many relatively small jobs that are I/O intensive. The synthetic observations have helped to interpret the real observations from the space missions (SDO, IRIS, Hinode), improving the data analysis pipeline, and providing the synthetic data for validation and verification of new data analysis methodologies.

Slide 4: Current HEC Application Characteristics
The StellarBox code (Wray et al., 2018) takes into account compressibility of the highly stratified plasma, 3D multi-group radiative energy transfer between the fluid elements, a real-gas equation of state, magnetic effects, and subgrid turbulent transport.

HEC Workshop "Key Questions"

- What are the **key factors today** that affect the load on computing assets and cause it to vary?

The key factors are large-scale, non-linear models, and numerous small-scale postprocessing jobs

Are there key functional gaps exist or are projected in the next 5 years or 10 years?

1. Jobs scheduling system
2. Data access system
3. The development of new codes or adaptation to modern computer architectures is near impossible for scientists because of a lack of funding and NASA restriction on code development within a grant.

4. Big data analysis and visualization tools

- Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?]

  Primary constraints are the long waiting queue and limited data storage. I could not submit a job for more than 2048 cores because of the extremely long queue waiting time. Sometimes even the 2048 core jobs stay in the queue for several weeks. Usually, I use the Cascade Lake nodes and Ivy Bridge nodes.

- Impact: [How does this affect fulfilling the mission need?]

These constrains decreases or delay the mission outcomes

Q&A:

Q from Daniel Duffy (Int) to Everyone: 2:10 PM
When you say more disk space, can you estimate how much more?
A: With helioseismology, a single snapshot is 80 gigabytes. I have to provide a whole data cube every 45 seconds. These data set can be provided to users by data transfer to a user's nobackup disk. If no NASA credentials I have to perform analysis by myself and send results to a collaborator.

Q from Brian Thomas (Int) to Everyone: 2:11 PM
Apologies if you mentioned and I missed it...what kind of ML for big data analysis?
A: We mostly apply the clustering analysis to synthetic data corresponding to IRIS and other instruments. Unfortunately, we have not capability to apply ML to big data. In particular, we are interested in developing DL to identify evolving features and track them in time. However, because of computational cost and model complexity, massive ML analysis needs to be developed. The development of ML algorithms for analysis of complex dynamical physical systems will be beneficial for a broad range problems in different fields.

Q from Robert Ciotti (Int) to Everyone: 2:11 PM
For the volumes of data you want to export outside, can you estimate urgency and volumes?
A: It is hard to estimate. The usual approach is to discuss with a collaborator to optimize the size of the data set by identifying that layers, parameters are needed for analysis or perform part of data analysis and transfer intermediate results. Usually, I transfer data to a user outside by small portions because of the user's disk limitations. Most people cannot process this data themselves. The typical approach to solving the problem of massive data set transfer of providing specific layers for selected variables. I have to process the same data set several times because usually during research, additional model parameters, slices etc. are needed.

4. Use Case: NASA Heliophysics Division (Note: No use case named, but this topic fits.)

Lika Guhathakurta
madhulika.guhathakurta@nasa.gov
NASA/SMD
Heliophysics Division

Main points:

P1: Heliophysics is a new science, as of just a few years ago. Billions of km of size in physical scope to study. Vast increase in data rates, archival volumes, and use of sophisticated instruments. ~2 TB/day from SDO, orders of magnitude more than prior instruments such as TRACE.

P2: It is critical to combine data analysis with computer models and data assimilation. We need AI.

P3: Traditional global models are not accurate or reliably predictive, which is needed esp. for space weather. We need exascale computing.

P4: Grand challenge: heliophysics are a frontier of knowledge…hard to predict several aspects.

P5: Build a NASA/NSF partnership on SW strategic capabilities: emphasize computer code development, data assimilating and mining, validation against observations.

P6: Enormous volume from sun to edge of solar system, complex spatial and temporal scales separated by magnitudes of 20 scales, in some cases, memory and data is huge. Integrate data science, computing, and implementation of AI.

Raw notes:

Slide 2: Heliophysics Past
Heliophysics reaches all the way from the Sun's interior to the outer reach of the solar system, billions of miles beyond the orbit of Pluto.
Space weather comes primarily from the interior of the Sun itself, and we have several missions trying to give us a 360-degree view of our side and the Sun side.

Slide 4: Multi-Scale Grand Challenge.
Heliophysics is also a frontier of knowledge; it pushes the boundaries of physics. How is the Sun's magnetic field generated? What causes the mysterious heating of the solar wind (corona?)?

Slide 5: Desired Progress Needed
We started with heavy use of HEC back in 2005.
The fundamental physics are embedded in all our problems.

Slide 6: Data Science, Computing & AI
With a bewildering array of agents, the heliophysics is highly complex over all scales, both spatial and temporal. It is interdisciplinary. Data science, computing, and AI are an absolute necessity.

Peter Williams provided this link (from presentation):
https://lwstrt.gsfc.nasa.gov/images/pdf/ACCEHS.pdf

Q&A:

Q from Brian Thomas (Int) to Everyone: 2:23 PM
Thoughts about AI/ML tools/tooling?

A: There is so much that is possible. We have only scratched the surface. You can take data (e.g., SDO), make models out of that, and extract features. The holy grail of solar physics is solar flares/X-flares. Add model value from solar observations.

---

5. Use Case: X-59 Sonic Boom Propagation Simulations

Will Doebler
william.j.doebler@nasa.gov
LaRC D321 Structural Acoustics Branch
Commercial Supersonic Technology Project
ARMD

Main points:

P1: Quiet sonic boom tech project, aircraft being built by Lockheed Martin, research to reduce sonic booms to sonic "thumps." Conduct test flights and survey the overflown communities. Research could possibly change sonic boom regulations. Physics-based simulation of X-59 sonic booms studying macro-level atmospheric effects.

P2: Use 100s of cores to plan booms for test planning. Pleaides helps with higher resolution.

P3: Security is key to public acceptance. Data loss would also not be good.

P4: Workflow summary-see slide. 5 TB of all output data.

P5: Large-scale, parallel computing required HEC clusters.

P6: Need more reduction is queue time, expanded MATLAB capability, Lustre storage space, and reduced turnaround for flight planning.

P7: Move from US to global-level modeling.

Raw notes:

Slide 2: Use Case: Mission Need

The aircraft is being built by Lockheed Martin. The goal is to create a sonic thump that will be acceptable to people on the ground to allow flying over land again.

Need to decide the flight plan, when to fly and where to fly.

Near the aircraft we have a signature that is complicated. The shocks come together and form a shock that is much quieter than past sonic booms. Come up with the loudness of sonic booms over various locations in the U.S.

Without HEC, we would have to run at lower resolution.

This was a one and done activity, but as we get different aircraft designs we could increase resolution and do global simulations instead of just over the U.S.

Slide 4 Key Questions

We used a Python wrapper, thus the Python request.

- Flight test campaigns, new/updated code inputs (CFD or atmospheric model), publication cycles, contract deadlines, unforeseen project needs

  - Constraints on Use Case: would appreciate reduced queue time. Would appreciate expanded MATLAB availability or support transcribing code. Would be nice to have more lustre storage space
  - Impact: reduced turnaround time would improve flight test planning and decision making
  - External Sources: None
- How will this Use Case evolve over the next decade?
  - Maturity: Likely will have on-going development for the next few years. X-59's flight test campaign ends within this decade
  - Growth: Move from USA-only to global resolution, increase resolution of inputs
  - Re-use: Possibly increase users, as others may want determine results for their aircraft designs. Possibly re-run the current use case as improved input data become available
  - Re-hosting: Need to maintain Python versions or support translation to other codes; may need to recompile executable for new architecture

Q&A:
Q from Robert Ciotti (Int) to Everyone: 2:37 PM
How to improve interoperability between Pleiades and K? (Note: K is a cluster at NASA Langley.)
I had few issues moving data; K is smaller. I have no complaints about the interoperability, but I may have too much MATLAB reliance, and there is not a lot of MATLAB availability on Pleiades.

Q from Robert Ciotti (Int) to Everyone: 2:38 PM
System configurations are similar enough to not be blocker?
A: Yes, configurations are similar enough not to be a blocker

Statement from Jon Jenkins (Int) to Everyone: 2:39 PM
If you have your own MATLAB licenses, you can use them on Pleiades and avoid license contention issues. (But it takes extra effort compiling MATLAB scripts into executables, which Will will look into, but is not something Will is familiar with)

---

6. Use Case: ISSM: Ice-Sheet and Sea-Level System Model

Eric Larour
eric.larour@jpl.nasa.gov
JPL

Main points:
P1: Ice Sheet and Sea Level model-fully unstructured, finite-element model system, which does not rely on grids—either meshes, or vertices on the Earth mesh, or ice sheet models that rely on simulated, unstructured models. Sparse connectivity.

P2: Demanding runs: altimetry data from ICESat-2 into ice sheet models. Lots of data dumped to disk, fast disk access required. Projections of sea levels. Very hard to use GEOS on unstructured meshes-no sign. Increase in perfection. We need very fast InfiniBand across clusters.

P3: Challenge: full-connectivity, dense operations from each CPU core to all other CPU cores.

P4: Very fast LFS requirements.

Raw notes:
ISSM is an unstructured finite element model system, which sets it apart from other Earth models. Nowhere do we rely on grids. The computations are done from vertices on the mesh to vertices anywhere else on the mesh. We repeat the solvers many times using the Dakota software. We typically run 5 days on 1,000 CPUs on Pleiades, with 10 samples of 100-CPU runs.
The other extremely demanding run is projection of sea-level. We end up dumping a lot of data while we are running, so very fast access to disk is required.

Q&A:
Q from Peter Williams (taken from key questions identified for Case Studies): What are the key factors that affect the workflows on the key systems?
A: It depends on whether we are doing adjoint runs or not, We almost never run on the NCCS systems because there is a one-day limit we can't get beyond. We want to run Monte Carlo samples; we split our big world into these mini-worlds. We are not doing embarrassingly parallel. We need 5 days of CPU time over 1,000 CPUs. We are very small in terms of footprint when we dump our result, except when we do adjoint runs and need 50 terabytes.

Q from Peter Williams (taken from key questions identified for Case Studies): How might the needs of this use case evolve over the next decade?
A: There is significant work to develop adjoint methods in C++. Argonne has a source -to-source transformation code (Fortran), which is used in NASA products including ECCO. We don't know if this is going to evolve into a stable or successful capability. For now we dump huge tapes on file. ICESat-2 is dumping huge datasets on us; NISAR (NASA-ISRO Synthetic Aperture Radar) will be even more data. We want to assimilate more and more of that data. Adjoint-based methods will be in huge demand and will we need to evolve from disk space and a stable LFS system
to STS transformations, computing heavily and not needing memory anymore?

Q from Ellen Salmon (Int) to Everyone: 3:14 PM
A late question for Eric L. - are there thoughts re getting data from NISAR, which (I understand) will be in the cloud? Or would you need to consider moving your HEC processing to the cloud due to large volume of NISAR data?
A: This is a case of growing by an order of magnitude. In 10 years, if we can compute Antarctica and Greenland over 10 million grids, that is a blip on the NISAR data. Will NISAR have Level-4 data? That is where I would want to operate. On the data assimilation and Earth system modeling with FEMS, If you are doing super-high-resolution grids globally, you are going to have a problem.

Q from Ellen Salmon (Int) to Everyone: 3:18 PM
Are others doing unstructured FEM grids? Do they have approaches to help re Eric's concern?
A: We're not the only ones running FEM for ice sheets. There are two other groups. We are all competitive with one another; there is not much difference between us.

Group Discussion of Key Questions

Mike Little: First time in a discussion like this…looking for conditions which may cause the evolution of your computing workflow on how much computational resource questions? **[Note: Any additional comments during the participant review would be most welcome.]**

Irina Kitiashvilli: I would like to have more support from computer science in terms of new tools to do more efficient data processing and organizing data. It is essential to move in this direction as datasets are getting larger and larger.

Eric Larour: It is more a fear than anything else. I notice more and more that discussions are geared towards GPUs and ML. In Earth system modeling, these things rarely come into play. I won't make that huge transformational shift if I don't see the performance gain. My fear is that I would find myself on the technological sidelines because of the computational structure that we have.

Arlindo da Silva: In computing history, moving from vector machines to MPI was a big shift. We make things work; they don't work by themselves.

Laura Carriere: What sort of data management support would be useful? Training? Software for search/discovery? Usage metrics? Other? **[Note: Any additional comments during the participant review would be most welcome.]**

Will Doebler: Would appreciate any training the HEC team is willing to put together. There was talk in the plenary session about achieving speedups by optimizing software for specific hardware. Perhaps showing an example of how to do this be a useful training. GPU use training, Intro to NASA HEC, etc.

Eric Larour: I find myself running MATLAB or Python again and again, reducing datasets to strictly what our models need. In ISSM we implemented multithreaded models, which are highly useful. Is there a way to run from the cloud using a multithreaded or even an MPI approach from a NISAR grid to an unstructured mesh or even a grid? I have five routines that I run all the time. It could be run on the cloud or on a cluster (e.g., Pleiades) as a service. We shouldn't have to run it ourselves.

Arlindo da Silva: Someone has to implement it. Python APIs. The service will require some configuration. The whole is interfaces.

Brian Thomas: What if you could deploy that service? Would that solve the issue? Are you willing as a user to deploy that service and allow it to be shared? It's whether they would be willing to solve the issue by creating a service for others. **[Note: Any additional comments during the participant review would be most welcome.]**

Eric Larour: I would be willing and have repeated tons of these algorithms. A lot of users will reinvent it again and again.

Arlindo da Silva: An unstructured mesh is probably something that has not been supported. Getting your data could be implemented server-side.

Eric Larour: When we initially look at the datasets that came from NSIDC, interpolation. In the data itself we have issues we have to untangle. There are other issues at play. ESMF has some of these mappings. Those should not be something that the user has to care about. If I request maps of SST.

Arlindo da Silva: You have to define this, it is not simple. You should have control of the interpolation.

Eric Larour: I had to develop my own interpolation to do mass conservation. If I specify a projection and want to conserve mass, I should be able to get that data without too much difficulty. DAAC centers should not do something [interpolation wise] that shouldn't be done.

Peter Williams: Irina, Luca, Grey, Will, anything to add to that conversation?

Robert Ciotti: RE: micro services. We are doing some work with JPL on a process pipeline that operates between remote DAACs, Amazon, and Pleiades. We then can publish those datasets.

Brian Thomas: Yes, I was hinting that if HEC provided microservices infrastructure support, the problem could then be solved (and shared) by your group **[Note: Any additional comments during the participant review would be most welcome.]**

Robert Ciotti: AI/ML make extensive use of SW middleware. Much of that can be ported to CPU. e.g., TensorFlow.

Lika Guhathakurta: Making modeling products available to the outside community is going to be a big need.

Grey Nearing: Going back to the comment about some applications working well on GPUs and others not, ML cannot be done on anything but GPUs. With other accelerators (TPUs, neuromorphic hardware), does NASA see itself keeping up with developments and seeing how they might be useful for different applications?[Grey Nearing: Full disclosure: I am asking here whether NASA has interest in funding work specifically to explore exotic hardware for science tasks.]

Piyush Mehrotra: We are always interested in looking at different technologies. There is an effort to do ML on CPUs.

Piyush Mehrotra: Has anyone looked at the cost of data in the cloud?
No answer. **[Note: Any comments during the participant review would be most welcome.]**

Will Doebler: No

Ellen Salmon: Anyone considering domain-specific languages yet?
No answer on WebEx but may need reflection from participants. **[Note: Any comments during the participant review would be most welcome.]**

Will Doebler: Not yet

Robert Ciotti: From the systems perspective, we are unfortunately somewhat disconnected from user requirements for I/O. Users show up and use what we have. So we often play catch up. HEC has focused mostly on CPU time required but very little on storage requirements. Then on storage requirements, there is not only volumes of data, there is the performance of the underlying technology that we provide. How can we better forecast needs here including the distribution requirements? **[Note: Any comments during the participant review would be most welcome.]**

Will Doebler: The use case I presented would benefit from increased Lustre storage space. The Lustre file number limit also limits throughput of my use case. Each core used creates several input and output folders. Some of them are deleted after execution, but during execution I can approach the file number and size limits, especially if I have more than one job running at the same time.

Kitiashvili: Agree. Because of the high computational cost of models, it is crucial to make them more accessible. Public release, even processed data sets, require a dedicated storage system with automatic backup. In present the time, publicly available data hosted by a data owner in the local nobackup directory.
Also, it is essential to have a procedure obtain allocation as project started. In a case, if the project finished, it is critical to be able to use the project allocation.

Arlindo da Silva: ESMF is capable to convert from lat-lon to scalable grids to whichever.

Mike Little: What is it that individual projects will need? We will have a separate survey to get that quantitative information. Do people have good ways of doing that kind of forecasts/making that sort of assessment? **[Note: Any comments during the participant review would be most welcome.]**

Irina Kitiashvilli: For Helio community, it is hard to have an accurate prediction in, for example, storage. With a new project I spend a lot of time on computation to get steady-state conditions. I do not remove these non-useful data until I start a production run. Sun is a highly dynamic object. It requires the development of different models that describe various activity conditions. Also, I keep old data for reanalysis if additional properties of plasma or observed phenomena need to be investigated.

Will Doebler: I agree it is hard to assess current and future usage. One reason I like using LaRC's K cluster is that I don't have to do any guessing/forecasting for my usage. It's just there for use even if I have to wait in the queue for a while.

William Thigpen: Is there anything that people need to do that the technology available today cannot address? **[Note: Any comments during the participant review would be most welcome.]**

Peter Williams: I think we have completed this discussion.

Mike Little: Thanks Peter. We will do a lot of follow-up via email. Thanks, everyone!

Appendix D2 Workshop Session 6-2 Inputs

Ashley Korzun	LaRC D205	Retropropulsion for Atmospheric Environments
Bil Kleb	LRC D305	Viscous Adaptive-Mesh CFD
Eric Nielsen	LaRC D302	Exascale Code Development for Emerging Architectures
Karl Edquist	LaRC D205	Supersonic Retropropulsion CFD
Melissa Carter	LaRC D301	Low Boom Flight Demonstrator Signature Validation
Craig Streett	LaRC D302	SLS Buffet / Aeroacoustics
Nettie Roozeboom	ARC AOX	Red Rover: Connecting Experimental and Computational Facilities and Communities

National Aeronautics and Space Administration

Workshop Session 6-2

2 June 2020

National Aeronautics and Space Administration

Use Case: Retropropulsion for Atmospheric Environments

Ashley Korzun
2 June 2020
NASA LaRC, Atmospheric Flight and Entry Systems Branch
STMD – Entry, Descent, and Landing (TA09)

Use Case: Mission Need

- Mission: Human-scale and high-mass Mars landers require propulsive descent and landing with potentially significant plume-induced aerodynamics and environments
  - Limitations in ground testing require increased reliance on computational modeling and simulation for conceptual design through flight implementation for mission infusion, vehicle performance evaluation, and risk characterization and reduction
- Criticality: NASA HEC has been the primary resource for PD production development work, but larger investigations and parametric studies cannot be completed with present HEC resources
- Mission Security Needs: Work is not open-source
  - Maturation to flight implementation and realistic propulsion systems require increasingly controlled access
  - External collaborators from industry and academia
- HEC Function: Post-processing and scientific visualization
  - Integral both during and after solutions completed
  - Data used to quantify aerosciences impacts to vehicle performance and definition of design requirements

- Impact of HEC: Isolated, under-resolved, low-fidelity simulations can be performed on HEC resources. Timely exploration of increasing fidelities (space, time, physical modeling) and eventual production database construction entirely infeasible with present HEC. Researchers currently rely on DOD systems and proposal-based DOE allocations.
- Mission Concept:

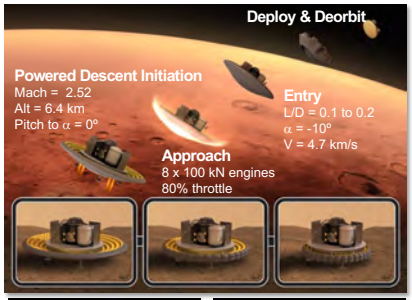

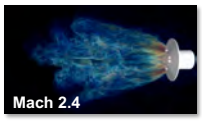

Use Case: Current HEC Application Characteristics

- Work flows for HEC simulations
  - Local geometry prep and gridding, then transferred to HEC
  - Solutions performed on HEC, results transferred back to local systems for analysis
- Work flows for Non-HEC simulations: $7.6M DOE allocations in 2019
  - Local geometry prep and initial gridding; transfer to DOD for grid completion; transfer to DOE for solutions
  - Solutions performed at DOE, basic on-site solution inspection; full datasets transferred to HEC at sustained rate of 40 TB/day for visualization; results transferred back to local systems for analysis

Mars Retropropulsion

- HEC: One run in ~9 mos on 5,000 SKL cores with 10-day waits for 5-day jobs
- Six runs done in 4.5 days on 3,312 GPUs
- Throughput of ~3 HEC facilities
- 985 TB of data generated

- Data
  - In: Computational grids of several GB to several TB
  - Out: Cutting planes, schlieren, iso-surfaces, contour maps at specified temporal frequencies (HEC); selected PDE solution variables for entire volume at specified temporal frequencies (non-HEC)
    - HEC solutions archived per researcher needs; non-HEC full solutions archived at HEC for long-term analysis
    - HEC solutions analyzed directly by researcher; non-HEC solutions require high-end visualization processing, followed by researcher analysis
  - Transfer: All outputs archived at HEC
- Computation
  - Type: HEC capacity solutions on CPUs; GPUs at capability scales are enabling technology for non-HEC solutions
- Evaluation: HEC will require substantial growth to accommodate these efforts in-house

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Desire for higher-fidelity simulations in space, time, and physical modeling; mission needs for more detailed characterization of design spaces and flight envelopes in lieu of ground/flight testing, uncertainty quantification
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Already limited in what can be done at HEC, in regards to both capacity and capability as well as timely turnaround
  - Impact: Lack of understanding of physical phenomena; ultimately increased cost and mission risk
  - External Sources: DOD, DOE
- How will this Use Case evolve over the next decade?
  - Maturity: Significant development work remains in the progression to an operation capability to support flight experiments and mission infusion with continued development
  - Growth: Use case will continue to grow in size and in number of solutions to be generated
  - Re-use: The number of users will continue to grow, branching across different activities (GT, FT, etc.)
  - Re-hosting: Have been working this for many years

Current Use Case (Survey)

- Achieving the mission objective
  - Aerodynamics of retropropulsion flows for atmospheric entry scenarios: time-dependent Navier-Stokes PDEs
  - Describe the application
    - Antecedents: Evolved from all major architectures of past 30 years
    - Programming languages used: 75% Fortran, 20% CUDA, 5% C/C++
    - Number of lines of code: 1M+
    - Workflow mechanism used to control code execution: shell script
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs: O(100)-O(1000) CPU cores on HEC; O(500)-O(5000) GPUs on non-HEC (100K-1M CPU cores equivalent)
    - Accelerator usage (type and memory requirement): NVIDIA Tesla P100/V100 (16/32), demonstrated on A100 (40)
    - Amount of memory required (overall or per process): Varies greatly by algorithm and user options
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other: MPI + OpenMP + CUDA
    - Run in batch, interactive, or near-real time: Batch

Current Use Case (Survey)

- Describe the data requirements for the application
  - Data flow: CAD to Grid to Solver to Visualization options and engineering analysis
  - Data loaded serial or parallel
  - O(1) GB to O(100) TB generated per run
  - Annual data – Input: 50 TB, Output: 3 PB
  - How does the output get analyzed: Visualization, time-series analysis
  - Describe the type of storage that this application needs: High-speed shared file systems, long-term archival systems
  - How is input data brought together for use? CAD clean-up
  - How much data needs to be curated (archived)? Several PB/year
  - How much data needs to be distributed and, generally, to whom? ~500 TB to general research community
- What are the key computing obstacles to effective execution of the current application? MPI latencies and synchronization
- Are there any analytic functions performed on the output of the application?
  - Time series analysis
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Yes (but never enough)

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals? Resources: Lack of early access to emerging hardware, timely hardware procurements of sufficient size. Application: Concurrency in temporal direction, multiphysics and associated load-balancing, increasingly disparate spatial and temporal scales.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components: higher-order, adaptive, chemistry, multiphysics
    - Increase in resolution and/or timesteps: Sky is the limit; constrained even on today's leadership class platforms
    - Rewrite to improve performance: Currently working with academia, other labs, vendors
    - Refactor for emerging platforms: Currently working with academia, other labs, vendors
    - Rehost onto cloud: Not likely at desired scales
    - Introduction of AI/ML/DL: Possibly
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores): Currently at 100K-1M cores equivalent; growth factor of 100-1000x in 10 years
    - Accelerator usage (type and memory requirement): Production use of accelerators today; currently preparing for further growth in heterogeneity of various types
    - Amount of memory required (overall or per process or per core): Will evolve in time as algorithms and data mapped to new hardware as needed
    - Parallelism of the application: Currently MPI at node and/or CPU core level, OpenMP threads at CPU core level, CUDA threads on GPU. PGAS implementation for GPUs implemented, currently optimizing with vendor. Exploring many types of programming model abstractions and comparing with optimized native implementations to inform long-term path forward.
    - Storage requirements: Likely to continue growing at fast pace
    - Run in batch, interactive, or near-real time: Mostly batch, near-real time for small engineering-class jobs

Use Case: Mission Need

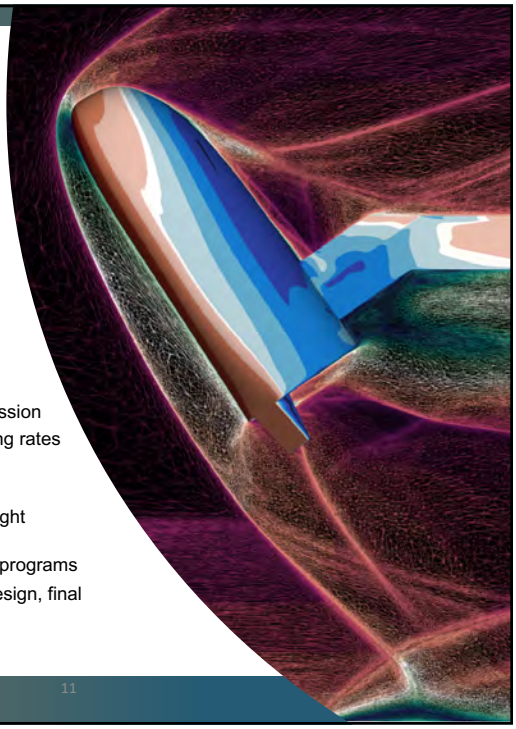

Mission: Aeroscience analysis and design of aerospace vehicles via computational science

HEC impact measured by portion of vehicle analysis and design conducted with error-informed HEC numerical modeling and simulation

Criticality: NASA's HEC is essential to accomplishing this work

If NASA HEC was no longer available to meet needs, would consider

- high-performance cloud computing,
- augmenting various center compute capabilities, and
- propose targeted studies to NSF, DoD, and DoE.

Mission Security Needs: (Confidentiality, Integrity, Availability)

Work ranges from open to SBU to classified, depending on phase and nature of mission

Uncertainty in HEC availability primarily causes schedule slips and impaired learning rates

Currently limited need for non-NASA-credentialed collaborators

HEC Function: HEC allows numerous CFD cases to be run for design exploration, flight database development, and/or uncertainty quantification

Simulation cases are typically run year-round to support various phases of various programs

Output is used to explore trade space during mission design, preliminary vehicle design, final design, mishap investigations, and "911" calls

Use Case: Current HEC Application Characteristics

- Workflow: AIAA 2019-2948: Case setup, launch simulation on coarse mesh, automatically adapt grid to final mesh, post-process results
  - Trigger: request for aeroscience analysis by program
  - Limited interaction by user during computing workflow; primarily only troubleshooting if automation fails
  - Infrequent requirement for quick reaction compute time except during mishap or specific mission needs (see slide note)
- Data:
  - In: Low-volume CAD geometry model and case-specific boundary and initial conditions
  - Out: Time series of vehicle flow field, which is typically distilled to reduced-order data.
    - High-volume flowfield data typically only retained until design analysis cycle complete or mission successfully flown.
    - Reduced to aerodynamic force and moments, pressure distributions, power spectral density distributions, or other reduced-order data via numerical integration, statistical analysis, and signal processing techniques.
    - Flowfield visualization often used to explore and understand physical phenomena
  - Transfer:
    - High-volume flowfield data temporarily stored on HEC storage system
    - Information necessary to create flowfield data and reduced-order outputs stored locally and/or in offsite version control repository.
- Computation:
  - Type: Highly-parallel batch CPU processing that grows from single node to thousands of nodes as mesh adaptation proceeds
  - Migrating to GPU processing and more frequent I/O (due to migration toward time-accurate modeling)
  - Ranges from capacity computing to capability computing as simulation mesh-size, physics, and other modeling requirements grow
- Evaluation: Currently not effective due to lack of support for growing job size, i.e., opposite of pbs_release_nodes, which sheds nodes

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Within a simulation: start with very coarse discretization and “adapt into” final mesh size, i.e., 1 to 1000s of nodes
  - Phase of project, e.g., pre-phase A, phase A, PDR, CDR, “911 call”, and so forth
  - AIAA summer and winter conference publication deadlines
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: current resources limit scope and quality of simulations attempted in terms of physics, time scales, volume, and uncertainty estimation
  - NAS currently divided into several machines, no provision for “next available” across machines
  - Impact: Throughput is highly uncertain due to queue wait time uncertainty
  - External Sources: No current use of non-NASA computing facilities
- How will this Use Case evolve over the next decade?
  - Maturity: Move from development state to operational capability within three years
  - Growth: Number of simulations, mesh resolution, number of time-steps, and physical model fidelity will all increase
  - Re-use: Number of people running this type of application with grow
  - Re-hosting: Strong need to migrate code to future computing architectures, especially mesh adaptation mechanics and high-volume, high-frequency output

National Aeronautics and Space Administration

Use Case: Exascale Code Development for Emerging Architectures

Eric Nielsen
June 2, 2020
NASA Langley Research Center

Principal Investigator, NASA OCE Effort on Emerging HPC Architectures
Technical Lead, NASA Langley HPC Incubator
Principal Investigator, DOE Leadership-Class HPC Awards
External panelist for DOE and DOD HPC Programs
Oak Ridge Leadership Computing Facility User Group Executive Board
AIAA CFD 2030 Integration Committee, HPC Roadmap Steward

- Code development targeting exascale CFD and multidisciplinary simulations across speed range, adjoint-based algorithms for error estimation, design, UQ, including application to chaotic systems.
  - More and more missions demanding higher fidelities in space, time, and physics modeling
  - Waiting months for an answer on a few thousand CPU cores is not feasible
  - Must scale up and also leverage emerging architectures
  - E.g., GPU-based simulation at capability scale is a game-changing technology for multiple missions
- HEC staff always very responsive and a pleasure to work with
- HEC has provided outstanding support for high-end visualization and high-capacity file transfers
- HEC is useful for CPU-based code development at small scales, but not amenable to larger scales, emerging architectures
  - Frequent, fast-paced access to large resources required to prepare applications to run at scale
  - Does not provide early access or assistance for emerging hardware technologies
  - Must frequently leverage other facilities, labs, vendors, and academia for these needs
- HEC would require substantial growth to accommodate these efforts

Mars Retropropulsion

- HEC: One run in ~9 mos on 5,000 SKL cores with 10-day waits for 5-day jobs
- Six runs done in 4.5 days on 3,312 GPUs
- Throughput of ~3 HEC facilities
- 985 TB of data generated

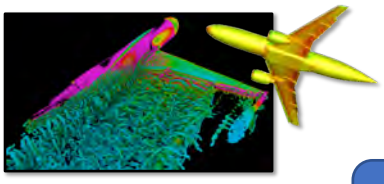

Transport High-Lift

- HEC: One run in ~6 mos on 5,000 SKL cores with 10-day waits for 5-day jobs
- Done in 31 hours on 828 GPUs
- 322 TB of data generated

JobID	Min/Max CPU	Wallclock Limit	Status	Last Event Time
4172092	85052/85052	28800	Finished	2020-05-14 03:52:58
4170592	12188/12188	28800	Finished	2020-05-13 11:59:41
4170478	12188/12188	28800	Finished	2020-05-13 11:44:14
4169485	12188/12188	28800	Finished	2020-05-11 10:11:49
4159993	12188/12188	28800	Finished	2020-05-10 22:16:25
4159969	12188/12188	28800	Finished	2020-05-10 21:41:55
4159396	48620/48620	28800	Finished	2020-05-10 17:23:05
	12188/12188	28800	Finished	2020-05-10 11:43:48
	10692/10692	14400	Finished	2020-05-10 01:00:01
	21340/21340	28800	Finished	2020-05-10 00:40:32
	10692/10692	14400	Finished	2020-05-09 23:54:45
	21340/21340	28800	Finished	2020-05-09 23:53:03

A Week on a DOD System

- 85,000-core job took about 4 hours to begin execution
- All others took less than an hour

- Size of runs: O(100)-O(1000) CPU cores on HEC; O(500)-O(5000) GPUs on non-HEC (100K-1M CPU cores equivalent)
  - FUN3D won Gordon Bell Prize 20 years ago for runs on 2,000 cores; this is still the norm on HEC
  - Expect growth of 100-1000x in next 10 years: trillions of DOFs
- Can fully solve PDEs using NVIDIA GPUs with no H/D data motion
  - Ten years of partnering with NVIDIA, DOD, DOE, academia
  - Demonstrated breakthrough performance on V100 in late 2017
  - Tesla P100/V100/A100 outperform Intel SKL/CSL by 2x/5x/8x
  - Approaching billion-way concurrency: scaled to 6,144 V100s; equivalent to 1.2M SKL/CSL cores
  - Hundreds of TB asynchronous output per run; no performance hit
  - Tremendous benefits for engineering-class jobs as well
  - Others using this capability for production work, outpacing our own use of the code
- 5-10 year key computing obstacles
  - Resources: Lack of early access to emerging hardware, timely hardware procurements of sufficient size
  - Application: Performance portability, V&V in face of asynchronous execution, concurrency in temporal direction, latency hiding, multiphysics, load balancing, Amdahl's Law at the disciplinary and workflow levels, increasingly disparate spatial and temporal scales, task-based vs bulk-synchronous execution, PGAS models, asynchronous and mixed-precision algorithms leveraging AI/ML-oriented hardware, in-situ processing, critical-path algorithms with little/no concurrency
- Working other GPUs, ARM, FPGA, etc…future will be very heterogeneous
- Currently comparing performance of wide range of programming models with optimized native implementations to inform path forward

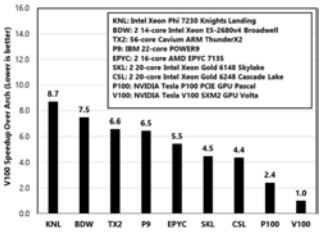

	Intel Skylake	ARM ThunderX2	NVIDIA V100	AMD MI-50	Intel PAC10 FPGA
OpenACC	TBD	TBD	1.1 Easy / Hard	TBD	No Plans
HIPSYCL	No Plans	No Plans	1.02 Moderate / Hard	TBD	No Plans
HIP	NA	NA	1 Moderate / Hard	1.32 Moderate / Hard	No Plans
SYCL For CUDA	NA	NA	1.02 Moderate / Hard	NA	No Plans
CUDA	NA	NA	1 Moderate / Hard	NA	No Plans
Vect. Intr. MPI	3.39 Easy / Hard	3.39 Easy / Hard	NA	NA	NA
OpenCL	TBD	TBD	TBD	TBD	200 Hard / Hard
OCCA	No Plans	No Plans	1 Moderate / Hard	1.32 Moderate / Hard	No Plans
Intel OneAPI	In Progress	No Plans	In Progress	In Progress	No Plans

- Concerned about long-term outlook for HPC at NASA
- Impacts long-term health of our science and engineering
  - Programmatic goals intentionally set low knowing compute is not available
  - Tough to attract top talent in computational science
- Hardware, software, workforce skills all need substantial, sustained growth
  - Very little engagement in broader HPC community
- **Would be happy to brainstorm paths forward**

November 2004 Top500

Rank	Machine	Site
1	BGL	DOE
2	Columbia	NASA
3	Earth Sim	Japan
4	MareNostrum	Spain
5	Thunder	DOE-LLNL
6	ASCI	DOE-LANL
7	System X	Virginia Tech
8	BGL	IBM
9		NAVO
10	Tungsten	NCSA

November 2019 Top500 + Future Systems

Rank	Machine	Site	Primary Compute	Perf (PF)	Power (MW)
	El Capitan (2023)	*DOE-LLNL*	*AMD GPU*	*~2000*	
	Frontier (2022)	*DOE-ORNL*	*AMD GPU*	*>1500*	
	Aurora (2021)	*DOE-ANL*	*Intel Xe GPU*	*~1000*	
	Fugaku (2020)	*Japan*	*Fujitsu ARM*		
1	Summit	DOE-ORNL	V100	144	9.8
2	Sierra	DOE-LLNL	V100	95	7.4
3	TaihuLight	China	Sunway	93	15.4
4	Tianhe-2a	China	Matrix-2000	61	18.5
5	Piz Daint	Switz	P100	21	2.4
6	Trinity	DOE-LANL	KNL	20	7.6
7	ABCI	Japan	V100	20	1.6
8	SuperMUC	Germany	Xeon	19	
9	Titan	DOE-ORNL	K20	18	8.2
10	Sequoia	DOE-LLNL	Power BGQ	17	7.9
32	Pleiades	NASA	Xeon	6	
39	Electra	NASA	Xeon	5	
117	Aitken	NASA	Xeon	2	

The following charts were used to stimulate a panel discussion where I was recently asked to speak to the state of NASA HPC along with senior leaders from DOE and DOD – thought these might be of interest. Please view in slideshow mode for dynamic content.

Audio/video from this forum is available at
https://livestream.com/aiaavideo/scitech2020/videos/200595144

High Performance Computing's Impact on Aerospace Prediction

Eric Nielsen
Computational AeroSciences Branch
NASA Langley Research Center

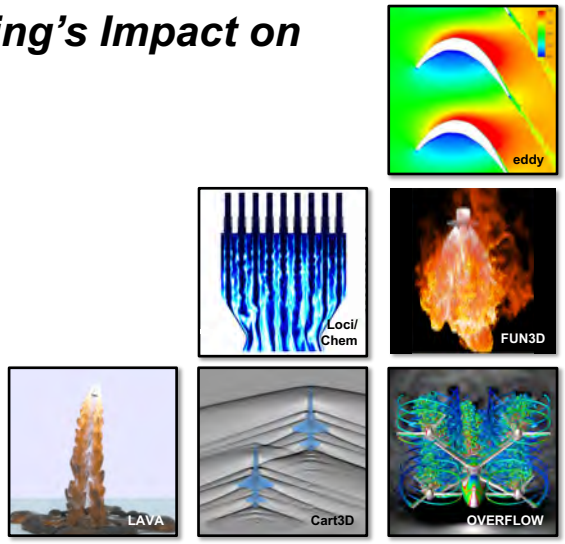

AIAA SciTech Forum 360
January 9, 2019

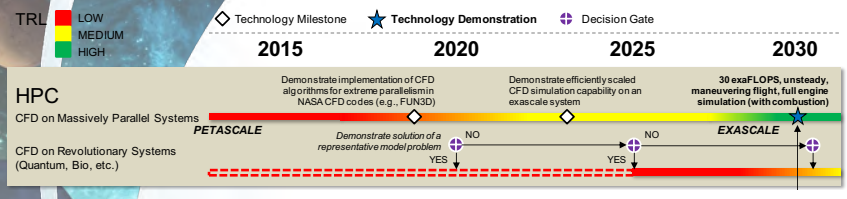

1. *"Current simulation software must be ported to evolving and emerging HPC architectures with a view toward efficiency and software maintainability."*
2. *"Investments must be made in the development of new algorithms, discretizations, and solvers that are well suited for the massive levels of parallelism and deep memory architectures anticipated in future HPC architectures."*
3. *"Increased access to the latest large-scale computer hardware must be provided and maintained, not only for production runs, but also for algorithmic research and software development projects…"*

"The aerospace CFD community is notoriously insular, publishing in AIAA or similar venues, with scant presence in computational science meetings hosted by SIAM, IEEE, and ACM…at a minimum, NASA should establish a presence at these meetings to keep abreast of developments in these areas…"

20

Euler Equations (inviscid) — *Reynolds-averaged Navier-Stokes: "RANS" (turbulence modeled)* — *Large-Eddy Simulations: "LES" (large scales resolved)* — *Direct Numerical Simulations: "DNS" (all scales resolved)*

Increasing physics, increasing cost →

- Today's analysis and design approaches typically rely on Euler and RANS simulations, each requiring $O(10^2)$-$O(10^4)$ CPU hours on moderate HPC resources
- CFD Vision 2030 calls for prediction of unsteady separated flows
- Current projection is full aircraft WMLES as a grand challenge problem in the 2060 timeframe using an entire leadership-class machine, with DNS following at infinity*
- For now, hybrid RANS-LES and canonical wall-modeled LES simulations are state of the art and may require $O(10^8)$ CPU hours on large HPC resources

Space Launch System Database Generation

Database	Code	Solutions (Grid Size)	Wall Time
Ascent	FUN3D	1,380 (60M)	2-4 weeks
Ascent	OVERFLOW	1,000 (500M)	2-3 months
F & M Wind Tunnel	FUN3D	600 (40M)	1 week
Booster Separation	FUN3D	13,780	3 months
Booster Separation	Cart3D	25,000	3 months

Derek Dalle, NASA ARC

* Spalart, P., Private communication.

An Histori

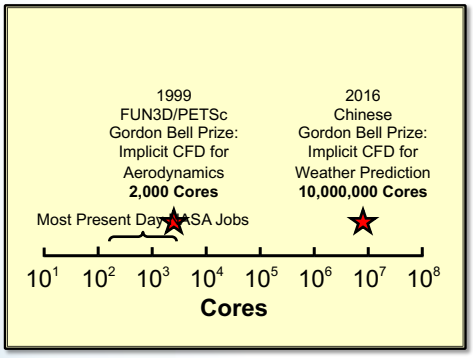

22

Software

- Strategic, long-term software development programs
- Project-driven environment leaves little time for skills development and exploring new paradigms
  - Heterogeneity is expected to grow dramatically over the next decade
  - Codes may run *slower* on new architectures if we stand pat
  - "I want the performance but I don't want to change my code"
- Proliferation of programming models
- Compilers can leave a bit to be desired

Algorithms

- Exposing vastly more concurrency
- Asynchronous, communication hiding, dynamic task-based schedulers
- Strong scaling, particularly for simulations with long time durations (HRLES, LES, …)
- V&V in the face of asynchronous execution
- Mixed-precision approaches leveraging specialized hardware driven by AI community

23

Hardware

- Job scheduling for capability development and applications
- Chicken and the egg, and the diversity of hardware in the data center
- Leveraging DOE hardware is very helpful, but not a silver bullet
  - Extremely high bar to qualify: equivalent of ~1 million cores now; ~10 million cores soon
  - Not enough compute available
  - Dependence on highly competitive, proposal-based systems not conducive to planning; high overhead for researcher
  - Sensitive data restrictions

Other

- Complex engineering workflows with items that have not been ported / scaled: Amdahl is a killer!
- Are we leveraging the potential of ML / AI / Big Data?
- Valuing computational science expertise on equal footing with traditional core competencies
  - Then attracting and hiring such workforce

24

Use of FUN3D on Summit for Mars Entry

- FUN3D shows 35x node-level speedup (6 NVIDIA Tesla V100s vs IBM POWER9)
- Scaled to 1,024 Summit nodes – performance equivalent of 1.1M Xeon Skylake cores
- Up to 200 TB of flowfield data stored per run
- Ensembles requiring years on capacity-managed Xeon systems are done in a workweek

NASA Langley HPC Incubator

- 3-year activity aimed at workforce development, HPC infusion into conventional projects
- Over 1700 participants in training courses; numerous guest speakers and collaborators
- Travel to HPC / computational science forums seldom supported by projects

FY20 New Start: "New CFD Algorithms Tailored to Emerging Computer Hardware Technology"

- 5-year project spanning 4 centers, funded by NASA Office of Chief Engineer
- Aimed at exploring workhorse aerosciences codes on new HPC architectures
- Workforce development through training, hackathons, boot camps, travel support
- Strengthening ties with OGAs, vendors, academia
- Procurement of testbed hardware
- Close coordination with NASA Advanced Supercomputing Division

25

National Aeronautics and Space Administration

Use Case: Supersonic Retropropulsion CFD

Karl Edquist, Bil Kleb, Ashley Korzun
2 June 2020
NASA Langley Research Center, D205
Descent Systems Study Project
Space Technology Mission Directorate, Game Changing Development Program

3/30/20 HEC Needs Assessment Plan B

Use Case: Mission Need

- Mission: Computational Fluid Dynamics (CFD) support of Supersonic Retropropulsion wind tunnel test
  - Uncertainty quantification of CFD for this problem is critical to the technology
- Criticality: NASA's HEC is critical to accomplishing this work
  - If NASA supercomputing was no longer available to meet your needs, we would seek federal HPC cloud computing and/or propose selected studies for DoE, DoD, and/or NSF resources
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - The current project is not ITAR
  - If HEC availability is interrupted regularly, major project milestones will be delayed or not completed
  - Currently there are no non-NASA users
- HEC Function: All CFD cases are being run on NASA HEC
  - Multiple jobs are run and post-processed daily
  - Output files are post-processed on NASA HEC or transferred to local computers
- Impact of HEC: NASA's HEC is critical to accomplishing this work
  - See above

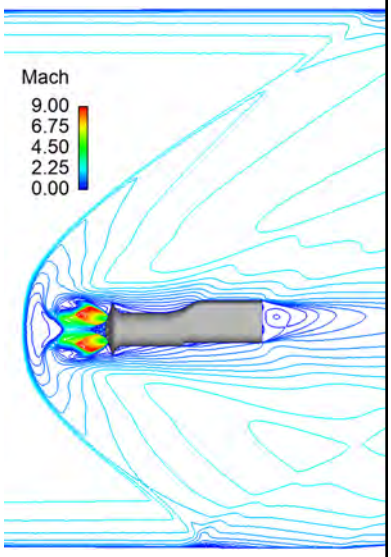

Use Case: Current HEC Application Characteristics

- Workflow:
  - CFD grid generation → CFD input deck creation → PBS job submission → post-processing by user → repeat as necessary
  - Is there a need for interaction by user, operator or other system in the computing workflow? Yes
  - Does this use case need quick reaction or near-real time computing? No
- Data:
  - In: User-supplied CFD grid, CFD input deck, and PBS script
  - Out: Binary CFD solution files, CFD code-specific files
    - Solution files must be retained
    - Solution files are visualized on local computers
  - Transfer:
    - Output files are transferred to the user's local system for post-processing
- Computation:
  - Type: CFD
    - Capacity computing for crippled milestones and for flight project milestones (e.g., Korzun's Summit work)
- Evaluation: The ability to regularly submit and complete jobs without significant delays (wait times can be days)

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Dozens of time-accurate CFD solutions are required on very large grids = significant memory and disk space requirements
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Reduced physics, time scales, design space, and uncertainty estimation due to lack of capacity and turnaround time
  - Impact: Cases need to be scaled back in order to be completed on the resources available in the required time
  - External Sources: No
- How will this Use Case evolve over the next decade?
  - Maturity: This Use Case will evolve into a NASA flight mission-critical function
  - Growth: The size and required number of this type of Use Case will grow
  - Re-use: As demand for the Use Case grows, more regular users will be needed
  - Re-hosting: Need workforce and software development to support migration to heterogeneous computing platforms and efficient I/O for time-accurate simulations

National Aeronautics and Space Administration

Use Case: Low Boom Flight Demonstrator Signature Validation

Melissa Carter
June 2, 2020
NASA LaRC, RD D301
ARMD CST

Use Case: Mission Need

- Mission: Low Boom Flight Demonstrator (LBFD) Signature Validation Analysis
  - Since FY19, the CST project has been working on improving the speed and accuracy of the codes needed for sonic boom signature analysis. NAS has been a critical part of testing the codes and baselining the speed improvements.
- Criticality:
  - NAS is considered **critical** as part of this mission. Without these computational resources, the CST project may have to expend significant operational resources to fine tune the flight operations of the X-59 aircraft through many more test flights than currently planned. Adjustments based on computational guidance will be much more efficient than additional flights of this experimental aircraft.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Due to the X-59 engine, data is EAR
  - Current plans have no external collaborators needing NAS access.
- HEC Function:
  - Current codes: LAVA, CART3D, FUN3D, USM3D, PCBOOM
  - Currently HEC primarily used for testing improvements to the code. However, during test flights, we will need to be able to submit runs, and have them completed overnight for analysis at the start of the workday.
  - The output of these data will show how well we are able to predict the near-field and ground boom signature, which is necessary prior to community acceptance testing

Use Case: Current HEC Application Characteristics

- Work flow:
  - Signatures from flight tests that don't match precalculated data sets (e.g., deflections different than predicted)
  - Interactions limited to ensuring that the proper computational model is uploaded and that jobs are submitted with the proper running conditions
  - Due to the nature of flight testing, the results will need to be turned around overnight
- Data:
  - In: 140-250 million cell grids, FUN3D, USM3D or LAVA input files
  - Out: Near field data, range depending on chase aircraft location, actual data needed for comparison will be a pressure line cut
    - Output data will be retained if proved to be beneficial to the analysis, retained on LFE until documented in final database
    - Statistical analysis versus chase aircraft and propagated to ground to compare statistically with ground measurements. There may be some visualization as well of shock structure in flight.
  - Transfer: to HEC storage system
- Computation:
  - Current run for LAVA: Wall Time: 13 hours, Number of CPUs: 1624, Pleiades systems
- Evaluation: Effectiveness of HEC will be measured by queue time and reliability

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Today, the loads are small and used for testing of code improvements.
  - Leading up to the flight tests, we will be building a database of potential flight conditions but this will be done without needing high priority
- Are there key functional gaps that exist or are projected in the next 5 to 10 years?
  - Constraints on Use Case: Biggest concern is queue time. Has been identified as a risk for the project.
  - Impact: The output of these data will show how well we are able to predict the near-field and ground boom signature, which is necessary prior to community acceptance testing
- How will this Use Case evolve over the next decade?
  - Maturity: Current developmental work will move to production for flights
  - Growth: HEC is a key part of the research portfolio. There will be ebbs and peaks in usage, but growth trend is anticipated to be upward.
  - Reuse: It is anticipated that the number of people will remain at or almost the same
  - Rehosting: Code refactoring, porting to GPUs and utilizing benefits offered by the latest hardware, may be needed to extract the maximum performance and throughput

National Aeronautics and Space Administration

Use Case: **SLS Buffet / Aeroacoustics**

Craig Streett
2 June 2020
NASA LaRC – D302 / NESC
SLS Program / NESC

Use Case: Mission Need

- Mission: **Simulation of turbulent flowfield about launch vehicle leading to unsteady forces on structure.**
  - **Data provided to structure-dynamics analysts leads to validation of design specs, identification of design exceedances, suggesting alternative structural concepts / optimal designs**
- Criticality / Impact of HEC: **Essential – NAS is only viable resource; requirements too large to obtain commercially**
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - **Program would likely not pursue this work at all, possibly leading to risk or suboptimal design**
    - **Could possibly pursue access to DoE systems, although data security is potential issue, and availability is difficult**
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - **National Security – LV technology is dual-use**
  - **SLS Program is milestone driven; inability to plan completion dates would force Program to consider alternatives**
  - **There is no need for external collaborators to directly access data on NAS**
- HEC Function: **Large-scale numerical simulation requires HEC**
  - **Large computational volume (cores x time); a given campaign might require 5-20 cases of 10-25M core-hrs <u>each</u>**
  - **Output data is extensively processed to provide inputs for structural dynamics analysis**
    - **Ex: Data at o(1M) locations x o(1M) time steps for each case**
- Mission Concept: **Accurate prediction of unsteady loads on LV, to provide data currently not possible to measure**
  - **See next slides for graphics**

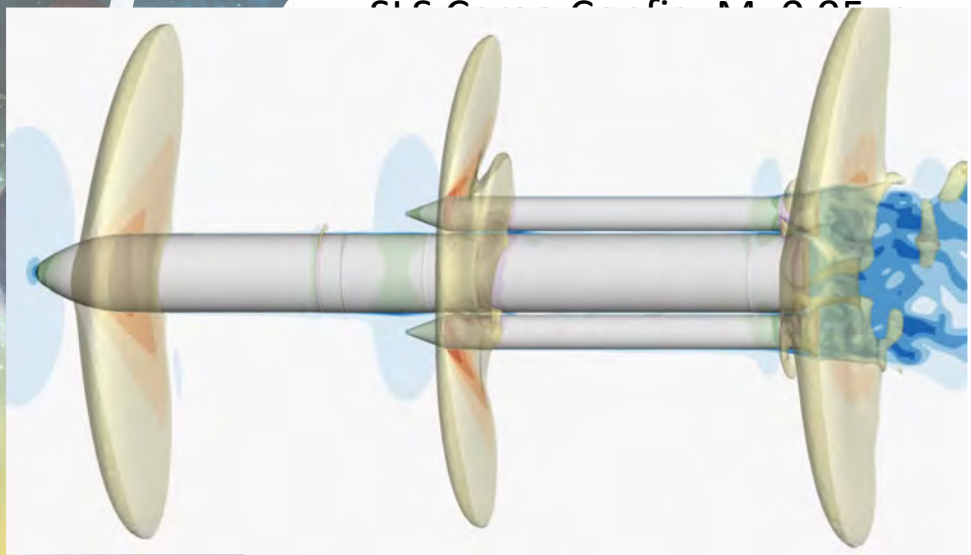

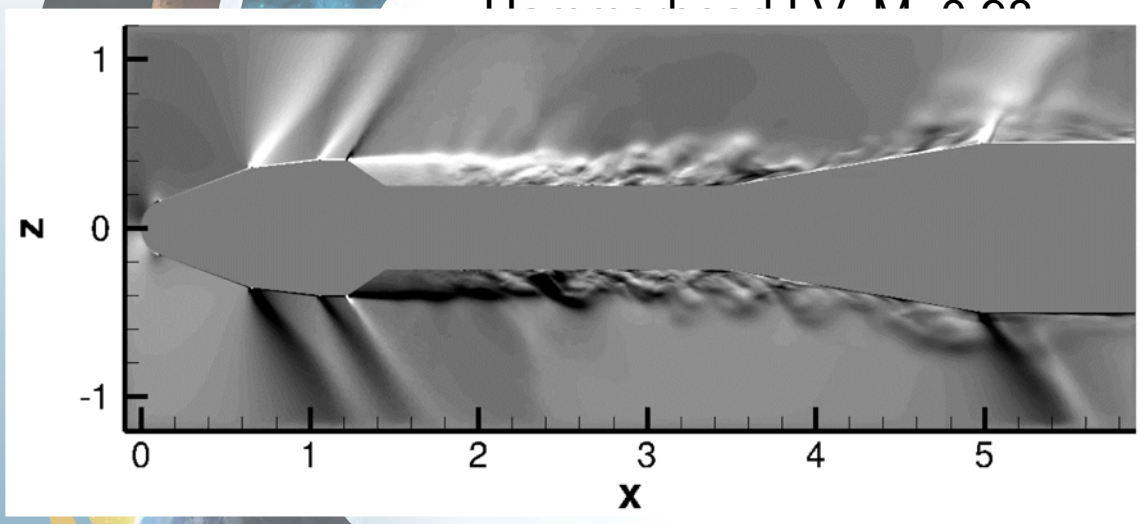

Use Case: Current HEC Application Characteristics

- Work flow: **OML CAD def'n, mesh gen (Pointwise®), trajectory def'n, flowfield initiation run sequence, data collection run(s), download/read/reformat raw surface pressure history data, time series analysis (rinse, repeat)**
  - **SLS Program triggers campaign by request through initiation of loads analysis cycle**
  - **HEC aspect of workflow is batch, with interaction only for data handling**
  - **Campaign is planned over O(months); quick response very rarely needed**
- Data:
  - In: **mesh – O(100GB), trajectory – O(10MB)**
  - Out: **Surface pressure histories collected – o(1M) pts x o(1M) time steps**
    - **Data retention desired for possible later re-analysis**
    - **Time-series analysis of surface pressures / small-region integrated loads**
  - Transfer:
    - **Data retained in HEC long-term storage, plus transferred to LaRC workstations for reduction/analysis**
- Computation:
  - **Sims are large-scale parallel numerical solution of PDE's – currently using either MPI on standard nodes or GPUs**
    - **Significant need for far more cores than NAS currently allows – (10-100x) – but data collection overhead will increase. Also problematic for GPU-based computation**
- Evaluation: **Adequate high-quality WTT and flight data exists against which test cases can be compared**

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - **Initiation of an SLS Program loads analysis cycle usually prompts many disciplines to each simultaneously begin computation campaigns – new reliance on asset reservations is not sustainable.**
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: **Capacity & turn-around are currently severely limiting**
  - Impact: **Program will not wait for results – will simply move this off critical path, and accept risk / suboptimal design.** ***We have the technology to do better, but infrastructure is severely limiting.***
  - External Sources: **DoE facilities are far more capable (1000x), but access is only available in special situations.**
- How will this Use Case evolve over the next decade?
  - Maturity: **Full acceptance of technology requires multiple successful applications, but opportunity limited**
  - Growth: **100-1000x increase in 4D resolution needed (now) for sims responsive to Program current needs – growth is to expand from select cases / conditions to generation of database**
  - Re-use: **This specific discipline will likely only have a handful of true practitioners for a while*, but wider acceptance by commercial LV designers is possible.**
  - Re-hosting: **Considerable effort already expended to port to GPUs; other architectures may require paradigm change.**

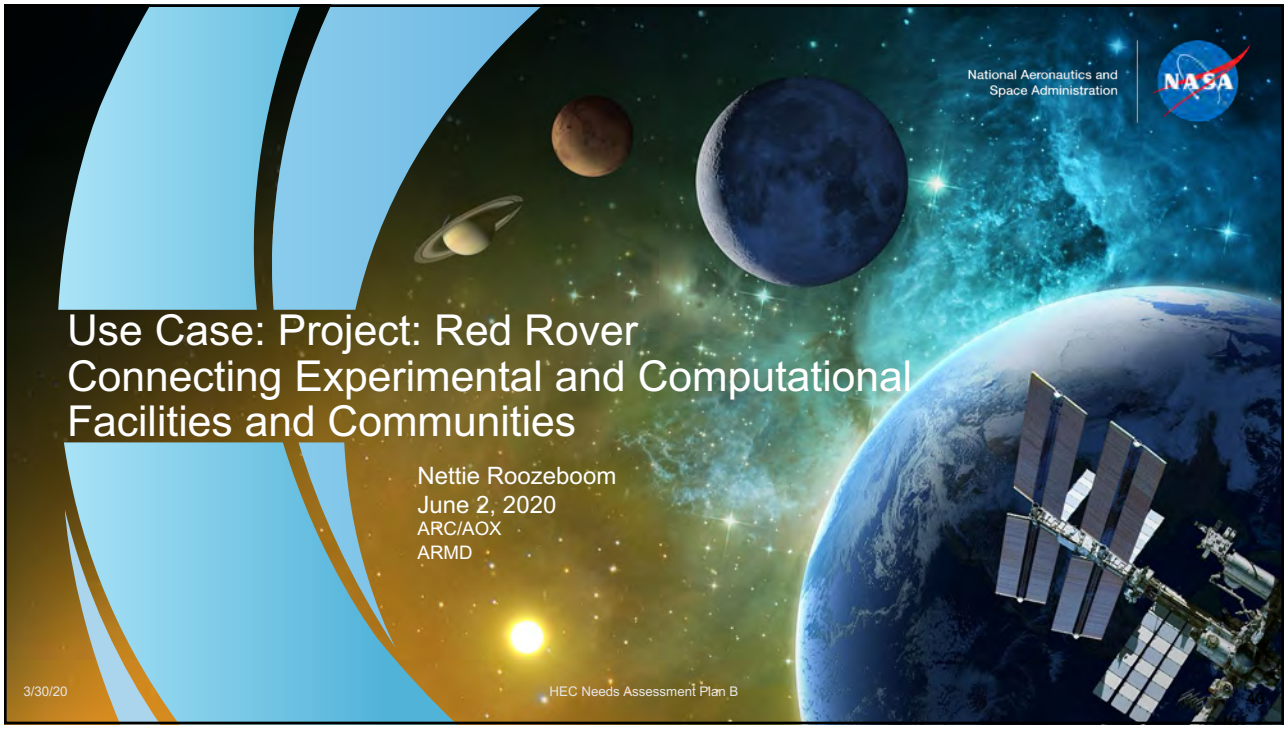

Use Case: Mission Need

- Mission: Connect the core competencies of high-end computing resources and knowledge to experimental facilities.
  - Can we run wind tunnel test more efficiently?
- Criticality: Experimental facility availability is a concern and computational fluid dynamics (CFD) tools are not sufficient for unsteady aerodynamic assessment
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    Spend $2Million in compute hardware and storage hardware that would need to be managed and replaced
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - National Security, Privacy (most wind tunnel data is either SBU/ITAR, company proprietary). This effort would not address Classified data.
  - If proper buffer storage is procured, then minimal impact. Would not want to decrease facility efficient
  - Currently, one individual from external company will request NAS access to pull data from NAS to their local company server.
- HEC Function: HEC was critical to designing, testing, implementing, and overseeing connection between facilities. Once the dat was at NAS, the processing software to process an optical data set was optimized (24 hrs → 5 min), and visualization team displayed results in near real-time.
  - Currently, this in development, but in future years, this connection could be used at least once per month.
  - Ultimately, deliver a better answer and improve efficiency of experimental facilities.
- Impact of HEC: NASA's HEC is highly essential. The overarching theme is combining core competencies of key NASA premier facilities.
  - Buy a system (not desirable)
- Mission Concept: From raw data to discovery and conversation.

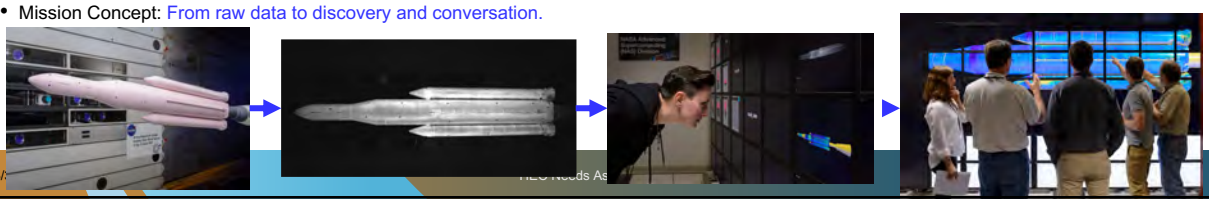

Use Case: Current HEC Application Characteristics

- Work flow:
  - A wind tunnel test is requested and planned
  - Is there a need for interaction by user, operator or other system in the computing workflow? It currently requires interaction, but automation (for the most part) could be introduced.
  - Does this use case need quick reaction or near-real time computing? Explain briefly. Yes, bringing high-end computational tools to assist in assessment of experimental data while the wind tunnel experiment is running is critical.
- Data:
  - In: September 2019, 150TB of uPSP optical data, from ARC Unitary Plan Wind Tunnel
  - Out: Current demo produced surface pressure-time histories for dense grid (~1Million nodes), Other data products are produced from there.
    - Unknown, currently writing results to nobackup (and raw data is on Lou)
    - Currently, data products and visualizations
  - Transfer:
    - With current demo, data is held at HEC local storage.
- Computation:
  - Type: [Summarize the type of computing used/needed and how HEC supports this today]
- Evaluation: HEC has enabled the transformation in how experimental tests are conducted and how to bring together core competencies of several organizations.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.]
  - uPSP technology is still being developed – grid refinement, understanding computations per frame (500k frames / data point * 500 data points) Maintain spatial resolution, but reducing increase signal.

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?] 3d visualization tools are slow – currently assessing how to share data with customer in interactive manner.
  - Impact: [How does this affect fulfilling the mission need?] Ideally have all functions in one place to avoid pulling data from NAS to local host.
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?] No
- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?] YES!!!
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?] Most likely
  - Re-use: [Will the number of people running this code/application grow?] YES!!!
  - Re-hosting: [What is needed to migrate code to future computing architectures?] To be honest, I don't know.

HEC Needs Assessment Workshop

Affinity Group Sessions
Tuesday 6/2/20 1pm – 4pm (AG6-2) CFD and EDL (10am – 1pm Pacific)
Facilitator: Marjorie Cole < marjorie.c.cole@nasa.gov>
Scribes: Jill Dunbar <jill.a.dunbar@nasa.gov> and Lindsay Ognoskie
Lindsay.t.ognoskie@nasa.gov

NASA's Goals that require High End Computing Capabilities:

Artemis: Returning astronauts to the Moon - Under Artemis, NASA will send new science instruments and technology demonstrations to study the Moon, accelerate plans to send astronauts to the Moon by 2024, and establish sustainable lunar exploration by 2028.

NASA, Boeing and SpaceX Launching Astronauts from American Soil - Boeing and SpaceX successfully launched American astronauts from the United States for the first time since the space shuttle's last mission 2011. Commercial spacecraft and rockets are delivering cargo to the space station, and more crew vehicles are planned.

International Space Station - The space station is a destination for next-generation spacecraft and a research and test environment that Astronauts use for exploration technology, advanced life support systems and human and robotic interfaces.

Mars and beyond - NASA's latest mission to Mars will include the most recent rover to head to the Red Planet, and to aid it in its exploration, it's carrying the first helicopter that will fly on another planet. Further out in the Solar System, OSIRIS-REx will visit the asteroid Bennu to grab a sample from its surface.

Quiet supersonic technology – NASA is continuing its research to make supersonic aircraft quieter and more amenable to flying in populated areas.

As NASA continues to lead the world in exploration and inspire the next generation it is vastly important to have in-house computational capabilities to reach these goals.

Key Point 1 – **Using Outside Resources** - Current HEC resources prevent NASA Mars and Human Space Travel; Deep Space, and Earth Science Research work from being completed and NASA currently turns to outside HEC resources that cannot be depended on for NASA requirements, especially regarding sensitive jobs.

- Ashley Korzunn – **NASA** Needs HEC for archiving data for long term analysis. These solutions cannot be generated on the present HEC systems as it would take years on current NASA resources where it takes a few weeks on DoD or DoE resources.
- HEC has been the primary resource and NASA relied on it to make significant leaps in progress however, it is now a struggle or nonviable to complete requirements as current investigations are at or beyond HEC capacity and investigations are continually growing in size.
- NASA team is forced to rely on external resources from DoE and DoD with no formal agreements set up to ensure NASA's requirements are met and It is not currently possible to run Sensitive But Unclassified jobs on external resources.
- Bill Klieb - HEC is critical to accomplishing this work; if there will be no in-house HEC there is no guarantee requirements and deliverables can be met. Recently tried to run at NAS and having trouble due to queue wait times; the lack of support for jobs that grow in size; and having to submit jobs to seven models across two classes of front ends.
- DoE and DoD are not long-term solutions:
    - they require proposals to be written and awarded, can't carry out a 5 – 10 year plan based on another agencies use of resources especially in a very competitive proposal award environment ;
    - require jobs to scale from 1 – 10 million cores used and not all NASA jobs are that large, DOE has a 7-year, $2 billion effort just for software development for their HEC jobs;
    - proposals are extremely time-consuming to formulate and prepare, months of background prep work goes into them and researchers don't necessarily have the time;
    - DOE systems available for such proposals are international, open-science resources. Cannot put ITAR or other classified data on them such as a lot of space-related work
    - Many NASA jobs do not fit the "shape" of the desired job topology at DOE. E.g., the 4D eddy-resolving CFD job, further complicating using outside resources.
    - DOE does not provide early-access to new architectures unless NASA is funded under their Exascale Computing Project. DoE funds/steers much of the technology development to vendors to get access to their capabilities earlier – hence a monopoly that effectively keeps others lagging several years behind.
- HEC is critical path technology on Mars EDL for 20 metric ton payloads for landing humans on Mars. . There is no current way to perform testing of this mission so need computational solutions. Currently NASA can do 4-5 metric tons at most; timing is important: landing in 2030 requires cargo landers with 20mt well in advance which means the models are needed in the 3-5 year timeframe to support that size payload.
- In aeroacoustics NASA is reaching the limits of NAS's capabilities and the work cannot use commercial computers or DoE or NSF due to ITAR. Work is planned months in advance and does not respond to changing needs of customer community. Need 100-1000 times more capacity to be able to respond to current needs. Large core jobs need reservations at NAS; one job has been waiting 23 days in the queue and still not run..

Key Point 2 – **<u>Incapable of Addressing Large NASA Quick Reaction Issues In-House</u>** - Current HEC can no longer support quick reaction runs when issues occur due to the size of the job in comparison to number of cores available and the queue times. NASA needs in-house computational capabilities to address mission issues especially as we move to human space flight and large constellations of science research systems.

- Bill Kleb – **NASA**s HEC no longer accommodates quick reaction runs due to job size as compared to the number of cores available and the queue times.
- Some projects are now running computational fluid dynamics jobs in support ARMD, HEOMD, STMD and NESC but are limited to the size job they can run in-house.
- As mesh refinement jobs continue to grow compute requirements keep doubling to eventually require hundreds to thousands of nodes, which NASA HEC cannot accommodate.
- Waiting 10+ day to try a tweak to the code when submitting 10-20 very brief runs at 10,000 cores a day to evaluate modifications/improvements is ineffective.

Key Point 3 – **<u>NASA's Outdated HEC is Eroding Science & Engineering Capacity</u>** - Current NASA HEC resources have not been modernized preventing advanced computational capabilities to be run in-house.

- Bill Kleb – NASA jobs require more frequent I/Os for time series runs, thus these jobs are migrating to GPU for greater efficiency but NAS cannot support these needs.
- As large computation jobs move towards using GPUs and related state-of-the-art (SOA) technologies NASA's HEC is not able to support these requirements due to a lack of modernization.
- Eric Nielsen - HEC would require substantial growth to accommodate efforts, expect 100 to 1000 times growth in next 10 years, trillions of DoFs. Jobs are around 1,000 cores on HEC, which was SOA about 20 years ago; current HPC SOA is equivalent to running on several million cores. Systems arriving in the next couple of years will be 10 times larger, around the equivalent of 50 million cores.
- Goals are currently set intentionally low because the current HEC system can't handle any more; Jobs NASA *should* be doing but cut back to what NASA *can* do instead.
- The long-term health of our science and engineering is at risk and the lack of compute power will gradually erode (is eroding) NASA's ability to perform cutting-edge science and engineering.
- NASA is losing good personnel and the opportunity to bring in new graduates because NASA's HEC is so far behind the state-of-the-art. NASA is forced to to get creative with computer science staff to come up with solutions, there is definitely a gap in our workforce.
- New architectures are being developed rapidly, NASA needs to adapt now, along with the industry and users.
- HEC does not provide the required rapid response to job turnaround that DoD does; there can be a long queue at HEC to get jobs run.

- Need programming models that permit code transfer to any platform; this is a tradeoff between portability and performance on a specific platform.
- Without access to resources, many additional flight tests will have to be performed than planned, increasing schedule and costs.
- CFD of aeroacoustics simulations are needed to obtain some solutions that can't be found in wind tunnel tests. If appropriate NAS capabilities are not available, Wind tunnel tests would probably sub-optimize design and increase risk.
- While STMD funds new technology development for EDL, they do not have robust funding for research, particularly if related to CFD simulations operating on advanced computer architectures. True engineering-oriented research in the space arena is a continuing issue for NASA. TRL<3 is difficult for any space-related problem, not just an STMD issue.

AG6.2 Presenters

1. **Craig Streett**
2. **Eric Nielsen**
3. **Karl Edquist**
4. **Melissa Carter**
5. **Nettie Roozeboom**
6. **William Kleb**

Attendees

7. Aaron Pina
8. **Ashley Korzunn- Updates provided**
9. **Craig Streett ** - No response to edits request**
10. Daniel Duffy
11. David Schuster
12. Elizabeth Hartman
13. Elizabeth Hook
14. Ellen Salmon
15. Emily Kuhse
16. **Eric Nielsen **Updates provided**
17. Francois Cadieux
18. Henry Jin
19. Jill Dunbar
20. Johnny Chang
21. Karen Petraska
22. **Karl Edquist - Updates provided**
23. Laura Carriere
24. Leigh Ann Tanner
25. Marjorie Cole

26. **Melissa Carter – Updates provided**
27. Michael Barad
28. **Nettie Roozeboom ** No response to edits request**
29. Piyush Mehrotra
30. Robert Ciotti
31. Robert Hood
32. Subhash Saini
33. **William Kleb – Updates provided**
34. William Thigpen
35. Yan-Tyng Chang

1. Use Case: Retropropulsion for Atmospheric Environments

Ashley Korzun
ashley.m.korzun@nasa.gov
NASA LaRC, Atmospheric Flight and Entry Systems Branch
STMD – Entry, Descent, and Landing (TA09)

Main points:

P1: Need HEC for archiving data for long term analysis. HEC role on this use case is limited because these solutions cannot be generated on the present HEC systems. Work completed on Summit took a few weeks, whereas the same solutions would take years on current HEC resources.

P2: HEC is the primary resource and we rely on it to make significant leaps in progress however, we struggle to complete requirements with current HEC system.

P3: An eventual database will require more resources, current investigations are at or beyond HEC capacity and investigations are continually growing in size.

P4: Project team is forced to rely on external resources for DoE and DoD (no formal agreements set up to ensure NASA's requirements are met). It is not currently possible to run Sensitive But Unclassified jobs on Summit.

Workload is year-round

Raw notes:

Slide 2: Mission Need
Applications for these vehicles – large systems- limitations for ground testing
Reliant on NASA HEC to test these landings
Increasingly controlled access to data and tools
Visualization of these large data sets needs to be done with HEC resources
Have partners for scientific visualization.
Use of Summit.

Slide 3: Current HEC Application Characteristics
Petabytes of data generated and transferred from DOE and back.
Archiving at HEC systems and stored there for future analysis
In some cases, full-volume solutions will take a while to dig into all data generated
Visualization done at NAS.

Slide 4: HEC Needs Assessment Workshop Key Questions
Doing work on this type of use case throughout the year- not projected to drop anytime in the future

Work from projects 2 years ago that was not completed because they can't get through the queues
Large parameter spaces become extremely computationally expensive.
Have partially completed work because can't get through queues at NAS.
Need to develop databases

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*
Q from Piyush (Int) to Everyone: time 1:25PM
mention of DOD systems. What scale of DOD systems are you using?
A: Use those systems to get around the 100,000 core capability.

Q from Mike Little(Int) to Everyone: time 1:27PM
Will this work be used for other planetary EDL?
A: Fairly specific to high-mass Mars applications. Not focused on plume-surface interaction in this use case but PSI has significant computational resource challenges as well and will be faced continually with insufficient resources and throughput over the next 4.5 years under the current project, as well as in providing requested support to HLS, CLPS, and MSR.

Q from Robert Ferraro (int) to Everyone: time 1:28PM
Are the problems the same/similar for the moon?
A: different due to presence of the atmosphere

Q :Mike Little (Int) to Everyone: time 1:28PM
Can you talk briefly for how you compete for DoE allocation? Is this handled through some kind of inter-agency agreement?
A: there is no interagency agreement. We built relationships over the years. This year, we had to submit competitive proposals and compete like everyone else.

Q Tsengdar Lee (int) to Everyone: time 1:29PM
Which mission directorate and what program funds mars retropropulsion research work?
A: STMD- all of this work is funding through STMD currently.

Q: Michael Little (Int) to Everyone: 1:30 PM

A: Michael Little (Int) to Everyone: 1:30 PM
100,000 cores

Q: Tsengdar Lee (Int) to Everyone: 1:34 PM
Does STMD expect your program to find your own computing resource?
A: We have leveraged as much as possible. STMD has been supportive of putting in proposals for external resources. Looking out to 5 years of resource requirements.

Additional info in chat: David Schuster (Int) to Everyone: 1:35 PM
While STMD funds new technology development for EDL, they do not have robust funding for research, particularly if you are talking about getting these types of simulations operating on advanced computer architectures. True engineering-oriented research in the space arena is a continuing issue for NASA. TRL<3 is difficult for any space-related problem, not just an STMD issue.

2. Use Case: Viscous Adaptive- Mesh CFD

Bil Kleb
bil.kleb@nasa.gov
Aerothermodynamics Branch
NASA Langley Research Center
ARMD, HEOMD, and STMD

Main points:

P1: Using adaptive mesh refinement that begins with very coarse grids that require only a single compute node; but as mesh refinement continues, compute requirements keep doubling to eventually require hundreds to thousands of nodes.

P2: Represents how some projects are now running CFD; Support ARMD, HEOMD, and STMD and NESC.

P3: HEC is critical to accomplishing this work; if no in-house HEC no guarantee requirements / deliverables can be met, Recently tried to run at NAS and having trouble due to queue wait times; the lack of support for jobs that grow in size; and having to submit jobs to seven models across two classes of front ends.

P4: HEC also supports quick reactions needs when issues occur but currently may not be capable of meeting these needs: size of the job compared to number of cores available and the queue times. Some load increases as part of preparation for summer and winter AIAA conferences.

P5: Require more frequent I/O needs for time series runs, migrating to GPU for greater efficiency.

Raw notes:

Slide 2: Mission Need
Question proposed: If you had a flight database that was developed by 50/50 wind tunnel and CFD, how would you measure that?
If NASA HEC was no longer available, they would have to figure out using cloud computing or something else outside- not guaranteed you're going to get anything.
All classified computing is done on secure clusters, not HEC

Most of the stuff we do is NASA mission related
Running cases year-round – different types of runs at different times
Function: Now encapsulating CFD into loop.

Slide 3: Current HEC Application Characteristics
Idea is it's supposed to be all automated, and only user involvement if there's a mishap
Viscous CFD cost increases appreciably over inviscid CFD like provided by Cart3D..
Also limited by the internal processor-class boundaries applied by NAS; would like to run across all intel processors, regardless of type.

Data Transfer--- info necessary to recreate the run and any reduced order outputs are what is stored on HEC
Do mesh generation as we go; informed by results.
Can use all the resources awarded.

Slide 4: Needs Assessment Workshop Key Questions
Phases of a project depend on how much we're doing
Key functional gaps: uncertainty estimation is currently very expensive /// We can throw jobs at you, but you have to line them up with particular machines
Issue: Can't just go run jobs. Queue wait times cause uncertainty; have to play games to submit on multiple systems.
External resources: idea is to bring this type of grid adaptation within three years
User Growth increases with more resources--- number of people has gone from about 3 to about 10

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*

Q from Piyush (Int) to Everyone: time 1:49PM
Is this a specific code or a generic description of an approach?

A: generic approach using FUN3D. need this to replace fixed grid. does grid convergence studies automatically.
Additional A: William Kleb (Int) to Everyone: 1:55 PM
Piyush: Generic description of approach (CFD growing their own grids via mesh-solution-adaptation from coarse to final); My specific approach (FUN3D-FV+refine+Engineering Sketch Pad) is documented in https://arc.aiaa.org/doi/abs/10.2514/6.2019-2948

4. Use Case: Exascale Code Development for Emerging Architectures
Eric Nielsen

eric.j.nielsen@nasa.gov
NASA Langley Research Center

Adjoint-based schemes.

Main points:

P1: HEC would require substantial growth to accommodate efforts, expect 100 to 1000 times growth in next 10 years trillions of DoFs. Most jobs are around 1,000 cores on HEC, which was SOA about 20 years ago. All of HEC is a few hundred thousand cores, though it is never used in that manner. Current HPC SOA is equivalent to running on several million cores. Systems arriving in next couple of years will be 10 times larger, around the equivalent of 50 million cores. This comment applies to both capacity and capability efforts – just need a lot more hardware either way.

P2: HEC has been helpful with things like visualization and high capacity file transfers, but runs are increasing, therefore the HEC resources need to increase as well. If not, the long-term health of our science and engineering is at risk. Lack of compute power will gradually erode (is eroding) NASA's ability to perform cutting-edge science and engineering.

P3: Losing good personnel and opportunity to bring in new graduates because NASA's HEC is so far behind the state-of-the-art. Need to get creative with computer science staff to come up with solutions, there is definitely a gap in our workforce. It is very hard to attract such skills going up against other national labs with "more toys" and long-term strategy and buy-in, as well as the Amazons, Googles, Facebooks, etc. Very many times we face well-qualified fresh-outs being offered 3-4 times salary above what the federal government can offer them (~maybe $75K).
I've been beating the bushes for years at job fairs, recruiting opportunities, career fairs at HPC/computational conferences, etc. On the rare occasion a hiring slot comes open, this issue comes up and NASA ends up falling back and just bringing in more engineers, etc because it is easier and that's what we've always done. The HPC Incubator at NASA Langley has worked hard to try to change this culture/mindset over the past 3 years.

P4: HEC staff is great but had to go elsewhere (DoE, DoD) for necessary core needs; new architectures being developed rapidly, need to adapt now but need significant funding to do so. DoE is not a good long term solution (Eric- please expand on why)

Many reasons DOE is not the solution to our problems. A few:

- Minimum bar for entry just to have a DOE proposal read is good scaling performance to 1 million cores equivalent. This will likely move to 10 million cores in a couple of years when their systems grow to 10 times larger. We have gotten to 1M cores with our code. I am not aware of any other NASA code that has pursued that size. Again, bear in mind NASA folks are doing jobs of about 1,000 cores on HEC. There is no mechanism for them to prepare their codes to jump 3 orders of magnitude – that is an *__enormous__* gap. And the lack of in-house workforce to bridge it. DOE has a 7-year, $2 billion effort just aimed at software development.

- Proposal-based system is very high-risk. Extremely competitive on an international scale. Only the best of the best receives allocations on those systems. One cannot carry out a long-term (5-10 year) research project with your primary (only?) resource on very thin ice from year to year. NASA is not given any special consideration; proposals go up against the top HPC groups in the world who are doing computing like this on a routine, everyday basis. More often than not, you will likely not win an award.

- These proposals are extremely time-consuming to formulate and prepare. Months of background prep work goes into them. Our researchers don't have that kind of white space in the conventional, what-have-you-done-for-me-lately projects.

- DOE systems available for such proposals are international, open-science resources. Cannot put ITAR or other classified data on them such as a lot of space-related work.

- Individual awards typically range from 500 million to 1 billion core-hours equivalent per year. As much as that is, that is intended for a single project. Not nearly enough compute available through DOE to meet NASA mission needs.

- Many NASA jobs do not fit the "shape" of the desired job topology at DOE. E.g., the 4D eddy-resolving CFD examples I mentioned towards the end of our AG discussion 2 weeks ago.

- DOE does not provide early-access to new architectures unless you are funded under their Exascale Computing Project. (And since they fund/steer much of the technology development, going to the vendors to get access is often met with the same reply – hence a monopoly that effectively keeps others several years lagged behind.)

P5: HEC does not provide the required rapid response to job turnaround that DoD does provide

As I'm sure many groups reported over the past couple of weeks, extended waits to run a job are a real issue. If I'm trying to get my code moved up from ~1,000 cores to ~10,000 cores, as a developer,
I need to run probably 10-20 very brief runs at 10,000 cores a day to evaluate modifications/improvements.
Waiting 10+ days to try a tweak to the code is a killer – I can't even remember what I was doing yesterday!
☺ And such waits are obviously not good on the application or 911 sides either.

P6: Goals are currently set intentionally low because current system can't handle any more

Lots of things NASA ***should*** be doing, but we cut back to what NASA
can do instead.

P8: Need programming models that permit code transfer to any platform; this is a tradeoff between portability and performance on a specific platform.

This is a tall order that the entire community is struggling with. Progress is being made, but lots of work yet to be done!

Need programming models that permit code transfer to any platform; this is a tradeoff between portability and performance on a specific platform.

Raw notes:

Slide 2: No Title

Have to be able to scale up. GPUs are a game-changing technology. This and other technologies are coming fast, and we have to be ready

Current HEC resources are not amenable to larger scales for 100,000 cores.

Developers need frequent fast-paced access --- typically have not had early access or received assistance and have had to look elsewhere for needs

HEC would need to require substantial growth in order to handle these needs

Comparison of turnaround times: 6 months to get Transport high-lift (Left-hand image). Right-hand image: Example of turnaround from DOD.

Slide 3: No Title

Hundreds of TB output per run- more than most will see in their entire career.

Programming models: propaganda out there about them but the key is portability versus performance portability

Will need trillions of degrees of freedom in 10 years.

Invested time to code on different architectures—caution that everything does not work as advertised.

Pre-selecting apps can be dangerous.

Slide 4: No Title

Concerned about long-term output for HEC at NASA

We have lost potential researchers due to not having the resources with HEC

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*

Q from Tsengdar (Int) to Everyone: time 2:04PM

is there a quantifiable statement what can be done with an exascale system?

A: that's problem dependent. high lift and aeronautics – hope is to cut 100 of millions of dollars out of aircraft development budget if we can do high lift successfully. Different exascale systems we can brainstorm through.

Q from Bill Kleb(Int) to Everyone: time 2:05PM
Nielsen: you mentioned need ready access to 100k cores for development; are there similar needs for applications?
A: I can't estimate that since I am not an app guy

Q from dan duffy (Int) to Everyone: time 2:05PM
What do you think about the variety of accelerator based systems, such as NVidia, Intel, or AMD? How would you recommend that we continue to write portable code while still getting as much performance out as possible?
A: open question for whole community. Programming models you can jump into and adopt. Remains to be seen how they may pan out on some of these architectures. Nvidia big one. AMD a thorn in Nvidia's side. We're working hard in our group to find the strategy that will help us find how to make portable code. Comes down to decorations on how you want to map that onto the hardware. Intel- you will need two code bases to get optimal performance on CPU and GPU.

5. Use Case: Supersonic Retropropulsion

Karl Edquist
karl.t.edquist@nasa.gov
LaRC D205
Descent Systems Study Project
Space Technology Mission Directorate, Game Changing Development Program

Main points:

P1: Long queue wait times are affecting project milestones

P2: Need to lean on the computational methods in order to make progress, and are running into issues due to resources.

P3: HEC is critical; this project only uses NAS resources. Critical path technology for Mars EDL for 20 metric ton payloads for landing human on Mars. . There is no current way to perform testing of this mission so need computational solutions. Currently can do 4-5 metric tons at most currently. Timing is important: landing in 2030 requires cargo landers with 20mt well in advance which means the models are needed in the 3-5 year timeframe to support that size payload.

P4: Have had to scale back fidelity to get solutions, in part due to time and resource limitations.

Raw notes:

Slide 2: Mission Need
Overarching goal- advancing the supersonic retro-propulsion via wind tunnel test and parallel CFD - ultimately for Mars

Image: depicting 5” diameter model where there is supersonic flow going from left to right and additional jets—unsteady, very complex flow field
HEC is critical to accomplish this- only using NASA resources for this fiscal year
Heavily depend on CFD models to build databases for Mars entry, descent, and landing.
Really long wait times- currently affecting project milestones for this year
Aiming to run many dozen cases for this year and somewhat behind- unsteady problems

Slide 3: Current HEC Application Characteristics
Problem: seeking to reach a position in the solution where you see the unsteadiness repeat itself so we can take averages of results to report- trying to get to the point to do a time average of an unsteady CFD solution
Part because of time and part because of limitations on computer resources, are having to back off on fidelity
Having to save more files than you would typically save due to uncertainty/ issues
Doing the solutions you can do with the resources you have and time you have rather than doing them how you would actually like to do solutions.
Queue wait times are multiple days; jobs are “outgunned” by HEOMD jobs running.

Slide 4: HEC Needs Assessment Workshop Key Questions
You never really get to a production mode for this problem the way things are at this TRL level right now.
Finding ourselves at step 1 where we are scaling back on our grids and codes to get some solutions on the table prior to this wind tunnel test happening--- critical path to Mars technology—no other way around building these aero databases--- this is something you cannot directly simulate in an earth based test facility --- very few facilities where we can test the way we want to, and they’re expensive, so we need to lean on the computational side and this will grow over the next 10 to 15 years as it’s needed for a flight mission.

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*
Q from Mike Little (Int) to Everyone: time 2:23PM
did you mean that you can't do 20 metric ton payloads with current capacity.
A: we can do a little over 1 metric ton with a rigid aero shell and a supersonic parachute. This is a technology that is replacing a parachute.

Q from Mike Little (Int) to Everyone: time 2:24PM
When does this capacity become necessary?
A. when we start brining pre-cursor hardware for humans. 2030 would be the actual mission. Cargo landers that go before the crew take much longer to get to mars. So we’re looking to have vehicles delivered to the cape by 2030s. anything bigger than the robotic landers, you’ll have to move past the chute and use retropropulsion.

Additional info from chat: Ashley Korzun (Int) to Everyone: 2:24 PM
~4-5t to the surface. At most. With today's modeling capacity

Q from Mike Little (Int) to Everyone: time 2:24PM
So, we need to be able to model the 20 metric ton model in 3 to 5 years in order to have flight ready hardware in 2030s?
A: yes- this is very integrated and affects design configurations ---all your phases of flight.

6. Use Case: Low Boom Flight Demonstrator Signature Validation

Melissa Carter
Melissa.B.Carter@nasa.gov
NASA LarC/D301
ARMD CST

Main points:

P1: No external sources to reach out to, so NAS is critical

P2: Need an overnight turnaround time

P3: Goal is validation: fly over communities at acceptable decibel levels. "Thump" not "boom

P4: NAS is critical to reach this goal: Without access to resources, many more flight tests will have to be done than planned

Raw notes:

Slide 2: Mission Need
There are 3 phases. We focus on phase 2- validate the boom signature from aircraft.
The goal is to hopefully have a low signature (thump thump) on the ground during phase 2, CFD is responsible for verifying the different decibel levels experienced on the ground.
NAS is critical because no outside source to go to- if we don't have access, looking at flying the aircraft many more times
Run time is about 12 to 13 hours, so can't wait in queues
Have to do overnight turnaround
CFD is near field solution
Quicker turn around, more answers we get, quicker we can do our work

Slide 3: Current HEC Application Characteristics
Once we get data, we can submit them for overnight running, and get them in the morning for signatures
May be some visualization needs.

Slide 4: HEC Needs Assessment Workshop Key Questions
Key factors: Not until in flight will we ask for high priority, and really need high priority

Current developmental work: might also be delayed due to COVID
Wait times identified as a risk.

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*
Q from Mike Little (Int) to Everyone: time 2:56PM
what would be the max acceptable queue wait time?
A: I like to see a solution within 24 hours but that's not a full solution. For some of the codes, as long as running with 24 hours, CFD'ers are happy. For the development part, won't be as detrimental, but when we get to the flight, I don't know if we will be able to wait more than 2 or 3 hours depending on how everything goes.

Q from Tsengdar Lee (Int) to Everyone: 2:57 PM
How often are the tests? If we set up a near real time queue for those test periods, would it be acceptable?
A: yes, probably going to need this time of queue. true goal is overnight. grid is made while flying. then processing in the morning to tell them what to do next.
Do not know the actual flight apps. I know they want to get multiple flights per day but not sure how many days in a row. Phase 2 flights go approximately 9 months- I don't have the schedule for that yet.

Q from Robert Ferraro (ext) to Everyone: time 2:58PM
Can you quantify the cost impact of more than 12 hour turnaround? During flight experiments.
A: not until we know how much these flights will cost us once we get going. Kind of a new world for me. Armstrong is used to flight tests, but I'm not familiar with how much exactly it will cost us.

Additional info from Piyush Mehrotra (Int) to Everyone: 3:02 PM
Melissa, we should be able to work these issues with support from HQ

6. Use Case: SLS Buffet/ Aeroacoustics
Craig Streett
craig.l.streett@nasa.gov
NASA LarC-D302/ NESC
SLS Program/ NESC

Main points:
P1: NAS is only viable resource; requirements too large to obtain commercially

P2: Reaching the limits with what's possible with NAS right now. Cannot use commercial computers or DoE or NSF due to ITAR. Work is planned months in advance, but does not respond to changing needs of customer community. Need 100-1000 times more capacity to be able to respond to current needs. Large core jobs need reservations; one job has been waiting 23 days in the queue and still not run. A full campaign would require selection of a limited number of cases.

P3: This work is a bit different; defining working on unsteadiness in flow fields.
The scale is quite broad and dependent on runtime and mesh size. It's a niche application because of the size of resources required.

P4: CFD simulations are needed to obtain some solutions that can't be found in wind tunnel tests. If NAS was not available, Wind tunnel tests would probably suboptimize design and increase risk.

P5: Need far more cores; need help from developers to mitigate overhead and data collection and storage.

Raw notes:

Slide 2: Mission Need

We're resolving unsteady fluid flow at 2 orders of magnitude or more.
Example: SLS Block-1 in forward skirt region, with electronic boxes mounted on equipment shelves—simulations being able to identify space-time with coherent and identify better equipment locations; the findings were not possible with wind-tunnel tests—only possible via simulation.
The Program "dodged a bullet" as a direct result of the equipment relocation.
NAS is the *only* viable resources. If these resources were not available, the program would not pursue such new techniques at all. The Program would fall back on old ways, leading to risk.
We could pursue DOE systems, but launch-vehicle technology is ITAR, so data security is an issue.
Cannot do space-time coherence in wind tunnel; only by simulation (Craig – please provide a reference that describes space-time coherence)
This capability may be used by commercial launch vehicle design teams.

Slide 3: Current HEC Application Characteristics

Reaching limits of what's possible at NAS in a practical sense.
The mesh generation for these types of runs is fundamentally different than simulation where you're only interested in time average flow field.
Once a set of simulations for a configuration is obtained, after analysis, sometimes have to re-mesh several times to get the spectrum resolved.
Need help from developers in order to mitigate overhead of data collection- have to be able to store surface data pressure every time step.
Typically, a batch runs take 20-30 days of wall-time, which doesn't include queue time.
Campaigns are planned over 6 months.

Problem: when the SLS program initiates the loads analysis cycle, each discipline requests asset reservations--- leading to only the way to get significant numbers of cores is through asset reservations--- leads to things sitting in cues, not being run
Capacity turn around is a severe constraint- if resources aren't available, program will have to be moved off of this critical path.

Slide 4: HEC Needs Assessment Workshop Key Questions
Reservations are not sustainable, but the only way to get lots of cores is through reservations.
Currently have jobs sitting in queue for 23 days –they probably won't run.
Maturity: Need multiple successful applications.

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*
Q from Bill (Int) to Everyone: time 3:21PM
Is there an optimal job size?
A: Currently, the practical size of jobs would be about 25,000 cores for 20 days wall clock time. That's the kind of thing that would make me happy in the near term. Optimal, I'd like to see 100 fold increase in computational volume in the near term. Would allow us to be sure of two orders of magnitude frequency range, accurate simulations…? Important benchmark.

Q from William Thigpen (Int) to Everyone: 3:23 PM
Do you have a preference CPU vs GPU?
A: The work Eric and his team have been doing shows that his code is very fast on the GPUs as long as the data collection issue can be mitigated. I'm sure with a short period of work, that team of experts could overcome. It doesn't matter to me at all, it's purely a matter of availability of the resources capable of handling the problem we need.

6. Use Case: Project: Red Rover Connecting Experimental and Computational Facilities and Communities
Nettie Roozeboom
nettie.roozeboom@nasa.gov
NASA/AOX
ARMD

Main points: Using AI to steer windtunnel experiments in realtime. This is useful in unsteady flow experiments to validate CFD with experimental results.

P1: Need to make sharing data between the wind tunnel and the NAS more efficient

P2: Have achieved success with diverse groups working together.

P3: We're connecting experimental and computational facilities; but not just facilities – experts and core competencies. Working with NAS was a critical partnership. Visualization support and the hyperwall were critical.

P4: Financial and socio-technical hurdles to be addressed– who has access to data while test is running

P5: TBD.

Raw notes:

Slide 2: Mission Need

Highlighted "share" because current situation we're living in, really highlights some efficiencies that could be gained here. What if we can't send our full analysis team? Can we work on this remotely? How can we share data? A lot of technical and financial hurdles that could be addressed by not having to send a team of folks to the wind tunnel.

All about getting the data in the right environment--- not in the right environment if keeping it at the wind tunnel.

If Supercomputing was no longer available, it would be a money issue and require a lot more time.

We are generating a lot of data, working with NAS folks, how do we get it from the wind tunnel to the NAS securely, but not create a back-up? Brings some risks.

Whole goal is to deliver a better answer. What's the difference between a decision and a bad decision? It's time.

Demonstration creates conversation.

Being able to see things up on the hyperwall allows discussion

Most recent test: generated lots of data; accomplished secure transfer from the wind tunnel to NAS.

Mission concept: UPSP was attempted for 10 years. Accomplished in 10 months.

Slide 3: Current HEC Application Characteristics

Workflow: A lot behind sending the data to NAS

The real time computing while that is important is fairly concentrated.

Transfer: ah hah moment--- there are improvements that can be made- security is very important, don't just download data and send it in an e-mail. Currently, downloaded to encrypted external drives, deliver it to the customer- now in their possession. What's the NASA solution? Procuring enough drives to store data and have a backup set became not feasible.

Slide 4: Needs Assessment Workshop Key Questions

Keep customer in mind so they don't have to pull data to their local machines to look at.

Q&A: * answers were given very quickly, so notes below are to be viewed as bits and pieces of the given statement, and can hopefully be better pieced together by the presenter*
Q from KAREN PETRASKA (Int) to Everyone: 3:34 PM
IF HEC was not available would it be possible to use the commercial cloud computing?
A: my knowledge about the cloud computing is somewhat limited, it seems like it would be. But I don't know enough about it. It needs to be fast, efficient, powerful, and secure. I do believe the security has been addressed with the cloud computing but I don't know about the tools or speed. So I would pause to have a definite answer there.

Q from Bob Ciotti (Int) to Everyone: time 3:49PM
camera count is doubling in the next test which grows the data volumes
A: yes, four now and eight cameras next. UPSP - fighting for as much light as possible. signal to noise problem - drives up the signal. 100GB per camera times 8

Additional info in chat from Bob Ciotti (Int) to Everyone: 3:51 PM
it is. we are working on that. it will double the data rate requirements.

Group Discussion of Key Questions
Unable to catch everything from the discussion, other than the chat comments. commentators should please edit their statements, as they will be a bit jumbled

Speaker: Eric Nielson
Let me summarize one of the challenges across different app spaces. The need for scale resolving CFD. Quite often with these apps you're interested in very long time durations to get statistics on unsteady dynamics. It becomes a 40 problem. Last year proposal put into DOE-requested 2 billion CPU hours (1 million/ year xeon core hours)- came back rejected because too small of a request. They give you a guidance on how many hours you should propose. Then figure out on that 4-dimension problem, what you can fit into that time. Need to be able to strong scale run a fairly large special problem, but when you do that the fourth dimension kills you. Can't afford to run those time dimensions. Finding over and over again we need our own machine to be able to perform this work. Going to other agencies isn't an option… scale resolving CFD, this is looming as a challenge that we need to crack internally.

Speaker: Piyush
See a lot of folks asking for huge resources, and that's fantastic. I'd like to point out that DoE has 10x budget more than we have. This Needs Assessment Workshop helps us by focusing on not just hardware we need but picking out codes—finding out where the gaps are, where the shortfalls are.

Speaker: Tsengdar Lee
What I find that what's needed is to communicate with decision makers. We set up something, we need an exascale computing system, but we also need in order to solve a problem, we need

to explain why we cannot solve it before, and what his new system, new capability can help us to solve. That's the kind of thing that will catch their eye. I've been asked to get that kind of information. That's why I asked what kind of statement, and quantifiable, and previously without that we have to prioritize the turbulence, and that may give us an arrow bar of 20%. And in order to narrow that, in order to do preciously edl, risk reductions, such and such, we need and exascale computing system. That's the kind of general comment that probably apply to many of the projects here.

Chat discussion:
from ERIC NIELSEN (Int) to Everyone: 4:01 PM
completely understand Piyush - please let us know how we can help you make that case going forward

from William Kleb (Int) to Everyone: 4:01 PM
NASA used to have one of the top supercomputers, how did we come to de-value supercomputing.

from William Kleb (Int) to Everyone: 4:04 PM
Are the DoE and DoD, percentage-wise of total budget, spending a comparable amount to NASA; or are they valuing supercomputing more than us?

from Robert Ciotti (Int) to Everyone: 4:04 PM
maybe people can also estimate the percentage of their work time is lost waiting for jobs to turnaround. workforce inefficiency metric?

from Robert Ciotti (Int) to Everyone: 4:05 PM
cost of programatic slips...

from Robert Ferraro (Ext) to Everyone: 4:06 PM
how will we do the human mars mission if we don't have the HEC resources in house?

from Piyush Mehrotra (Int) to Everyone: 4:06 PM
Eric: Like Tsengdar just said, specific examples that can be done with larger systems would help. Specific examples that the managers at HQ can relate to.
systems would help. Specific examples that the managers at HQ can relate to.

from William Thigpen (Int) to Everyone: 4:07 PM
DOE Has twice the budget and spends more than 10 times on HEC

Speaker: Mike Little
Plan for this is to use this to justify increasing the HEC resources/ support programs- whole point for producing this report is look at what we need in the future, how do we increase the

budget and what does that buy the Agency in terms of mission accomplishment, risk reduction, design iterations.

Appendix D3 Workshop Session 6-3 Inputs

David Giles	GSFC 618	AERONET NRT and Historical Inversion Processing
Phil Partain and Heather Cronk	CIRA/Colorado State University	GeoCarb Data Operations Center
Jennifer Dungan, Andrew Michaelis, Ramakrishna Nemani	ARC SGE	NASA Earth Exchange
Peter Tenenbaum, Jon Jenkins	ARC SGE ARC TN	TESS Science Processing Operations Center
Paul von Allmen and Trinh Vo	JPL 398	First-Principles Modeling of Thermoelectric Materials
Thomas Vandal	ARC SGE (BAERI)	Deep Learning for Geostationary Data
Bob Ciotti	ARC TNC	HECC Project NASA Advanced Supercomputing

National Aeronautics and Space Administration

Workshop Session 6-3

3 June 2020

National Aeronautics and Space Administration

Use Case: AERONET NRT and Historical Inversion Processing

David Giles
June 2, 2020
GSFC/618
Science Mission Directorate/AERONET]

Use Case: Mission Need

- Mission: Perform inversions of ground-based remote sensing measurements to obtain atmospheric aerosol characteristics
  - The HEC impact can be measured based on how long a data set can be processed; millions of retrievals can be performed in a fraction of the time than computed on the AERONET server.
- Criticality: HEC is critical to providing near real-time and historical reprocessing of this data set in a timely manner.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Additional hardware would need to be purchased to provide the processing power and storage to compute these retrievals.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - No impact on inadvertent data release.
  - HEC availability interruption would significantly delay historical reprocessing of retrievals
  - No current mission need for external collaborators.
- HEC Function: HEC provides thousands of processing cores to run serial job requests.
  - Near real-time (NRT) data requests run between 1 and 3 times a day, while historical reprocessing may run quarterly for evaluating scientific questions.
  - NRT retrievals of aerosol characteristics are made available on the AERONET web site and historical reprocessing advances understanding of the capabilities and limits of the retrieval code.
- Impact of HEC: HEC is essential to completing this work since it would take several months to complete historical reprocessing on the local system.
  - Additional hardware would need to be purchased to provide the processing power and storage to compute these retrievals.
- Mission Concept: See supplementary slide 1

Use Case: Current HEC Application Characteristics

- Work flow:
  - The AERONET system generates SLURM jobs batches and input data to run on HEC system.
  - The scripts are automated for near-real time data processing but user interaction is partly needed for historical reprocessing of data to monitor job completion.
  - This application is used in NRT and can have an issue when processing or disk resources are not available.
- Data:
  - In: Measurement data are received at least daily from the local AERONET system in NRT between 1MB and 50MB; historical data submission can be quarterly and up to 50GB
  - Out: Aerosol characteristics are provided for each hourly day time measurement input from hundreds of sites worldwide; data are analyzed on the AERONET system using local visualization tools and AERONET web site
  - Transfer:
    - Data are transferred to the AERONET system (aeronet.gsfc.nasa.gov) for retention and dissemination
- Computation:
  - Type: Multiprocessor/core environment is needed to run many serial tasks
    - The capacity is available now but limited to maximum number of processes that can be run by one user
- Evaluation: The number of processed retrievals is the success of the HEC resources

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Science data reprocessing – data are often reprocessed to answer scientific questions, associated with new product versions, or publications
  - Routine measurements and field campaigns demand near real time processing
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Increased processor speed with reduce processing time; an increase in the availability of processing nodes for serial job requests will reduce job latency.
  - Impact: The measurement database is expected to grow so the need to perform faster computations will grow.
  - External Sources: Retrieval processing does not run at a non-NASA processing facility
- How will this Use Case evolve over the next decade?
  - Maturity: Operational capability with continued development
  - Growth: New measurement and ancillary data base inputs will require continue historical reprocessing
  - Re-use: The number of staff running code will not change significantly.
  - Re-hosting: Code migration requires necessary compilers (C and FORTRAN)

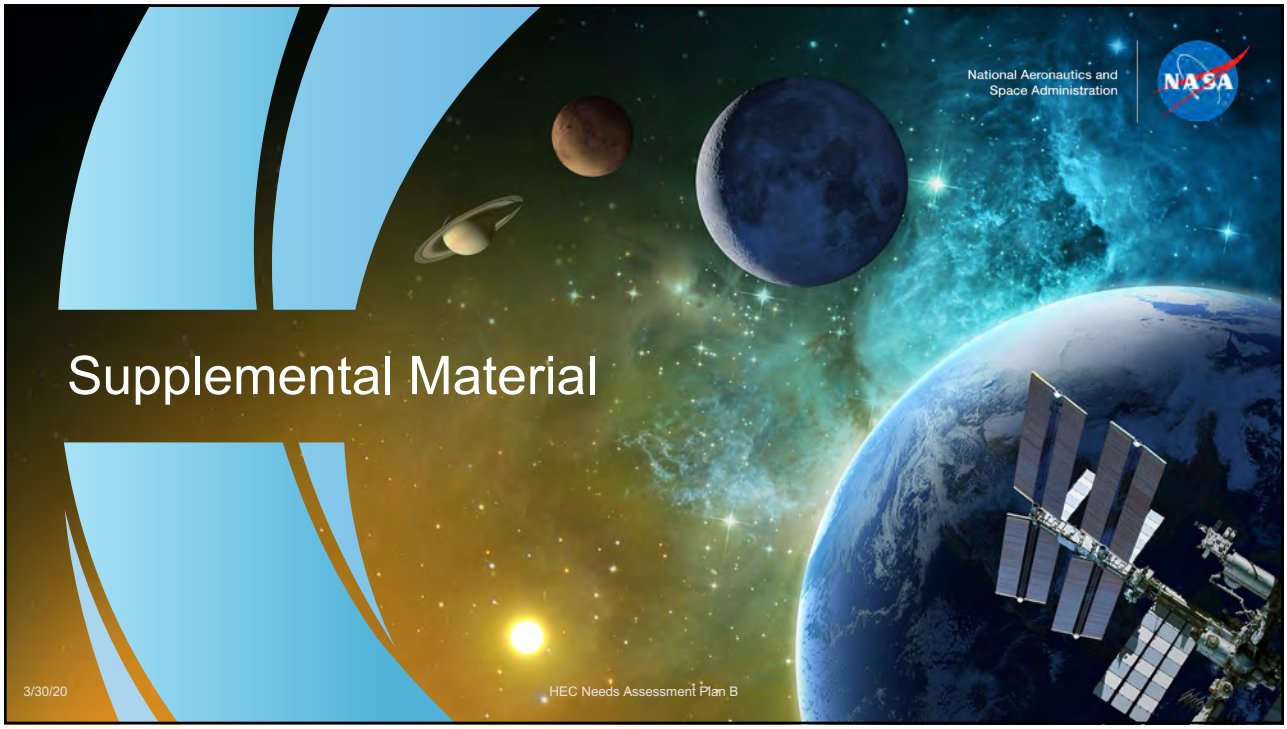

Mission Concept

HEC facilitated a recent publication on the Version 3 AERONET aerosol retrievals and uncertainty estimates

Sinyuk, A., Holben, B. N., Eck, T. F., Giles, D. M., Slutsker, I., Korkin, S., Schafer, J. S., Smirnov, A., Sorokin, M., and Lyapustin, A.: The AERONET Version 3 aerosol retrieval algorithm, associated uncertainties and comparisons to Version 2, Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2019-474, in review, 2020.

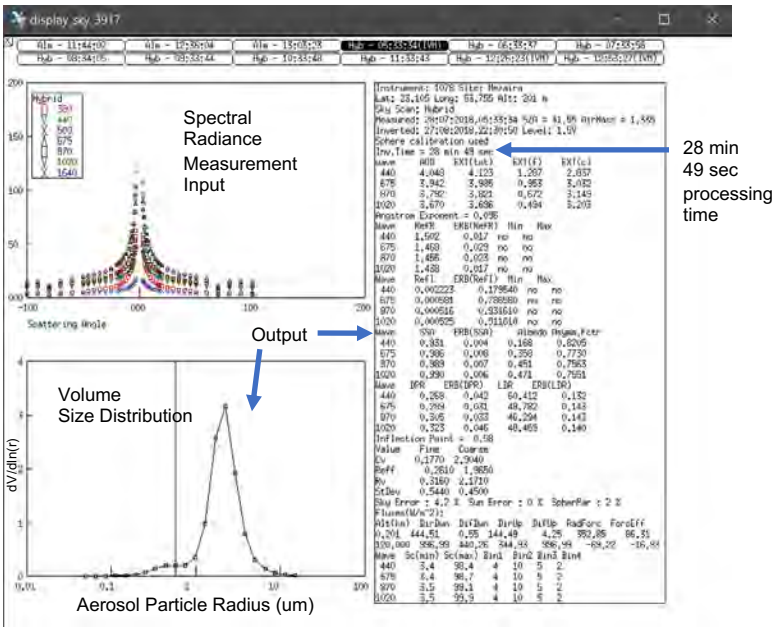

Information about the Survey

- The following slides describe the kind of data that the Needs Assessment will collect via a survey
- This material is needed to supplement the description of the Use Case presented to the Affinity Group but will not be presented.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: performs inverse computations of sky radiance and aerosol optical depth measurements to obtain atmospheric aerosol characteristics
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones: Runs on AERONET system
    - Programming languages used (95% FORTRAN, 5% C.)
    - Number of lines of code: ~45,000
    - Workflow mechanism used to control code execution (shell script, SLURM,PoDS)
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): 12 nodes with 28 per core
    - Accelerator usage (type and memory requirement):
    - Amount of memory required (overall or per process): 1.5GB per process
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other: No, parallelism, runs serial
    - Storage is covered below
    - Run in batch, interactive, or near-real time: Batch mode

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output: Read input file->Process Data->Output file
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other
    - Load input file and read files line by line
  - How much input and output are generated per run
    - No input is generated during run; one output file is generated per input
  - Summary of the total amount of input and output data per year: 150GB output per year (estimated)
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? Visualization, statistics
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other: Needs fast read/write for tar and gzip (sometimes batches do not finish due to slow disk access)
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used. Linked to prepared SLURM batch jobs
  - How much data needs to be curated (archived)? Only temporarily on HEC system
  - How much data needs to be distributed and, generally, to whom? All of the output is distributed to AERONET
- What are the key computing obstacles to effective execution of the current application? Disk storage to run historical processing and number of processing nodes per user
- Are there any analytic functions performed on the output of the application? Timing is computed
- Do you have available the workforce necessary to re-host your code on any new platform necessary? Yes

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals? Availability of processing time and disk storage
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.) Yes
    - Increase in resolution and/or timesteps: Not likely
    - Rewrite to improve performance: Yes
    - Refactor for emerging platforms: Not expected
    - Rehost onto cloud: Possible, but no currently planned
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components: Possible but not currently planned
  - If possible, describe how the computational requirements will evolve over time: Probably the same as above but depends on improve performance versus new components
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements
    - Run in batch, interactive, or near-real time

Use Case: Mission Need

- Mission:
  - GeoCarb is a NASA Earth Ventures Mission that requires significant data processing resources to accomplish its mission requirements
  - Mission success depends on meeting data processing requirements
- Criticality/Impact of HEC:
  - Without HEC resources, the GeoCarb mission would incur significant cost by purchasing a large hardware cluster to be located at Colorado State University and supplementing with commercial cloud computing resources
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Inadvertent data release could cause public and science community confusion
  - Due to the large processing workload (8200 SBU/day max), frequent interruptions could result in the inability to meet product delivery requirements
  - The workflow operators are Colorado State University employees
- HEC Function:
  - The L2 Full Physics products will be run daily on Pleiades as input arrives
  - The output contains CO2, CO, and CH4 gas concentrations, the retrieval of which are key mission objectives

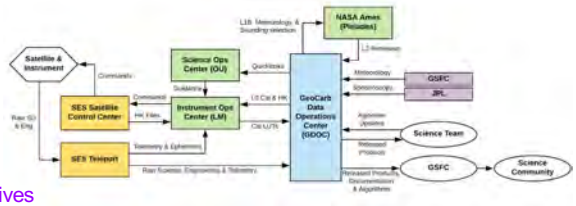

Mission Concept

Use Case: Current HEC Application Characteristics

- Workflow:
  - The workflow is triggered by existence of L1B input data
  - Ideally, no people in the operational workflow
    - Needs: scripts/dashboards to monitor current processing status, identify job success or failure, and notify CSU operators in the event of failures
  - 6% of the data will be processed in near real time to support identification of instrument issues. 100% of the data will be processed with a ~1 week lag and the data delivery latency requirement is 60 days.
- Data:
  - In: A rolling 2.5 TB comprised of three dynamic products and a set of static data. This estimate includes a 30 day buffer for data not affected by retrospective processing and a two week buffer of other inputs in case an issue is found. All data will be coming from CSU.
  - Out: The main algorithm produces per-sounding retrieval files which must be aggregated into a granule-level file. Further post-processing, packaging, and analysis will happen at CSU
  - Transfer: Data is written to HEC storage, aggregated into granule-level files, and transferred to CSU
- Computation: Distributed, independent processing
- Evaluation: HEC is a partner in meeting our mission requirements

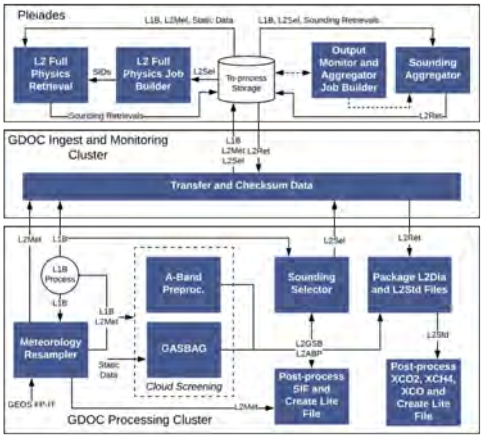

GeoCarb Level 2 Processing Workflow

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Large volume of observations (up to 1 million per day)
  - Data pipeline bottlenecks (downlink issues, lower-level product issues) causing the need to catch up
  - Annual science data reprocessing campaigns
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - The current capacity is not enough to fully support science data reprocessing within reasonable timeframe for OCO-2, which is a very similar use case. GeoCarb will have considerably more data to process during reprocessing.
  - Impact:
    Without adequate resources to support reprocessing, GeoCarb would likely need to use public cloud resources to supplement, which would present a budget issue for the cost-capped mission.
  - External Sources:
    The majority of the GeoCarb data pipeline components will run at Colorado State University, with the most computationally expensive algorithm running on HEC resources (Pleiades).
- How will this Use Case evolve over the next decade?
  - Maturity: GeoCarb is currently in pre-launch development and will move to testing this year and operations nominally in early 2023
  - Growth: None (beyond increased data for reprocessing during each campaign)
  - Re-use: None (the number of people running this application will not grow)
  - Re-hosting: Ideally no need to refactor code for future architectures if we are running in containers

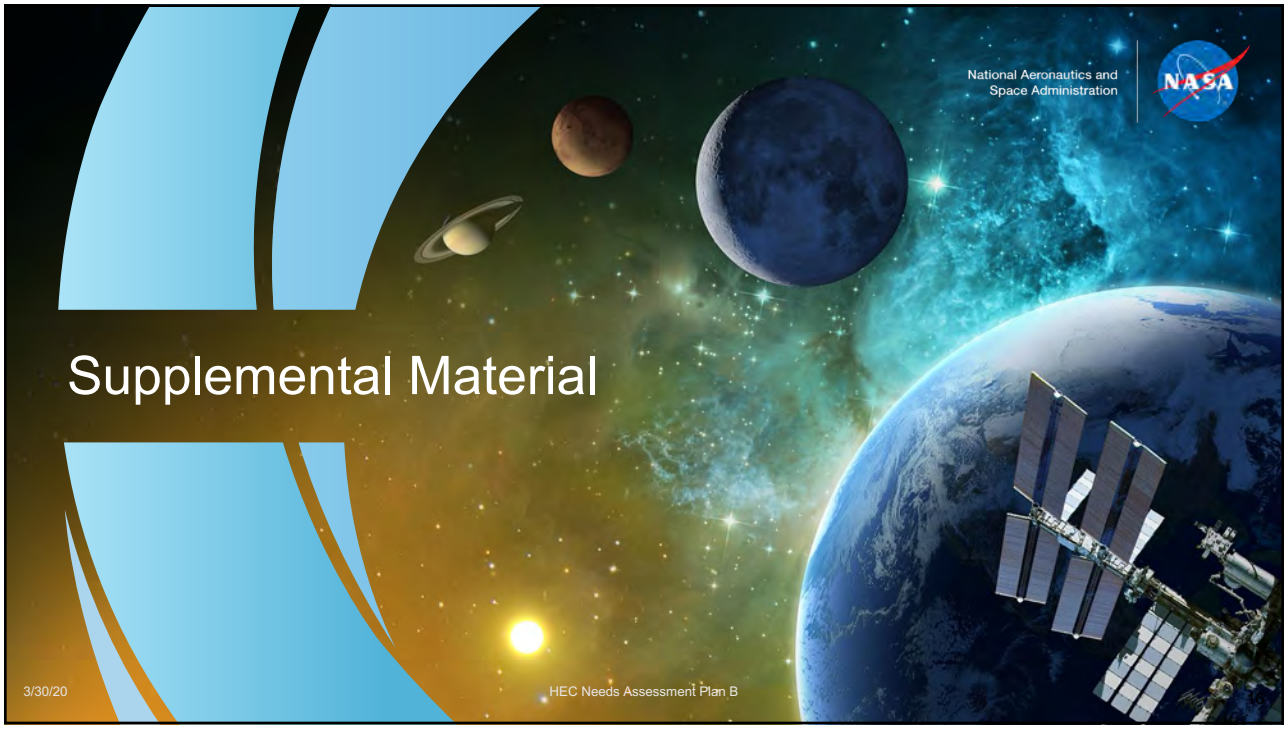

Information about the Survey

- The following slides describe the kind of data that the Needs Assessment will collect via a survey
- This material is needed to supplement the description of the Use Case presented to the Affinity Group but will not be presented.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones
    - Programming languages used (If multiple languages used, provide estimate percentage of the code in each languages.)
    - Number of lines of code
    - Workflow mechanism used to control code execution (perl, jupyter, shell script, etc.)
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage is covered below
    - Run in batch, interactive, or near-real time

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other)
  - How much input and output are generated per run
  - Summary of the total amount of input and output data per year
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena?
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used.
  - How much data needs to be curated (archived)?
  - How much data needs to be distributed and, generally, to whom?
- What are the key computing obstacles to effective execution of the current application?
- Are there any analytic functions performed on the output of the application?
  - Timing, computing load, workflow
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.)
    - Increase in resolution and/or timesteps
    - Rewrite to improve performance
    - Refactor for emerging platforms
    - Rehost onto cloud
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements
    - Run in batch, interactive, or near-real time

Use Case: Mission Need

- Mission: Enable research and applications with regional, continental and global Earth observations within NASA's Earth Science Division strategic directions
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Inadvertent release of immature results could lead to public confusion and/or professional embarrassment
  - HEC availability interruptions delay the achievement of objectives.
  - The Earth Science Division model is based largely on "external" collaborators (without NASA credentials).
- HEC Function
  - Supports and maintains hardware, software and governs job scheduling, provides user support for system-relevant activities (e.g. login, installing software, develops and maintains documentation, provides runtime troubleshooting)
  - Codes run by NEX users are in a variety of languages and often get revised and re-run
  - The output of product generation jobs and analyses support mission objectives by being archived at a DAAC and by being published in the conference and peer-reviewed literature
- Impact of HEC:
  - If NASA supercomputing were no longer available to meet NEX needs, each supported scientific team would be "on their own" to decide how to perform their work, including buying systems and using commercial, university or NSF/DOE cloud resources
  - . Complex costing of cloud services, scoping & procuring individual compute clusters with large storage capabilities, obtaining an ATO - HECC significantly reduces project level budgets and overall efforts in aggregate for NASA programs, enabling more science.
- Mission Concept: Using NASA's private cloud computing to achieve what is otherwise not feasible

3/30/20 22

NEX as a Data Analytic Center

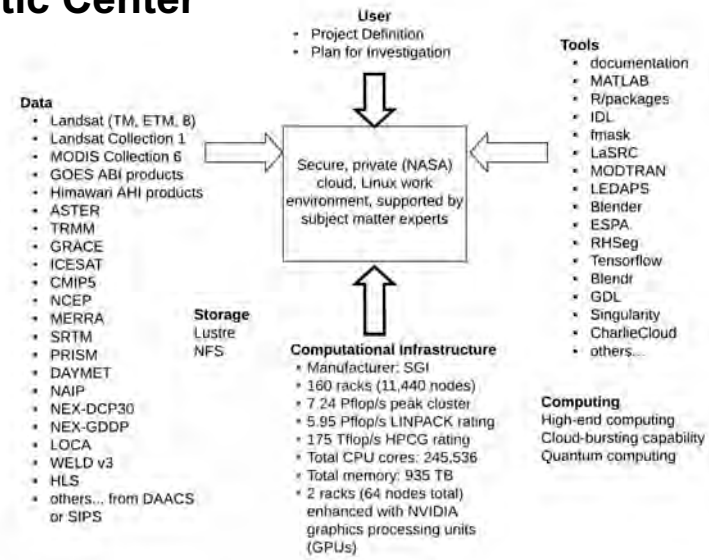

Use Case: A major current thrust of NEX is to support geostationary products (GeoNEX)

- Work flow:
  - As new data becomes available from NOAA, JMA, and KMA, it is acquired and set of workflows are triggered
  - Workflow(s) filter, transform, and package consumable, unified geostationary data products
  - Results are placed on the Lustre file system(s) for NEX users
  - If newly consumable product(s) are "approved", post or expose product(s) on the NAS data portal
  - Goal is to have limited to no technician/user interaction with any of the workflow(s) producing products or posting results
- Data:
  - In: GOES 16 and 17, Himawari Level 1a, KMA level 1a data inputs (roughly 30 TB/year per platform, full disk)
  - Out: Both input and output data are retained; input for ease-of-access and reproducibility, output for use by researchers (roughly 80 TB/year per platform)
  - Transfer: Generated products are stored locally and or exported for public use on the NAS Data Portal
- Computation/Storage:
  - Currently, CPUs are used by recruiting thousands of cores in an embarrassingly parallel fashion (process per granule or granule set).The special HEC capabilities that are critical for this specific use case are high performance file systems and with fast interconnects, networking (in and out of NAS) and storage management.
  - Roughly 1 hour core per GB input (Broadwell architecture) = 1 SBU per GB
- Evaluation: The effectiveness of HEC in supporting this work is currently judged qualitatively by feedback from NEX users/technicians.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - For NEX use cases, the most important factor is the readiness of users to utilize the computing assets. This is affected by their familiarity with HEC computing environment/Linux/PBS/MPI but more importantly time in the academic year (PIs, postdocs and other users have competing priorities)

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Storage is our main limit, user education and limited messaging services are also constraints
  - Impact: Storage may limit the ability to achieve global geostationary product generation or other big-data challenges; limited services (messaging) impose more burden on technicians and developers for semi-resilient processing pipelines
  - External Sources NEX team members have had experience running processes on AWS and Google Earth Engine, which provide many turnkey services but costs are hard to project

- How will this Use Case evolve over the next decade?
  - Maturity: GeoNEX product generation has potential to move from a development state to an operational capability with continued development
  - Growth: GeoNEX products will increase both spatial and temporal resolution to support evolving future mission needs
  - Re-use: The number of people running GeoNEX will likely not grow; the number of NEX users may grow
  - Re-hosting: Most code by NEX users would need to be released or open-sourced and then possibly refactored if needed

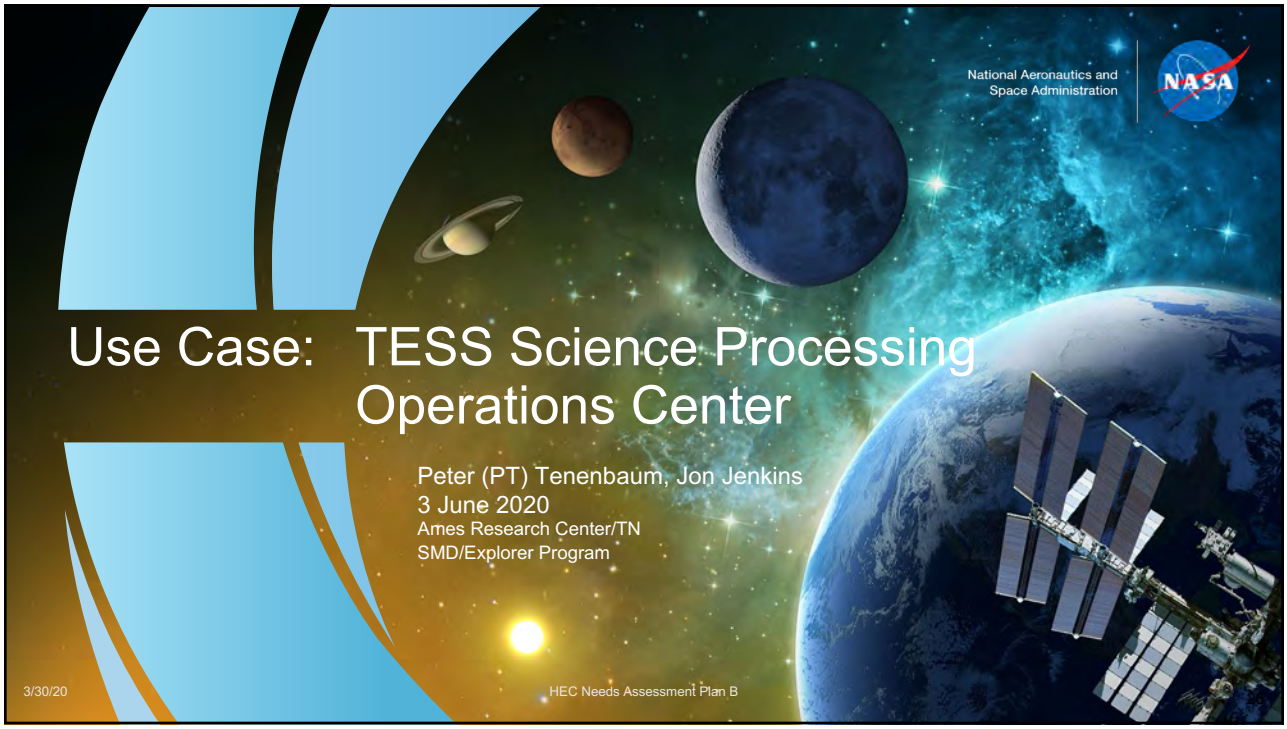

Use Case: Mission Need

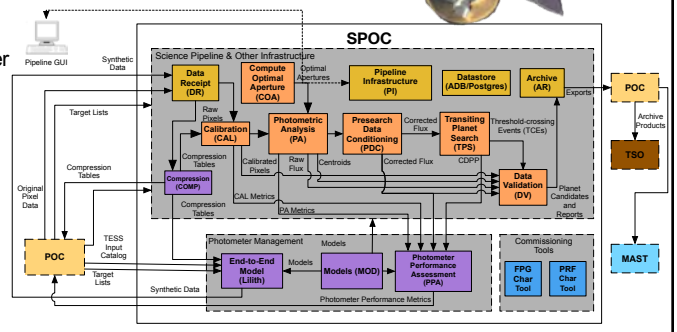

- Mission: TESS Mission's Science Processing Operations Center (SPOC)
- Criticality: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - We know of no viable alternative to utilizing the HEC for processing the science data from TESS, which is a cost-capped mission. Our concept of operations was predicated on access to the NAS Pleiades Supercomputer. All other options would have incurred significant costs and schedule delays
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Inadvertent release of data is low risk as our data products are hosted by the Mikulski Archive for Space Telescopes (MAST)
  - If HEC availability is interrupted regularly, we would not be able to meet our data processing and export schedule
  - What is mission need for external collaborators (without NASA credentials) to access system or to run jobs? We have two pipeline scientists at MIT who require remote access to our task files to conduct their work to analyze the data quality
- HEC Function: We run jobs on Pleiades nearly continuously as TESS downlinks ~1 TB every two weeks
  - We provide all the official archival data products based on the image data collected by TESS
- Impact of HEC: The HEC is indispensable to achieve our objectives and have a successful mission

Use Case: Current HEC Application Characteristics

- Work flow:
  - Processing begins with notification of data downlink from spacecraft
  - Pipeline operator pulls over data and metadata, imports into datastore, and starts pipeline
  - Need to respond within 1-2 days of data availability to meet throughput requirements
- Data:
  - In: Approx. 1 TB / month of raw pixel data, plus descriptions of SC pointing, target star lists, etc.
  - Out: Approx. 1 TB / month FITS files delivered to MAST for use by science community
    - To date, all raw data, data products, and intermediates ("task files") preserved since start of mission
      - Do not expect to be able to do this in the extended mission (begins in July)
  - Transfer:
    - We store the results in our HP 3PAR data storage system, then export the archival files to the MAST
      - Bandwidth to both destinations critical
- Computation:
  - We divide data into individual targets (stars), run multiple processing steps with 1 core / target
    - No communication between cores/nodes
- Evaluation: Without the NAS we could never have proposed this project much less executed it

Original Flux

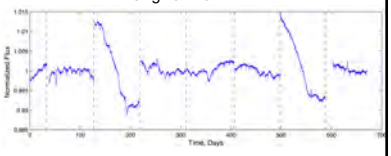

Systematic Error-Corrected Flux

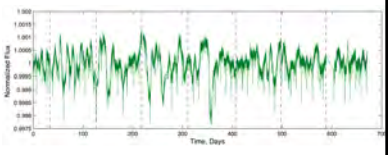

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Data drops are 2 per month (so idle periods and then super-busy)
  - Data volume determined by number of targets to process
    - Recently added a capability that boosted # of targets roughly 8x
  - Computations consume ~113,500 SBU2s per year
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Nothing major, occasional wait periods during heavy Pleiades use or down times
  - Impact: We do not anticipate any constraints from NAS in the extended mission, though data volume will increase.
  - External Sources: All processing is performed within the NAS enclave.

- How will this Use Case evolve over the next decade?
  - Maturity: TESS SPOC is mature (currently performing V & V on release 5.0)
  - Growth: In extended mission, will add an additional data type and increase # of images/month by a factor of 3
  - Re-use: Pipeline infrastructure will be used by additional missions
  - Re-hosting: N/A

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: Data analysis for transiting planet detection / photometry instrument in Earth orbit
  - Describe the application
    - Antecedents: Kepler Science Operations Center processing, also at NAS
    - Programming languages used:
      - Major: Java, MATLAB, C++
      - Minor: Python, BASH scripts
    - Number of lines of code approx. 500,000
    - Workflow mechanism used to control code execution: Spiffy pipeline infrastructure application
  - List the functional computational and storage requirements for the application
    - Size of runs 2000 cores typical
    - Accelerator usage None
    - Amount of memory required 4-12 GB per core
    - Describe the parallelism of the application: No communication between cores, each core performs an independent analysis of 1 target object (usually a star or CCD)
    - Storage is covered below
    - Batch jobs plus some interactive (Pipeline operator starts and monitors runs, looks for smoke)

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output: Data is marshaled from a datastore, provided in binary format on a working disk array, algorithm places results in same format on same array, they are persisted to the datastore
  - Each core loads its own data from the disk and stores its own results to disk
  - How much input and output are generated per run? Over 1 TB / run (1 run per month)
  - Summary of the total amount of input and output data per year: Order 15-20 TB / year inputs and outputs
  - Results go to MAST and are analyzed by an army of researchers all over the world
  - RDBMS and datastore (netCDF) on dedicated, high-performance (sort of) disk array from 3Par (now HPE); we also use nobackupp15 for binary files used and produced by algorithm (algorithms have no direct access to database or datastore)
  - How much data needs to be curated (archived)? All results archived at MAST
  - How much data needs to be distributed and, generally, to whom? All results archived at MAST
- What are the key computing obstacles to effective execution of the current application? netCDF and HDF5 do not support concurrency so marshaling and persisting are inefficient.
- Are there any analytic functions performed on the output of the application? No.
- Do you have available the workforce necessary to re-host your code on any new platform necessary? No.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - Extended mission: 3rd data type, plus increase in volume of existing data types
  - What are the key computing obstacles that could effect these goals? Lack of concurrency support in netCDF and HDF5
  - Describe how the application and workflow is expected to evolve over that time frame
    - Data volume will increase by factor of approx. 2 in the extended mission
  - If possible, describe how the computational requirements will evolve over time
    - Aside from higher data volumes, we expect no changes to computational requirements

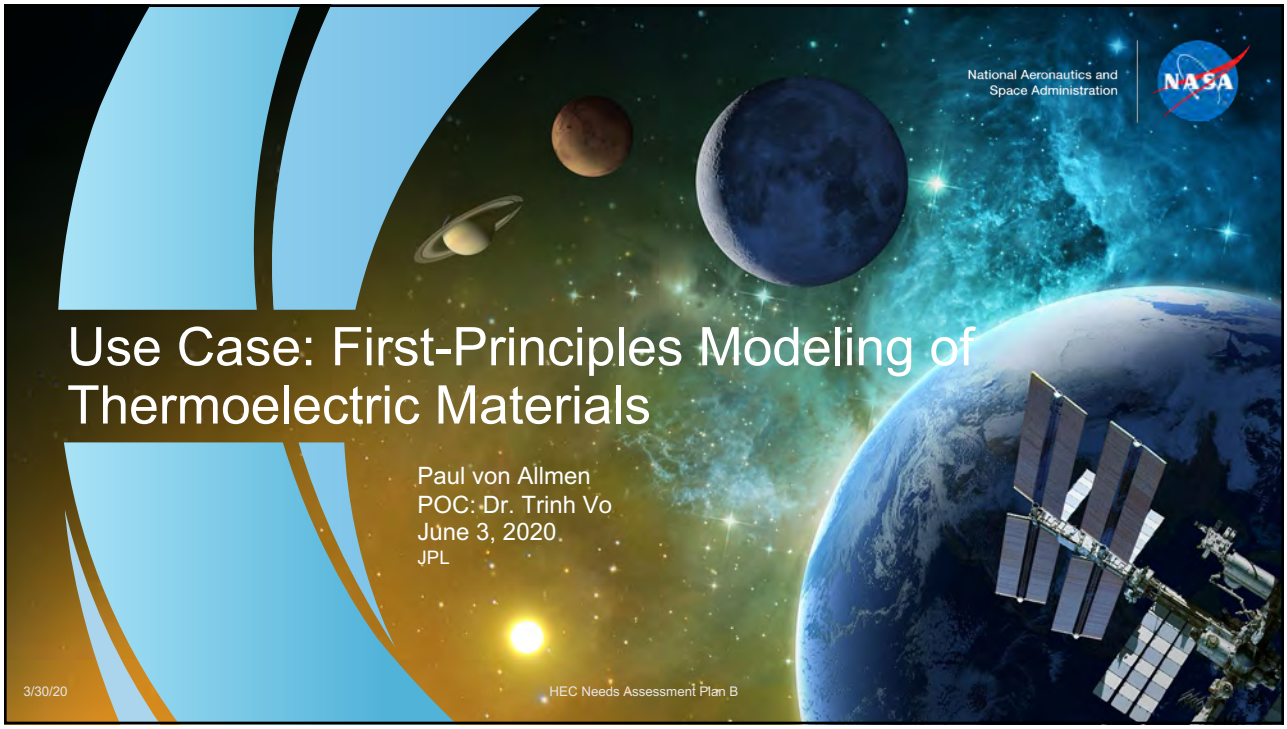

Use Case: Mission Need

- Mission:
  - This effort supports the design of new materials for Radioisotope Power Generators in deep space missions.
  - The number of runs and hence the exploration of the parameter space depends on the availability of HEC resources.
- Criticality:
  - If NASA supercomputing was no longer available, we would need to apply for resources at other HPC center (UCSD, TAMU, commercial). The potential cost of these resources would reduce the scope of your research.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Data is low TRL scientific research grade. However, crucial new results may impact the competitiveness of our program.
  - Regular interruption of HEC can be detrimental to our calculations, depending on the frequency, relative the check points in the simulation code.
  - No interaction with entities outside of NASA is needed.
- HEC Function:
  - HEC is used on a continuous basis to run our simulations.
  - The output of the HEC calculations (electronic band structures) are used for local processing at JPL.
- Impact of HEC:
  - If NASA supercomputing was no longer available, we would rent time on a commercial facility.
- Mission Concept:

14-1-11 zintl band structure

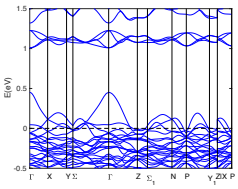

Use Case: Current HEC Application Characteristics

- Work flow:
  - A run is submitted with an input file, and potentially large restart files.
  - There is no need for interaction by user, operator or other system in the computing workflow.
  - No near-real time computing is required.
- Data:
  - In: small input file with parameters; potentially 100's GB restart files.
  - Out:
    - The output data requires long-term retention for future analysis.
    - The electronic band structures are used to compute density of states and thermoelectric transport coefficients.
  - Transfer:
    - The output is transferred to a local server at JPL.
- Computation:
  - Type:
    - Standard MPI on CPU computation is sufficient.
- Evaluation: NASA HEC (Pleiades) has been essential for accomplishing our work.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Publication cycles affect our use of NASA HEC.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Storage capacity on Pleiades is a limiting factor.
  - Impact: We spend time transferring large files between JPL and Pleiades.
  - External Sources: We run our HEC codes exclusively on Pleiades.
- How will this Use Case evolve over the next decade?
  - Maturity: The project needs for HEC will remain unchanged or grow slightly (20-30%).
  - Growth: A larger number of runs may be needed. A new HEC code may be used.
  - Re-use: The number of people running the codes may grow from 1 to 2 or maybe 3.
  - Re-hosting: Migration involves recompiling with the proper environment (Intel compilers and numerical libraries)

Use Case: Mission Need

- Mission: Develop computationally efficient machine learning models for processing and extracting information from petabyte scale Earth observations from geostationary and low earth orbit satellites.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Confidentiality is not of a large concern
  - HEC availability of GPU resources is critical
  - High-performance read capability to move data from storage to make best use of GPUs.
- HEC Function
  - Supports and maintains hardware and software for GPU computing which enables training and inference of deep learning models. Anaconda environments for deep learning frameworks allows new users to quickly apply deep learning models to their research. Packages used include python, PyTorch, TensorFlow, Nvidia Cuda drivers, and Jupyter notebooks.
  - Models developed support mission objectives by generating new approaches and datasets leading to peer-reviewed publications and software tools.
- Impact of HEC:
  - Without GPU availability, training deep neural networks is not feasible, and researchers would be forced to find external GPU resources from cloud providers or NSF/DOE resources.
- Mission Concept: Using NASA's computing power to learn from large data resources provided by the NASA Earth Exchange.

Use Case: Geostationary Virtual Sensors with Deep Learning

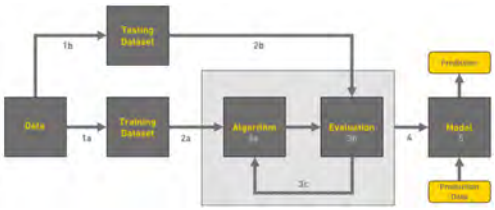

- Workflow:
  - Collect training and testing datasets, build data pipelines by preprocessing raw data sources.
  - Iteratively refine architecture with multi-GPU experiments
  - Select final model and perform inference (GPUs for large models)
- Data:
  - In: High-dimensional imagery from GOES 16/17 and Himawari-8 Level 1a data are input.
  - Out: Inputs and labels are used to learn the neural network. The resulting network is then used to performance inference and process generate new datasets. Networks are selected by validating on test data.
  - Transfer: Software and data stays within NEX
- Computation
  - GPU usage with MPI capabilities and fast storage read times enable reduced training and quicker iterations of experiments.
- Evaluation: The effectiveness of HEC in supporting this work is judged quantitively by the availability of resources for GPU computing and the training examples/second seen by the model.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Most of the computing power is used while running parameters searches and experimenting with data inputs. Distributed training is necessary when networks can take multiple days to train to converge. Experiments are executed sporadically during the project lifecycle.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: The number of GPUs is a constraint.
  - Impact: More users training deep learning models will reduce research productivity.
  - External Sources: Cloud computing (AWS, Google Cloud, Azure) provides unlimited GPU access with state-of-the-art hardware, instances are optimized for deep learning applications. Cloud computing at a large scale provides hardware when needed but at a large cost.
- How will this Use Case evolve over the next decade?
  - Maturity: Multiple deep learning networks will move to production data pipelines for efficient processing.
  - Growth: Increase in data resolution, more users, larger networks.
  - Re-use: The number of users requiring GPUS will grow
  - Re-hosting: Minor refactoring required, code is open-sourced

Information about the Survey

- The following slides describe the kind of data that the Needs Assessment will collect via a survey
- This material is needed to supplement the description of the Use Case presented to the Affinity Group but will not be presented.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs
    - High-dimensional matrix multiplication, gradient estimates
  - Describe the application
    - Antecedents: Open-source software
    - Programming languages used: Python and PyTorch
    - Thousands
    - Python scripts, Vim, Jupyter, Tensorboard
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Dozens of GPUs, multiple terabytes of storage for training data
    - GPU accelerators
    - CPU memory can be limited, GPU memory places a limit on training samples and model size
    - Intranode training with multiple-gpus, MPI jobs to starting training runs.
    - Interactive for development and PBS for large jobs

Current Use Case (Survey)

- Describe the data requirements for the application
  - Data streamed to node, GPU computes gradients and updates model, returns learned weights
  - Streamed via a petastorm distributed database based on Apache Parquet
  - GBs
  - Terabytes of training data, inference could generate petabyte scale data
  - Statistics computed using MPI scripts, visualization using Tensorboard and matplotlib
  - High speed shared file systems and GPU availability
  - High-quality examples are selected from data sources, cleaned, resampled, and stored in petastorm
  - Millions of examples taking single terabytes of storage
  - Outputs are distributed at NEX
- What are the key computing obstacles to effective execution of the current application?
  - Availability of GPUs
- Are there any analytic functions performed on the output of the application?
  - All the functions are analytical, robust statistics are computed on outputs
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Background in machine learning, statistics, python, and PyTorch

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
    - Too few GPUs per ML researcher will reduce availability and limit experimental decisions
  - Describe how the application and workflow is expected to evolve over that time frame
    - Networks will increase in size, use different learning algorithms, and produce more data products.
    - Resolution of datasets will continue to increase,
    - Models will become more efficient over time
    - Refactor for emerging platforms
    - Rehost onto cloud
    - Introduction of a general set of tools developed for high-dimensional datasets
  - If possible, describe how the computational requirements will evolve over time
    - Grid searches and experimental designs will be selected based on availability
    - GPU usage will increase
    - Memory is not a large concern
    - MPI
    - Terabyte scale
    - Batch and interactive (development)

National Aeronautics and
Space Administration
NASA
Use Case: NASA Advanced Supercomputing
Bob Ciotti
June 2, 2020
ARC TN
3/30/20
HEC Needs Assessment Plan B

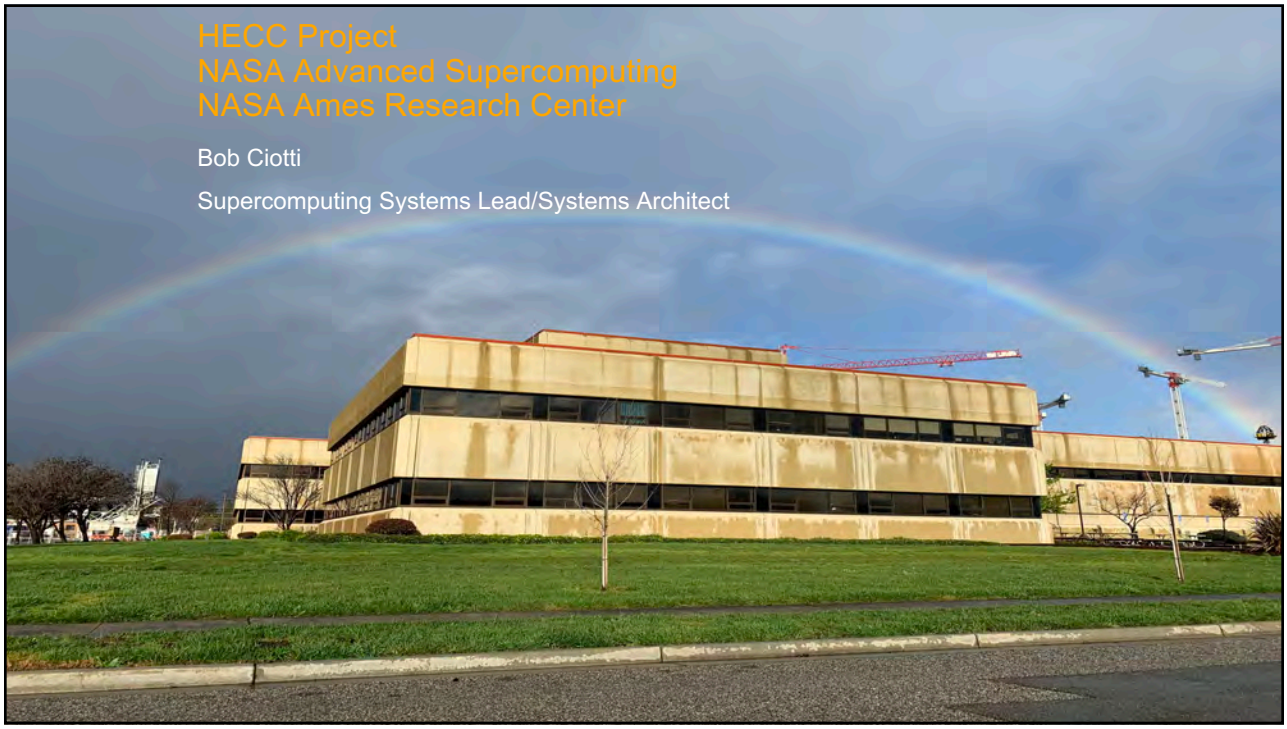
HECC Project
NASA Advanced Supercomputing
NASA Ames Research Center
Bob Ciotti
Supercomputing Systems Lead/Systems Architect

HECC Computational Lan pe

Compute Bound

Internal to HECC

Well Understood Computations

$x + y = z$

Embarrassingly Parallel

Systems Integration Team
24x7x365 Ops
Evolving HW Architecture
Evolving SW Services

Tightly Coupled

Complex and Evolving Computations

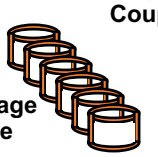

Data/Storage Intensive

$$\frac{\beta}{\gamma_w^2}\int_0^{2\pi} h_w(s')\frac{\partial G_w(s,s')}{\partial \tau}L(s')\,ds' + \frac{\varepsilon_w}{\gamma_w^2}\left(\frac{e_w(s)}{2}+\int_0^{2\pi} e_w(s')\frac{\partial G_w(s,s')}{\partial n}L(s')\,ds'\right) - \frac{\beta}{\gamma_c^2}\int_0^{2\pi} h_c(s')\frac{\partial G_c(s,s')}{\partial \tau}L(s')\,ds' - \frac{\varepsilon_c}{\gamma_c^2}\left(-\frac{e_c(s)}{2}+\int_0^{2\pi} e_c(s')\frac{\partial G_c(s,s')}{\partial n}L(s')\,ds'\right) = 0$$

External Network Intensive

Systems Evolution

Classic view of HPC
- large/sometimes singular computing problem
- standalone/self contained - few external dependencies
- detailed construction of static inputs
- results interpreted locally
- compute the more critical resource

Major shift to data analysis and its distribution
- increases in computing capabilities - 'Moore's law'
- increases in volumes of data produced
- wide geographic distribution of ***compute providers*** and ***users*** as well as and ***generators*** and ***consumers*** of data products
- moved from only compute being most critical resource to include data analytics, its understanding and distribution.

How can HECC adapt to ingest, compute, preserve, and distribute information to a de-centralized collection of people, systems and data?

Program/PI Challenges

Data Processing
- Ground operations/processing
- Data source, mid-points, destinations
- Tools and Databases
- Scheduling/Backlog/Cost

Data Distribution, Access Control and Sovereignty
- How to share
- Where is the data needed
- Who can access
- Internal or external
- Anonymous or Credentialled
- Where to store various portions
- Ability to move around to different providers

Data preservation
- Where to save and for how long
- Disaster recovery

Interoperability
- Often run on many platforms may need to make cost tradeoffs
- How to project costs and remove barriers to competition?

HECC Development

- Data Processing
  - Dedicated systems attached to supercomputing
  - Dedicated purpose-built filesystems
  - Specific user SW and Access requirements
  - Dedicated Databases
  - Facilitate scheduling needs for timely processing (e.g. satellite/observational data streams)
- Data Distribution, Access Control and Sovereignty
  - Filesystems can be exported externally
  - Data Accessible both inside and outside
  - Either credentialled or anonymous
- Data Preservation
  - Automated Tape Archiving
  - Indexing/Dedicated Databases
- Interoperability
  - Provide compatibility and interoperability between multiple environments through singularity containers

Project Coordination Activities

Dedicated systems/capabilities – Projects: TESS/SDO/HySDS/ECCO/NEX/uPSP
- Customized configurations
- 'Front End' Servers
- Databases
- Specific user SW and Access requirements
- Dedicated filesystems

TESS – Transiting Exoplanet Survey Satellite
https://www.nasa.gov/tess-transiting-exoplanet-survey-satellite
SDO – Solar Dynamics Observatory
https://sdo.gsfc.nasa.gov/
HySDS – Hybrid Science Data System
https://hysds-core.atlassian.net/wiki/spaces/HYS/overview
ECCO – Estimating the Circulation and Climate of the Ocean
https://www.ecco-group.org/
NEX – NASA Earth Exchange
https://www.nasa.gov/nex
UWT - Unitary Plan Windtunnel
https://www.nasa.gov/centers/ames/orgs/aeronautics/windtunnels/index.html

Automated Tape Archive - Projects: SDO
- Investigate moving from a separate user driven tape archive to coupled high performance disk/tape archive.

User Configurable Software Environment - Projects: HySDS, AERO, Machine Learning/Data Analytics
- singularity containers

Data Distribution - Projects: HySDS/SDO/ECCO/NEX/uPSP
- Data products generated inside HECC can be accessed both internally as POSIX files or externally through HTTPS via a GUI web portal or simple programmatic interfaces like wget and or file range requests.

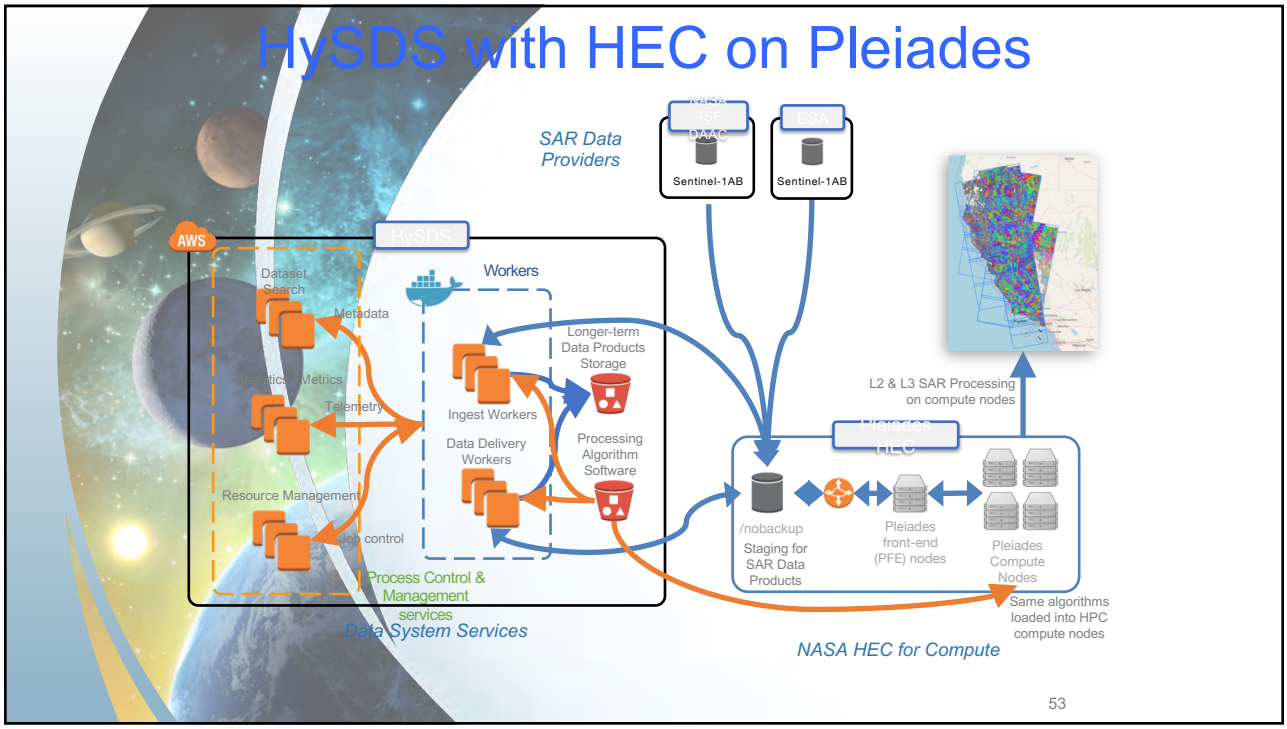
HySDS with HEC on Pleiades
SAR Data Providers
NASA ASF DAAC
Sentinel-1AB
ESA
Sentinel-1AB
AWS
HySDS
Workers
Dataset Search
Metadata
Analytics Metrics
Telemetry
Resource Management
Job control
Process Control & Management services
Data System Services
Longer-term Data Products Storage
Ingest Workers
Data Delivery Workers
Processing Algorithm Software
L2 & L3 SAR Processing on compute nodes
Pleiades HEC
/nobackup
Staging for SAR Data Products
Pleiades front-end (PFE) nodes
Pleiades Compute Nodes
Same algorithms loaded into HPC compute nodes
NASA HEC for Compute
53

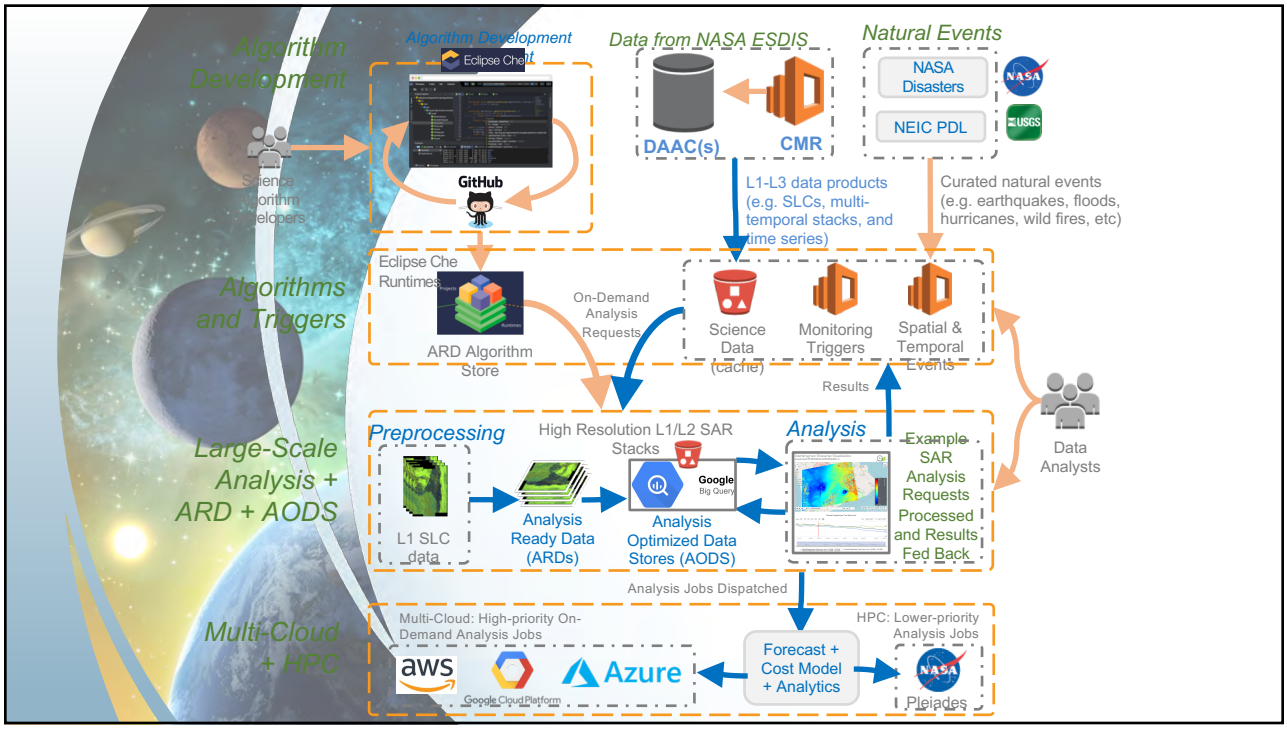
Algorithm Development
Science Algorithm Developers
Algorithm Development
Eclipse Che
GitHub
Data from NASA ESDIS
DAAC(s)
CMR
Natural Events
NASA Disasters
NEIC PDL
L1-L3 data products (e.g. SLCs, multi-temporal stacks, and time series)
Curated natural events (e.g. earthquakes, floods, hurricanes, wild fires, etc)
Algorithms and Triggers
Eclipse Che Runtimes
ARD Algorithm Store
On-Demand Analysis Requests
Science Data (cache)
Monitoring Triggers
Spatial & Temporal Events
Results
Large-Scale Analysis + ARD + AODS
Preprocessing
L1 SLC data
Analysis Ready Data (ARDs)
High Resolution L1/L2 SAR Stacks
Google Big Query
Analysis Optimized Data Stores (AODS)
Analysis
Example SAR Analysis Requests Processed and Results Fed Back
Data Analysts
Analysis Jobs Dispatched
Multi-Cloud + HPC
Multi-Cloud: High-priority On-Demand Analysis Jobs
aws
Google Cloud Platform
Azure
Forecast + Cost Model + Analytics
HPC: Lower-priority Analysis Jobs
Pleiades

Infusion Examples

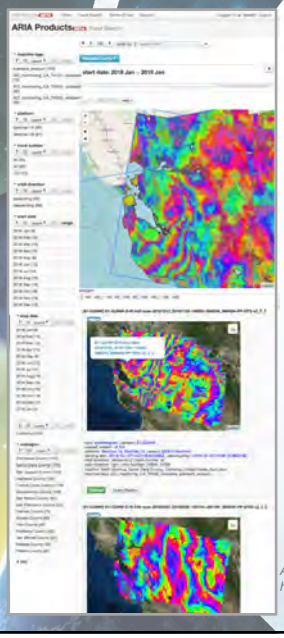

Infusions

- *Advance Rapid Imaging & Analysis (ARIA)*
- *MEaSUREs water vapor & cloud classification*
- *OCO-2 L2 full physical sounding retrievals*
- *Getting Ready for NISAR (GRFN)*
- *Joint ESA-NASA MAAP*
- *SMAP-in-the-Cloud (in development)*
- *SWOT (in development)*
- *NISAR (in development)*

Processing and Delivery of L2 and L3 Data Products in cloud

Multiple SDS Processing Modes

- Bulk (re)processing
- Forward "keep-up" processing (running 24/7)
- On-demand processing
- Urgent Response

NISAR 5X (430TB+ per day)
~8,300 compute nodes
260K+ cores
2PB+ of scratch disk

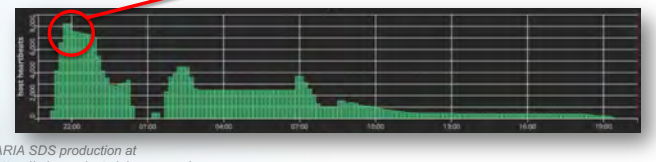

ARIA SDS production at https://aria-products.jpl.nasa.gov/

Common Needs and Gaps

Projects often pull in large volumes of existing data from many different external sources including DAACs, NASA mission data, International sources, even ***air-gapped*** windtunnels driving the need for high performance wide area networks to bridge the connections from HECC to the outside world including the cloud. HECC limited to a 10Gb WAN connection.

Projects create large volumes of content totaling 10's of petabytes for the limited number of projects identified to either distribute the content – and/or continue to analyze it on HECC resources over months or years.

Projects often need a database to index or otherwise describe that content and a consistent way to access the data from both inside and outside of HECC.

Projects may be required to share results either by policy or specific mission objectives.

Use Case: Current HEC Application Characteristics

- Work flow: Workflow can be triggered frequently for data produced by observational platforms or periodically for campaigns associated with focused upper level product analysis.
  - Since these workflows can operate over long periods of time, operators are often involved in monitoring/shepherding
  - Not a real time requirement, but needs the ability to start and ramp up fairly quickly, often a node at a time to some reasonable level like several hundred nodes.
- Data:
  - In: Varies in size greatly. Often data comes from multiple sources including NASA DAACs, university, or international efforts. Some are quite large to multiple petabytes.
  - Out:
    - Retention varies. Typically analyzed and shared with other teams, some with access to HECC and some outside NASA. Subsequent analysis often desired.
    - Desire to retain data locally and to publish it externally or to copy it elsewhere.
- Computation:
  - Type:
    - Typically capacity including a mix of CPU and GPU. But this characterization may underestimate the effort required to effectively process multi petabyte data sets is also a challenging problem that benefits from the robust supercomputing infrastructure that can effectively exploit parallelism.
- Evaluation:

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Availability of data inputs
  - Commitments to processing timelines required to make data products available.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case:
    - WAN network performance for input and distribution
    - Sufficient storage space for computational results and the ability to maintain access to that data both internally and externally for months or years. Multiple petabytes.
  - Impact: Varies – see project listings
  - External Sources: Codes or portions of the processing pipelines are run on other NASA machines, external systems, and/or on cloud systems. Some run internally mainly for cost reasons, others run externally on cloud
- How will this Use Case evolve over the next decade?
  - Maturity: use cases will continue to mature driven by the continued growth of data products
  - Growth: Data drive growth - significant
  - Re-use: More teams likely to adopt some sort of processing pipelines.
  - Re-hosting: Important to run on multiple different platforms and possibly even to reproduce earlier results.

HEC Needs Assessment Workshop

Affinity Group Sessions

Wednesday 6/3/20 1pm – 4pm (AG6-1) (nominal affinity: Science Data Processing)
(10am – 1pm Pacific)

Facilitator: Peter Williams <peter.williams6850@gmail.com>
Scribes: Jarrett Cohen <jarrett.s.cohen@nasa.gov> and Sean Keefe rosemary.e.keefe@nasa.gov

KEY POINTS

Mission Use Cases covered during HEC Need Assessment

AERONET atmospheric aerosols, GEOcarb, Earth Science, TESS (Transiting Exoplanet Survey Satellite) science processing, thermoelectric material design, Geostationary data and Deep Learning

Impacts of HEC on Mission Needs: Range of Impacts and Significance

- Efficiency: calculation time is a fraction of that of alternatives, duplication of computing resources would occur if HEC unavailable and programs purchase separate systems, Data Analytic Centers assume efficiency and effectiveness increase as multiple data sets and mission areas utilize a shared HEC environment, NASA Earth Exchange efficiency dependent upon shared resources and consistent training that accelerates user readiness and preparation to leverage HEC resources
- Dependency: Earth Ventures missions require significant data processing that HEC provides
- Dependency: TESS science processing is HEC dependent and replacing the workflow, while possible, exceeds budget projections;
- Mission dependency: Deep space missions depend on thermoelectric power generation and that depends on HEC resources to design and simulate with near-constant, low-demand runs
- Earth observations depend on Machine Learning and Deep Learning to gain efficiency and effectiveness in the models and in the methods for extracting and processing observation data; ML/DL requires GPU HEC to accomplish the training in a cost-effective manner within NASA span-of-control (Vandal)
- Climate models, weather models, wildfire models, space weather, geologic hazards, and aeronautical engineering, all require HEC for basic research, simulation, and system design and testing.

Mission Security Needs relevant to HEC: Range and Significance

- Confidentiality
- Integrity
- Availability: Varies, with some missions dependent and others less so
- Public Confusion/Controversy: Yes, especially if immature data is released
- National Security
- Intellectual Property

Criticality of HEC to Mission Needs

- Financial consequences: Cost of Alternatives compromise program efficiency and effectiveness
- Mission Security
- NASA HEC-GPU is essential for Machine Learning and Deep Learning required for Earth Science observations and to leverage NASA Earth Exchange data resources further

Workflow Issues and Needs

- NASA Earth Exchange efficiency dependent upon shared resources and consistent training that accelerates user readiness and preparation to leverage HEC resources
- Limited services (messaging) and 1-off services (poor sharing): HEC could centralize this service across agency and across collaborators
- Alignment and sequencing of project/program demands is challenging for some needs, not all.
- Moving large datasets from NASA platforms to HEC and back is already an issue; moving those datasets to non-NASA HEC systems seems likely to compound the challenge.
- As HEC application includes more analytics, not just computation, demand on resources grows
- Workflow affected by distributed workforce, decentralized systems, and decentralized data, all of which are increasing

Data Input, Output, and Transfer: Issues and Needs

- Possibly keep this section focused on processing speed (FLOPS) and data volume (BYTES) as a specific look at these two elements of workflow issues or needs
- *Input, output, and transfer* : Need fast interconnects, networking (in and out of NAS), and high-performance file management systems; throughput demands of volume and time require HEC
- *Computation and Storage:* Storage is a primary technological limit

Evaluation: HEC Effectiveness and Value Measures

- Might roll this category under the earlier "Impacts" section [or cross-walk impacts and evaluation to ensure each reinforces the other – evaluate based on critical impacts]
- Computational efficiency, available disk storage, time in queue, work of support team
- Tools for the data analysis, visualization and expertise support

Key Factors Affecting Current Load on Computing Assets

- Publication cycles and University calendars
- Distributed training needed for ML and DL while in research and experiment stages

Anticipated Evolution of Use Case and Mission Need

- Maturity
- Growth: increasing data resolution means more data input with concurrent expectation that it will be incorporated into models and simulations
- Re-Use: Increasing expectations that available data from NEX and other sources will be leveraged, often in ways that depend on HEC
- Re-Hosting
- Other(s): Decentralization of workforce/colleagues, systems, and data resources, as well as data repositories

Functional Gaps: Existing and Anticipated

- Constraints: big data analysis (including ML), disk space, computational allocation, queue, network, I/O, #GPU's (ML and DL)
- Impact of Constraint or Gap: insufficient performance to obtain the required models and data analysis, mission failure, ML/DL is expanding which means the training expands on HEC

resources also needed for research production; more ML/DL training demands on same HEC resources reduces research productivity because competition for same computing resources

- External Sources

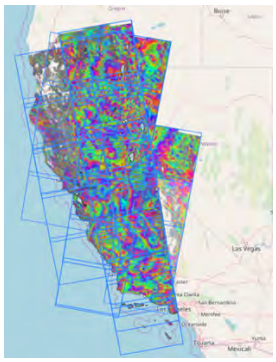

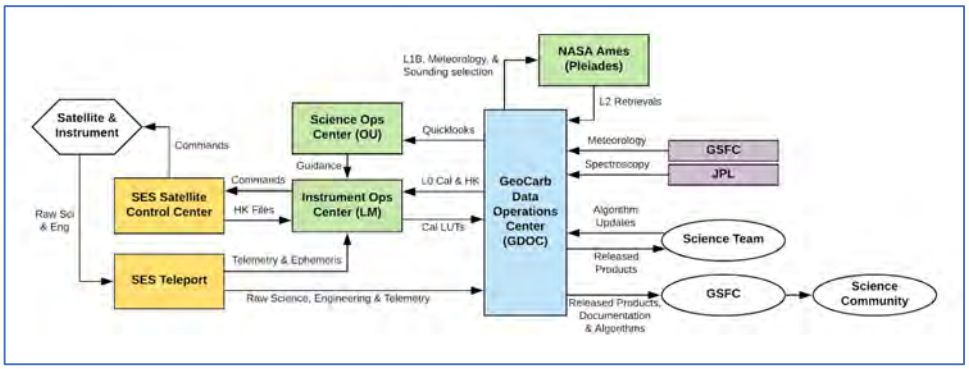

Figure 1. Example of distributed mission (Cronk)

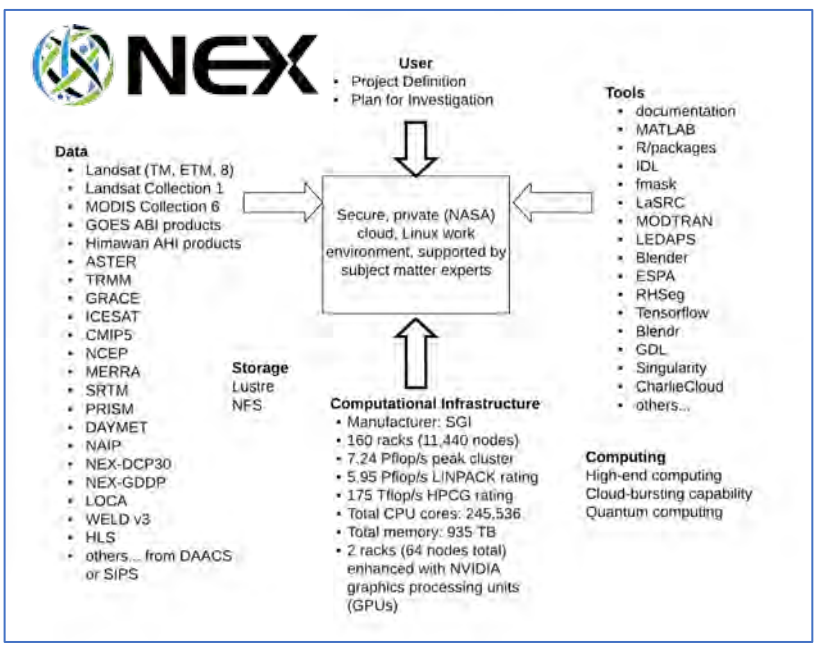

Figure 2. Example of datasets and tools interacting with HEC (Ciotti)

Attendees

1. Mike Little, NASA Goddard, ESTO and HEC Program, Convenor
2. Peter Williams, Facilitator
3. **Presenter #1: David Giles, GSFC/618**
4. **Presenter #2: Heather Cronk, CIRA/CSU**
5. **Presenter #2: Philip Partain, CIRA/CSU**
6. **Presenter #3: Jennifer Dungan, ARC, Code SG**
7. **Presenter #3: Ramakrishna Nemani, ARC/SGE**
8. **Presenter #3: Andrew Michaelis, ARC/SGE**
9. **Presenter #4: Peter Tenenbaum, ARC**
10. **Presenter #4: Jon Jenkins, NASA Ames, Kepler/TESS/Pathfinder**
11. **Presenter #5: Paul von Allmen, JPL**
12. **Presenter #6: Thomas Vandal, AMES/BAERI**
13. **Presenter #7: Robert Ciotti, NAS**
14. Aaron Pina, NASA HQ, Weather Strategy Plan
15. Asri Vasiri
16. Dan Duffy, NCCS
17. Daniel Chrichton
18. Ellen Salmon, NCCS
19. Elizabeth Hartman, NAS
20. Elizabeth Hook, NASA HQ
21. Ellen Salmon, NAS
22. Emily Kuhse, , ARC/TN
23. Haoquiang Jin, ARC/TN
24. Jarrett Cohen, NASA Goddard, HEC Program
25. Johnny Chang, ARC/TN
26. Laura Carriere, NCCS
27. Leigh Ann Tanner, NASA Ames/NAS HEC Program
28. Michelle Moyer, NAS
29. Piyush Mehrotra, NAS
30. Robert Hood, ARC/TN
31. Sean Keefe, NCCS
32. Sharad Gavali, ARC/TN
33. Subhash Saini, ARC/TN
34. Tsengdar Lee, NASA HQ, HEC Program
35. Vidya Bobbiligama, NAS
36. William Thigpen, NAS, HECC Manager
37. Yan-Tyng Chang, CSRA

1. Use Case: AERONET NRT and Historical Inversion Processing

David Giles
david.m.giles@nasa.gov
GSFC/Code 618
SMD

Raw Notes 1:
Mission Concept description. Important for satellite remote sensing. Many users of this data, important that we provide the data to them on a continuous basis. Historical processing can sometime lead to slowness in data files being saved or extracted. Cause issues ending up in against-the-wall time. We can have issues with disk resources not available due to maintenance.

Data transfer only used for moving input and output files, on a temp basis, local system storage. Capacity is limited to number of SLURM jobs. Model does not have any parallel, multi-core processing units. HEC impact can be measured based on how long a data set can be processed; millions of retrievals can be performed in a fraction of the time than computed on the AERONET server. HEC is essential to completing this work since it would take several months to complete historical reprocessing on the local system.

The number of processed retrievals is the success of the HEC resources.

Science data reprocessing – data are often reprocessed to answer scientific questions, associated with new product versions, or publications. Routine measurements and field campaigns demand near real time processing Try to adjust to new changes to improve systems, HW & SW architecture, compilers. We

will change from Fortran to C to help make SW more streamlined. Use NCCS Discover supercomputer. Model to obtain aerosol characteristics.

Importantly, we would not be able to have released our latest product version (Version 3) in a timely manner without the use of NCCS discover resources. In this regard, HEC also facilitated AERONET research and a recent publication on the Version 3 AERONET aerosol retrievals and uncertainty estimates:

> Sinyuk, A., Holben, B. N., Eck, T. F., Giles, D. M., Slutsker, I., Korkin, S., Schafer, J. S., Smirnov, A., Sorokin, M., and Lyapustin, A.: The AERONET Version 3 aerosol retrieval algorithm, associated uncertainties and comparisons to Version 2, Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2019-474, accepted, 2020.

Slide 2: Use Case: Mission Need
600 sites worldwide. Measurements taken hourly. Program can take 5 to 45 minutes to run. We can process on Discover on the order of 6,000 of these retrievals at any one particular time. On our previous server, we could only do 100. We have 60-times increased capability and performed twice as fast. Making things faster would be beneficial (improving code, hardware, storage, etc.). Because of a lot of input/output, the storage can be slow and allow the jobs to run up against the wall time. It would take 6 to 8 months to process on our previous system; now we can do it in a matter of a few weeks.

Slide 3: Use Case: Current HEC Application Characteristics
We ready the information on our local server and then send it to NCCS. Data transfer is only used for moving the input and output files between our server and NCCS. We run perhaps 6,000 jobs at once. The capacity is currently limited to that number. The model does not currently have parallel multiprocessing in it to use more cores at a time.

Slide 4: HEC Needs Assessment Workshop Key Questions
Over the next decade, we will continue what we're doing and try to adjust to all the changes going on. There are several changes to the hardware and software. We hope to change from Fortran to C to help with making the software more streamlined.

Slide 6
There are many users of this data. It is important for us to maintain access to the data and continually improve it.

Q&A:

Q from Laura Carriere (Int) to Everyone: 1:23 PM
How often, per day, would you like to run your NRT runs?

David Giles:
Every 3 hours would be optimal.

Q from Michael Little (Int) to Everyone: 1:25 PM
Is Aeronet output used for operational requirements of other systems, like weather forecasts or NOAA air quality forecasting

David Giles:
As far as operationally, to some extent. There is some latency in measurements itself. If we can make it more available in a shorter amount of time, it could be used operationally.

Q from Michael Little (Int) to Everyone: 1:26 PM
That is, are there agreements to deliver data to other agencies?

David Giles:
Not as far as I know, nothing formal, NASA has mandates for using ground-based measurements for validation of satellite retrievals (e.g., MODIS, VIIRS, MISR, OMI). We also work closely with GMAO, MERRA-2 and so forth. We are participating with other entities, nationally and internationally.

We work collaboratively with NOAA, and they obtain aerosol products routinely from the AERONET web site for weather forecasts,air quality, and geostationary satellite algorithm development and validation.

Q from Peter Tenenbaum (Int) to Everyone: 1:27 PM
Bottlenecking of data reads -- what format is being used for the data that is read? HDF5? netCDF? Something else?

ASCII text formatted files

Q from JENNIFER DUNGAN (Int) to Everyone: 1:30 PM
How many SBUs does it take to do the latest complete historical reprocessing?
~20000 (old type SBU)

Q from Laura Carriere (Int) to Everyone: 1:31 PM
Can I set up a call with our HPC team to discuss ways we can help you? I have a number of thoughts.

Yes, we had a meeting the next day. We are attempting to restructure processing jobs to use SLURM array and utilize additional resources on NCCS.

Q from David Giles (Int) to Everyone: 1:38 PM
Currently, data input and output are in ASCII format. These files a tar.gz for transfer between systems.

---

2. Use Case: GeoCarb Data Operations Center

Heather Cronk, CIRA/Colorado State Univ.
heather.cronk@colostate.edu
Philip Partain, CIRA/Colorado State Univ.
philip.partain@colostate.edu

Raw notes:
Using Pleiades to generate level 2 product. Criticality: w/out Pleiades, we would be building a large HW cluster. We can't afford that—less expensive than using commercial cloud. Security: Confidentiality is important until products are released. We've been allocated up to 8,200 SBU/day, but we expect to typically use half of that.

HEC Function: We process a huge workload. Intend to set up scripts & dashboards. 6% in real time, meant to support ID of instrumentation issues, with one-week lag we retrospectively process the data (uses updated calibration). Level 2 data products go to the Goddard DAAC in 60 days. Rolling repo of data on Pleaides.. Input data will be resident on the system for 2 weeks. All input data comes from CSU.

Algorithm runs per sounding. Trying to understand and develop our workflow. Partner with HEC to meet our level 2 processing requirements.

Slide 3: Use Case: Current HEC Application Characteristics
Level-2 processing split up between Pleiades and CSU. The retrieval is quite expensive computationally; we are averaging 8 minutes per sounding, and we process up to 1 million soundings per day. We have partnered with HEC to meet our mission requirements in terms of mission processing and delivery.

Large volume of observations. If we need to catch up due to downlink, processing, or transfer issues, could require a burst of horsepower required. Expecting annual data science reprocessing campaigns.

Slide 4: HEC Needs Assessment Workshop Key Questions
We can run into variations each day in terms of what kind of processing we need.
The algorithm is based on the one for OCO-2 (although OCO-2 is a polar orbiter rather than a geostationary instrument). Splitting between CSU and Pleiades has some risk with passing data back and forth. What we will be doing day to day is pretty well known.

Level 2 algorithms built on OCO-2, which has used Pleiades for reprocessing campaigns, 5 years of data, smaller amount due to polar orbit and narrow swath. Resources available were not enough to complete their reprocessing. OCO-2 split their reprocessing between the commercial cloud, Pleiades, and on-premises. Concerns about Pleiades capacity for reprocessing. Using the commercial cloud helps with available resources at additional expense.

Logistical concerns with passing data back and forth. 2023: target date for nominal ops. Scope not intended to grow. The number of folks running algorithms on Pleiades not anticipated to grow. Hope to run in containers, which would simplify our ability to distribute processing and would support different architectures.

Q&A:

Q from Tsengdar Lee (Int) to Everyone: 1:41 PM
In the past, when HEC is supporting data processing, there was a requirement not to update/upgrade the computing systems for the upgrade may change the product. Is this still a concern?

A: We will want to keep track of system changes, but will be running regression tests with each update of the algorithm or hardware. It should be hardware- and OS agnostic, particularly if we are able to run our software in containers.

Q from Brian Thomas (Int) to Everyone: 1:47 PM
Not sure I understood the bottlenecks, what are you getting out of commercial cloud which you need (and cannot get from Pleiades), containers? Transfer speed between data and algorithm? What downlink speeds do you need vs what exist?

A: There is nothing we're getting from commercial clouds that we're not getting from Pleiades. The only thing we could get is unlimited resources so we can scale up, but there is a cost associated with that. We are hoping to do as much as possible with HEC resources.

Q from Brian Thomas (Int) to Everyone: 1:47 PM
Shouldn't the move to containers ameliorate (lessen the impact) of OS-based issues?

A: Yes, I would hope that it would. We are not having OS issues, but it does protect us from changes to the underlying architecture.

Q from Robert Ciotti (Int) to Everyone: 1:48 PM
Do you want or plan to save the data locally at Pleiades to support reprocessing?

A: That is something we are considering. Originally we were planning not to, but after talking to Bob Hood's team, that is a possibility we are thinking through.

Q from Brian Thomas (Int) to Everyone: 1:49 PM
So, scalability is the issue with Pleiades?

A: Yes; it is over-subscribed now, and I feel it will get more and more popular.

Q from Robert Ciotti (Int) to Everyone: 10:50 AM
How much data?

A: Data provided as follow-up:
Current estimates for L2 input products that could be stored on Pleiades to support reprocessing:

File Type	Estimated Daily Volume (GB)	Estimated Mission Volume (GB; nominal 3-year mission)
L1B	91.3	99847
Resampled Meteorology	3.6	3950
Selected Soundings File	0.01	10.7

Output products (not for long-term storage on Pleiades):

File Type	Estimated Daily Volume (GB)	Outcome
L2 Retrieval Output (per sounding)	1500	Deleted after aggregation
L2 Retrieval Output (per granule)	1500	Transferred after aggregation
L2 Retrieval Logs	30	Transferred after aggregation

3. Use Case: NASA Earth Exchange

Jennifer Dungan, ARC/Code SGE
Jennifer.L.Dungan@nasa.gov
Andrew Michaelis, ARC/Code SGE
andrew.r.michaelis@nasa.gov
Ramakrishna Nemani, ARC/Code SGE
rama.nemani@nasa.gov

Main points:
P1: Team of excellent folks. Unusual use case: many use cases in exchange. Work on regional global continental scale observations. Can't be performed at university clusters. We also support a # of smaller use cases. Use case: geostationary data.

P2: Security. Public release of emails long ago in use cases, climate projection, sensitive topic. Most use cases aren't publicly sensitive. Happy with current security level, but there are security hurdles. Avail of

system can impact objectives. Projects often go into cost extension due to other factors. ESD model based on external collaborators with no NAS credentials. Bring on remote users. HEC supports hardware/software job scheduling user support, which is fantastic, help day of night. Variety of languages revised or rerun. Output of those projects goes to a DAAC or published. Impact: no longer avail: all teams w project at these scales would have to use cloud resources, a severe funding hurdle. Costing model complex for individual PIs to scope without long.

Researchers are using NADA's private cloud do what otherwise is not possible.

P3: NEX as a data analytic slide

P4: Major thrust: support geostationary products such as GOES 16 etc. Not real-rime data, though we have experimented with commercial cloud-AWS. 30 TB per year/platform. 90 TB from 4 satellites/year. We can rerun our processes. Thousands of cores workflow. 1 SBU/GB. Intensive data use, not compute. Users provide qualitative feedback.

P5: NEX, most important factor: how ready are users to use the environment/how familiar with Linux, etc. More importantly: PIs on system have many competing priorities. Sudden, sporadic use can result. User availability issue.

Use cases drive by storage and use case. We don't have and could use messaging services- a constraint.

Raw notes:
Slide 2: Use Case: Mission Need
NEX involves a lot of use cases. It supports groups that get grants.

I will go through one use case: geostationary data.

Public embarrassment could result because NEX deals with climate data. Remember the case where climate change scientists' emails were leaked? Climate projection is a publicly sensitive topic.

HEC user support is fantastic. Users are happy that they can call any time of day or night and get help.

Slide 3: NEX as a Data Analytic Center
NEX is one instantiation of this concept.

Slide 4: Use Case: A major current thrust of NEX is to support geostationary products (GeoNEX)

Single use case. Triggered when data becomes available from the satellites (see list under Data).

Main computation is done on Pleiades.

Compute is not particularly intensive, but the data use is intensive: 1 SBU per GB.

Slide 5: HEC Needs Assessment Workshop Key Questions
How ready are the users to use the environment? Are they familiar with the computing environment? Also, with other priorities, PIs and others end up being sporadic users.

Q&A:

Q from Tsengdar Lee (Int) to Everyone: 2:06 PM
Are you aware of any HPC platforms that offers web services? You mentioned services when you were talking about the pipeline. HPC is basically a platform-as-a-service and provides limited software and messaging services. Those Big Data platforms often have other services.

Jennifer Dungan:
By web services, I think you mean messaging services.

Andrew Michaelis:
I think we snuck that in there as a nice-to-have. We are looking for asynchronous messaging integrating with LDAP services to build our pipeline, rapid nq, for example. NOAA is using it for some of their stuff. I can compile a list for you.

from Brian Thomas (Int) to Everyone: 2:07 PM
What turnkey services in GCP (Google) or AWS are of interest? What about data architectures built around other technologies like Apache Spark (which can work well with highly parallel problems)

[Questions to be addressed later]

Q from Piyush Mehrotra: 2:07 PM
Jennifer mentioned AWS. **What would be useful coming out of cloud in an HPC environment?**

Q from Brian Thomas (Int) to Everyone: 2:08 PM
So...microservices built around LDAP/storage would be useful.

Q from Brian Thomas (Int) to Everyone: 2:09 PM
Your architecture is more of SOA (Service-Oriented Architecture) vs microservices.

Q from Robert Ciotti (Int) to Everyone: 2:09 PM
We do have some use cases with celery and rabbitMQ in development.

Q from Robert Ciotti (Int) to Everyone: 2:09 PM
Yes. [to a prior question]

Q from Robert Ciotti (Int) to Everyone: 2:09 PM
No. [to a prior question]

Q from Robert Ciotti (Int) to Everyone: 2:10 PM
Let's talk.

Q from Brian Thomas (Int) to Everyone: 2:10 PM
Messaging services are typically built for a given data architecture (you develop).

Q from Piyush Mehrotra (Int) to Everyone: 2:10 PM
Jennifer mentioned AWS services - which ones are useful for your workflow?

Q from ANDREW MICHAELIS (Int) to Everyone: 2:10 PM
Let's talk, Bob, when you have some time.

Andrew Michaelis: **Can anyone just join and link up with a queue?** If so, we would happy to be beta testers.

4. Use Case: TESS Science Processing Operations Center

Peter Tenenbaum, ARC
peter.tenenbaum@nasa.gov
Jon Jenkins, NASA/TNC
jon.jenkins@nasa.gov

Raw notes:
Slide 2: Use Case: Mission Need

Diagram: Gold boxes are the infrastructure. Orange boxes are the science pipeline.

We take data at the pixel level for each star.

We planned from the start to architect and run our pipeline in this manner It is based on Kepler. It took us 10 months to process 2 years of data before. After moving to Pleiades, it would take a few months to reprocess all of our data.

There are several layers of review and approvals before data gets released.

It takes 5 to 7 days to process 1 month of data.

Slide 3: Use Case: Current HEC Application Characteristics

Graphics show a Kepler example; processing is dominated by instrumental effects that we have to remove.

Data volume is increasing, and storage capacity is not increasing at the same rate.

We have several thousand cores, each chewing away on its own chunk of data happily and not caring what its siblings are doing.

Slide 4: HEC Needs Assessment Workshop Key Questions
We are not using a constant number of cores in steady state.

The allocation of >100,00 SBUs includes increase in targets and data coming from the extended mission.

Re-hosting: Not a chance.

Q&A:

Peter Tenenbaum:
Lots of discussion in chat box about messaging software (MQs). We would like to use something that is written and supported by somebody else, not us.

Q from Peter Williams:
Anything else from the chat box that you would speak to, beyond what you just did?

Peter Tenenbaum:
In terms of scalability, we had a similar experience to one of the discussions, our ability to get data into and out of pipeline is much more of a constraint than the availability of cores for processing. We are spending all of our time/are too slow marshaling and persisting, so more cores would not help.
For various reasons, the mission developed its own message system, Tiny MQ (TMQ). We would like to move to a third-party MQ that will work well on the NAS.

5. Use Case: Use Case: First-Principles Modeling of Thermoelectric Materials

Paul von Allmen
pva@jpl.nasa.gov
JPL/3980
Directorate (if NASA)

Main points:
P1: We are mostly using AWS, processing of data, research projects, ocean works, cloud computing and comp'l needs & processes that are embarrassingly parallel.

Small project going on for 10 years, records the devt of new material for RPG to support the nuclear program. Large calcs that can take weeks to months on Pleiades.

P2: Cost prohibitive, so we are critically dependent on Pleiades.

P3: Interruptions can be a problem, but not a problem any longer.

P4: We run as long as possible, transfer large files (GB) to JPL files.

P5: We have encountered difficulties with getting storage.

Raw notes 1: Use MPI to code Fortran. Expending sustainment funding from nuclear power program. We may run different sets of codes which would increase workloads. If we had to migrate, not too hard. Without HEC, some runs would take a year, not practical.

Raw notes 2:
Most of the work we are doing is using mostly AWS. We do use some of the OCO-2 and OCO-3 data, research projects like ocean work. The computational processes are embarrassingly parallel. I want to focus on the type of work where a computer like Pleiades is essential.

Slide 2: Use Case: Mission Need
Supports NASA's nuclear power problem. We do first principles calculations for highly complex materials that have a very large number of atoms. The calculations can take weeks to months running on several hundred cores on Pleiades.
We submit a run with a input file, and it runs as long as it can, given the rules. The output has to be transferred to JPL; it is a few gigabytes.

Slide 3: Use Case: Current HEC Application Characteristics
The restart files are large (hundreds of gigabytes) and come from hundreds of runs that we do on Pleiades.
The aim is better radioisotope generators.

Slide 4: HEC Needs Assessment Workshop Key Questions
We many run a different set of codes in the future. Runs on a single server would take a year or more, which is not practical.

Q&A:

Q from Michael Little (Int) to Everyone: 2:50 PM
NASA advertises that the RPG program is being performed in conjunction with Dept of Energy. Do they actually work with you or only provide oversight at a high level?

Paul von Allmen:
It must be at a high level, because I have no contact with anyone at DOE. We get our grant directly from NASA. DOE does have high resources available.

Q from Brian Thomas (Int) to Everyone: 2:54 PM
What would lessen the need to transfer large files from JPL to Pleiades? Backed up storage? Moving some processing to Pleiades?

Paul von Allmen:
Transfer is not a huge problem, unless we transfer them from JPL to Pleiades to JPL. It would be nice to have larger storage and a homesite at Pleiades; it would make it unnecessary to move a lot of data from JPL to Pleiades.

Q from Brian Thomas (Int) to Everyone: 2:56 PM
Can you quantify much larger storage?

Q from Piyush Mehrotra (Int) to Everyone: 2:56 PM
How much storage is needed?
Paul von Allmen: The technical point of contact is Dr. Trinh Vo; she can give exact numbers. It is in the hundreds of gigabytes.

---

6. Use Case: Use Case: Deep Learning for Geostationary Data

Thomas Vandal
thomas.vandal@nasa.gov
AMES/BAERI

Main points:
P1: HEC availability of GPUs is critical. Analytic functions. Overall go to speed up processing on high-D data sets. Without HEC we would not be working on deep learning. Use V100s for training.

P2: We try to keep data and computing close together. Multiple GPUs used to train models: much faster. One constraint how fast we can get models onto GPUs.

Raw notes:
Slide 2: Use Case: Mission Need
We are currently focusing on geostationary satellites, but there are applications for low earth orbit satellites as well.

We use GPUs constantly.

Our overall goal is to speed up the processing of high-dimensional datasets.

We do this pipeline; it runs directly from NEX data to GPUs to our training so we can do evaluation of data on NEX.

Slide 3: Use Case: Current HEC Application Characteristics

The current application I am looking at is virtual sensing, predicting a synthetic image. The majority of the work is generating these training datasets and is where we use a lot of CPU time. We are testing a bunch of parameterizations. After we select the model, we can use it to perform inference on our datasets. You usually want to do inference on the GPUs. We mostly use the V100s for training and the K40s for inference (K40s are not sufficient for training).

One model per GPU: analyze more parameterizations.

Multiple GPUs on one model: train model faster.

How fast can we get the data onto the GPU? Spinning disk does not take advantage of the GPU. Flash disk does.

Slide 4: HEC Needs Assessment Workshop Key Questions
I have many combinations of models to train and test.

If there are more than a couple of users doing deep learning experiments, we are going to get pretty clogged up.

Cloud computing: Our last project would have cost about $10,000 on AWS.

Q&A

Q from Tsengdar Lee (Int) to Everyone: 3:07 PM
Can you train one model across multiple GPUs across multiple nodes? I am thinking of a typical Summit nodes with 2-4 GPUs per node but there are thousands of nodes.

Thomas Vandal:
Yes. That is currently supported by PyTorch and TensorFlow. NVIDIA drivers need to be compiled with some sort of MPI flags.

Q from William Thigpen (Int) to Everyone: 3:07 PM
What would an optimum number of GPUs be? Are there any preferences on the number of GPUs per node?

Q from Brian Thomas (Int) to Everyone: 3:09 PM
Do you limit your search of hyperparameters and/or algorithms/models because of GPU constraints? (Sounds like you are worried that more folks using GPUs will do so in future as well.)

Thomas Vandal:
Yes, but there are ways around this. Say you have two hyperparameters, if I add a third, you end up with 2^3, and you end up with many more possibilities. These hyperparameters may correspond to anything from learning rates, regularization terms, and weighting factors. In the future, I believe more users will be doing deep learning at NAS and these sort of experiments could be common practice.

Q from William Thigpen (Int) to Everyone: 3:09 PM
There is an order in process that will raise the number of 4-GPU nodes to 48.

Thomas Vandal:
I don't have a big preference between 4 and 8 GPUs. Four seems to be sensible. When you get to eight, the question is how fast can you get the data from the node to the GPU? Too much time for data transfer. As applications continue to become accelerated with GPUs, they may be needed on every node.

Q from Brian Thomas (Int) to Everyone: 3:09 PM

Are there ML services (which exist in AWS/GCP) that you wish existed for Pleiades?

Thomas Vandal:
I don't think so, not for my personal use case. Maybe for other users need a node or very simply environment with everything there. Anaconda works reasonably well, and I use that on the cloud too.

Q from ANDREW MICHAELIS (Int) to Everyone: 3:10 PM
To add to Bill's question, host memory per GPU, memory onboard GPU? What's good for what you are doing specifically? Just curious...

Thomas Vandal:
Memory per GPU, what that ends up meaning is how big a network do you want to put on the GPU? The activations within the GPU blows up memory. A full GOES image with 16 bands would blow up the memory instantly. There is a high memory requirement, but all these things can be worked around.

[Questions to be addressed later]

Q from Daniel Kokron (Int) to Everyone: 3:12 PM
How much storage would you need on node to hold the model and training data?

Thomas Vandal:
Training data is typically on the order of terabytes and models between 0.5-5 GB.

Q from Robert Ciotti (Int) to Everyone: 3:16 PM
Users can build their own containers. **[NOTE: This may be in response to an earlier question. Bob, if you remember which one, can you just move this text to it as a response? PW]**

Q from Daniel Kokron (Int) to everyone:
Thomas, can you quantify if or how much Lustre is impacting your ability to train. Put another way, how much would you benefit from high-speed on-node storage?

Thomas Vandal:
Tested with 2GB of training data on node and Lustre.
Data: Batch size of 8 per GPU of arrays of size 64x64x16
Database reading from Apache Parquet through the Petastorm library
8 multi-processing threads for data loading

1-GPU On-node (/tmp directory): 27.3 samples/second
1-GPU Lustre (/nobackupp10 directory): 27.0 samples/second

(Copy model onto each GPU)
4-GPU On-node (/tmp directory): 41.6 samples/second
4-GPU Lustre (/nobackupp10 directory): 41.3 samples/second

Found very little speed up by training on node for single and multi-gpu setups.

---

7. Use Case: HECC Project, NASA Advanced Supercomputing, NASA Ames Research Center

Robert Ciotti

bob.ciotti@nasa.gov
NAS/ARC
Directorate (if NASA)

Main points:
P1: Run systems team at NASA for quite awhile. Often meeting w Pis to see what they need and adapt system, a continuous evolution. Issues addressed: security, policy, how we manage system, something durable developed by us and given back to user cmty. Hopefully these new features or tools are used for diff projects. We want to develop a generalized capacity at as a low level. Match comml capabilities as much as possible without protection.

P2: SC becoming a wide geographic distribution of users, providers and consumers. Shift to data distribution. Common issues heard from presenters (data processing, distribution, preservation, interoperability) drive our development.

P3: Multi PB database, specialized internal systems, users want them avail on outside as well. Give users containers to build their own customized env't. Diversity of work is too much for us to manager. Cloud has that advantage, although there are security issues there in container envts.

P4: Integration between us and DAACs and cloud.

P5: Use Case: 5 PN GEOS/ECCO Global ocean model data: first Data Portal use case.

How we interface with teams (via messaging) helps, we are integrating web technologies to help partitioning of work, integrate with scheduling system. Nodes are pulling work and code from containers.

Raw notes:
Runs the systems team at NAS.

Slide 2: HECC Computational Landscape
Run 24x7x365 but also continuously evolving the systems. Create a system that is durable for the user community. Want to develop tools and features that can be used by multiple projects. Develop a generalized capability so that more people can use it.

Slide 3: Mission Concept – Systems Evolution
We are trying to adapt and change.
Over the past several years, we have seen a shift to creating these very large datasets that need analysis and distribution. System consumes 10 megawatts and 17,000 nodes.

Slide 4: Program/PI Challenges and HECC Development
The challenges drive the development that we do internally. Have machines for what they are doing. Example: TESS within supercomputing environment. We do hardware acquisition, maintenance, etc. to allow the scientists to focus on the work they are doing.
Common theme: These groups want the data to be available both inside and outside the environment. Indexing and databases: You typically want a database for metadata; we have specialized systems to do that.
We have a particular OS build that drives our local configuration. We want to allow our users to build their own customized environments within singularity containers.

Slide 5: Use Case #1 – GEOS/ECCO

They will want to have 50 petabytes of data available both internally and externally.

Slide 7: Use Case #2- HySDS with HEC on Pleiades
HySDS component runs in AWS. Built a front-end system where we interface with the components on AWS and pull the data into Pleiades. Hua Hook/JPL could not do this on his own. Web-centric technologies are things we are integrating to facilitate the partitioning of work that gets sent to Pleiades. Pulling work from AWS (containers and code).

Slide 8
Meta-function that enables a project or campaign to do the work they need to do. Pull data from different sites, do data analysis, etc. Take inputs (machine learning algorithm) to decide what is important.
HPC was viewed as a standalone capability. Now it is being integrated with cloud computing and other things out in the world.

Q&A

Brian Thomas: Arguably, cost is orthogonal to high-/low-priority for some jobs the cost will never work for users and on premise (Pleiades/HEC) will be required.

Robert Ciotti: Understanding what costs are is important for senior management, what are we spending in various activities, how to optimize that. **How do we share tremendous amount of data on cloud, instead of copying out of cloud, copy to next better cloud provider? How to inject competition with these cloud providers? In a changing environment, let's provide competition. Questions; how to we do cost containment and maximize opportunity?**

Brian Thomas: Agree on understanding the costs, should be transparent to end user so they can make informed decisions.

Group Discussion of Key Questions

Peter Williams: For each HEC Needs Assessment Workshop Key Questions**,** are there any key factors that have not been mentioned or that you want to underscore?

Key Question: What are the key factors today that affect the load on computing assets and cause it to vary?

Peter Tenenbaum?: Several different classes: bulk processing on continuous satellite streams, disaster-related activities. **How are you going to facilitate that when you have a spike in demand?** Have an infrastructure that allows you to pick between the two.

Jennifer Dungan: Pleiades usage is 95% because the number of users is high; there are no spikes. If all academic NEX users decided to do their work in summer, it would affect the load on NEX.

Heather Cronk: Underscore: messaging/notification systems, higher storage quotas -- not necessarily in just in bytes but also in number of files.

Paul von Allmen: Should we distinguish between the needs of large projects like ECCO, HySDS, NEX, ... and small research projects needing large computing resources but not having the budget to run on the cloud?

Brian Thomas: Key Factors: seems like many are looking for some data architecture support from HEC so that they can leverage standard services for identity management/authentication, data access, finding services, containers, etc. Support for machine-to-machine processing vs humans mediating as much as they do today.

Key Question: Are there key functional gaps that exist or are projected in the next 5 years or 10 years?

Robert Ciotti: It is tens of petabytes. The way satellites might collect exabytes. Computations might produce exabytes too. There is growing demand, and it is going to continue.

Laura Carriere: One underlying theme I've heard that is not hardware related is a growing need to support more complex workflows, often spanning multiple organizations. Within NASA, we need to support this in accordance with security policies and this makes us much less agile than commercial providers. We can solve these problems, but they take time and human resources. **Have presenters experienced this? If so, can we quantify loss to science because of the time it takes us to provide solutions?**

Brian Thomas: Key Gaps? Reusability of products of compute for downstream folks (models, processed data, containers which may be built upon, etc.)

Brian Thomas: Concur with Laura C on complex workflows.

Jon Jenkins: The Surface, Biology and Geology Directed Observables Mission will collect ~15 TB/day and generate ~75 TB/day of data products. Key areas include data transfer into/out of NAS and processing capability.

Key Question: How will this Use Case evolve over the next decade?

Jennifer Dungan: We haven't talked about how confidentiality is maintained now for HEC. External users find one-size-fits-all security to be too burdensome. Some say "Forget this" when confronted with what it takes to get an account.

Brian Thomas: guest.nasa.gov could be used for id management

Brian Thomas: We have a monolithic approach, and we don't have to. In the past, we had more is more. In the Big Data picture, size matters. In some areas, it is about the heterogeneity of the environment and the data interacting with that. How do we share, find, leverage resources? If somebody comes up with a great way to access data in the HEC environment, how do people find out about it?

Peter T.: BOBL NES does this: several diff classes: bulk reprocessing actty, going back to reprocess large amounts of data, emergency-specific disaster-related activities: run it internally or externally due to spike in demand? Need infrastructure that slows you to choose between the two?

Jennifer: The reason that Pleiades usages are high 95% # of users is high there are no spikes because it is always full. From a NEX standpoint, if all res did their res at same time, or use a lot of cycles, scheduling on top of each other, that could affect the load.

Jenkins: cue gets short here…

Peter W.: Key functional gaps that exist or projected, ones not covered or to underscore?

Brian Thomas: Agree on understanding the costs, should be transparent to end user so they can make informed decisions

Brian Thomas: Key Factors: seems like many are looking for some data architecture support from HEC so that they can leverage standard services for identity management/authentication, data access, finding services, containers, etc. Support for machine to machine processing vs humans mediating as much as they do today

Bob Ciotti: dozens of PB problem. Satellites & simulation may produce exabytes of data. How to distribute, archive it, move it around.

Heather Cronk: Agree with Brian's points about knowledge sharing -- this was extremely useful for me to see what has already been done and what we could leverage for GeoCarb and I'm not sure how I would have learned all this outside of this forum

Jennifer Dungan: I'll second Heather's comment. I was happy to learn about HySDS

Robert Ciotti: Data Portal capability is available to everybody. Several projects use that. It is pretty early days for that. Data Portal capability access to anyone on web, attach to portal, things that PIs tell us they want exported.

Piyush Mehrotra: Reflecting on Brian's comment., develop capabilities for specific missions, then make them available for other missions/HEC users.

Andrew Michaelis: Please consider reviewing something like Confluence (wiki), that would be available to all NAS/NASA users. A simple confluence wiki not refined would be helpful for other users. Low-cost mechanism to help other users, accessible within NASA network.

Brian Thomas: The evolution of computing can be seen as constantly developing a layer of abstraction over the last generation of compute. Higher levels of abstraction make it easier to use the compute. **What are the next levels of abstraction to give me the extra oomph that I need to carry out my project?**

Mike: Thanks, Peter, for ending us on time. Thanks to presenters and participants who shared a lot with us, important to understand how to evolve HEC systems to continue to be responsive to users' needs. Interesting discussion that came out of this. Presenters will be sent a note within 24 hours to react to notes to consolidate to session report. Please email updated slides to me, most current for report. Thanks!

Appendix D4 Workshop Session 6-4 Inputs

David Wagner	LaRC D309	Femera
James Warner	LaRC D309	Uncertainty Quantification for Expensive, High-Fidelity Models
Jeffrey Hill	ARC TSA	Aerothermal Analysis of Planetary Entry Systems
Mehdi R. Khorrami	LaRC D302	High-fidelity, System-level Airframe Noise Simulations for Civil Transports

National Aeronautics and Space Administration

Workshop Session 6-4

4 June 2020

National Aeronautics and Space Administration

Use Case: Femera

David Wagner
June 4th, 2020
Langley Research Center
Research Directorate,
HPC Incubator
Transformative Tools & Technologies

Use Case: High-performance Finite Elements for Elasticity

- Mission: Higher-fidelity elasticity simulation requires **new algorithms** (no Moore's Law)
- Impact: **HEC is needed for the expertise** required to develop HPC algorithms
  - 85% mini-app kernel speedup **+ continuing improvement + full app development**
- Criticality: HEC is required for large production jobs
  - LaRC mid-range K-cluster is intended for development and small production jobs
- Mission Security Needs: Minimal (mini-app is open source)
  - Academic collaborators: research and development (open science)
  - Industry and government collaborators: flight qualification guidelines (open standards)
- HEC Function: HPC collaborators
  - Finite element (FE) analysis is a routine mission need.
  - Results can be used to **qualify flight components**.
- Mission Concept: Additive manufacturing analysis (top right)
  - Optimization progress (left)
  - Comparison to traditional FE methods. (bottom right)

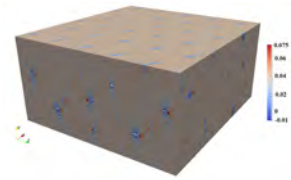

Lack-of-fusion pore model (26 MDOF), courtesy of Dr. Saikumar Yeratapally, NIA

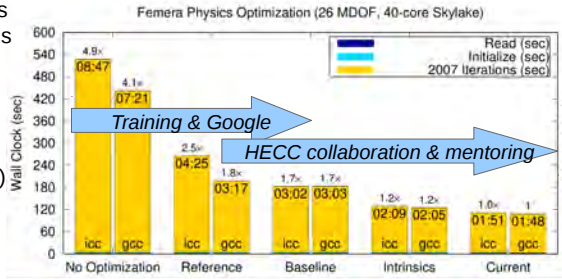

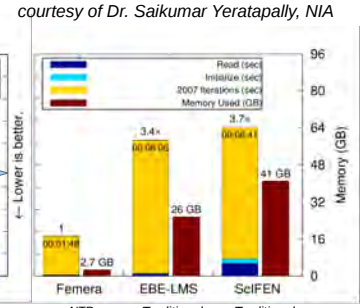

Optimized in collaboration with Daniel Kokron and Gabrielle Jost, NASA AMES High-end Computing Capability

NTR LAR-19531-1 | *Traditional Element-by-element* | *Traditional Sparse Matrix*

Use Case: Current HEC Application Characteristics

- Work flow: Mesh, partition, physics, solve, compare results
  - Femera **mini-app** is used for performance optimization
  - User interaction limited to a few simulation parameters
- Data:
  - In: Mini-app reads in parallel (<35 GB/model, Lustre)
  - Out: Timing data, residual, and correctness check
    - Data retained for open science publication
    - Retention also required for NASA Modeling and Simulation Standard compliance
  - Transfer:
    - Data archive location TBD.
- Computation: Femera mini-app is single-node
  - Capacity and special capability requirements expected.
- Evaluation of HECC: **Collaboration with HPC experts**
  - Performance optimization: Speedup = after / before
  - Mentoring: Continuing after-effort capability development and improvement

Isotropic Linear-elastic Quadratic Tetrahedra Per-core Performance

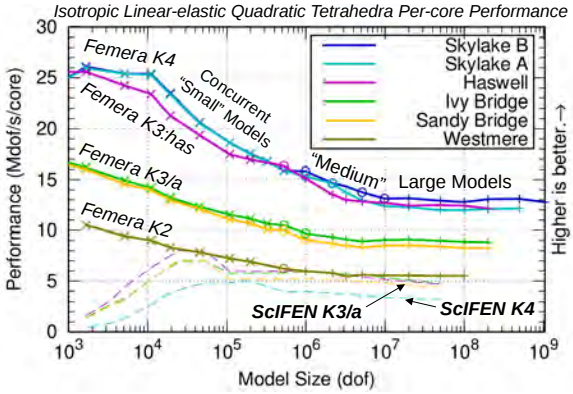

- *Traditional finite element application (ScIFEN) limited by main memory bandwidth*
- *Femera limited by cache bandwidth*

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Large-scale, non-linear, multi-scale, multi-physics, high-order models for: uncertainty quantification, optimization, inverse problems, digital twin, training data; **performance testing**, publication and conference cycles, project time lines (**milestones**), emergency **forensic analysis**.

- Are there key functional gaps, existing or projected in the next 5 years or 10 years?
  - Constraints on Use Case: **Long-term NASA support** for software capability development and maintenance.
  - Impact: Current **mission needs are scaled** proportional to available capabilities.
  - External Sources: Academic, government, and industrial partners are expected to develop and run this code.

- How will this Use Case evolve over the next decade?
  - Maturity: Femera (NTR LAR-19531-1) will move from a development state (mini-app) to an operational capability with continued development.
  - Growth: Femera will increase resolution, time-steps, multi-physics, and multi-model capabilities to enable evolving future mission needs.
  - Re-use: We expect the number of people running the Femera library and application to grow.
  - Re-hosting: Training, domain and **HPC expert collaborators**, high-performance supporting codes, and access to **rapid development cycle hardware** are needed to migrate code to future computing architectures.

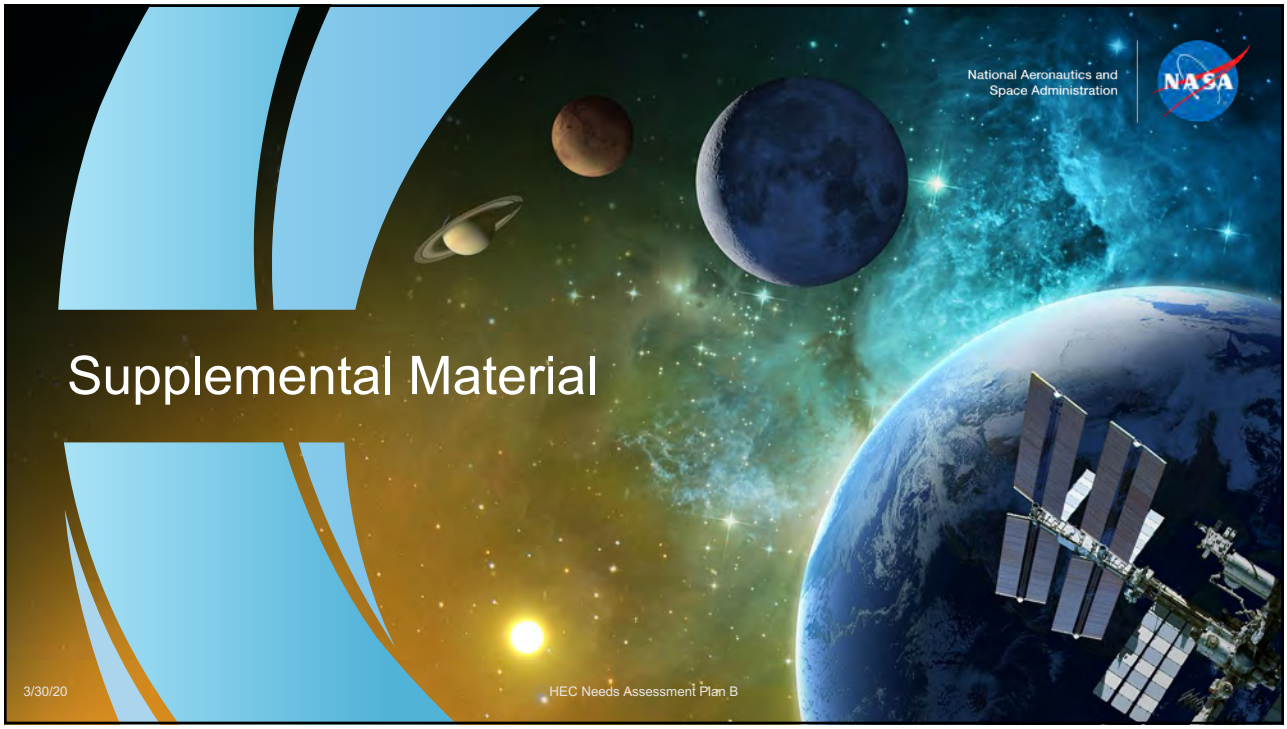

Current Use Case (Survey)

- Achieving the mission objective
  - Computation: tensor-based, matrix-free implementation of the standard (low-order) finite element method for elasticity
  - Describe the application
    - Antecedents: Mini-app kernels may be integrated into ScIFEN (NASA finite element application)
    - Programming languages:
      - Mini-app: C/C++11, shell scripts
      - Full app: C/C++11 (library, unit tests) Python 3.6 (bindings + unit, integration, performance tests)
    - Number of lines of code: Mini-app 16k (full app 1k–it's only a few weeks old)
    - Workflow mechanism used to control code execution: Python, shell script, command line
  - Computational and storage requirements
    - Size of runs: Mini-app is single node only
    - Accelerator usage: Proof-of-concept only
    - Amount of memory required: Varies with use, 100 Mbytes/MDOF double precision (<10% traditional FE application).
    - Parallelism of the mini-app: OpenMP, nested for concurrent "small" (<1 MDOF) model analyses.
    - Run in batch (Yes), interactive (Limited), or near-real time (No)

Current Use Case (Survey)

- Data requirements for the Femera mini-app
  - Data flow: Mesh and physical parameters in, computations done, timing and residuals out.
  - Data loaded in parallel at start of each model, one file per (over-decomposed) partition .
  - Input up to 30GB, outputs only timing data and residuals.
  - Performance profiling is eating up several terabytes of input per year (output negligible in comparison)
  - Output analysis: See charts in this presentation for examples.
  - Type of storage: high-speed shared file system
  - Input data requires some re-ordering for internal use.
  - Very little data needs to be curated (archived). The input meshes for performance tests are just cubes.
  - Performance data needs to be distributed to interested academic and research parties (open science).
- Key computing obstacles to effective CPU execution: mostly cache bandwidth; some routines compute-limited.
- Timing and performance analysis is done on the output.
  - Timing, computing load, workflow
- Do you have available the workforce necessary to re-host your code on any new platform necessary? No
  - Long-term NASA support is needed for development and maintenance of software capabilities.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - Access to new hardware for rapid-cycle development is expected to be the key obstacle.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity: comprehensive user plugin capability, multi-physics, hypercomplex algebra, on-demand in-memory (only) grid generation and manipulation, external libraries (partitioning, multigrid, etc.)
    - Increase in resolution and timesteps: Yes
    - Rewrite to improve performance: Yes, continually
    - Refactor for emerging platforms: Yes
    - Rehost onto cloud: Not expected
    - Introduction of AI/ML/DL components: Yes
      - Replace non-linear material response computation with a physics-informed neural network
      - Integrate with external uncertainty quantification applications (Python API)
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs: single-node (current mini-app), to hundreds of cores (existing models), to thousands+ cores (expected)
    - Accelerator usage: Identified as a requirement, but not yet scheduled for development
    - Amount of memory required: Varies with use, 100 Mbytes/MDOF double precision (<10% traditional FE application).
    - Application parallelism: OpenMP, nested for concurrent simulations (current mini-app) to MPI+OpenMP (in development)
    - Storage requirements: Varies with use
    - Run in batch (Yes), interactive (Limited), or near-real time (Yes)

Use Case: High-performance Finite Elements for Elasticity

- Mission: Higher-fidelity elasticity simulation requires **new algorithms** (no Moore's Law)
- Impact: **HEC is needed for the expertise** required to develop HPC algorithms
  - 85% mini-app kernel speedup **+ continuing improvement + full app development**
- Criticality: HEC is required for large production jobs
  - LaRC mid-range K-cluster is intended for development and small production jobs
- Mission Security Needs: Minimal (mini-app is open source)
  - Academic collaborators: research and development (open science)
  - Industry and government collaborators: flight qualification guidelines (open standards)
- HEC Function: HPC collaborators
  - Finite element (FE) analysis is a routine mission need.
  - Results can be used to **qualify flight components**.
- Mission Concept: Additive manufacturing analysis (top right)
  - Optimization progress (left)
  - Comparison to traditional FE methods. (bottom right)

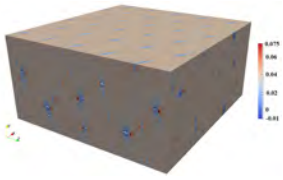

Lack-of-fusion pore model (26 MDOF), courtesy of Dr. Saikumar Yeratapally, NIA

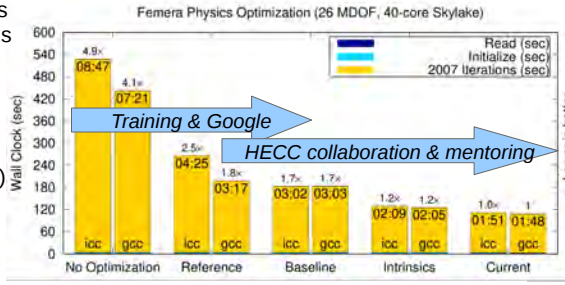

Optimized in collaboration with Daniel Kokron and Gabrielle Jost, NASA AMES High-end Computing Capability

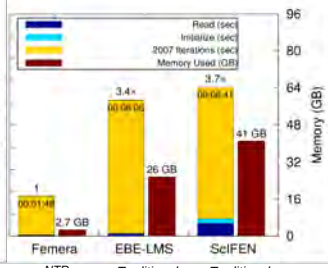

NTR LAR-19531-1 *Traditional Element-by-element* *Traditional Sparse Matrix*

National Aeronautics and Space Administration

Use Case: Uncertainty Quantification for Expensive, High-Fidelity Models

James Warner
June 4th, 2020
Langley Research Center
Durability, Damage Tolerance, & Reliability Branch D309

Use Case: Mission Need

- **Mission:** Uncertainty quantification (UQ) quantifies mission risk, decreases margins/costs, provides confidence to decision makers
  - Emerging paradigms like Digital Twin (DT) and Certification by Analysis (CBA) require rigorous UQ; Importance across disciplines
  - UQ with high-fidelity, physics-based models is computationally intensive. **HEC: More compute = more confidence with UQ**
- **Criticality:**
  - Performing high-fidelity UQ is **infeasible** without HEC since it's adding an outer loop to already intensive computational model
    1) Traditional (Monte Carlo simulation-based) approaches require $\gtrsim O(10^4)$ model evaluations
    2) Emerging data-driven (scientific machine learning) approaches require training of deep neural networks with large datasets
  - Without NASA supercomputing, this type of work would likely be turned over to partner (e.g., Sandia, Academia, …)
- **Mission Security Needs:** (Confidentiality, Integrity, Availability)
  - Interrupted HEC availability could result in unquantified/poorly quantified risks for missions, missed deadlines for research when analysis is particularly intensive (computing rare event probabilities, hyperparameter tuning for scientific ML)
- **HEC Function:** UQ on HEC is required for any mission with probabilistic requirements, CBA, reliability-based design optimization
- **Impact of HEC:** [See criticality]
- **Mission Concept:** UQ as an outer-loop application for a general physics-based simulation

Traditional Sampling-Based approach

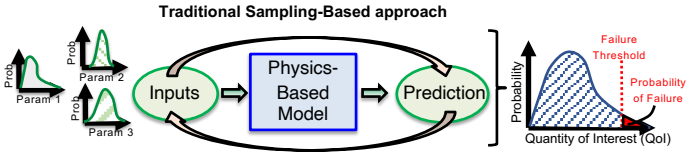

Scientific Machine Learning Approach

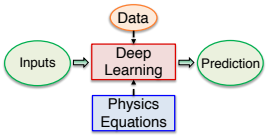

Use Case: Current HEC Application Characteristics

- **Work flow:**
  - Sample-based UQ methods require an outer loop: 1) draw random inputs, 2) evaluate model, 3) store outputs
  - Scientific machine learning requires solution of an expensive, nonconvex optimization problem to train model
    - User must often do manual hyperparameter tuning through trial-and-error for accurate UQ
- **Data:**
  - In: All required data for physics-based simulation + probabilistic info for inputs (e.g., random number generator)
  - Out: The outputs/QoIs from the model must be stored for each iteration of the outer loop (for each sample)
    - Ex. 1: a probability of failure analysis where the output is a Boolean for pass or fail (small amount of data)
    - Ex: 2: case where probabilistic information is requested for a spatially varying quantity at all nodes in mesh (potentially huge amount of data)
  - Transfer:
    - Visualize remotely on HEC system (e.g., Paraview) or transfer to local machine for small output data
- **Computation:**
  - Sample-based UQ: batch evaluations of many moderately-sized models with distributed memory parallelism
  - Scientific machine learning UQ: GPUs for training deep learning models
- **Evaluation:**
  - Overall, both approaches to UQ are highly amenable to HEC. More resources generally means faster, more accurate analysis. Sample-based UQ depends mainly on the efficiency of the model used

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - Sample-based UQ run times depend primarily on the complexity of underlying model used
    - The number of model evaluations required for UQ increase substantially for rare event calculations (e.g., probability of failure)
  - Scientific machine learning run time depends on amount of training data and complexity of governing physics.
    - Training of deep learning networks constrained by physics is difficult, requiring trial-and-error approach to tune hyperparameters
- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - Constraints on Use Case: Top-end GPU hardware for training deep learning is in high demand; can lead to lengthier queue times, longer turn around time.
  - Impact: Future missions that push performance envelope with narrower margins will need high-fidelity UQ
  - External Sources: None currently.
- **How will this Use Case evolve over the next decade?**
  - Maturity: Scientific ML methods will continue to mature/develop; Sample-based UQ is currently relatively mature
  - Growth: Growing compute power/resources means more confidence in modeling & simulation
  - Re-use: Prevalence of UQ and data-driven methods is expected to increase with time
  - Re-hosting: Viewing UQ as an outer loop application – migration to new hardware will be dependent on the underlying models themselves being migrated

Information about the Survey

- The following slides describe the kind of data that the Needs Assessment will collect via a survey
- This material is needed to supplement the description of the Use Case presented to the Affinity Group but will not be presented.

Current Use Case (Survey)

- Achieving the mission objective
  - Sample-based UQ:
    - Typically applications are developed in Python programming language
    - Python controls code execution; manages the evaluation of the computational model in outer loop
    - Resource/storage requirements highly dependent on the computational model used; but output storage can be huge when the number of samples is large and the quantity of interest is a spatially varying parameter defined at all nodes in a computational mesh (storage/memory O(# nodes x # samples))
    - MPI used at the outer loop level (running multiple model evaluations in parallel) while each single model evaluation can be running in parallel (shared memory)
    - Typically run in batch mode. Perform all required model evaluations for UQ, then aggregate/postprocess results in Python driver to obtain estimates of statistics, probability distributions, probability of failure, etc.
  - Scientific machine learning UQ:
    - Implemented in Python using Tensorflow2 deep learning library; O(1000) lines of code
    - Main computation is the training of deep neural network; typically performed on GPUs (V100s, K40s)
    - Thus far limited to single GPUs but extension to multiple GPUs is possible (and likely in future)
    - Often multiple trials are run in batch on several GPUs during hyperparameter tuning (different set of parameters per GPU)
    - Debugging often done interactively

Current Use Case (Survey)

- Describe the data requirements for the application
  - Data is typically loaded in serial and then distributed across processors. For sample-based UQ, the is parallel execution of multiple physics-based models (e.g., finite element analysis) that each have different input parameters.
  - The amount of input data depends on the nature of the uncertainty in the computational model – if the uncertain parameters are mesh dependent, then each model executed by the UQ outer loop will have its own computational mesh and as such the total amount of input data can be quite large (size of one computational grid x the number of samples needed for UQ)
  - The amount of output data depends on the UQ analysis, but again can be quite large. If probabilistic information is requested for an entire spatially varying quantity defined at all nodes of a computational grid, then there total amount of output data will be O(# nodes x # samples).
  - The output for UQ gets analyzed by computing statistics, histograms, probability of failure, etc. across the samples. If the quantity of interest is spatially varying, then the computed mean/variance fields can be visualized with Paraview.
  - Typically, only the computed statistics / postprocessed output data is archived and distributed. So the larger raw output sample dataset usually does not need to be stored long term.
- What are the key computing obstacles to effective execution of the current application?
  - For sample based UQ, the coordination of parallel model execution from Python/PBS is key for effective execution. Care must be taken, for example, to ensure consistency among MPI communicators when the models being executed in parallel are running in parallel themselves with MPI.
- Are there any analytic functions performed on the output of the application?
  - The models used with sample-based UQ can be analyzed/profiled independently of the UQ driver code itself
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Sample-based UQ is application/model specific. We need to work with domain experts to integrate their models into a UQ outer loop

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - As computing power continues to increase, the calculation of expensive UQ quantities (rare events, failure probabilities) for high-fidelity, physics-based models will be feasible without having to resort to less accurate surrogate models
    - An increase in resolution for UQ analysis means higher confidence / more precision in predicted quantities
    - The evolution of sample-based UQ applications that act as outer loops to computational models will depend on the migration of those models to new platforms themselves.
    - We expect a continued evolution of scientific machine learning approaches to UQ that capitalize on advancements in deep learning and GPU hardware; possibly gaining equal footing to more traditional sample-based UQ methods
  - If possible, describe how the computational requirements will evolve over time
    - As computational models grow in complexity over time, UQ algorithms will need to improve to keep pace. Approaches to increase convergence (get similar precision with fewer model evaluations), multifidelity methods that combine predictions from low fidelity models with the original high-fidelity model, and scientific machine learning approaches will be important to make UQ feasible for expensive models.
    - Sample-based UQ methods will continue to rely on MPI, distribution of individual model evaluations across processors. Scientific machine learning approaches will rely on GPU hardware for training the underlying deep learning networks.

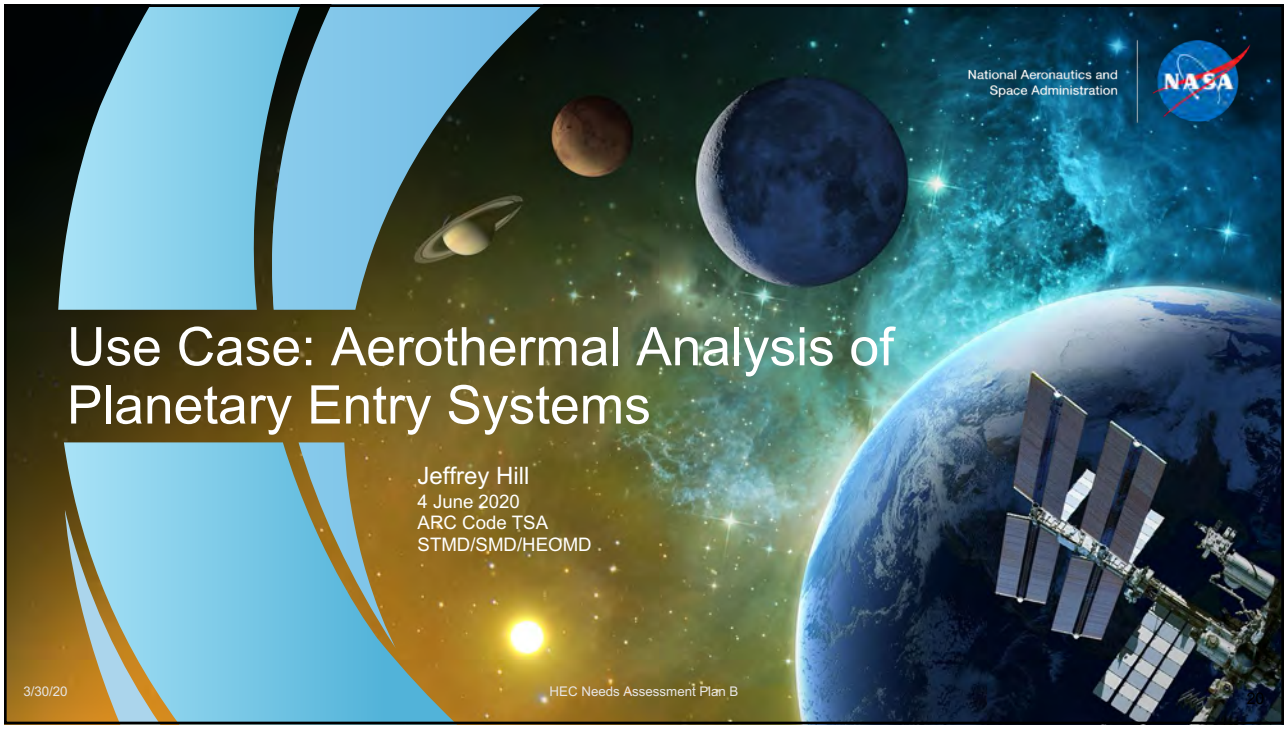

Use Case: Mission Need

- **Mission:**
  - NASA's planetary entry systems (Orion, MSL, etc.) encounter extreme heating when entering the atmosphere.
  - Aerothermodynamics is the estimation of this heat transfer so suitable thermal protection systems may be designed.
  - Primary tool is computational fluid dynamic (CFD) simulation, specialized to address thermochemical non-equilibrium.

- **Criticality:**
  - Aerothermodynamic CFD is **mission-enabling** for human space exploration and many planetary space missions.
  - If NASA supercomputing was no longer available:
    - Large projects (Orion, M2020) would likely buy and maintain their own compute resources.
    - Small project (Discovery, New Frontiers) would seek to acquire resources to meet mission analysis needs.
    - Code TSA would purchase and maintain for internal R&D and support of small missions.

- **Mission Security Needs:**
  - Confidentiality: Inadvertent release has National Security implication (EAR99, ITAR, SBU).
  - Integrity: Calculations are expensive (compute and labor). High-integrity, long-term storage is extremely valuable.
  - Availability: Typically not critical path; short outages (< ~2 days) can be accommodated ~monthly.
  - Access: Typically all collaborators will have active NASA affiliation and credentials.

Use Case: Mission Need

- **HEC Function:**
  - All simulation work performed on the Pleiades Supercomputer.
  - Code TSA consist of ~20 full time engineers, all of which are running multiple cases per week on NAS.
  - Simulation results compiled into data packages passed to adjacent engineering disciplines (TPS, Structures, Flight Mechanics)

- **Impact of HEC:**
  - HEC is a critical resource that enhances mission outcomes by reducing non-recurring engineering costs.
  - Aerothermal CFD simulation would be significant program cost if required to purchase time or systems.
  - Smaller programs would likely carry extra TPS mass to mitigate risks due to reduced analysis.

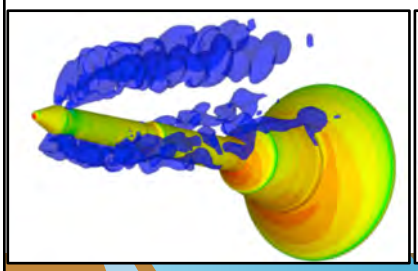
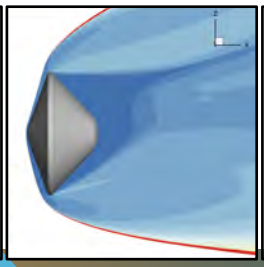
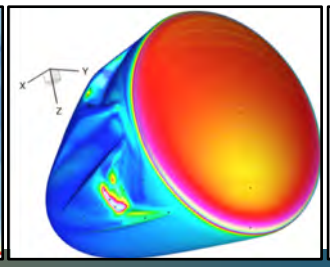
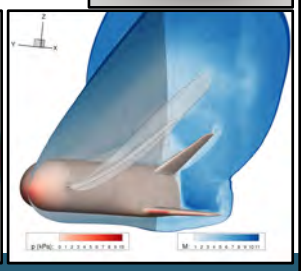

Use Case: Current HEC Application Characteristics

- **Work flow:**
  - Project defines entry conditions (airspeed, entry angle) and candidate vehicle design.
  - Critical point along the entry trajectory are identified, pre-processed, submitted for analysis on NAS.
  - Analysts reduces simulation results to engineering quantities and transfers to project for further engineering.
  - *Is there a need for interaction by user, operator or other system in the computing workflow? Yes*
    - *Some limited user interaction require to assess simulation progress and make corrections.*
  - *Does this use case need quick reaction or near-real time computing? Generally, no.*

- **Data:**
  - In: Simulation grid on which the CFD equations will be solved, 100MB-1GB per simulation.
  - Out: Flow variables (densities, velocities, temperatures) for each grid point, 1GB-10GB per simulation.
  - Transfer:
    - Engineering data extracted from simulations on NAS and scp'd to workstations for further processing.
    - Reduced engineering data typically archived in project-specific data servers.
    - Typically, O(100) output simulation retained on Lou during engineering phase (3-5 years).
    - After that, simulation outputs deleted, with inputs (grids, config, scripts) saved for future re-use.

Use Case: Current HEC Application Characteristics

- **Computation:**
  - Massively parallel solution of non-linear PDE by domain decomposition.
  - Multiple Program, Multiple Data (MPMD) model for parallel programming.
  - Currently we only target traditional CPU architectures.
  - Primarily capacity computing, but application is sensitive to inter-node communication performance.

- **Evaluation:**
  - Queue Wait Times: lower is better for rapid iteration and debug of simulation issues.
  - Case throughput: Faster turn around times improved engineering effectiveness.

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - Program Needs: Orion says run cases, we run cases.
  - Publication Cycle: Primarily target AIAA Scitech (Jan) and AIAA Aviation (June).
  - Summer interns.

- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - *Constraints on Use Case:* **Current resources do not allow large, chemically reacting flow fields to be simulated in a time accurate fashion on engineering-relevant timescales.**
  - Impact: Some engineering (dynamic stability, retro-propulsion) relies on experiment or low-fidelity modeling.
  - External Sources: No

- **How will this Use Case evolve over the next decade?**
  - Maturity: Capability is mature; no major changes expected.
  - Growth: Anticipate strong increase in demand for time-accurate, high-resolution, and multi-physics simulation.
  - Re-use: Expect user base to remain stable or slightly increase in the next 10 years (Artemis, Moon-to-Mars)
  - Re-hosting: **Currently taking steps to re-architect software to be GPU-capable.**

Code TSA Simulation Tools

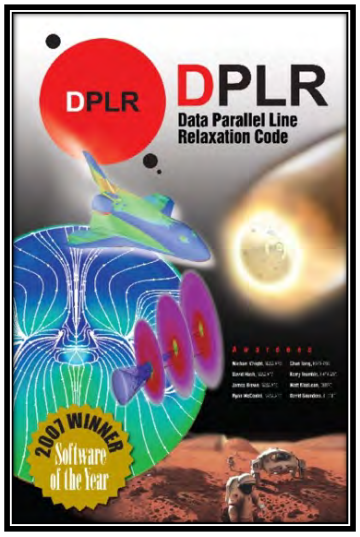

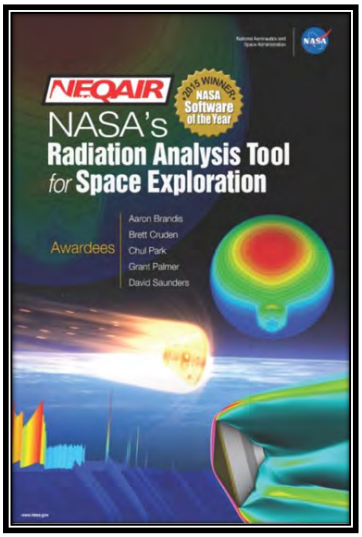

Background

- Core competency of ARC/TSA and LaRC/D305 is the estimation of heat transfer to planetary entry vehicles via computational simulation.
  - Convective Heating (reacting flow CFD)
  - Radiative Heating (radiation transport simulation)
- Workhorse tools for aerothermal CFD at NASA are LAURA (LaRC), DPLR (TSA), and US3D (TSA/U. Minnesota).
  - Second order accurate finite volume method
  - Euler backward / BDF2 time integration
  - Fully coupled finite rate chemistry
  - Multi-temperature thermal non-equilibrium
  - Semi-automated grid alignment to the bow shock

We must model a unique set of gas physics due to the extreme velocities associated with planetary entry, O(10) km/s.

Afterbody Calorimeter Placement

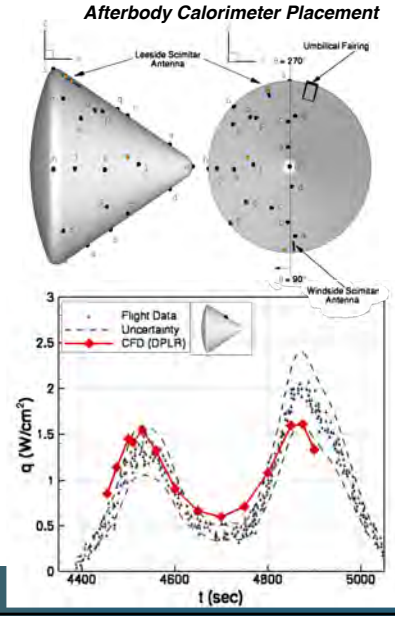

Core Method

- The core computational algorithm for aerothermal CFD is the repeated assembly and solve of the following block-sparse linear system:

$$\left(\frac{\boldsymbol{I}}{\Delta t} + \boldsymbol{A_i} - \boldsymbol{B_i} - \boldsymbol{C_i}\right)\delta \boldsymbol{u_i} + \underbrace{\sum_j (\boldsymbol{A_j} - \boldsymbol{B_j})\delta \boldsymbol{u_j}}_{\text{Face neighbor flux linearization}} = \boldsymbol{s_i} - \sum_j (\boldsymbol{f_j} - \boldsymbol{f_j^v})$$

Inv. Flux Linearization
Vis. Flux Linearization
Source Term Linearization

- Algorithm is broken into four independent kernels:
  - Assemble inviscid flux + linearization (i-, j-, k- faces as separate loops)
  - Assemble viscous flux + linearization (i-, j-, k- faces as separate loops)
  - Assemble source term + linearization (requires $\vec{\nabla} \cdot \vec{u}$ from viscous flux)
  - Solve linear system (line relaxation)
  - Update

Data Parallel Line Relaxation

- Line solvers are a critical technology for efficient aerothermal simulation at planetary entry conditions.
  - High cell aspect ratios common in BL and at the bow shock
  - Accurate wall gradients => stringent convergence criteria
- Current state-of-the-art is the Data Parallel Line Relaxation algorithm for which the DPLR code is named.
  - Iterative method to get approximate solution to <u>global</u> linear system.
  - Efficient if lines of strong coupling exists, e.g. normal directions through the boundary layer or the shock front.
  - Highly scalable on convention CPU+MPI+Infiniband systems.

Can DPLR be implemented efficiently on GPU?

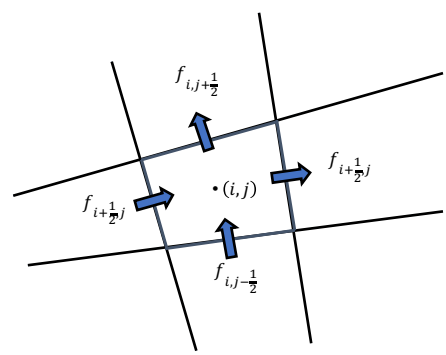

Given:

$$R_{ij} = f_{i+\frac{1}{2},j} - f_{i+\frac{1}{2},j} + f_{i,j+\frac{1}{2}} - f_{i,j-\frac{1}{2}}$$

Require:

$$\rho\left(\frac{\partial R_{ij}}{\partial U_{i,j\pm\frac{1}{2}}}\right) \gg \rho\left(\frac{\partial R_{ij}}{\partial U_{i\pm\frac{1}{2},j}}\right)$$

Data Parallel Line Relaxation

- Algorithm:

  Given a block-banded linear system $Ax = r$ (first-neighbor stencil on structured grid)

  Assign each cell to a line and order/partition $A = T + O$, such that:

  T is block tridiagonal containing flux Jacobians for tightly-coupled neighbors

  O is block banded containing flux Jacobians for loosely-coupled neighbors

  Perform in-place LU-factorization of T

  Initialize $x_0 = 0$

  Solve $x_i = T^{-1}(r - Ox_{i-1})$, which requires

  High-value work hides communication:
  - Block sparse matrix-vector product
  - Block banded forward elimination
  - Block banded back substitution

  Asynchronously communicate overlapping components of x_i

  Iterate ~5 times until convergence

Use Case: Mission Need

- **Mission:** Accurate prediction of airframe noise to help mitigate its impact on communities near airports
  - HEC is required to perform high-fidelity, system-level airframe noise simulations and to develop/evaluate noise reduction technologies prior to expensive flight tests
- **Criticality:** Essential
  - Unavailability of NASA supercomputing voids the mission unless time is rented on large commercial facilities or Cloud Computing is procured
- **Mission Security Needs:** Confidentiality, Integrity, Availability
  - Breach of proprietary data (SBU) clauses would result in professional embarrassment for self and agency
  - Interruption of HEC availability severely hinders timely achievement of project milestones
  - There is no need for external collaborators
- **HEC Function:** Use case depends entirely on NAS supercomputers and data visualization
  - Simulations are generated often
  - Output is used to determine airframe noise signatures and certification metrics
- **Impact of HEC:** See Criticality item above
- **Mission Concept:** Development of airframe noise reduction technologies for regional jet class (Fig. 1), large civil transports (Fig. 2), and supersonic designs (Fig. 3)

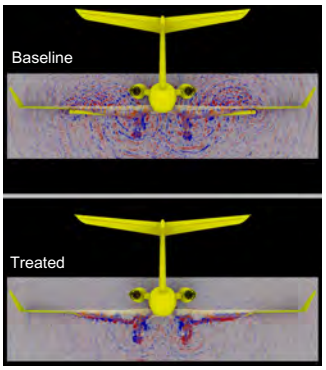

Fig. 1 G-III landing gear and flap noise

Fig. 3 Generic supersonic transport

Fig. 2 Surface pressure on full-scale B777-300ER

Use Case: Current HEC Application Characteristics

- **Work flow:** Discretization of computational volume, simulation execution, data extraction and post-processing
  - Case setup directs discretization, I/O, and simulation
  - Repeated user interaction caused by walltime limitations and crashes due to network and file system problems
  - Near-real time computing is not necessary
- **Data:**
  - In: simulation setup input (~10 GB), volume discretization (~0.5 TB)
  - Out: time-dependent aerodynamic data on surfaces and flow field
    - 2-3 large files, ~130 TB total, ideal output rate (IOR) ~1TB/h; 1500 small files, ~30 TB total, IOR ~0.5 TB/h
    - Retention is required for all data
    - Data post-processed in-cluster for visualization and generation of input for acoustic propagation code
  - Transfer**:** Output kept in local HEC storage until post-processing completion; extracted data transferred off-site
- **Computation:**
  - Pre-processing: small node-count, large shared memory (2+ TB)
  - Simulation: single control process (1 TB local memory), 4000-10000 cores/simulation, high-performance I/O
- **Evaluation:** HEC currently offers a unique capability to perform high-capacity simulations of unprecedented scale

HEC Needs Assessment Workshop Key Questions

- **Key factors today** that affect the load on computing assets
  - Use of large-scale, non linear model to perform various initial, potential debugging, and grid resolution studies
- **Key functional gaps today** or projected in the next 5-10 years
  - Constraints on Use Case:
    - Capacity/mission share that limits turn-around times, development, and debugging
    - Availability of high-memory nodes for control process
    - Lack of reliable, high bandwidth, and fast I/O
    - 5-day walltime limit
  - Impact: slow turn-around (up to 30 days in queue, simulation time several weeks), extremely slow debugging
  - External Sources: not required
- **Use Case evolution over the next decade**
  - Maturity: Already operational
  - Growth: increase spatial resolution and model complexity
  - Re-use: code users outside of case may increase, but their simulation sizes may not surpass current case
  - Re-hosting: commercial application requires support of software vendor

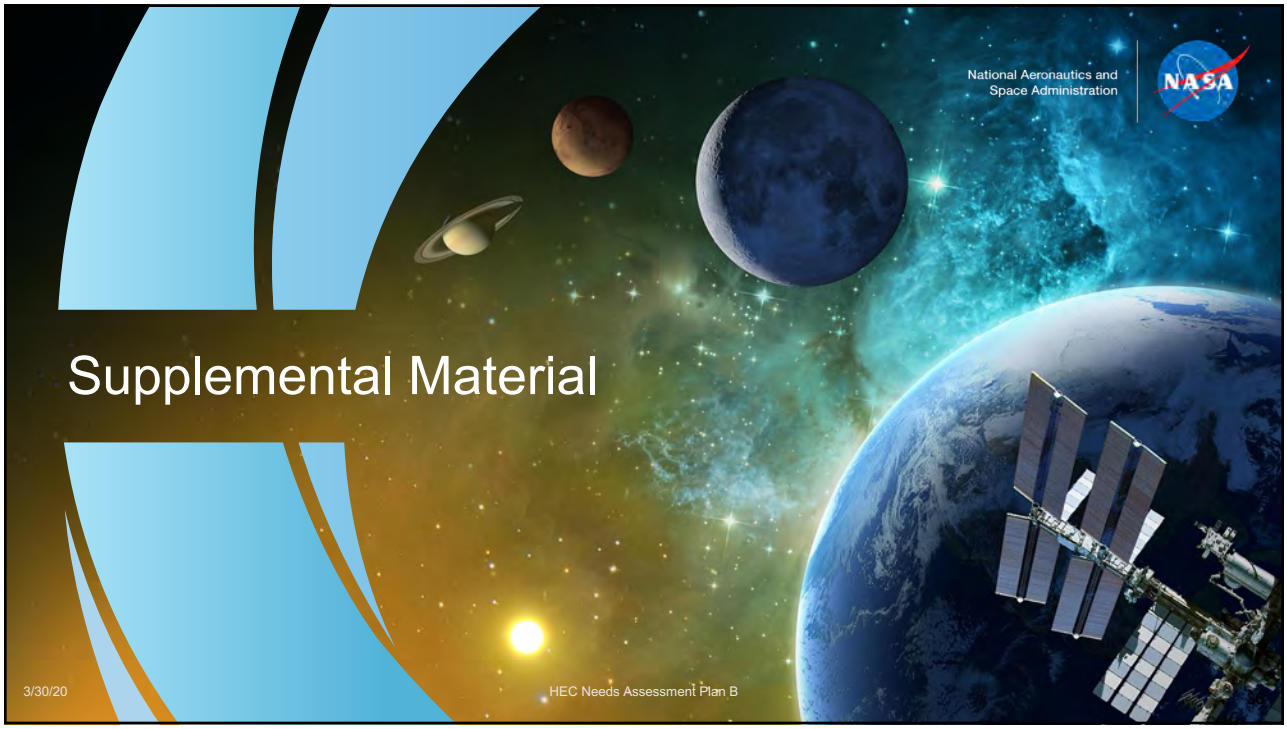

Current Use Case (Survey)

- **Achieving the mission objective**
  - Computational fluid dynamics (CFD) simulation to obtain acoustic response of flow from interaction with solid surfaces
  - Solver based on lattice-Boltzmann equation
    - Antecedents: none
    - Programming languages used: Not specified
    - Number of lines of code: Not specified
    - Workflow mechanism used to control code execution: shell script
  - Functional computational and storage requirements for the application
    - Size of runs: 4,000 – 10,000 cores
    - Accelerator usage: none
    - Amount of memory required: overall ~ 20 TB, ~1 TB for control process, 2-3 GB per simulation process
    - Parallelism of the application:
      - Discretization/pre-processing: shared memory, OpenMP
      - Simulation: MPI
    - Simulations run in batch mode

Current Use Case (Survey)

- **Data requirements for the application**
  - Describe the data flow of the application from input to output
    - Initial simulation setup: geometrical model representation and simulation characteristics
    - Discretization: generates simulation mesh following prescribed resolution distribution
    - Simulation
      - Reads discretized mesh (+ potentially previous solution for initialization)
      - Diagnostic information during initial transient phase
      - After initial transient, output requested measurement data:
        - High-frequency, high-resolution, short-signal transient data surface/volume for visualization
        - Time-averaged, long-signal data on model surface and volume for validation/verification
        - High-frequency, long-signal surface data for acoustic post-processing
    - Post-processing
      - Visualization data → create animations
      - Validation/verification data → extract low-fidelity data to compare to experimental measurements
      - Acoustic data → use acoustic post-processing techniques to extract synthetic microphone data

Current Use Case (Survey)

- **Data requirements for the application**
  - Data loading into application: serial, all at start
  - Input and output generated per run
    - Input: 0.5 – 2 TB
    - Output: 150 TB
  - Total amount of output data per year: 500 – 1,000 TB
  - Output data analysis
    - Visualization
    - Far-field propagation via acoustic integral method (Ffawcs-Williams and Hawkings)
    - Direct solution interrogation and plotting
  - Storage needed by application
    - High-speed local serial storage (parallel I/O not implemented)
  - Input data do not require any pre-treatment
  - Archival data requirements: 20 – 150 TB per simulation
  - Data distribution requirements: 10 – 50 TB to NAS visualization group

Current Use Case (Survey)

- **Key computing obstacles to effective execution**
  - Availability of compute resources (node availability, mission CPU share)
  - Network/file system stability for large-scale, long-running applications
  - Availability of shared-memory node for discretization step
  - Availability of high-memory node for control process of simulation
  - Serial file I/O performance
  - Available disk storage
- **Analytic functions performed on output**
  - Regular assessment of actual vs. theoretical performance
- **Workforce to re-host code**
  - Workforce available for re-hosting to standard industrial platforms only (commercial code)

Future of the Current Use Case (Survey)

- **Ability to predict total aircraft noise for landing and takeoff operations of any civil transport**
  - Goal achievable in the mid-term (5 – 10 years) if the following conditions are met:
    - Sufficient computing capacity and robust file system with large I/O bandwidth
    - Software improvements
      - Parallel I/O
      - Ability to handle an extremely large number of mesh cells (>50 B)
      - Reduce memory requirements of head node
  - Application and workflow evolution
    - Increase in scope and complexity by considering airframe and propulsion components to predict total aircraft noise
    - Increased spatio-temporal resolution to extend noise prediction to higher frequencies
    - Rewrite to improve performance: dependent on software vendor's cost/benefit analysis
    - Refactor for emerging platforms: dependent on software vendor's cost/benefit analysis
    - Rehost onto cloud: dependent on software vendor and possible if data security and availability can be assured

Future of the Current Use Case (Survey)

- **Computing needs of the application as it evolves over the next 5-10 years**
  - Estimated increase in computer requirements
    - Jobs running routinely on > 50,000 cores
    - Amount of memory required: 2 to 3 times current Use Case needs
    - Parallelism of the application:
      - Discretization/pre-processing: shared memory, OpenMP
      - Simulation: MPI
    - Storage requirements: 2 to 3 times current Use Case needs
    - Simulations run in batch mode

HEC Needs Assessment Workshop

Affinity Group Sessions

Thursday 6/4/20 1pm – 4pm (AG6-4) (nominal affinity: Computational Materials and Structures)
(10am – 1pm Pacific)
Facilitator: Chris Dwan <chris@dwan.org>
Scribes: Michelle Moyer <michelle.c.moyer@nasa.gov> and Emily Kuhse <emily.kuhse@nasa.gov>

Attendees

1. Mike Little, Convener, NASA Goddard Space Flight Center, Earth Science Technology Office
2. Chris Dwan (facilitator)
3. Presenter #1: David Wagner, NASA Langley Research Center, D309
4. Presenter #2: James Warner, NASA Langley Research Center, D309
5. Presenter #3: Jeffrey Hill, NASA Ames Research Center, Aerothermodynamics Branch
6. Presenter #4: Mehdi Khorrami, NASA Langley Research Center, D302
7. Angela Butcher (part of session)
8. Arsi Vaziri, NASA Ames Research Center, NAS Division
9. Blaise Hartman, NASA Ames Research Center, NAS Division
10. Brian Thomas, NASA Headquarters
11. Cetin Kiris, NASA Ames Research Center, NAS Division
12. Daniel Duffy, NASA Goddard Space Flight Center, NCCS
13. Daniel Hedin (NASA HQ/HEOMD)
14. Dan Kokron, NASA Ames Research Center, NAS Division (part of session)
15. Elizabeth Hook, NASA Headquarters, HEC Program
16. Emily Kuhse, NASA Ames Research Center, NAS Division
17. Henry Jin, NASA Ames Research Center, NAS Division
18. Irina Kitiashvilli, NASA Ames Research Center, NAS Division
19. Johnny Chang, NASA Ames Research Center, NAS Division
20. Karen Petraska, NASA Headquarters, OCIO
21. Leigh Ann Tanner, NASA Ames Research Center, NAS Division
22. Michelle Moyer, NASA Ames Research Center, NAS Division
23. Piyush Mehrotra, NASA Ames Research Center, NAS Division
24. Robert Ciotti, NASA Ames Research Center, NAS Division
25. Robert Ferraro, NASA Jet Propulsion Laboratory, Science Research and Concept Development
26. Robert Hood, NASA Ames Research Center, NAS Division
27. Sharad Gavali, NASA Ames Research Center, Advanced Computing Branch
28. Sherry Chang, NASA Ames Research Center, NAS Division
29. Subhash Saini, NASA Ames Research Center, Advanced Computing Branch
30. Ted Manning, NASA Ames Research Center, Advanced Computing Branch
31. William Thigpen, NASA Ames Research Center, NAS Division

Key Points

Mission Use Cases covered during HEC Need Assessment

Femera, Uncertainty Quantification (UQ), Aerothermal analyses of planetary reentry systems, Airframe Noise Simulations for Civil Transports

Impacts of HEC on Mission Needs: Range of Impacts and Significance

- Dependent: Higher-fidelity elasticity simulation requires new algorithms, as opposed to advancement of existing ones (Wagner)
- Dependent: Deep Learning and Machine Learning depend on large datasets upon which the training can occur (deep neural networks) and the speed upon which the learning relies (Warner)
- Dependent: aerothermodynamic simulations and modeling depends on HEC because the computational fluid dynamics necessary to achieve the simulations and models (Hill)

Mission Security Needs relevant to HEC: Range and Significance

Security needs discussed confirm the five areas mentioned in the Use Case outline (bullets) and adds the importance of stability (system instability caused by security threats compromises mission integrity) and protection of intellectual property (of particular concern in public-private partnerships)

- Confidentiality
- Integrity
- Availability: Varies, with some missions dependent and others less so; downtime increases risk of missing transient events (rare and significant);
- Public Confusion/Controversy
- National Security: Export control requirements could be compromised by an inadvertent leak
- Financial value (integrity?): Calculations and resulting models are valuable as investments and to private industry seeking competitive advantage (Hill)

Criticality of HEC to Mission Needs

- Note: Similar to “Impacts of HEC to Mission” (above)
- HEC can allow qualification of flight components prior to physical manufacturing
- Uncertainty quantification and digital twin for material development and testing (Wagner)
- Computational speed (FLOPS) and power (BYTES)
- Data storage capacity
- Allocation time (SBU)
- Financial consequences
- Mission Security
- Massively parallel and Embarrassingly parallel
- Risk reduction: manned missions (safety) and modeling uncertainty/instability
- Knowledge Growth: inability to mature/test breakthrough concepts, such as reduced fuel burn (Choudhari)
- Deep Learning and scientific machine learning require training of deep neural networks using large datasets, meaning HEC is necessary and essential (Warner)
- Aerothermodynamic computational fluid dynamics (CFD) is mission-enabling for human space exploration and planetary space missions (Hill)
- Supported missions would need to buy and maintain their own compute resources or acquire those on the open market (Hill)

Workflow Issues and Needs

- This section and the subsequent one about "key factors affecting load" might be integrated

Data Input, Output, and Transfer: Issues and Needs

- Possibly keep this section focused on processing speed (FLOPS) and data volume (BYTES) as a specific look at these two elements of workflow issues or needs
- *Input, output, and transfer* seem generally associated with--if not dependent upon--computation and storage. Could combine this section with the following section or explain why these are distinctly important given the purposes of this needs assessment (The idea is that either lumping or splitting are fine as long as understandably explained based on meaningful considerations)
- *Computation and Storage:* A principle aspect of this challenge is that output from some work becomes input for other work, thus storage between runs becomes central to workflow. Storage on or as part of HEC system has costs, as does moving off/on (transfer) because the volume of data is significant.

Evaluation: HEC Effectiveness and Value Measures

- Might roll this category under the earlier "Impacts" section [or cross-walk impacts and evaluation to ensure each reinforces the other – evaluate based on critical impacts]
- Performance of computational efficiency, available disk storage, time in queue, work of the support team, code/compiler/system stability,
- Tools for the data analysis, visualization and expertise support
- Number of observation trades/strategies/architecture possible to evaluate in near-real-time (NOS-LeMoigne)
- Researcher demand and impact of resulting science (Owen)
- Computational power, including FLOPS and Bytes, sufficient for desired confidence in models and simulations (Warner)
- Volume of high-capacity simulations (Khorrami)

Key Factors Affecting Current Load on Computing Assets

- Publication cycles and University calendars
- Revisions to major models
- Focal area of selected projects (temporal and geographic scales up to planet)
- Reprocessing of science data (preprocessing?)
- Scalability: Effect(s) of adding nodes, increasing amount of data to process
- Availability of adequate cores to complete job(s) in desired timeframes (Jacobson)
- Physics requiring small step times combined with low frequency structures requiring long simulations (tension?) (Jacobson)
- Practical considerations, like availability and queuing, drive decisions about data resolution and sampling frequency in models, as opposed to research needs or project needs (Chouhdari)
    - Data transfer speeds are increasingly a major constraint
- Run times depend on complexity of the underlying model(s) (Warner)
- Needs of supported missions, like Orion (Hill)
- De-bugging (Khorrami)

Anticipated Evolution of Use Case and Mission Need

- Maturity: Eddy resolving Computational Fluid Dynamics and nonlinear structures will mature; scientific machine learning is likely to see significant evolution/maturity
- Growth: Increases in time-step frequency and spatial resolution, data assimilation, ensembles (multi-model, multi-parameter, seasonal forecasts, climate scenarios), model coupling, AI/ML and calibration needed for training, CFD expects smaller timesteps and more of those, NISAR launch likely to increase demand for InSAR analysis runs.
- Re-Use: Advances in CFD likely to attract more users and decrease the required expertise; CFD likely to become more common as part of Multi-Discipline Analysis (MDAO)
- Re-Hosting: Anticipate need for access to emerging hardware to test and develop application; challenge of migrating models, not just availability of new hardware, shouldn't be overlooked (Warner); re-architecture of software to work on G
- Other(s): Onboard processing, integration of mission management, anomaly detection, and multiple S/C and multiple nodes, fast integration,

Functional Gaps: Existing and Anticipated

- Constraints: big data analysis (including ML), disk space/data storage, computational allocation, queue, network, I/O, restrictions on working with non-NASA credentialed partners in academia (Choudhari), emergency forensic analysis (Wagner), multi-discipline and multi-physics involved in flow-fields are constrained by capacity (speed and volume) of current resources (Hill)
    - Limited data storage capacity forces scientists to perform mundane tasks of shuffling data just to keep projects moving forward (Peters-Lidard)
    - "Extrapolations from the most resource intensive jobs we were able to run 5-10 years ago would paint a bleak picture in terms of meeting the widening gaps in the future." (Chouhdari)
- Communication speed between ground and S/C and S/C-S/C (spacecraft)
- Logarithmic relationship between HEC speed and mesh resolution for some modeling means an order of magnitude increase in HEC speed is needed to double resolution (Choudhari)
- Impact of Constraint or Gap: insufficient performance to obtain the required models and data analysis, mission failure, systems considered advanced today will be considered moderate in the near future as anticipated evolution of Use Cases and Mission Needs occur
- External Sources

Figure 3. Thermodynamics of reentry (Hill)

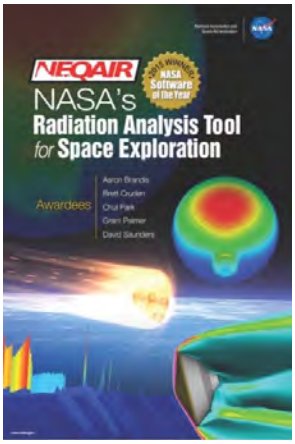

Figure 2. Simulation Tool example that relies on HEC (Hill)

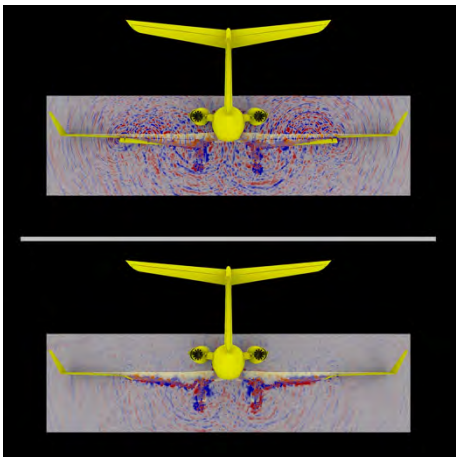

Figure 1. G-III landing gear and flap simulation (Khorrami)

1. Use Case: Femera

David Wagner
david.wagner-1@nasa.gov
LaRC, D309

Main points:
P1: Higher-fidelity elasticity simulation requires new algorithms (no more Moore's Law); HEC is needed for the expertise required to develop HPC algorithms. In collaboration with HECC, Femera was able to obtain a 90-95% speed up.

You have to have a team to do this kind of work. It requires domain knowledge to know what kinds of things you can do algebraically, as well as HPC expertise. Often these people don't talk to each other.

P2: New methods take a long time to be adopted. Most people using the existing methodology (traditional sparse matrix) just accept how much memory it requires and how slow it is.

P3: Project is only funded for three years. There is a lot of concern about what will happen after that, since it doesn't make sense for engineers to invest time in adopting a tool that may become unstable and is no longer supported. A commitment to a level of long-term support is needed; this is missing at NASA.

P4: Biggest spike in demand would be seen when/if there is an accident and forensic analysis is required.

Telework has also caused a spike in usage, especially on the local cluster at Langley; which is intended for development, not production runs. HEC resources are required for large production jobs.

P5: Needs can be quantified based on limit of the biggest problem we can feasibly solve. Doing concurrent model analysis is a big benefit.

Raw notes:

Slide 1: Use Case: High-performance Finite Elements for Elasticity
Femera: Open-source finite-element mini-app developed under developed under LaRC's High Performance Computing Incubator; Gabrielle Jost and Dan Kokron @ NAS provided significant help to get the core kernel performing, and also mentoring for refactoring into (new technologies?). New methods take a long time to be adopted.

Managed to get a code running which thought would be faster using a bunch of tiny little operations. Compilers don't like this, so needed help to get performance up to where it should be. Just about everyone uses traditional sparse matrix and just accept how much memory it takes and how slow it is.

Improved performance - 85% kernel speedup - to where it should be in comparison to SciFEN Sparse Matrix. They were able to get the code running faster than traditional finite element analysis.

General finite element analysis code for elasticity and thermo elasticity lags behind codes for fluid analysis.

Slide 2: Use Case: Current HEC Application Characteristics
Currently developing mini-app into full application.

SciFEN - seeing decreasing performance w/ newer processors; Femora doing better. Trying to scale up to a scale not seen before in FE w/ elasticity.

Slide 3: HEC Needs Assessment Workshop Key Questions

Forensic analysis – when there is an accident, will cause an immediate need for lots of compute to figure out what's going on.

Telework spiked usage; LaRC cluster had to limit size of runs; cluster not meant for production. New efforts may create immediate need for a lot of performance.

Project only funded for 3 years. Doesn't make sense to develop a full application only to abandon it.

Missing at NASA: Long term support for software capability and maintenance; getting support for three-year project; doesn't make sense to develop a full application only to abandon it.

Project will work with NAS experts (Gabrielle Jost and Dan Kokron) again.

Q&A:
from Chris Dwan (Ext) to Everyone: 10:16 AM

Do you mean Google as in research, or Google as in Google Cloud? Text?

A: I mean, for example, I get a weird compiler error and I Google it. Have difficulty accessing cloud resources. Slow getting the legal aspect in place.

from Brian Thomas (Int) to Everyone: 10:20 AM

You should consider putting your code in the agency source code repository, developer.nasa.gov

from Brian Thomas (Int) to Everyone: 10:21 AM

that would promote sharing, and hopefully, longer term support (via community) If you don't have funded support.

from Brian Thomas (Int) to Everyone: 10:21 AM

It is ITAR/EAR compliant

from Brian Thomas (Int) to Everyone: 10:21 AM
No, developer.nasa.gov is now capable of hosting internal, sensitive code.

from Jeffrey Hill (Int) to Everyone: 10:22 AM

If opensource, agency has a public GitHub as well.

A: We have a local GitLab and there is a NASA GitHub for open source; we released three days before govt shutdown so could not put it in the NASA github. Putting it somewhere will it remain is a good idea.

from Jeffrey Hill (Int) to Everyone: 10:24 AM

I think code.nasa.gov describes how to get your code there

from Piyush Mehrotra (Int) to Everyone: 10:24 AM

NAS can help you get access to AWS cloud if you want to run your benchmarks on it.

from Brian Thomas (Int) to Everyone: 10:25 AM

Yes, code.nasa.gov is the right resource

from Brian Thomas (Int) to Everyone: 10:25 AM

Basically, you will need to pass the SRA process.

from Chris Dwan (Ext) to Everyone: 10:24 AM

Question: Are there metrics that we can use to quantify the observation that we scale our ambitions to the available compute resources.

A: Getting metrics indirectly, 500k deg of freedom; doing concurrent model analysis is of benefit (missed this); want to go 10x in problem size (5M deg of freedom); solve bigger problems in same amount of time; main problem is memory limitations.

Can quantify based on limit of the biggest problem they can feasibly solve. Doing concurrent model analysis is a big benefit.

from Chris Dwan (Ext) to Everyone: 10:26 AM

Have you seen good examples of the sort of expert resource you're asking for in terms of HPC expertise?

A: You have to have a team to do this kind of work. Needs domain knowledge to know what kinds of things you can do algebraically to design new algorithms; need HPC expertise. Often these folks don't talk to each other very much.

Warner: Models exist in academia - universities have groups of research staff computer scientists to help develop code, to help professors and students.

from Piyush Mehrotra (Int) to Everyone: 10:28 AM

As I said during an earlier session, DOE with 10x NASA's budget for HPC has projects to optimize the whole software stack from applications to systems software. -Something we might aspire to.

from Brian Thomas (Int) to Everyone: 10:29 AM

there are data science groups around the agency which can help with this kind of computing software. ARC, LaRC both have teams like this

from Brian Thomas (Int) to Everyone: 10:29 AM

POC ARC : Nikunj Oza, LaRC: Ed Mclarney

the fluids branch at LaRC has not quite that level but more than just a lot of other branches where they don't have computing expertise and long-term software support.

2. Use Case: Uncertainty Quantification for Expensive, High-Fidelity Models

James Warner
james.e.warner@nasa.gov
LaRC, D309
Directorate (TBD)

Main points:
P1: Without HEC, this work would not be possible and would likely be handed over to a partner.

P2: For UQ, more compute gives decision makers more confidence.

P3: Outputs can be quite large and potentially huge datasets can be a problem. Both UQ approaches are highly amenable to HEC, and scientific machine learning UQ scales well on GPUs.

P4: There can be lengthy queue times. This impedes ability to run the application more freely.

P5: For rare-event calculations, the cost goes up significantly.

Raw notes: May be a unique case – going to talk more generally about UQ, not a particular app.

Slide 1: Use Case: Mission Need
Uncertainty Quantification (UQ) is needed to support emerging paradigms like Digital Twin and Certification by Analysis. UQ can be looked at as an outer-loop application (as described by Prof. Willcox in the plenary). More compute = more confidence in UQ. More confidence for decision makers.

Traditional (Monte Carlo) sample-based approaches: Many advanced UQ methods still require this process (outer loop); general physics-based model; run thousands of times to generate probability; critical use case is reliability or probability of failure; resolving rare events is more expensive (see first chart in slide)

Emerging data driven approaches: Use deep learning; integrate physics constraints into deep learning to do UQ. Embed physics directly into learning and training process, predictions are physically admissible; this method is gaining more traction; requires a different HPC approach (see second chart in slide)

Edge case UQ predictions even more expensive to solve.

Consequences of not having HEC: deadlines for research, ML hyper-parameter tuning is expensive. Work would not be possible and would likely be handed over to partner if no HEC. Sandia is active in the machine learning area, with lots of tools.

Slide 2: Mission Need
For sample-based approach, outputs can be quite large, multiplied by number of samples, and that can be a problem. Potentially huge dataset.
Both approaches highly amenable to HEC. Parallelization scheme depends on the approach; python drivers are used for sample-based approach, embarrassingly parallel; Scientific Machine Learning UQ scales well on GPU (the more you have, the more efficient)

Slide 3:
Rare event calculations – cost goes up quite high; typical use case at NASA.

Running a lot of these on their mid-size cluster on v100s and K-40s. High demand, can experience lengthy queue times there.

Difficult to train ML networks; batch training with different hyperparameters

ML UQ methods gaining a lot of popularity. Expect them to mature. Sample-based approach is pretty much mature.

Interesting sell: for UQ, more compute = more confidence rather than, for example, higher fidelity

from Michael Little (Int) to Everyone: 10:38 AM

Ref to the note that without HEC you'd have to turn it over to Sandia, do you collaborate with them today?

Yes, working with some people in optimization and design at Sandia; also, some connections with UQ group that develops Dakota software, but no official collaborations

from David Wagner (Int) to Everyone: 10:40 AM

I'll comment that I'm working on getting a formal collaboration with Sandia, but it's progressing slowly. Also, Gabrielle, Dan, and I took my mini-app to a DoE-sponsored GPU hackathon.

from Chris Dwan (Ext) to Everyone: 10:43 AM

Question: Do you have opinions on the various generations / models of GPU - in particular the ones that aggregate intra-board traffic on a dedicated PCI or even DRAM bus?

from Brian Thomas (Int) to Everyone: 10:44 AM

Q: could TPU hardware be of interest ?

from Brian Thomas (Int) to Everyone: 10:44 AM

TPU = Tensor Processing Unit (Google hardware)

A: Thus far, limited in a lot of this. Work being done under IRAD. Limited to GPU nodes at LaRC. So far, v100 is 6-7 times faster than k40s. Allocated for new V100 node we'd have restricted access to for the next year. Kind of impeded in studying this method without means to run application more freely.

---

3. Use Case: Aerothermal Analysis of Planetary Entry Systems

Jeffrey Hill
jeffrey.p.hill@nasa.gov
ARC, TSA
STMD/SMD/HEOMD

Main points:
P1: This is not a "vanilla" CFD solver. It requires tracking more degrees of freedom per grid node and is much more computationally intensive than traditional low-speed CFD code.

P2: ARC/TSA discontinued in-house cluster to support smaller missions after cost-benefit analysis. Work was moved to HEC. However, if things moved off-site and there was a latency issue, TSA would reconsider a local cluster.

P3: Would be able to use an infinite amount of time if available, but a 10X increase would be a very meaningful improvement. Right now, turnaround is too slow to include certain types of aero/aerothermal analysis results (e.g. damping simulations) in CDR/PDR because of queue times. Ideally, there would be tighter integration with the design cycle. Adding simulations for radiation transfer would require another 10X increase.

P4: There is interest in moving to GPUs and this is already being done at a low level but the effort is restricted by funding. A 2 to 3-year program to support the development work would be a good step.

P5: Projected demand over time is seen as consistent, if not growing. Orion will shift to more of an operational domain, with less engineering required, so computational needs will decrease. However, other projects would likely pick up that slack. The Mars Sample Return Program, which has Earth and Mars entry vehicles, is an example. There is also a series of science missions going to Titan.
If Moon to Mars becomes a real project, will be huge step up and require time-accurate simulation and a database of time-accurate simulations the likes of which we haven't done before. In near term, number of runs increases in preparation for design reviews.

Raw notes:

Slide 1: Use Case: Mission Need
Major computational use cases in TSA: Mars 2020, Orion, small planetary missions. extreme heat transfer and heat loads – 10,000 K; running simulations to provide environment definitions to engineers building thermal protection system (TPS).

Planetary probes could not carry much payload without the risk analysis.

Not a vanilla CFD solver. More degrees of freedom per grid node that needs tracing with flow equations. Much more computationally intensive (more math per cell) than a traditional low-speed CFD code.

Code TSA discontinued in-house cluster to support smaller missions and moved to NAS because of cost-benefit analysis; going back would shift cost back to branch.

Need high confidentiality; need long-term reliable storage; a primary task with their own cluster was managing backups and such.

Simulations are done early, asynchronous with other design tasks, can handle short outages but not much beyond that.

Slide 2: Use Case: Mission Need
Large workload on Pleiades, lot of people interacting.
HEC is a critical enabler for these projects.
Run gamut of different vehicle classes (capsule geometries; space plane; launch abort, etc), but always interested in heat transfer.
Expect to see more time-accurate simulations. Not just chemistry but evolving that time accurately = order of magnitude more expensive computation.

Slide 3: Current HEC Application Characteristics

Shape, mass – run pre-NAS; get critical flight envelope, construct grid, then batch out to NAS.

Limited interaction, text based from command terminal. Check every couple of hours, review convergence history, can make adjustments to time steps or performance if necessary; not high-frequency interactions.

There are occasional needs for quick reaction or near-real time computing, though generally not needed. For example, Shuttle Return-to-Flight, need to grid and run high-fidelity CFD in 24 hours. May develop something like this for Orion, but no official request from mission yet.

Input is a grid that defines geometry and fluid volumes 100 MB – 1 GB; computation cost of solving at each grid point; builds up over time as project progresses.

Transfer output from HEC back to local workstation, may archive on project specific server (for example, at Johnson or JPL). May hold data 3-5 years on lou mass storage at NAS, depending on how project manages its data. Often save inputs, configuration scripts for future use.

Slide 4: Current HEC Application Characteristics

Being able to move to GPU is something we're exploring.
If you give us more time, we are going to use it – no upper limit. Will add more cases if it becomes faster to run them
Queue time is an issue for larger simulations. Usually 1:1 simulation vs wait time.
Usually start with coarse grid; reserve private nodes at NAS.
Case throughput – gotten good at this and expanded capability has helped this. We'd like to be more tightly integrated with design cycle.

Slide 5: HEC Needs Assessment Workshop Key Questions

Usage is not spread equally throughout the year; have publication cycle, with spikes in January and June. Also, summer interns (20-30% increase with 5 or 10 interns).
Constraints: Not enough time to include in CDR or PDR, especially for retropropulsion for human landing on Mars, and other larger more time accurate cases.

Need to increase fidelity and radiation transport into main solver loop, factor of ten increase on simulation.

Biggest paradigm shift: want to move to GPU architectures; FUN3D and Sandia SPPARKS codes are making good progress w/ GPUs.

from David Wagner (Int) to Everyone: 10:46 AM
Getting access to hardware for day-long interactive development is something that I meant to highlight as a bit of a roadblock.

from Chris Dwan (Ext) to Everyone: 10:47 AM
My deferred question: Is there any benefit to transfer learning? This is the "take and bake" of machine learning - with a model mostly trained on a large public corpus of data and then fine-tuned with local parameters.

from James Warner (Int) to Everyone: 10:53 AM

We are relying solely on Tensorflow for our physics-informed ML work and have seen that that's common in the literature right now. Tensorflow auto-differentiation is effective/fast and is needed for incorporating PDEs into training. It was also really easy to get our code to run efficiently on GPUs with TF.

from Chris Dwan (Ext) to Everyone: 10:56 AM

Question about reproducibility. What, if anything, do you do to support reproducible research for your results?

from Jeffrey Hill (Int) to Everyone: 11:14 AM

On the topic of reproducibility: We archive inputs and maintain version control of the tools. That said, we're not very good at this. While we can often reproduce the solution, we regularly lose the data reduction methodology because this is still mostly a manual process

from Michael Little (Int) to Everyone: 10:59 AM

for Jeff - have you found an knee on the performance curve where inter-node comms is the bottleneck?

from Michael Little (Int) to Everyone: 11:01 AM

how big a computing resource do you need to release the constraint limitation on slide 6

A: Factor of 10 would be a big one for us to allow for time accurate simulations, whether CPU or GPU. Run in batch mode with lots of cases.

from Brian Thomas (Int) to Everyone: 11:03 AM

You mentioned that reduced data were important and saved 3-5 years. If you could would you save longer? Is data transfer an issue? Do you spend significant resources doing your own data management?

No problem with long-term storage; keeping all simulations for Orion; lou mass storage at NAS has been a good solution in terms of archive with relatively low overhead.

---

4. Use Case: High-fidelity, System-level Airframe Noise Simulations for Civil Transports

Mehdi R. Khorrami
mehdi.r.khorrami@nasa.gov
LaRC, D302
ARMD

Main points:

P1: These simulations allow flight testing to be done after the legwork for technology development has been put in and permit analysis of some of the finer details that you might not be able to look at during flight testing or wind tunnel tests.

P2: With urban air mobility (UAM) coming into the picture, there will be increasing demand for high-fidelity, system-level aeroacoustic simulations. While the computational needs may be slightly smaller for these types of vehicles, they present new challenges such as fleet presence or how sound bounces off buildings in a surrounding urban area. Turnaround times will also become more crucial to answer specific questions.

P3: Because of their success and accuracy, these airframe noise simulations are also being applied to supersonic aircraft, particularly commercial SST concepts. For these aircraft, we need to determine their acoustic signature on the ground and establish the proper noise metrics to assess their impact on community noise during landing and takeoff.

P4: High-fidelity airframe simulations generate vast amounts of data and are very I/O dependent. Reliable, high-bandwidth, fast I/O would go a long way to expedite turn-around times. Simulations have been self-limited on the number of cores requested because of excessively long queue wait times.

P5: Our computational demands are steady throughout the year because we are milestone-driven. In the next 5-10 years, we can expect a significant push from industry for the FAA to open airspace for UAM operations. NASA will be expected to assist federal regulatory bodies in the development of noise metrics by mapping out the acoustic signature of these vehicles.

Raw notes:

Within the Computational AeroSciences Branch, for the past 12-13 years we have been pushing the envelope in high-fidelity airframe noise analysis. Now looking at system level, e.g. entire aircraft, vs component level; needed to get the accurate noise signature.

Slide 1: Use Case: Mission Need
Aircraft noise reduction is a critical goal for NASA. Airframe noise is a bigger component of overall noise than the engines during landing.

All our airframe noise simulations are time-accurate by definition. To generate noise, need to have unsteady flows.

When you can predict airframe noise and evaluate performance of new noise-reduction technologies, you can cut down expensive flight testing; allows flight testing to be done at the end of the technology development process, after you have put in the legwork.

Ten years ago, our biggest simulations were components such as wing or landing gear. Only in the last few years have we started doing them routinely. Now, we are considering full-scale aircraft geometries in landing configuration – pushing the envelope of what is possible. Have been able to prove to HQ the unique value that these simulations bring to the table. Doing noise reduction technology development in virtual space postpones costly flights to the end of the evaluation process; also, can examine all of the finer details you might not be able to consider during flight testing or wind tunnel tests.

Use visualization to gain global insight on near field, then figure out how to reduce noise-producing flow fluctuations caused by complex geometry.

Group typically execute simulations year-round; not dependent on conference cycle.

Noise is the number 1 complaint that the FAA receives. Need to predict and understand how much noise an aircraft generates in order to expand air traffic. With urban air mobility (UAM) coming into the picture, you will see a lot more of this type of noise simulations to develop proper metrics for FAA certification. Turnaround time becomes more of an issue for the UAM community due to the vast number of configurations and situations that must be examined.

Grand challenge of airframe noise simulation is a large civil transport during landing; five-year collaboration with Boeing. Because of success and accuracy, migrating to supersonic class of aircraft to get noise signature for take-off and landing; we do have not have experimental data, or some of the empirical models.

To get mid-to-high frequency noise, we need grids with small cells to resolve the finer flow structures produced by complex components such as landing gear. These grids contain billions of cells. For aircraft certification, FAA requires 50 hertz to 20 kilohertz. Resolution of this frequency range necessitates long-time simulations with minute time steps.

Slide 2: Use Case: Current HEC Application Characteristics
Eventually want to bring aircraft certification into digital/virtual space.

Airframe noise simulations generate vast amounts of data. I/O is very important for us (very I/O dependent)
Preprocessing – small node count, large memory. For simulations, we limit the number of cores to avoid very long queue wait times (can be 3-4 months).

Slide 3: HEC Needs Assessment Workshop Key Questions
Visionary - what we would like to do in 5-10 years: improve spatial resolution, bring in propulsion noise for total aircraft noise signature; 2-3 times the size of current simulations. Usage of code is growing.

Reliable, high bandwidth, fast I/O would go a long way (for example, to not need to repeat simulations due to filesystem failures or other issues).

Side Discussion in Chat (not covered in following Group Discussion notes)

NOTE: Notes clock is PT, while this section of chat clock times are ET (3 hrs.)

from Chris Dwan (Ext) to Everyone: 2:37 PM

Do we have a feel for the ratio between that optimization capacity as opposed to capacity for production runs?n

from David Wagner (Int) to Everyone: 2:39 PM

I'll ballpark it at 2/7, as one colleague doing this has shifted those loads to the weekend.

from David Wagner (Int) to Everyone: 2:40 PM

...his is a code that runs full-scale production on Pleides

from David Wagner (Int) to Everyone: 2:44 PM

Problem sizes for computational materials are just starting to grow as more researchers start to get access to codes that allow running large physics-based (rather than empirically-tuned) models.

from David Wagner (Int) to Everyone: 2:58 PM

I concur with the gentleman re. my code

from Robert Ciotti (Int) to Everyone: 3:01 PM

I'd characterize NASA as code heavy. porting code is labor heavy and when your done 5 or 10 years later, the architecture you started where on is gone. At this point time, architecture are in a period of flux, where architecture are evolving more rapidly. hit a moving target.

from Chris Dwan (Ext) to Everyone: 3:03 PM

I cannot overstate the -qualitative- change for a research team in reducing runs from weeks to hours.

from Chris Dwan (Ext) to Everyone: 3:03 PM

I've seen it a couple of times in my career, and it's like a different kind of science.

from David Wagner (Int) to Everyone: 3:03 PM

concur

from Brian Thomas (Int) to Everyone: 3:07 PM

the infrastructure for 'code mapping' should be by HEC. But missions then do particular mapping. Ex containers: HEC needs to host the container registry and provide base, vetted images

from Brian Thomas (Int) to Everyone: 3:07 PM

as well as examples for the rest of the agency to template off of

from David Wagner (Int) to Everyone: 3:08 PM
It would help to do some divide-and-concur strategy. For example, do we need a different grid generator/partitioner/etc. for every different application?

from William Thigpen (Int) to Everyone: 3:09 PM

Would they be interested if it was a funded development?

from Brian Thomas (Int) to Everyone: 3:09 PM

promoting infrastructure which lowers impedance on code sharing would answer that question

from David Wagner (Int) to Everyone: 3:09 PM

even a 10x scale-up in the core solver, for example, will likely require an entire toolchain scale-up

from Robert Ciotti (Int) to Everyone: 3:09 PM

development for the future capabilities is traided for operational near-term deliverables,

from David Wagner (Int) to Everyone: 3:09 PM

good point

from Brian Thomas (Int) to Everyone: 3:10 PM

I think a gathering (such as this) for use cases and an analysis of common ones would provide targets for infrastructure development

from Robert Ciotti (Int) to Everyone: 3:16 PM

is NASA losing their leadership position in aeronautics because of how HEC investment are structured?

from Robert Ciotti (Int) to Everyone: 3:17 PM

maybe as that relate to how other countries invest in R&D in HEC.

Group Discussion of Key Questions

Note: Lots of side discussion in the chat (previous section); mostly not covered here in the notes.

Mike Little: Are there key factors that affect load and variabilities other than conf. publications and interns?

Jeff Hill: Project needs/milestones drive the most increase. Latest analysis shows more cases running; not uncommon to try and juggle 20-30 simulations in a week. In order of importance: project milestones, conf. season, interns.

Mike Little: Projecting over next decade, are these programs your supporting moving towards some kind of convergence? Moving towards crescendo vs annual cycle.

Mehdi Khorrami: Our simulations are milestone driven. Within the next 5-10 years, there will be a significant push from UAM industry for FAA to open airspace. NASA should really map out the acoustic signature of these vehicles. Massive, very rapid need for answers. Supersonic aircraft – push from industry for FAA to develop noise certification guidelines. Currently, there no semi-empirical models nor flight data exist, so role of simulation is increasing. These two example of things to come will push the boundaries of simulation-based noise prediction.

Mike Little: Urban air mobility – flight regime and behavior are different than some of the larger aircraft. Do you see that it reduces demand on computational resources?

Mehdi Khorrami: The UAM simulations could be slightly smaller maybe, because don't have to deal with system of systems like landing gear, but as you stated, they have their own issues. We are starting to look at noise footprint from a fleet and interaction with the surrounding environment, how does the noise bounce off of buildings in urban areas with multiple vehicles operating. Also, because the noise signatures span a broad frequency range, no free lunch here – have to do some of the heavy simulations.

Jeff Hill: Anticipate Orion and more of an operational domain, less engineering. Expect computational needs to decrease. Have other projects that would likely pick up that slack. Mars entry, Mars sample return, sample retrieval lander; Earth return vehicle (conservative risk posture, incl. planetary protection issues); series of science missions going to Titan. These will keep load at least constant. Moon-to-Mars w/ human landing systems needing supersonic retropropulsion would be a huge step up with time-accurate simulation, database of time-accurate simulations the likes of which we haven't done before; would dramatically increase HPC need.

Chris Dwan: The more we can define the features and their impact on the business, the more likely we are to get funding. Do end stakeholders understand the HPC?

Jeff Hill: Programs are more speculative; researchers are trying to speculate; Ashley Korzun, LaRC for example: supersonic retropropulsion. (She presented at yesterday's session.)

Little: What are the limits imposed by current assets; trying to explain so stakeholder will understand. Do you tell program managers what you can't do because you don't have resources?

Chris Dwan: Schedule driven – what we can deliver. No program managers are responsible for 2025-30 timeframe so cannot scope projects for 2030 at a program level.

Mehdi Khorrami: At every turn, at every opportunity, I always hammer home to project managers that either due to turnaround, more capacity, too many users, we need more compute power. We need more capacity, the current situation is not sustainable. Value of having HPC at our disposal. Always mention this and the value of HPC when presenting at HQ.

Chris Dwan: Non-NASA computing resources. Which ones and why?

Mehdi Khorrami: Cloud computing not realistic. There are benefits to interacting with NAS experts. Can't just run on cloud turnkey. Need people who know how to work with the system you have. NAS experts are essential; value of human experts for the specific systems is essential.

Chris Dwan: Sounds like no one is using non-NASA resources on this call.

Jeff Hill: ITAR data precludes us from doing so.

Mehdi Khorrami: Also, SBU issues; one has to be very careful.

Chris Dwan: Evolution over next decade of the use case. All seem to support on-going development. No particular maturation scale. Anything that will qualitatively change from development to operations?

David Wagner: Mini-app will become a full application. Just started. Worried about support issues going forward (continual testing, maintain integrity); fully functional version is a huge package. Would like to see level of commitment for long term support. Application intended to be built by user. Requires some level of sophistication.

Chris Dwan: Several of you talked about moving to new architectures (GPU). What would be needed to move forward and migrate?

Jeff Hill: Fundamental challenges. There will be some development work. Chief engineer sponsored a study to move some of these CFD codes. Starting to explore moving to v100 platform (Eric Nielsen's HPC RNA project). I think we need more things like that. Have the people, but finding the funding and someone to organize the training is a challenge. Workload outstrips R&D budget. Needs to be agency or MD level and have some real funding behind it.

Mike Little: Funding needed for each?

Jeff Hill: Enough resources at NAS to keep us busy for next few years and enough projects. Training could be paid for, but we're talking about man-years of effort. It's a labor issue for the developers.

David Wagner: On-going cost.

Jeff Hill: 2-3-year program would be a good step but needs ongoing funding.

Chris Dwan: Assume shifting from one location to another with the same architecture. Less work or comparable?

Jeff Hill: Less work. If things moved offsite and there was a latency issue, we'd have to reconsider local cluster.

Irina Kitiashvili: Warning about big changes. Typically, projects are grant restricted; code development is restricted. No major changes possible for code.

Chris Dwan: Unless it's part of the funded work to migrate codes, it's very difficult to find the time.

Irina Kitiashvili: If new architecture, will require significant work. Don't know whether it will be an improvement to migrate. Providing new architecture without migration help can leave project stranded.

Chris Dwan: What are the questions we missed?

David Wagner: NAS doesn't have any plans for exascale computing. Is anyone going to need that?

Jeffrey Hill: We'd love to have access to that capability. It would fundamentally change the analysis we can do. Even 10x improvement would be very meaningful, maybe don't need to get to exascale. DOE exascale computers are using advanced architectures, which takes a long time to match codes. Exascale will lead to needing investment in advancing codes.

Mehdi Khorrami: Same here. Lower turnaround time; additional time would allow for more analysis and understanding of what we need to do next.

Piyush Mehrotra: Specific quantitative examples of what you could do with more resources (10x larger, 100x larger) would be very useful.

Jeff Hill: For landers, we are trying to characterize hypersonic thru subsonic capsule stability. Capsules dynamically unstable descending in Mach number. Parachutes bring stability; need to decide when and if parachute is used. To assess, pre-parachute descent needs time-accurate highly resolved wake simulation. We can't turn them around fast enough now, 10x increase would be enabling. Also, supersonic retropropulsion.

Jeff Hill: Will pull together a chart to provide.

Mehdi Khorrami: Running on 100k cores reduces month to days to speed up analysis.

Tsengdar Lee: If we build such a system, who should be the one to fund mapping the code to the new architecture?

Jeff Hill: Hard question. General experience is that flight projects are not going to do it. Need investment at the R&D level. Entry systems modeling has existing connection to R&D. Need agency level or MD level.

Bill Thigpen: If the funding was there, would people in the groups be interested in working on those teams to do the improvement?

Jeff Hill: Yes, definitely. Have a couple of people already doing this at a low level but restricted by funding limitations.

Mehdi Khorrami: Using 3rd party code (commercial code) with primarily automotive customers; they would have to do cost/benefit; has to work for them economically.

Historically, Aero has footed most of the bill for expansion. Does the space side contribute/take any ownership?

Tsengdar Lee: My understanding talking to HEOMD side, very focused on near-term mission. If they don't have enough capacity building an aero DB, they would invest. Solving very specific problems for very near term (<5 years). Not planning for future.

Mehdi Khorrami: For a lot of the capability, ARMD has to spend the money, but Space takes priority in hour of need and there must be a balance. Someone has to say this is not sustainable.

Jeff Hill: STMD is coming in to do more investment, but still mission-focused and technology-focused.

Mike Little: We are doing six sessions; special sessions for large groups doing the same kind of things; by the time we get done we should be able to identify everyone in NASA who needs HPC. We will have a comprehensive survey of what needs exist.

Irina Kitiashvilli: App development - can we request NASA to include code development? Slowing down SMD.

Mike Little: This will be one of the elements we talk about in final report. Send email with specific suggestions or add to notes to include in input.

Mehdi Khorrami: In 2017, while we were developing landing gear noise reduction technologies, we performed a fit-check a few months before the flight test. A conflict between tire and fairing was discovered during the check and they had to cut a big chunk of the fairing. All of the work we had done was on the line if the aeroacoustic effects of this last-minute modification could not be evaluated. Asked the change process to stop, contacted NAS to get priority. We were able to run and analyze simulations to determine that cutting the fairings would have a minimal impact. Took 4 weeks. More resources would have allowed a more cost-effective evaluation and faster turnaround, to get multi-million-dollar flight test right. The cost of flight-testing the modified fairing without having assessed its impact via simulations would have been significant if it turned out that the change was detrimental to the performance of the landing gear.

Appendix D5 Workshop Session 6-8 Inputs

Brian Allan	LaRC/D303	RVLT High Fidelity Modeling - Lift+Cruise UAM Concept Vehicle
Gaetan Kenway, Jeffrey Housman, Emre Sozer, James Jensen, Jared Duensing and Cetin Kiris	ARC/TNA	High Efficiency Database Generation with LAVA
Francois Cadieux, Michael Barad, Cetin Kiris	ARC/TNA	Scale-Resolving Simulations with LAVA Cartesian AMR
Li Wang	LaRC/D302	Multidisciplinary Analysis and Optimization for Rotorcraft Aeromechanics
Jim DeBonis	GRC/LTN	Large Eddy Simulation for Propulsion Flows
Reynaldo J. Gomez	JSC/EG3	HEO/Launch and Entry Aerosciences
Mujeeb Malik	LaRC/D302	TTT-RCA Technical Challenge of Accurate CL_{max} Prediction
Michael Barnhardt	ARC/TSA	Modeling Parachute Dynamics for Entry Systems
Michael Rogers	ARC/TNA	Transformational Tools and Technologies (T^3) Project HEC usage
John Lawson	ARC/TI	Computational Materials & Molecular Simulations

National Aeronautics and Space Administration

Workshop Session 6-8

8 June 2020

National Aeronautics and Space Administration

Use Case: RVLT High Fidelity Modeling - Lift+Cruise UAM Concept Vehicle

Dr. Brian Allan
June 1, 2020
LaRC/D303
ARMD

Use Case: Mission Need

- Mission: Perform high fidelity Computational Fluid Dynamic (CFD) simulations of NASA RVLT Urban Air Mobility (UAM) concept vehicles for development and validation of conceptual design tools.
- Criticality: HEC is critical to accomplishing this work as it is the only NASA resource capable of meeting our simulation requirements.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - RVLT would reduce high fidelity modeling and shift focus to lower fidelity models.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Research performed on HEC is typically open, but some work has been done on ITAR related work.
  - Most simulations take multiple queue submissions. If these jobs get interrupted, it can result in significant delays since the job would need to be cleaned up and resubmitted back into the queue.
  - Currently there is no need for those without NASA credentials to access HEC.
- HEC Function:
  - Applications run all year to meet mission objectives.
  - High Fidelity modeling simulations of UAM vehicles can capture critical aerodynamic interactions between the vehicle fuselage, wing, rotors, and rotor wakes that are not validated in NASA lower fidelity design tools. Without experimental data, high fidelity modeling plays a critical role in assessing and validating design tools.
- Impact of HEC:
  - Currently we do not have another option if NASA supercomputer was no longer available and we would need to adjust our planned work to available resources.

NASA RVLT Lift+Cruise Concept Vehicle CFD Simulation

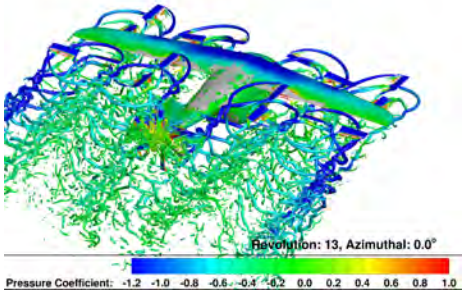

Use Case: Current HEC Application Characteristics

- Work flow:
  - Compute stage starts after creation of a computational grid and preparation of input files (from local resources)
  - Each simulation may take multiple job submissions; results are checked daily (faster turnaround would encourage more automation here).
  - Near-real-time response is not needed, but queue wait times > 1-3 days discourage production use.
- Data:
  - In: Computational (volume) grid and text input file defining run conditions.
  - Out: Final (volume) flow solution and time-history of selected volume or surface subsets.
    - All results are typically saved to mass store, though running significantly more cases would make this prohibitive. If compute resources were reliably available, simply rerunning cases would be an alternative.
    - Post-processing includes convergence checking and flow visualization.
  - Transfer:
    - Output goes to local compute resources for post-processing and to HEC mass store.
- Computation:
  - Type: Computational Fluid Dynamics (CFD)
    - This includes capacity computing but can benefit from advanced (GPU) HW options in the future.
- Evaluation: Effectiveness of HEC is judged on turnaround (queue wait plus execution) time for each run.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Project goals and milestones
  - Support for wind tunnel test planning
  - AIAA and AHS conferences
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Queue turn around time, capacity, queue max time
  - Impact: Increased turn around time limits the number of cases that can be run narrowing the scope of the investigation.
  - External Sources: None
- How will this Use Case evolve over the next decade?
  - Maturity: This code and application are currently used in industry and will continue to be developed as part of NASA strategic goals.
  - Growth: Increased resolution with faster turn around time to increase number of cases run
  - Re-use: Number of researcher running this code/application is expected to increase.
  - Re-hosting: A significant research effort is needed to move code to GPU clusters. Currently there is a research effort to explore porting the code/application to GPU clusters.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: Computational Fluid Dynamics
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones: None
    - Programming languages used: Fortran (99.9%), C(0.1%)
    - Number of lines of code: 350,000
    - Workflow mechanism used to control code execution: Shell scripts
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing): (estimating for 250M point grid case)
    - Size of runs (nodes or cores): ~2000 cores
    - Accelerator usage (type and memory requirement): currently none
    - Amount of memory required (overall or per process): 500GB
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other: both MPI and OpenMP for each case, and multiple cases
    - Storage is covered below
    - Run in batch, interactive, or near-real time: Batch

Current Use Case (Survey)

- Describe the data requirements for the application (sized for 250M point grid, 2000 cores)
  - Describe the data flow of the application from input to output
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other)
    - Input is serial ~24GB restart read by manager, distributed via MPI to workers
    - Output *can* be written in parallel (but then needs post-processing to reassemble), usually just serial
  - How much input and output are generated per run: 24GB input, 200GB to 1TB output per run
  - Summary of the total amount of input and output data per year above, times 60 runs/year
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? Statistics, convergence checks, visualization, integrated (force/moment) values, acoustic analysis of surface pressures
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other reasonably fast shared scratch space, longer term mass store
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used.
  - How much data needs to be curated (archived)? 50-100GB/run
  - How much data needs to be distributed and, generally, to whom? 20GB/run to user's workstation
- What are the key computing obstacles to effective execution of the current application? Queue wait time
- Are there any analytic functions performed on the output of the application?
  - Timing, computing load, workflow timing, load balance, integrated flow quantities, acoustic analysis
- Do you have available the workforce necessary to re-host your code on any new platform necessary? 25%
  - Describe any skills or capacity needed. Expertise in GPU and future HW architectures, current code structure, and aerodynamic applications.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could affect these goals? Mainly availability of adequate resources.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.): New methods for prediction of turbulence transition. Adjoint and real gas (chemistry) are possible extensions.
    - Increase in resolution and/or timesteps: Needed for transition, improved acoustics, and for other applications.
    - Rewrite to improve performance: This is ongoing.
    - Refactor for emerging platforms: This is ongoing.
    - Rehost onto cloud: No (due to NASA IT security restrictions).
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components: No
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores) Could increase 10-100x (trade off with wall time), but only if nodes are available.
    - Accelerator usage (type and memory requirement) Could use GPUs or other HW, again only if available.
    - Amount of memory required (overall or per process or per core) Steady or decline (based on using more cores).
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other Better parallelization, but with more communication.
    - Storage requirements Directly related to available compute resources.
    - Run in batch, interactive, or near-real time Batch for production (but development environment needs interactive).

National Aeronautics and Space Administration

Use Case: High Efficiency Database Generation with LAVA

Gaetan Kenway, Jeffrey Housman, Emre Sozer, James Jensen, Jared Duensing and Cetin Kiris

June 8, 2020

Computational Aerosciences Branch (TNA), NASA Advanced Supercomputing division, NASA Ames Research Center
ARMD and HEOMD

Use Case: Mission Need

- Mission: Perform RANS database generation for ARMD projects. Current projects include the X57 Maxwell, X59 LBFD and Novel Propulsion Airframe Integration (NPAI)
  - How to measure impact: Validation against wind tunnel and flight test data; how much CFD data is incorporated into design decisions/revisions; cost reductions from CFD lowering the number of required tests
- Criticality: NASA HEC is critical to accomplishing this work
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Most likely the quantity of simulations would be curtailed, and a much smaller set of simulations performed on Commercial Cloud Computing. Projects would have to accept additional risk and possibly perform more physical experiments
- Mission Security Needs: Ranges from pubic to restricted
  - What is impact of inadvertent release of data? Depends on classification, potentially severe for restricted data
  - We rely on HEC availability to meet project goals and deadlines. Currently database generation time is measured in months
  - Industry partnerships are important and sometimes require non-NASA access to both systems and data
- HEC Function: Use HEC computing to generate large-scale databases for ARMD projects. HEC storage resources are used for data generated
  - Performed O(100) simulations of the X57 Maxwell in a wide variety of flight configurations
  - Performed O(100) simulation of the NASA CRM BLI wind-tunnel model ahead of upcoming NTF tunnel entry
- Impact of HEC: HEC provides ARMD projects with "Numerical Wind Tunnel" capability. With fewer NASA operated wind tunnels than in the past, CFD analysis is more important than ever
- Mission Concept: Aero-database can serve both aero and space vehicle concepts

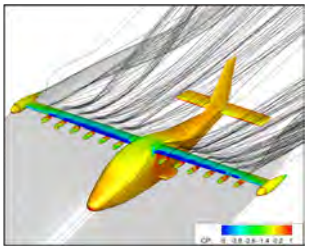

X57 Maxwell CFD simulation using actuator regions for high-lift propellers

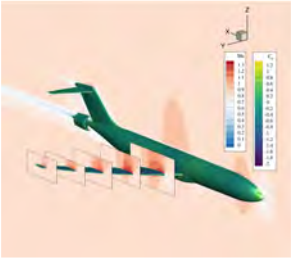

Simulation of the NASA CRM with a flow-through Boundary Layer Ingestion (BLI) nacelle. Wind tunnel-scale model simulations performed before entry into the NTF wind tunnel. NPAI application.

Use Case: Current HEC Application Characteristics

- Work flow:
  - What triggers this workflow? Specific projects decide on the CFD requirements in order to fulfill the mission objectives. Database generation is often used to reduce the risk associated with a wind-tunnel or flight test. It can provide much needed guidance to the physical test which reduces the time and costs associated with real-work testing
  - A CAD geometry is usually provided. Standard preprocessing tasks are CAD clean, surface and volume mesh generation, input file preparation, job submission, and visualization of results on HEC or workstations
  - Typically does not need quick reaction or near-real time computing.
- Data:
  - In: Volume Grid Files (1 GB-10 GB). Optionally additional triangulations in areas of interest (~10 MB)
  - Out: Volume data (1-20 GB) typically one for ever steady state simulation
    - Surface Data (~10 MB) Sample surfaces or point probes
    - Volumetric data is usually retained for post processing analysis and long-term storage
    - Volumes are used to generate visualizations to help understand flow physics
  - Transfer:
    - Mostly stays on HEC system, or is downloaded to local Linux workstation for faster visualization
- Computation:
  - Type: Capacity computing: Large number, O(100-1000) simulations each of which use between 20M and 200M grids nodes and requiring between 1 and 10 hours of computing time. Jobs can use anywhere from 100 to 1500 cores.
- Evaluation: HEC is essential for providing the required compute capacity for large scale database generation. Computational improvements over the past two years have resulted in decreased computer usage between 5x and 30x. This has allowed much better coverage of mission critical conditions to be evaluated, but more computer resources are still required

LMC608 from the 3rd Low Boom Prediction Workshop. Design is comparable to the final LBFD configuration

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Project deadlines for databases
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Available capacity on the HEC limits the system. The long queue times limit the number of cases that can be computed. Results in database generation times that are measured in months.
  - Impact: With an increased availability of computing resources and increases in the computation efficiency of the curvilinear solver, we would like to see database generation time drop from months to weeks.
  - External Sources: None
- How will this Use Case evolve over the next decade?
  - Maturity: Routine database generation of many ARMD projects and a reduced reliance on wind tunnel testing. Using computational resources to advance concepts further along the TRL scale before physical prototyping and tests begins.
  - Growth: We foresee an increase in the fidelity of simulations to scale-resolving simulations. We expect these simulations to be used to complement database generation to gain a deeper understanding into the flow physics behind the design. Scale-resolving simulations are ~30X more expensive than our steady-state RANS simulations for a highly optimized code
  - Re-use: We expect code to be released to selected alpha/beta testers
  - Re-hosting: We expect the existing code base to be well positioned for future CPU architectures. However, ongoing efforts will be required to port significant portions of the code for future accelerator hardware

Scale resolving simulation of the NASA Juncture Flow Model computed using LAVA Curvilinear. Passive particles colored by stream-wise velocity.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: LAVA Curvilinear performs both steady-state RANS simulations (implicit time-integration) as well as time-accurate scale-resolving simulations using explicit time-stepping.
  - Describe the application
    - Antecedents: LAVA framework has been under active development at ARC for 10+ years
    - Programming languages used: primarily Fortran 2008
    - Number of lines of code: 50,000+
    - Workflow mechanism used to control code execution: input file or Python script
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): 40-4000 core CPU cores
    - Accelerator usage (type and memory requirement): None for production runs currently
    - Amount of memory required (overall or per process): 1-3GB per core, Implicit <1GB per core, Explicit
    - Describe the parallelism of the application: Hybrid MPI/OpenMP: 1-2 MPI rank per node, OpenMP for on-node parallelism. SMT is beneficial for explicit time stepping. Dynamic load balancing using OpenMP.
    - Storage is covered below
    - Run in batch for most jobs, interactive for development and debugging

3/30/20 HEC Needs Assessment Plan B 14

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output: Pre-generated volume grid and control file. Optionally surface triangulations and point probe locations. All input is read at the start and thereafter only results are written.
  - How are the data loaded in the application: At startup: Grid and optional check-point file loaded using MPI I/O. Small I/O is read only by master
  - How much input: Grid files 1-25 GB, Checkpoint files (1-100 GB) and up to 100k files of ~10 MB for scale resolving simulations
  - Summary of the total amount of input and output data per year: input ~ 100 GB, output ~ 1 Petabytes
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? For steady-state simulations commercial software (Tecplot) is commonly used.
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other: need large files systems to handle MPI I/O parallel writes. Code has built in capability for buffering surface probe outputs to reduce writing frequency and allow for a large number of files to be written simultaneously, greatly increasing the effective write bandwidth.
  - How is input data brought together for use? For database generation, preprocessing tasks are performed mostly on workstations.
  - How much data needs to be curated (archived)? Only post-processed data and input data < 100 GB
  - How much data needs to be distributed and, generally, to whom? Post-processed data, to project partners: NASA or private partners
- What are the key computing obstacles to effective execution of the current application? Exceedingly long queue times for ARMD. Ratio of queue time to run time can easily exceed 10 for jobs under 500 cores which is completely unacceptable.
- Are there any analytic functions performed on the output of the application? No
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed: We have some of the workforce necessary, but we could certainly use more skilled programmers and software engineers with experience coding for the latest CPU hardware (e.g. Fujitsu ARM) and GPUs (e.g. NVIDIA V100/A100).

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could affect these goals? The main obstacles for database generation is the sufficient timely availability of capacity-level computational resources. We expect steady-state analysis to continue to provide sufficiently accurate engineering-level simulations for a wide variety of ARMD application. Without increase in availability of capacity computing (hundreds jobs from 100-1000 cores for ~8hours) this will be impossible to achieve.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components: Increasing use of scale-resolving simulations to complement steady-state RANS simulations
    - Increase in spatial resolution: This will lead to lower numerical error and higher confidence in database predictions
    - Rewrite to improve performance: LAVA Curvilinear has been completed refactored and efforts are still ongoing to further improve computational performance on existing computer architectures.
    - Refactor for emerging platforms: As new platforms become available, we will be investigating OpenACC for effective accelerator parallelism.
    - Rehost onto cloud: No plans
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components: No plans yet
  - If possible, describe how the computational requirements will evolve over time
    - Data base generation from O(100) simulations to O(1000) simulations
    - Accelerator usage (type and memory requirement): most likely Nvidia GPU, but Fujitsu, NEC, AMD, and Intel all possible provided human resources are available for code porting and optimization
    - Amount of memory required (overall or per process or per core): 3-5GB/core is optimal for Implicit, <1GB per core is sufficient for Explicit time stepping
    - Describe the parallelism of the application: MPI for internode comm, OpenMP in-node CPU threads, OpenACC for in-GPU threads (future)
    - Storage requirements: from 1 Petabytes/year to 10 Petabytes/year
    - Run in batch, interactive: no change, mostly batch, some interactive for setup and debugging

National Aeronautics and Space Administration

Use Case: Scale-Resolving Simulations with LAVA Cartesian AMR

Francois Cadieux, Michael Barad,
Cetin Kiris (POC: cetin.c.kiris@nasa.gov)
June 8, 2020
Computational Aerosciences Branch (TNA), NASA Advanced Supercomputing Division, NASA Ames Research Center
ARMD and HEOMD

Use Case: Mission Need

- **Mission**: Predict ignition over-pressure (IOP) and vibro-acoustic loads on launch vehicles (HEOMD) and far-field acoustics for aircraft and rotorcraft (ARMD) with CFD within short enough turnaround to help with vehicle design safety and noise reduction – Accomplished with LAVA Cartesian AMR Navier-Stokes or Lattice-Boltzmann
  - **How to measure impact**: validation against wind tunnel, ground test, and flight test data; how much CFD data is incorporated into design decisions/revisions; cost reductions from CFD replacing some tests; direct impact on design and construction for KSC Main Flame Deflector (MFD);
- **Criticality**: NASA HEC is critical to accomplishing this work
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Rent time on a secure commercial facility, agreement with DOE/DOD, or secure commercial cloud computing;
  - LAVA development is application-driven from many different mission directorates and partnerships
  - *Convenient & responsive experience with capacity computing and storing data on NAS systems*
- **Mission Security Needs**: ranges from open-to-public to ITAR
  - What is impact of inadvertent release of data? Depends on classification, severe if data is ITAR protected
  - We rely on HEC availability to meet project goals and deadlines, any projects that are in final design and are starting construction have strict deadlines
  - Industry partnerships are important and sometimes require non-NASA access to both systems and data
- **HEC Function**: Use HEC computing to perform scale-resolving simulations and HEC storage resources for data generated
  - Performed 12 Orion Launch Abort simulations with LAVA Cartesian AMR (350-800M cells, ~2 Million SBUs over 3 years)
  - Performed 5 LAVA Lattice-Boltzmann simulations of SUI Quadcopter (100-600M cells ~ 200k SBU over 3 months)
- **Impact of HEC**: LAVA simulation results were directly used to decide on re-design of KSC MFD; KSC decided to stop construction when contractors veered off design until new LAVA simulations helped determine course-correction and resume construction; CFD was not used in determining vibro-acoustic environment before, now it is used along test data to drive down uncertainty margins; demonstrated capability to predict high-frequency broadband noise for small rotorcraft
- **Mission Concept**: See videos to the right for Orion Launch Abort System

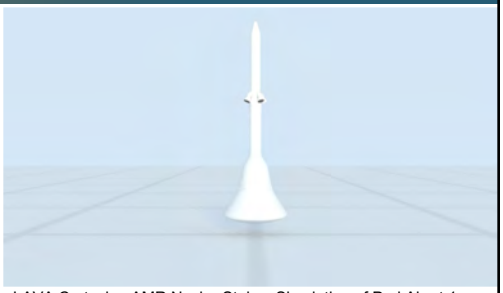

LAVA Cartesian AMR Navier-Stokes Simulation of Pad Abort 1 flight test where Orion LAS accelerates from rest to 10x Earth's gravity: passive particles seeded at the nozzle colored by velocity magnitude: white is fast, dark orange is slow.

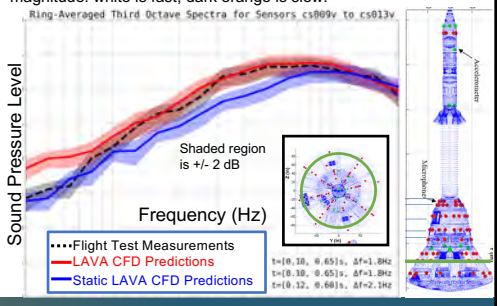

Use Case: App Characteristics

- **Work flow**:
  - Triggers: KSC approached the LAVA group about infrastructure re-design of Launch Pad 39B for SLS, and returned when construction was different than design to ensure safety; Orion Loads & Dynamics team from JSC expressed interest in capability from previous work on KSC launch infrastructure; Advanced Air Mobility market is extremely interested in new methodologies like Lattice-Boltzmann to better address noise concerns with multi-rotor aircraft
  - Create or receive CAD design, cleanup and prepare CAD for CFD, prepare inputs, submit job, post-process results in parallel on HEC to obtain surface and volume statistics, spectra, and space-time cross-correlations
  - Typically does not need quick reaction or near-real time computing.
- **Data**:
  - In: Closed surface triangulations of the geometry and any visualization cut planes (typically < 1GB, read once)
  - Out: Volume data (1-200 GB, every 1-10+ min, ~200 TB total per run), surface data (<1GB, every 1-5 seconds, ~600 GB total per run), point data (1MB, every 0.1-2 seconds, ~ 1 GB total per run)
    - Post-processed data requires retention: Statistics (200 GB), Surface spectra (<10GB), point probes (1GB)
    - Volumes are used to generate visualizations to help understand flow physics, surface and point spectra and cross-correlations are generated using parallelized python post-processing tools on Pleiades
  - Transfer:
    - Mostly stays on HEC system, or is downloaded to local Linux workstation
- **Computation**:
  - Type: Capability computing necessitating large resources over long period of time to obtain converged statistics and spectra at high space-time resolution
    - Typically run on 4,000 to 16,000 cores for 1 week to 1.5 months
- **Evaluation**: Unsustainable because reservation required to get started with jobs above 4,000 cores

Simulation of launch ignition of SLS at KSC. Temperature: white is hotter & brown is cooler. Plume contoured by exhaust-mass fraction (compared to air). Small green people are shown for scale.

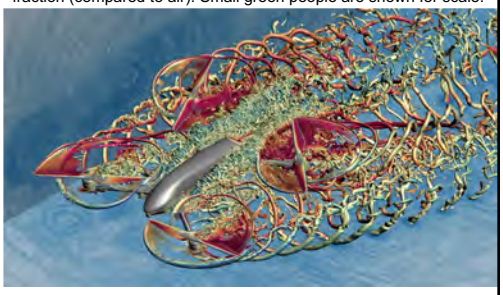

LAVA Lattice-Boltzmann SUI Quadcopter 5% tip chord simulation: Q-criterion colored by vertical velocity; cut planes: log of pressure gradient mag

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Project deadlines (design reviews, etc), publication cycles, algorithmic and numerical improvements
- Are there key **functional gaps** exist or are projected in the next 5 years or 10 years?
  - **Constraints** on Use Case: Can only run a handful of Orion Launch Abort Scenarios, or KSC launches in a year, primarily due to lack of concurrent capability computing resources (hard enough to get one 16,000 core reservation for 1 month): goal is to cover a database of scenarios in under 1 year.
  - **Impact**: More scale-resolving simulations could…
    - Reduce risk and costs for KSC infrastructure improvements for SLS and commercial launches (SpaceX)
    - Drive down uncertainty in Orion LAS vibro-acoustic environment and potentially reduce structural weight if we could cover more scenarios before it is built
    - Provide understanding to reduce noise of Advance Air Mobility aircraft designs
  - **External Sources**: None
- **Future**?
  - **Maturity**: Scale-resolving simulations for near-field and far-field acoustics will become routine within next decade
  - **Growth**: Foresee increase in frequency of 1+ Billion cell computations that require months to obtain converged statistics and spectra, increase in computational complexity due to coupling with other physics (fluid-structure interaction, multi-phase liquid-gas, particle physics for contaminants, combustion chemistry, etc)
  - **Re-use**: LAVA team continues to grow its users and developers, and will be released to selected alpha testers in the next 1-2 years, and beta version later to NASA and partner community
  - **Re-hosting**: Ongoing efforts to port LAVA to GPU in collaboration with NAS, foresee increase in need for skilled developers to migrate code to future computing architectures.

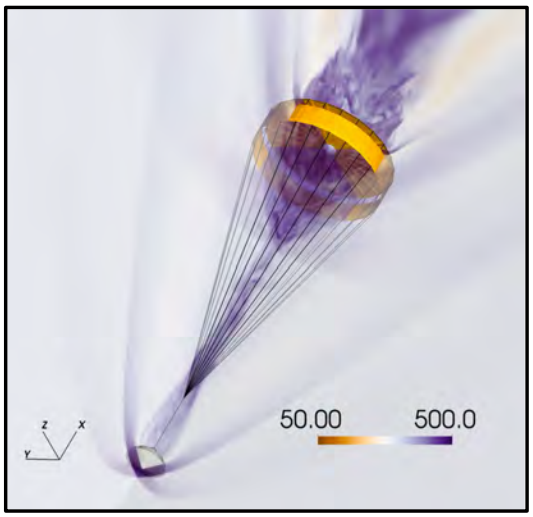

Coupled CFD + non-linear structural dynamics of Sub-scale MSL/Viking DGB Parachute: Temperature (K) color contours.

National Aeronautics and Space Administration

Supplemental Material

INSERT ANYTHING OVER 3 SLIDES HERE

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: LAVA Cartesian AMR performs time-accurate (explicit) scale-resolving simulations either by discretizing compressible Navier-Stokes or Lattice-Boltzmann (low Mach) equations
  - Describe the application
    - Antecedents: LAVA framework has been under active development at ARC for 10+ years
    - Programming languages used: primarily C++
    - Number of lines of code: 10,000+
    - Workflow mechanism used to control code execution: python and shell scripts
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): 4-20k CPU cores
    - Accelerator usage (type and memory requirement): None for production runs at the moment
    - Amount of memory required (overall or per process): 1 GB – 80 GB per MPI rank
    - Describe the parallelism of the application: Hybrid MPI/OpenMP: 1-2 MPI rank per node, OpenMP for on-node operations, extreme scaling via oversubscription and dynamic OpenMP loop-balancing
    - Storage is covered below
    - Run in batch for most jobs, interactive for development and debugging

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output: surface triangulations, point probe locations, and input file are read at start, no reading after job starts, only results output
  - How are the data loaded in the application: all at start, small files read by master and communicated, large files like checkpoint files are read in parallel where each rank reads its owned chunck
  - How much input (5-6 files of a few MB) and output (50-100k files of 100 MB, 10-50k files of up to 200GB) are generated per run
  - Summary of the total amount of input and output data per year: input ~ 100 GB, output ~ 1-2 Petabytes
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? We compute and store on-the-fly statistics in the volume output, surface output is analyzed via spectra (FFT) and cross-correlations, but we would like to do it for volume output too, volume is mostly used for visualization with help from NAS viz team with passive particle, computational schlieren, and volume rendering
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other: need large files systems (amount of data produced is large) that are able to handle MPI parallel writes efficiently because asynchronous is not always an option given the frequency requirements for output for certain projects (especially for far-field acoustics); currently writing high-frequency volume output is the most time consuming aspect of simulations, and surface output can be nearly 30% of run-time cost when output frequency is high (every 1-4 time step).
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used.
  - How much data needs to be curated (archived)? Only post-processed data and input data < 1 TB
  - How much data needs to be distributed and, generally, to whom? Post-processed data, to project partners: NASA or private partners
- What are the key computing obstacles to effective execution of the current application? Getting large jobs > 4,000 cores for more than 5 days requires reservation. We generally severely limit volume output to avoid it taking up precious run-time.
- Are there any analytic functions performed on the output of the application? No, but acoustic propagation to far-field, spectra, cross-correlations for post-processing
  - Timing: 2-30 min, computing load: 20-100 nodes, workflow: mpi-parallel code to post-process volume/surface data using a devel job on Pleiades
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed: We have some of the workforce necessary, but we could certainly use more skilled programmers and software engineers with experience coding for the latest CPU hardware (e.g. Fujitsu ARM) and GPUs (e.g. NVIDIA V100/A100).

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could affect these goals? We want to increase the use of predictive computational aero-acoustic (CAA) scale-resolving simulations in the design of future aircraft and spacecraft. Without increase in availability of capability computing (jobs from 4,000 – 20,000+ cores for ~1 month) this will be impossible to achieve.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.): we are adding in new physics to perform fluid-structure interaction, multi-phase liquid-gas simulations, coupled Lagrangian particle physics to model contaminants, which will increase # of flops/time step and code complexity
    - Increase in resolution and/or timesteps: Both will be drivers for acoustics as we need longer duration for converged spectra across wider frequency range
    - Rewrite to improve performance: Ongoing refactor of LAVA Cartesian infrastructure to reduce complexity/dependencies and improve asynchronous communication
    - Refactor for emerging platforms: Participated in GPU Hackathon in Boulder, CO (2018) for Lattice-Boltzmann single box mini-app, will participate in NASA Ames GPU hackathon (Fall 2020) to port LAVA LBM multi-box mini-app to evaluate outer loop threading strategies
    - Rehost onto cloud: No plans
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components: No plans yet
  - If possible, describe how the computational requirements will evolve over time
    - From current largest: O(1e4) to O(1e6) way concurrency if available
    - Accelerator usage (type and memory requirement): most likely Nvidia GPU, but Fujitsu, NEC, AMD, and Intel all possible, as much memory as possible
    - Amount of memory required (overall or per process or per core): from 80 GB per MPI rank to 384 GB per MPI rank (2 per node)
    - Describe the parallelism of the application: MPI for internode comm, OpenMP in-node CPU threads, OpenACC for in-GPU threads
    - Storage requirements: from 2 Petabytes/year to 100 Petabytes/year
    - Run in batch, interactive: no change, mostly batch, some interactive for setup and debugging

National Aeronautics and Space Administration

Use Case: Multidisciplinary Analysis and Optimization for Rotorcraft Aeromechanics

Li Wang
June 8, 2020
Research Aerospace Engineer
Computational AeroSciences Branch/NASA LaRC
Revolutionary Vertical Lift Technology Project

Use Case: Multidisciplinary Rotorcraft Optimization

- **Mission**: Development of a high-fidelity multidisciplinary analysis and optimization (MDAO) tool coupling high-fidelity computational aerodynamics, rotorcraft comprehensive analysis with structural dynamics, and noise prediction models for rotorcraft aeromechanics
  - The mission objective requires high-fidelity analysis and sensitivity analysis of fluids, structures, dynamics, propulsion, acoustics, and complex interactions among various disciplines. Accurate resolving of all the physics involved is a technical challenge requiring orders of magnitude more HEC resources than what can be currently accessed.
- **Criticality**: HEC is critically important for completion of the mission
  - If NASA supercomputing was no longer available, then that would create a real barrier to the mission considering large computing resources required and restricted software installation. An option could be to write proposals to access DoD resources or buy a new system.
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - Inadvertent release of data causes exposure of protected geometry/structural models to unauthorized sources, which is prohibited.
  - If HEC availability is interrupted regularly, work cannot be made toward timely completion of the mission
  - The current mission does not require access of the MDAO tool from non-NASA credentials.
- **HEC Function**: HEC enables MDAO tool development based on RANS simulations and low-frequency rotor noise predictions. Optimization limited to grids containing less than 10-million grid points considering resource availability and run time limit.
  - Lots of coupled analysis/design runs will be needed in next 5 years with orders of magnitude more resolutions to resolve turbulence for higher fidelities, blade/engine/airframe acoustics, and interactions
  - Coupled solutions are compared with experiments for model validation. Optimization performance is assessed by computational time to a viable design solution
- **Impact of HEC**: The MDAO mission cannot be met without HEC resources, which need to be expanded
- **Mission Concept**: Development of high-fidelity multidisciplinary analysis and design capability advances rotorcraft optimization for improved rotor performance and reduced noise generation

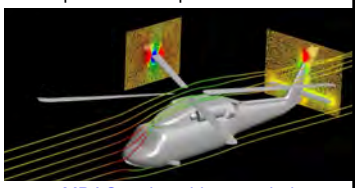

MDAO tool enables coupled aero/structural/acoustic optimization

Use Case: Current HEC Application Characteristics

- **Work flow**: Multiple discipline models perform coupled analysis, sensitivity analysis, and design solutions are optimized in a sequence of design iterations.
  - Optimization driver initiates a design process and manipulates workflow and data transfer in subsystems
  - The optimization process is automatic without user interactions. A recent aero/structure optimization used 122 wall-clock hours on 2,200 CPUs for a two-point (hover and forward flight) Blackhawk rotor-blade optimization. The demonstration was performed with isolated rotors on coarse grids (no aeroacoustics). Much larger computing resources and advanced architectures needed for full geometry and acoustic optimization.
  - Computing time to achieve a viable design solution should target on a few hours or less to satisfy engineering needs
- **Data:**
  - In: Rotorcraft MDAO computations involve asynchronous operations on large amount of time-dependent data for overset-grid assembly, mesh, flow, structural, and acoustic sensitivity analysis, and multipoint design considerations. MDAO on a 10-million-point coarse grid, 720 timesteps, and 2 design points requires about 2 TB time-dependent data written to disk as inputs of the adjoint modules to compute sensitivities.
  - Out: 2 TB data are written to disk at every design cycle; 25 GB data for restart optimization
    - Produced data retained for postprocessing and data analysis
    - Visualization of flow fields and assessments of blade deflections, optimized blade geometry, acoustic characteristics, convergence of design objective and constraints are used to produce mission results
  - Transfer:
    - Data from discipline models transferred to local for postprocessing and visualization; restart data archived in NAS mass storage servers
- **Computation**:
  - Type: Primary CPU computing; the current time to solution requires revolutionary enhancement
    - Cost will increase due to needs of *higher* mesh/time resolutions, *accurate* acoustic prediction, *longer* time duration, and *more* design points
- **Evaluation**: Timely progress toward mission completion requires that HEC expands the hardware capacity by 100-500x in next 5 years

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Many factors affect the load: mesh/time resolutions, fidelities in resolving flow turbulence, acoustic source prediction, noise metric evaluation, and multipoint design considerations, etc.
  - Available resources are not sufficient to meet the mission requirements. Our optimization demonstration has to perform with simpler geometry, reduced resolutions and time integration, targeting only low-frequency acoustic noise because of limited resources.
  - Some coupled analysis/design computations were delayed due to extendedly long wait time (a 36-day wait for 5,000 cores)

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Computing resource capacity, long wait time that affects the total time to solution. Emerging compute architectures should be supported.
  - Impact: Insufficient resources, long wait time, and interruptions in middle of computations will cause mission failure
  - External Sources: Run primarily on HEC

- How will this Use Case evolve over the next decade?
  - Maturity: The high-fidelity MDAO tool is expected to mature and used for production work with substantially more computing resources
  - Growth: Mesh/time resolutions will increase together with longer time duration to achieve future mission objectives (maneuvering rotorcraft design, aeroacoustic optimization for high-frequency noise reduction)
  - Re-use: This high-fidelity MDAO tool offers unique capabilities; its application will grow significantly in government and industries.
  - Re-hosting: The CFD codes have been ported to GPUs and will port to any emerging hardware for improved performance.

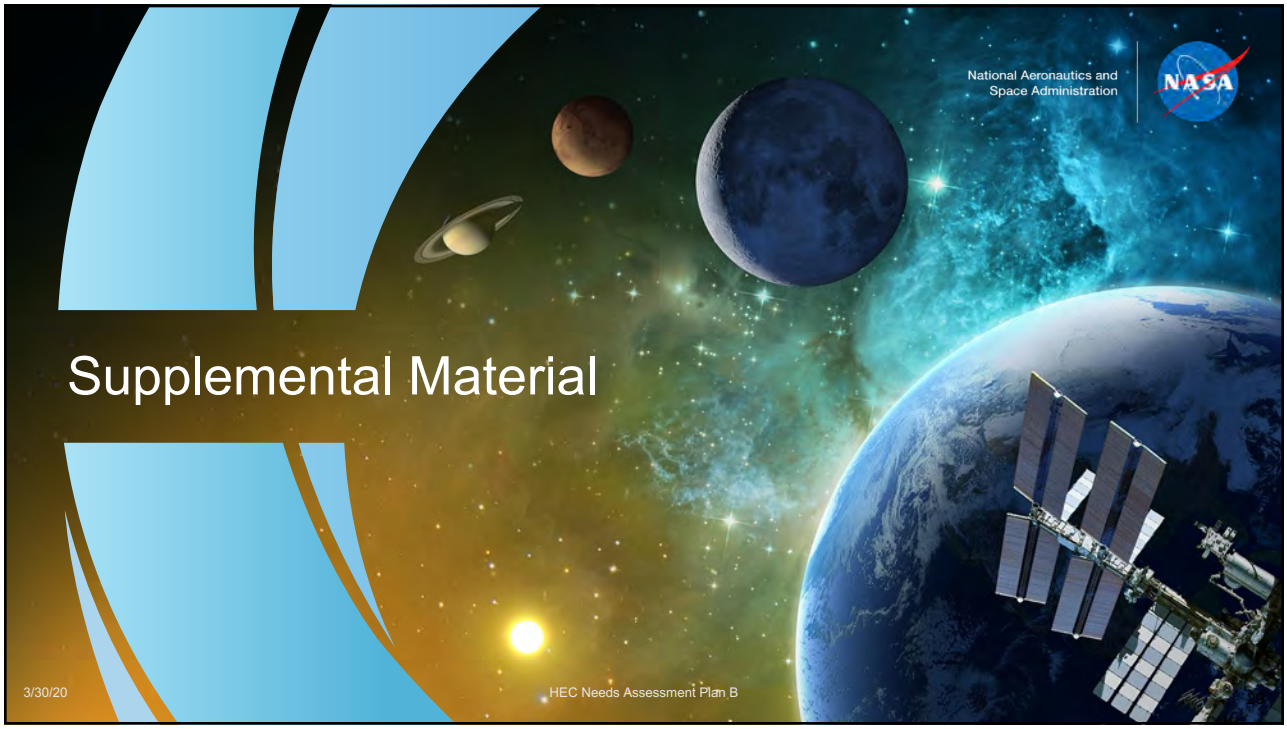

Current Use Case (Survey)

- Development of high-fidelity multidisciplinary analysis and optimization tool for rotorcraft aeromechanics
  - Multiple discipline models perform coupled rotorcraft analysis and sensitivity analysis. Gradient-based optimization is processed by a quadratic programming optimizer to minimize objective function subject to design constraints
  - The MDAO tool aims at rotorcraft shape optimization with thousands of design parameters to improve vehicle's performance and to reduce rotor noise
    - Existing models and capability: overset grid assembly, RANS simulations on dynamic, deforming, overset grids, multibody structural dynamic analysis, acoustic-noise prediction and propagation, adjoint-based sensitivity analysis and optimization
    - Fortran (80%) / C++ (20%)
    - 1,100,000 lines of code
    - Operations of various discipline models are executed either by a Fortran driver with ISO-C-Binding interfaces or shell script
  - List the functional computational and storage requirements for the application
    - 5,000 processing cores
    - MPI parallelization
    - Overall about 5,600~11,000GB memory usage
    - The application primarily uses MPI parallelization in forward analysis and sensitivity analysis
    - A single-point design optimization on a CFD grid with only 10-million points requires 1 TB storage for adjoint sensitivity computations; the storage requirement grows linearly with the number of design points and the number of timesteps.
    - Run in batch

3/30/20 HEC Needs Assessment Plan B 30

Current Use Case (Survey)

- Data requirements for the application
  - Input data include ASCII/binary/HDF5 formatted data and output formats include ASCII/HDF5/binary/Tecplot (plt, szplt) formats
  - Each discipline model (FUN3D/DYMORE/ANOPP2) loads the input data when executed for initialization
  - Input/output data are approximately 1 TB for a single-point optimization based on RANS turbulence modeling
  - Total about 20 TB of input and output data per year (some archived, some discarded)
  - Output data is analyzed through Tecplot tools for flow visualization and assessments of design functional convergence, optimized blade geometry, blade deflections, and acoustic metrics.
  - Data storage relies on local file systems in run time and high-speed shared file servers for archive
  - Input data for MDAO arranged to satisfy requirements of individual discipline models
  - Useful data is identified and then archived from local to the Lou Mass storage systems
  - The data is not currently distributed
- Key obstacles to effective execution of the current application: limited resource capacity, availability, and long wait time
- Are there any analytic functions performed on the output of the application?
  - Time to forward and adjoint solutions for individual discipline models, optimization functional convergence
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - The CFD model can run on GPUs, while the structural and acoustic models run on CPUs. Rehosting the entire multidisciplinary system to advanced computing architectures can be done for improved performance.

Future of the Current Use Case (Survey)

- Multipoint urban air mobility (UAM) aero/structural/acoustic optimization involving takeoff/landing, maneuvering, and cruise flight conditions: if running on 20,000 CPUs, we still need 15-20 days of run time; we must seek emerging architectures to scale up and reduce run-time cost
  - Key computing obstacles: computational resources, capacity, total computing time, and wait time
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in code performance in forward analysis and sensitivity analysis (scalability, full adjoint formulations)
    - Increase in mesh/time resolutions and timesteps to resolve high-frequency blade/engine/airframe noise, turbulence, and long-time integration for maneuvering-flight simulations
    - Rewriting the code or running the application on other computing architectures for improved performance can be made
    - Refactor for emerging platforms
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs can be 100-500x
    - Total memory usage can increase to 600,000 GB requiring faster network communications
    - MPI/CPU and GPU computing is still dominant; rehosting to other architectures also possible
    - Storage requirements become high (e.g., 250+ TB) due to longer time durations needed for maneuvering optimization and higher mesh resolutions for better resolving turbulence and aeroacoustics
    - Run in batch, short wait time (less than several hours but not days/weeks/months), and increased maximum computing time

Use Case: Aerodynamics & Aerothermal Databases

- Mission: Development of aerodynamic/aerothermal databases for Commercial Crew, MPCV, SLS vehicles
  - HEC resources have a critical role in the extension of databases to flight conditions. Without HEC, the vehicles would likely be designed with reduced performance to allow for greater uncertainty and costs would increase.
- Criticality: Essential
  - Without agency HEC these products would be replaced by a combination of local resources, cloud computing, additional wind tunnel testing, and additional flight testing if possible.
- Mission Security Needs:
  - The SLS/MPCV data is considered NASA SBU, CCP TPPI. Some of the data is subject to ITAR restrictions, and SLS has works with some limited-rights data.
  - The SLS aero analysts are constantly using the NAS resources. If the current availability is further decreased, schedules will be delayed, with the possibility of effecting SLS launch schedules.
  - Access for external collaborators is not generally required. Occasionally access may be requested for university professors and non-credentialed contractors
- HEC Function: Computational Fluid Dynamics (CFD) simulations
  - Simulations are run daily by multiple users
  - CFD results are combined into databases of aerodynamic forces and moments and heating distributions This data is used to design the vehicle and to determine when the vehicle can and cannot fly.

Use Case: Current HEC Application Characteristics

- Workflow: Databases are developed to support design cycles, verification cycles, and flight-readiness analysis cycles
  - A new work-flow is triggered when of the above cycles begins for one of the SLS family of vehicles or Orion has new geometry or trajectory.
  - The work-flow is a combination of automated and manual processes with user interaction on a daily basis.
  - None of this use is considered real-time analysis, but it is time critical
- Data:
  - In: Vehicle geometry, flight conditions enveloping all possible flight scenarios
  - Out: The output data is distilled into aerodynamic force and moment, and aerothermal databases. Database are delivered to the Programs and used by multiple engineering groups in the program
    - Typical data retention includes computational flow-fields and convergence history files that are stored on NAS long-term storage facilities
  - Transfer: Large solution and grid files are saved on HEC long-term storage archives and/or transferred to local storage. Databases are stored external to NAS
- Computation:
  - Type: Analyses encompasses a wide range of computational jobs from tens of thousands of jobs that each require 1000 to 10,000 core hours, to a few dozen much larger jobs that each require 100,000 to 1,000,000 core hours
  - The higher priority work requires special reservations
  - Current CFD analyses are constrained by the current HEC capacity. Given more resources, the teams would run more simulations and/or higher fidelity simulations
- Evaluation: HEC has been very effective at producing key data needed by the HEO Programs. However, the growth of capacity has been over past few years has been modest. A more aggressive approach that ensured fully populating the entire NAS expansion facility by FY2022 would be more consistent with the goal of returning astronauts to the moon by 2024.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - The key factor in large variations of SLS usage is the cyclic nature of the vehicle design cycles. While we have constant utilization of HEC, surge is necessary during beginning of Orion design cycles or for mission pre- and post-analysis. Increasing need for unsteady, reacting flow simulations.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: The CFD analysis performed by SLS is constrained by the current HEC capacity. Turn around time is a large limiter for some of the desired moving-body and 6-DoF cases. More processors and longer run times are the only current mitigations.
  - Impact: We often cannot afford to perform some of these analyses due to the perceived cost and tax they would place on resources. More accurate, higher-fidelity analysis in the preliminary design cycles could reduce future design cycle time, or even eliminate entire design cycle
  - External Sources: No
- How will this Use Case evolve over the next decade?
  - Maturity: Actual SLS flights will provide feedback and maturity to the vehicles and databases
  - Growth: The SLS family is schedule to evolve into more a more powerful vehicle that will continue to need aerodynamic analysis for the next 10 years. Orion work will move from design to sustained engineering.
  - Re-use: The number of users supporting Orion/SLS is expected to be stable over next 10-15 years unless new Program objectives are introduced.
  - Re-hosting: HEOMD has not invested, help will be needed along the way

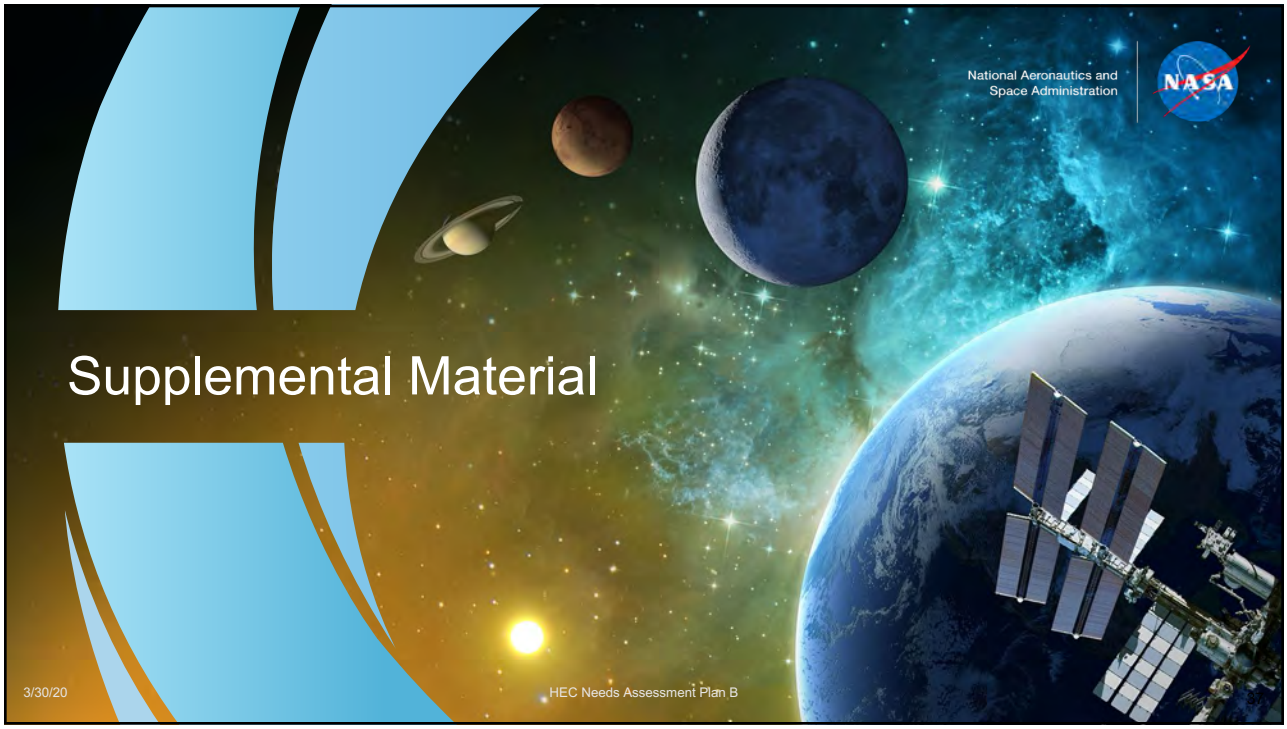
National Aeronautics and
Space Administration
NASA
Supplemental Material
3/30/20
HEC Needs Assessment Plan B

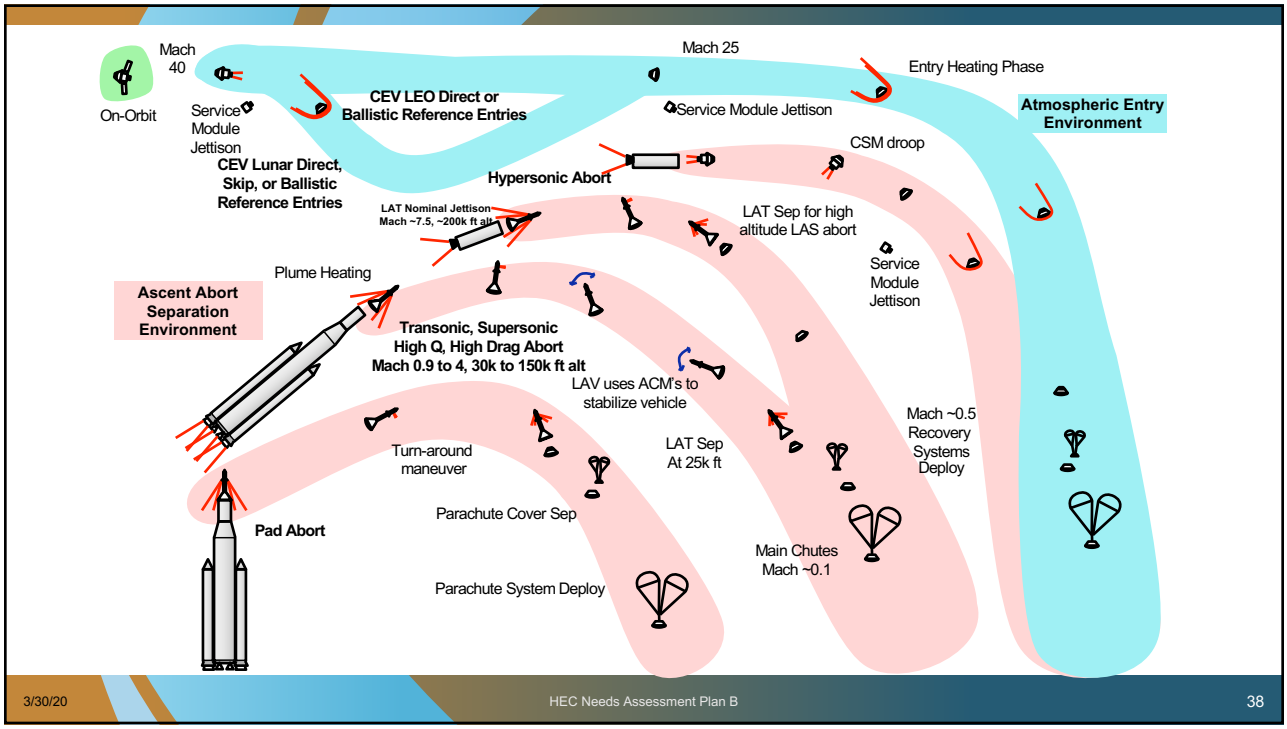
Mach 40
On-Orbit
Service Module Jettison
CEV LEO Direct or Ballistic Reference Entries
CEV Lunar Direct, Skip, or Ballistic Reference Entries
Mach 25
Service Module Jettison
Entry Heating Phase
Atmospheric Entry Environment
CSM droop
Hypersonic Abort
LAT Nominal Jettison Mach ~7.5, ~200k ft alt
LAT Sep for high altitude LAS abort
Service Module Jettison
Plume Heating
Ascent Abort Separation Environment
Transonic, Supersonic High Q, High Drag Abort Mach 0.9 to 4, 30k to 150k ft alt
LAV uses ACM's to stabilize vehicle
Mach ~0.5 Recovery Systems Deploy
Turn-around maneuver
LAT Sep At 25k ft
Parachute Cover Sep
Pad Abort
Main Chutes Mach ~0.1
Parachute System Deploy
3/30/20
HEC Needs Assessment Plan B
38

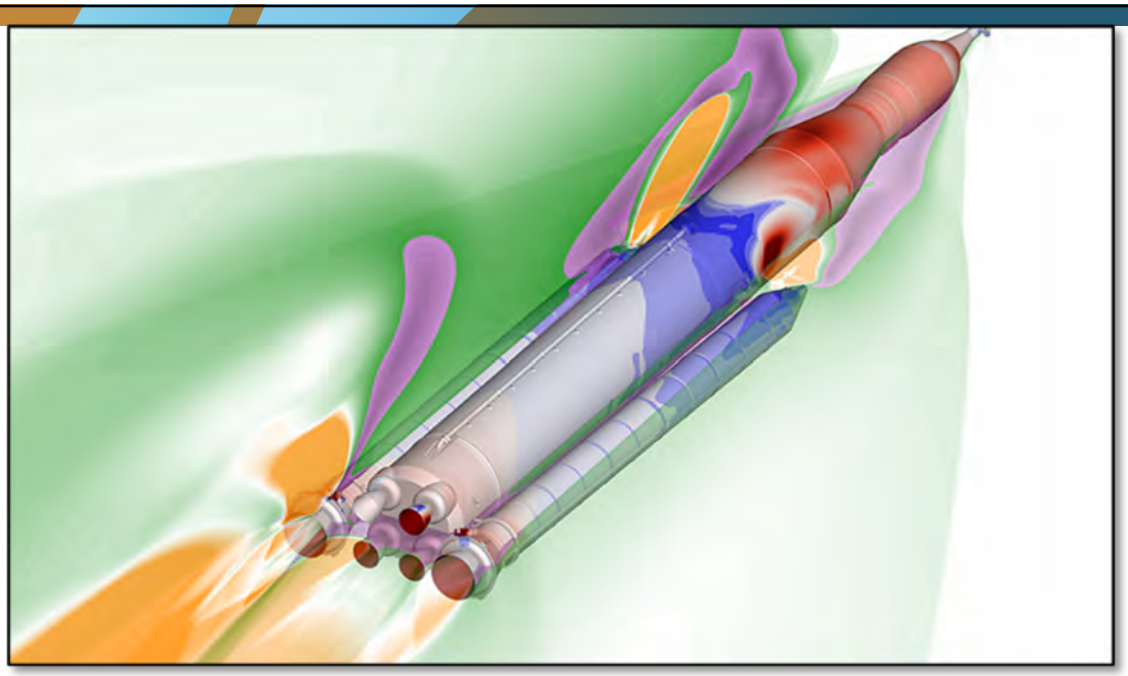

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones
    - FUN3D 970768 lines F95, OVERFLOW 288485 lines F90/889 lines of C, DPLR, LAURA, Loci-CHEM, DAC
    - Combination of shell scripts, Python, Perl depending on application
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - 1 node to 10,000 cores depending on application
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage is covered below
    - Run in batch, interactive, or near-real time

National Aeronautics and Space Administration

Use Case: TTT-RCA Technical Challenge of Accurate CL_{max} Prediction

Mujeeb R. Malik
June 8, 2020
Langley Research Center/D302
ARMD/TACP/TTT

Use Case: Accurate Prediction of Aircraft CL_{max}

- **Mission:** Accurate Prediction of CL_{max} to Enable Aircraft Certification by Analysis is a Technical Challenge (TC) under TACP
  - The TC requires development/assessment of eddy resolving methods that require orders of magnitude more HEC resources than the current RANS based methodology, which is unable to predict CL_{max} with the required accuracy.
- **Criticality:** HEC is *critically* important for successful completion of the TC. Without compute resources, mission cannot be accomplished.
- **Mission Security Needs:** (Confidentiality, Integrity, Availability)
  - Inadvertent release of premature data denies NASA officials to follow established procedures for the export control review.
- **HECC Function:** Enables solution of the nonlinear partial differential equations with billions of grid (discretization) points
  - There will be a need for a large number of simulations required during next 5 years, employing varying fidelities (DNS, LES and WMLES)
    - WMLES, LES and DNS require 10x, 100x and 1000x compute resources compared to the current practice of prediction using RANS
    - The TC goal is to mature the WMLES technology for practical configurations, but DNS and LES required for canonical configurations to assess lower fidelity (RANS and WMLES) tools and develop reduced cost physical models
    - The output data is compared with ongoing experiments to establish efficacy of computational tools, develop more efficient methods, and establish best practices for performing industrial
- **Mission Concept:** Development of accurate computational tools will enable aircraft certification by analysis, saving hundreds of million dollars by reducing flight certification test points.

Euler Equations (inviscid)	*Reynolds-averaged Navier-Stokes: "RANS" (turbulence modeled)*	*Large-Eddy Simulations: "LES" (large scales resolved)*	*Direct Numerical Simulations: "DNS" (all scales resolved)*
1	10	100	1000

Increasing physics, increasing cost →

Courtesy, Boeing

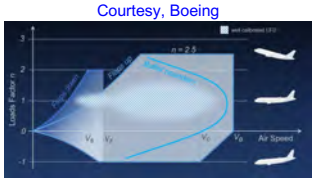

Use of computational tools across the flight envelope will enable certification by analysis

Use Case: Current HEC Application Characteristics

- **Work flow:**
  - Simulations initiated by defining the geometry, grid and flow conditions, which determine the run time. No user interaction except restarts.
  - A recent DNS for a small fraction of the flow separation configuration shown below involved 2.3 billion grid points, which used 48 GPUs for 2 months. Much larger compute times needed for simulations involving complex flow physics. Figure on the left shows the difference between the DNS and the lower fidelity RANS results. Increase in Reynolds number by a factor of 2 would require 8x resources used here.
- **Data:**
  - In: In the above example, 248GB grid and 644GB "flow restart" from scratch; latency determined by the connection speed.
  - Out: 644 GB restart flow data are written to the disk every hour. Latency based on PCI-express connection.
    - Statistical data computed on the fly as the simulation runs. Instantaneous visualization from the restart runs. Need to write much larger unsteady data files, but not possible on GPUs as the process slows down.
    - Statistical data and flow visualization used to provide mission results.
  - Transfer:
    - Restart data to the NAS mass storage server; statistical data transferred to local systems.
- **Computation:**
  - Type: Both the CPU and GPU computing is used, based on availability.
  - The present HECC system is designed to allow "capacity" computing, our need is "capability" computing
- **Evaluation:** Timely simulations would require that HECC **expand the hardware capability by 100x in 5 years**.

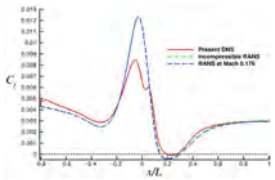

Current CFD prediction technology, RANS, not able to predict flow characteristics to the desired level of accuracy.

Configuration of a flow separation experiment. In this example, simulation of only a small spanwise section was performed. DNS at the desired Reynolds number would require 8x compute resources.

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - Available resources are not sufficient and demand far exceeds the capacity. Some simulations are simply not attempted because of the wait time. Our simulations involve solution of nonlinear PDEs, which require hundreds of thousands to millions of CPU hours (or equivalent GPU capability) for each run.
- **Are there key functional gaps that exist or are projected in the next 5 years or 10 years?**
  - Constraints on Use Case: The wait time is too long. Also, need to embrace the emerging hardware.
  - Impact: We will fail to accomplish the mission, if HECC does not significantly expand compute capability.
  - External Sources: We have used DOE's Edison and Summit supercomputers, with time allotted via competitive proposals. On Edison, for example, we could use 46,000 cores routinely, with little wait in the queue, which is something not available at HECC. Our recent proposals to DOE were not successful.
- **How will this Use Case evolve over the next decade?**
  - Maturity: The ongoing research will mature and be implemented in one or more CFD codes that will be used for production runs to support the mission needs.
  - Growth: Resolution and time-step requirements will increase to support future mission needs.
  - Re-use: Number of people running this application will grow using various codes under development.
  - Re-hosting: Already porting the code(s) to GPUs and will port to any new hardware for improved performance.

Use Case: Mission Need

- Mission: Commercial Crew, Artemis, Mars 2020. All entry missions of past decade have faced parachute qualification issues.
  - Congressional support called out in NASA Authorization bill
  - Current capabilities are restricted in fidelity. Only run one-off cases, can't run parametric analysis for parachute opening loads / failures. HEC is therefore enabling to make this a production capability in the hands of engineers.
- Impact of HEC / Criticality: Enabling – HEC is critical path in parachute simulations
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Without NASA or another agency stepping up, missions would be forced to depend on expensive, long-duration test campaigns while accepting greater residual risk
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Confidentiality: Low risk; Proprietary data of commercial partners
  - Availability: No effect at current level of interruption.
  - External collaborators: Limited academic support
- HEC Function: Provides compute resources: Nodes, compilers, libraries, etc.
  - Frequency: Currently a few simulations per year, aim for weekly cadence.
  - Output: Risk assessment for parachute design, inform design or manufacturing improvements

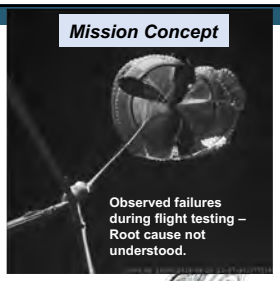

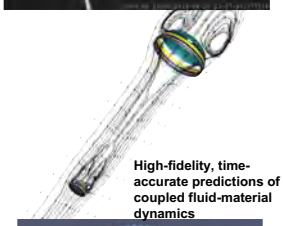

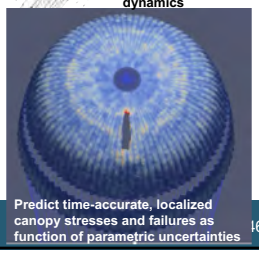

Use Case: Current HEC Application Characteristics

- Work flow: Geometry definition, grid generation, code config, running, post-processing (rate limiting step on the current system).
  - Triggers: Mission design cycles, qualification, Failed parachute flight tests
  - Interaction: Heavy user interaction, both for problem set up, quality assessment, and data interpretation
  - Does this use case need quick reaction or near-real time computing? Explain briefly. No
- Data:
  - In: Mesh of aeroshell, parachute, cables. Volume: Order(GB), both simulation data from disk and node memory
  - Out: Volume: Order(TB). Data files of flow variables for post-processing
    - Retention: Yes, regular instantaneous solution files for analysis of dynamics. Set by mission traceability requirements.
    - Analysis: Statistical and visualization. Parameter identification of dynamical systems.
  - Transfer: Raw data is local to HEC storage system; Post-processed or sub-set data transferred to missions for analysis
- Computation:
  - Type: Distributed parallel computing on Pleiades. High-memory, non-COTS visualization and data manipulation packages
  - Near-term, this is a capacity computing problem. GPUs could enable more rapid, sophisticated analysis in the future
- Evaluation: For current R&D posture, we are making good progress with what we have. However, it is insufficient to meet program requirements for eventual production capability. Computer access and power are rate-limiting steps.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Mission design cycles, field campaigns, increase in commercial space

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Lack of capacity leading to low throughput. Forced to reduce simulation fidelity (resolution) in order to make calculations manageable. No dedicated software + hardware solution for data mining/visualization TB data
  - Impact: Increased costs and risks to programs. Forced to rely on sparse, non-representative parachute flight tests. No insight into driving physics makes design process reactive.
  - External Sources: No

- How will this Use Case evolve over the next decade?
  - Maturity: Program demand for operational capability but only achievable with increased HEC support.
  - Growth: Capability for increased resolution already exists, only limited by compute capacity.
  - Re-use: If successfully transitioned to operational capability, application use will grow with it.
  - Re-hosting: [What is needed to migrate code to future computing architectures?] Substantial code re-write

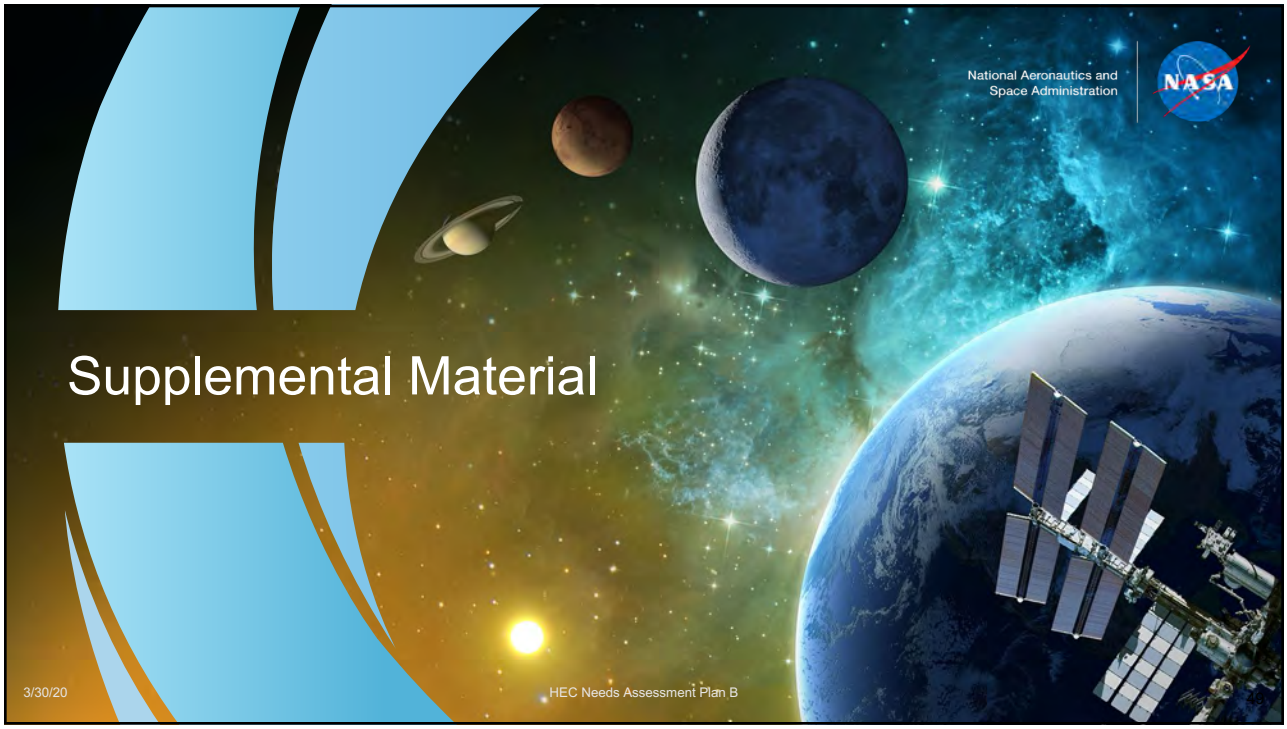

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: Solves systems of partial differential equations to predict system dynamics
  - Describe the application
    - Antecedents: Application has been developed on Pleiades and supplemental support on university-owned clusters
    - Programming languages used: Primarily C++ with some scripting
    - Number of lines of code: Multiple applications, each on the order of 10s of thousands of lines.
    - Workflow mechanism used to control code execution: Shell scripts
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): 1000 cores today at limited resolution. 10,000 cores desired scale
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process): 2-16 GB per process for data generation, depending on parallelization. 64+ GB for data manipulation/visualization
    - Describe the parallelism of the application: Distributed memory application built on MPI
    - Run in batch, interactive, or near-real time: Batch. Interactive for code development

3/30/20 HEC Needs Assessment Plan B 50

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.)
    - Increase in resolution and/or timesteps
    - Rewrite to improve performance
    - Refactor for emerging platforms
    - Rehost onto cloud
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements
    - Run in batch, interactive, or near-real time

T^3 HEC usage

The T^3 Project has used 10.6M SBUs in FY20 (thru May 31), which is 55% of the ARMD usage
- foundational tool development, physics-based modeling

The big users include CFD (including transition prediction), Combustion, and Materials & Structures (M&S)
- in the future M&S and Multidisciplinary Design, Analysis, and Optimization job sizes will grow

All Systems NOP SBUs	PI	PI Company	Title	
1,373,532	Moder, Jeffrey P	Glenn Research Center	National Jet Fuel Combustion Program CFD	Combustion
1,294,328	Murman, Scott M.	NASA Ames Research Center	ARC for RCA	
1,223,994	Debonis, James Raymond	NASA Glenn Research Center at Lewis Field	Large-Eddy Simulations for Propulsion Flows	
931,288	Malik, Dr. Mujeeb R.	NASA Langley Research Center	LES of High RE Flows	
670,414	Georgiadis, Nicholas J.	NASA Glenn Research Center at Lewis Field	Advanced Propulsion Turbulent Predictions Methods	
602,430	Kiris, Cetin	Ames Research Center	TTT Workshop participations for V&V	
489,270	Glaessgen, Edward	NASA Langley Research Center	Computational Materials for Design, Certification and Sustainment	M&S
401,988	Balakumar, Ponnampalam	NASA Langley Research Center	DNS and LES of Turbulence	
223,920	Moder, Jeffrey P	Glenn Research Center	Study of combustor-turbine interactions	Combustion
194,534	Choudhari, Dr. Meelan M.	NASA Langley Research Center	Transition Control in Swept Wing Boundary Layers	
192,736	Borghi, Michael Raymond	NASA Glenn Research Center at Lewis Field	Numerical Simulations of Turbulent Heat Transfer	
176,635	Good, Brian S	NASA Glenn Research Center at Lewis Field	Computational Materials Modeling for Environmental Barrier Coatings	M&S

T^3 HEC usage supports critical ARMD milestones

Program / Agency Milestone Number	Technical Challenges (TCs) and Agency Annual Performance Indicators (APIs)
TACP06 **RCA**	Develop and demonstrate computationally efficient, eddy-resolving modeling tools that predict maximum lift coefficient (CL*max*) for transport aircraft with the same accuracy as certification flight tests. (2025) Supporting US industry to bring improved aircraft products to market faster and with greater confidence in performance.
TACP04 **Combustion Modeling**	Predict the sensitivity of lean blowout (and soot emissions) to changes in fuel composition occurring with the use of alternative fuels (or blends) where the relative difference in fuel sensitivity between simulations and experiments is less than 20%. (2021) Enables early assessment of promising candidate alternative fuels without costly and time-consuming experimental testing.
API 20-3.2.3.3 **RCA**	Complete detailed analysis of turbulent heat flux data obtained from NASA's Turbulent Heat Flux (THX) experiment to enable better computational tools for prediction and design of future air vehicle propulsion systems. (2020) Provide tools that enable accurate prediction of film cooling used for protection of surfaces from hot combusting gases in propulsion systems.

T^3 Combustion: Mission Need

- Mission: Computational tools for advanced propulsion systems and alternative fuels assessment
  - HEC critical to computational tool development and model assessment
- Criticality: HEC provides critical support for this foundational research. Without it progress stops.
- Mission Security Needs: Foundational research databases are not generally sensitive, although some geometries from our partners may be proprietary. Work is generally not time critical and can be interrupted for limited periods.
- HEC Function: Ongoing research not tied to specific events. New capabilities published, as well as transferred to partners and stakeholder community.
- Impact of HEC: NASA's HEC is critical to accomplishing this work.
- Mission Concept:

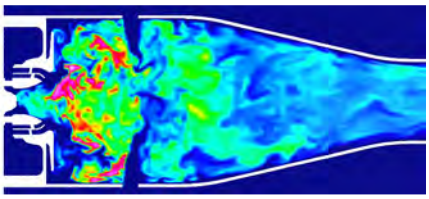

C1 = Gevo Alcohol-to-Jet fuel used in National Jet Fuels Combustion Program (NJFCP) experiments

Temperature contours in Combined Combustor-High Pressure Turbine simulation

Including wall heat transfer and cooling holes (at least in the turbine), would at least double the CPU-hrs (to 2.4M CPU-hrs). Adding soot would double again (4.8M CPU-hrs). Adding radiative heat transfer would put this at 10M CPU-hrs or more.

T^3 Combustion: Current HEC Application Characteristics

- Work flow: Ongoing foundational research. Adjust parameters between runs. Real time not required.
- Data:
  - In: Geometry / grid / parameters
  - Out: Restart files saved for further analysis, statistics files, flow visualization
  - Transfer: local HEC storage, post-processed files transferred to other computers
- Computation:
  - Type: Currently using CPUs. Large computations pushing capacity computing limits
- Evaluation: Rate of completing new cases. Ability to add additional physics.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Ongoing research (fairly uniform load), except perhaps to complete runs for conference publications
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Ability to add realistic geometry details and additional physics is limited
  - Impact: Slows progress on simulating relevant problems
  - External Sources: No, Combustion effort currently not using other compute resources
- How will this Use Case evolve over the next decade?
  - Maturity: Additional functionality / physics will continue to be added
  - Growth: More complex models with additional physics will increase the size of the problems tackled
  - Re-use: Transitioned computational tools and models will be used by others
  - Re-hosting: Code would need to be refactored for efficient GPU utilization

National Aeronautics and Space Administration

Use Case: Computational Materials & Molecular Simulations

John Lawson
June 4, 2020
ARC/TI
ARMD/TTT

Use Case: Mission Need

- **Mission**: *Computational modeling/design accelerates development of high-performance, application-specific materials for a wide range of aerospace applications, e.g. metal alloys, composites, batteries, coatings, biosystems, nanosystems, etc.*
  - *These computational tools dependent critically on HEC. NASA HEC is essential to use and further develop these tools*
- **Criticality**: **If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?** *Buy our own cluster*
- **Mission Security Needs: (Confidentiality, Integrity, Availability)**
  - *Inadvertent data release will result in public confusion, professional embarrassment, National Security, Privacy, etc*
  - *When HEC availability is interrupted regularly, it reduces our rate of progress*
  - *Currently, no significant mission need for external collaborators (without NASA credentials) to access system or to run jobs*
- **HEC Function**: *Significant number of cycles required to use these applications*
  - *We regularly run HEC applications to support mission objectives. It is the major part of our work.*
  - *Computational simulations are important part of integrated technology development (experiments + modeling)*
- **Impact of HEC***: NASA's HEC is increasing essential to accomplish this work. This HEC application is major growth area.*
  - *If NASA supercomputing was no longer available to meet your needs, we would buy our own system or try Commercial Cloud Computing*
- **Mission Concept**: *Shape memory alloys (SMAs) have many benefits for morphing aircraft structures. HEC materials simulation can be used to design SMAs with application specific properties. Use case predicts phase transition between crystal structures (B2->B19')*

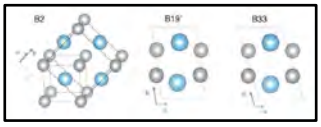

Use Case: Current HEC Application Characteristics

- **Workflow:** preprocessing, simulation run, postprocessing
  - *What triggers this workflow? Desire to know SMA properties for specified composition*
  - *Is there a need for interaction by user, operator or other system in the computing workflow? Convergence checks*
  - *Does this use case need quick reaction or near-real time computing? No*
- **Data:**
  - **In:** *Minimal input data: chemical composition, crystal structure, simulation parameters*
  - **Out:** *Computed properties, simulation diagnostics, restart files*
    - *Properties data files are retained, restarts are not*
    - *Postprocessing, analytics performed on NAS system*
  - **Transfer:**
    - *nobackup directory for short term retention and analysis*
- **Computation:**
  - **Type:** *Electronic structure/DFT computations and molecular dynamics simulations*
    - *Utilize open source or commercial software. Main HEC needs are lots of fast cycles*
- **Evaluation:** *Can we get the cycles we need in a timely manner?*

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - *Varies depending on workload, research objectives, milestones*

- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - **Constraints on Use Case:** *Faster and More is always better*
  - **Impact:** *Unclear. We are running simulations near the limit of current technology in terms of simulation size and simulation time*
  - **External Sources:** *No*

- **How will this Use Case evolve over the next decade?**
  - **Maturity:** *Use case will move from development to operational. Always new use cases.*
  - **Growth:** *It will be applied to more complex systems*
  - **Re-use:** *Possibly*
  - **Re-hosting:** *Depends on the architecture*

HEC Needs Assessment Workshop

Affinity Group Sessions

Monday 6/8/20 1pm – 4pm (AG6-8) (nominal affinity: CFD (10am – 1pm Pacific)
Facilitator: Peter Williams <peter.williams6850@gmail.com>
Scribes: Katie Pitta <katharine.pitta@nasa.gov> and Emily Kuhse <emily.kuhse@nasa.gov>

KEY POINTS

Mission Use Cases covered during HEC Need Assessment

Computational Fluid Dynamic simulations, High-fidelity Revolutionary Vertical Lift Technology (RVLT) Urban Air Mobility (UAM), RANS Turbulence Modeling data generation, rotorcraft aeromechanics, scale-resolving aeropropulsion CFD simulation, aerodynamic and aerothermal databases, maximum coefficient of lift (CL_{max}), parachute dynamics for entry/re-entry systems, Transformational Tools and Technologies (T^3). Computational material simulations.

Impacts of HEC on Mission Needs: Range of Impacts and Significance

- Dependent: High-fidelity modeling and simulations not possible through non-HEC means because little experimental data from emerging technology (Allan).
- Non-HEC turbulence modeling is takes longer and costs more, affecting quality of simulations, increasing risk and lowering confidence, and increasing need for physical experiments, models, and prototyping (Kenway, reinforced by Gomez and most presentations)
- Multi-disciplinary analysis and optimization mission objectives require high-fidelity analysis across multiple disciplines and, therefore, depend on HEC to meet mission
- Dependent: Aeropropulsion needs HEC to conduct simulations and modeling, no other way to conduct computations
- "Without HEC, the vehicles would likely be designed with reduced performance to allow for greater uncertainty and costs would increase." (Gomez)
- Dependence: Transformative aeronautics concepts program depends on HEC resources to fulfill mission (Malik)
- Parachute qualification is critical for all entry systems, yet flight test methods with little other analysis produces little to no data about failure source or causes. Non-HEC qualification leads to high-cost, high-risk designs.
- As alternative fuel technology is developed and introduced, HEC is essential for simulating and modeling dynamics of fuel burn and propulsion systems
- High-performance, application-specific aerospace materials require computational modeling and design using HEC resources because no other way to do it (Lawson)

Mission Security Needs relevant to HEC: Range and Significance

- Confidentiality
- Integrity
- Availability: control factor for project goals and timelines
- Public Confusion/Controversy: Yes, especially if immature data is released
- National Security: some simulations have significant security classifications; SBU (sensitive but unclassified) if common;
- Intellectual Property

Criticality of HEC to Mission Needs

- Financial consequences: Cost of Alternatives compromise program efficiency and effectiveness
- SLS simulations and modeling demand HEC and launch scheduled depend on meeting that demand (Gomez)
- Mission schedule for SLS and Orion integrated with simulation and modeling of aerodynamics
- Mission Security
- HEC is essential to conduct simulations necessary to reduce risk to advanced wind tunnels and flight tests
- Reduced availability of wind tunnels today increases criticality of HEC to CFD missions, especially at high-end (speed, physics, aerodynamics, etc.)
- HEC simulations and modeling allow great confidence in vehicle design and performance, which has long-term operational implications and production-cost implications.
- Projects may require rescoping if HEC availability changes

Workflow Issues and Needs

- Queue wait times discourage production use for CFD simulations: Lower wait times likely to encourage more production use and result in accelerated knowledge production (Allan)
- Results go to mass storage in HEC, but more cases would make this prohibitive (Allan)
- Interaction/quick reaction needed to correct or adjust coding before run is lost
- Choice, selection, or assignment of processor greatly affects runtime and queue time: need to better understand and address
- Changes to SLS or Orion geometry drives workflow of CDF aerosciences (Gomez)
- Changes in parachute design or system geometry, as well as test failure, drive HEC cycles

Data Input, Output, and Transfer: Issues and Needs

- Possibly keep this section focused on processing speed (FLOPS) and data volume (BYTES) as a specific look at these two elements of workflow issues or needs
- *Input, output, and transfer*: output post-processing converted to visualization also requires HEC
- *Computation and Storage:* Storage is a primary technological limit; coupled analyses needed for multidisciplinary simulations and models, same for sensitivity analysis and design solutions

Evaluation: HEC Effectiveness and Value Measures

- Most critical measure is **turn-around time**: HEC effectiveness drops inversely with turn-around time
- Computational time to a viable design solution is critical measure (multiple presenters--Wang
- Full-multidisciplinary simulation and modeling (e.g., rotor blade, atmosphere, acoustics)

Key Factors Affecting Current Load on Computing Assets

- Demand exceed capacity, so some simulations are not attempted (millions of CPU hours or equivalent GPU capacity)
- Demand for better, more robust/higher-fidelity understanding requires great capability to model and simulate physics within the fluid/gas and propulsion systems, leading to greater HEC needs
- Publication cycles and University calendars
- Mission design cycles and growth of commercial space missions

- Wind tunnel test planning: simulation work can support wind tunnel planning/design
- "Available resources are not sufficient to meet the mission requirements. Our optimization demonstration has to perform with simpler geometry, reduced resolutions and time integration, targeting only low-frequency acoustic noise because of limited resources." (Wang)
- Cost and turn-around time limits adoption of scale-resolving simulations

Anticipated Evolution of Use Case and Mission Need

- Maturity: transition from development to operational capacity limited by system capacity, so smaller, less complicated problems are predominant.
- Growth: increasing data resolution means more data input with concurrent expectation that it will be incorporated into models and simulations, significant MDAO application growth expected in industrial and government sectors as multidisciplinary concerns turn to multidisciplinary models
- Re-Use: Increasing expectations that available data from NEX and other sources will be leveraged, often in ways that depend on HEC;
- Re-Hosting: Moving code to GPU clusters needs study/research; movement to future architectures anticipated to require substantial code rewrite.
- Other(s): Decentralization of workforce/colleagues, systems, and data resources, as well as data repositories; reduced reliance and availability of wind tunnels is expected to remain or fall further, leading to growing dependency and need for CFD simulation and modeling.

Functional Gaps: Existing and Anticipated

- Constraints: big data analysis (including ML), disk space, computational allocation, queue (wait time), network, I/O, #GPU's (ML and DL), embracing emerging hardware
- Impact of Constraint or Gap: insufficient performance to obtain the required models and data analysis, mission failure, ML/DL is expanding which means the training expands on HEC resources also needed for research production; more ML/DL training demands on same HEC resources reduces research productivity because competition for same computing resources
- External Sources

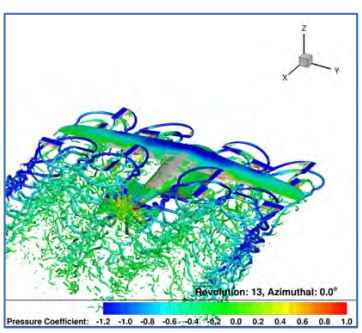

Figure 1. NASA RVLT Lift+Cruise Concept Vehicle simulation (Allan)

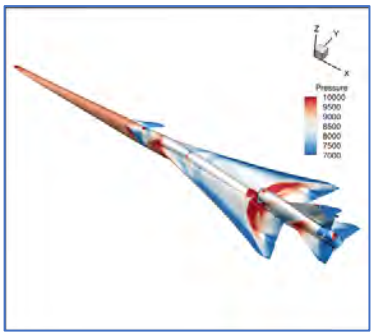

Figure 2. LMC608 from the 3rd Low Boom Prediction Workshop. Design is comparable to the final LBFD configuration (Kenway)

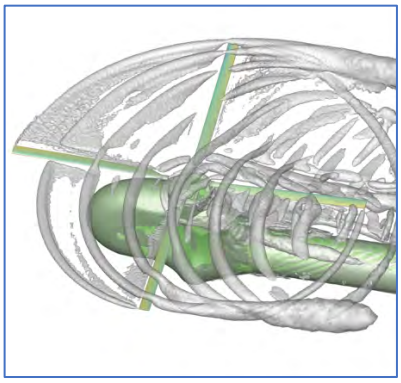

Figure 3. Multidisciplinary Computational Fluid Dynamics (Wang)

Figure 4. Including wall heat transfer and cooling holes (at least in the turbine), would at least double the CPU-hrs (to 2.4M CPU-hrs). Adding soot would double again (4.8M CPU-hrs). Adding radiative heat transfer would put this at 10M CPU-hrs or more.

Attendees

1. Mike Little, Convener, NASA Goddard, ESTO and HEC Program
2. Peter Williams, facilitator
3. Presenter #1: Brian Allan, NASA Langley Research Center, D303
4. Presenter #2: Gaetan Kenway, NASA Ames Research Center, Computational Aerosciences Branch
5. Presenter #3: Francois Cadieux, NASA Ames Research Center, Computational Aerosciences Branch
6. Presenter #4: Li Wang, NASA Langley Research Center, D302
7. Presenter #5: James DeBonis, NASA Glenn Research Center, LTN0
8. Presenter #6: Reynaldo Gomez, NASA Johnson Space Center, EG311
9. Presenter #7: Mujeeb Malik, NASA Langley Research Center, D302
10. Presenter #8: Michael Barnhardt, NASA Ames Research Center, Aerothermodynamics Branch
11. Presenter #9: Michael Rogers, NASA Ames Research Center, Computational Aerosciences Branch
12. Presenter #10: John Lawson, NASA Ames Research Center, Intelligent Systems Division
13. Aaron Pina, NASA Headquarters, DK000
14. Daniel Duffy, NASA Goddard Space Flight Center, NCCS
15. Elizabeth Hartman, NASA Ames Research Center, NAS Division
16. Elizabeth Hook, NASA Headquarters, DK000, HEC program
17. Elizabeth Lee-Rausch, NASA Langley Research Center, D302
18. Emily Kuhse, NASA Ames Research Center, NAS Division
19. Henry Jin, NASA Ames Research Center, NAS Division
20. Irina Kitiashvili, NASA Ames Research Center, Advanced Computing Branch
21. Johnny Chang, NASA Ames Research Center, NAS Division
22. Katherine Pitta, NASA Ames Research Center, NAS Division
23. Piyush Mehrotra, NASA Ames Research Center, NAS Division
24. Robert Ciotti, NASA Ames Research Center, NAS Division
25. Robert Ferraro, NASA Jet Propulsion Laboratory, 8310
26. Robert Hood, NASA Ames Research Center, NAS Division
27. Sherry Chang, NASA Ames Research Center, NAS Division
28. Subhash Saini, NASA Ames Research Center, Advanced Computing Branch
29. Ted Manning, NASA Ames Research Center, Advanced Computing Branch
30. William Thigpen, NASA Ames Research Center, NAS Division

1. Use Case: RVLT High Fidelity Modeling- Lift+Cruise UAM Concept Vehicle

Brian Allan
brian.g.allan@nasa.gov
LaRC, D303

Main points:
P1: High-fidelity modeling plays a critical role in this work, especially since there is very little experimental data for this type of vehicle.

P2: Without access to HEC resources, work would have to be de-scoped greatly.

P3: Queue capacity and turnaround time are the biggest limiting factors of this work.

P4: In the future, the project would like to look into more acoustics simulations and more high-resolution cases.

P5: Work can benefit from GPU/accelerator hardware. There is a limited effort looking at porting OVERFLOW to run on GPU clusters.

Raw notes:
Slide 1:
High fidelity modeling of rotorcraft to validate and compare to conceptual design tools.

Not having a lot of experimental flight data, using CFD to test and design multiple concept vehicles.

HEC is critical: no other platforms or resources at NASA to compute these kinds of problems due to large size; smaller computers could not support needs, would have to downscale the goal of the project without the resources available.

Cloud computing not an option at the moment, using ITAR software (OVERFLOW)

Biggest problem: turnaround time, waiting in queues 4-9 days recently

Slide 2:
Performing flow visualization on supercomputers, rather than local machines; storing data on Lou and Mass Storage. This greatly reduced the need to download large amounts of data to local machines from NAS.

Can benefit from GPU hardware in the future. Would like to start looking at it.

Slide 3:
Some of the work is done in support for wind-tunnel tests, which can increase number of runs needed during a given period.

Reduced queue times would increase number of use cases possible, currently limited to how many cases can be run. Capacity and max queue time is a big limiting factor.

Future: Want to start looking at high-resolution acoustic simulations. Some effort toward looking to run OVERFLOW on GPU clusters.

Q&A:
Piyush Mehrotra: Do you have need for capability computing, that is large core count jobs?
Brian Allan: Using between 1k-2k cores right now. It sometimes it doesn't scale. The high-resolution cases will dictate how many cores we can go to.

Piyush Mehrotra: Any codes other than OVERFLOW?
Brian Allan: We do use a rotorcraft analysis code, but right now, OVERFLOW is the heavy code that we're using.

Note: For rotorcraft applications, RVLT mainly uses OVERFLOW and FUN3D but for this Lift+Cruise UAM use case we are currently only using OVERFLOW and are working on future plans to use FUN3D.

---

2. Use Case: High Efficiency Database Generation with LAVA

Gaetan Kenway
gaeten.k.kenway@nasa.gov
ARC/TNA

Main points:
P1: What do we want HECC resources to be, capacity or capability? Goals for each can be somewhat conflicting. Can we increase the overall capacity so that we can run the multitudes of jobs needed?

P2: Goal of work is to reduce reliance on wind tunnel and physical testing. There are fewer wind tunnels available, so improvement in CFD computations will be critical going forward.

P3: Current time to generate databases are being measured in months. Would like to get that down to weeks and believe that this improvement will be critical going forward.

P4: Throughput is key to efficient work: Any one job of many waiting in queues is fine, but if most or all are waiting, it can greatly impact efficiency and deadlines.

P5: Project is looking at GPUs and is hopeful that these simulations can be efficiently ported over to the next generation of hardware.

Raw notes:
Slide 1:
Capacity computing side of HEC. Perform RANS simulations to generate large databases for ARMD projects, including NPAI, X-57, X-59, Low Boom Flight Simulator.

Numerical Wind Tunnel Capability; reduce the amount of reliance on wind tunnel and physical tests, also help to target physical tests. Fewer wind tunnels than we had in the past, so being able to do these CFD runs in databases is more important than ever.

Without HECC, smaller numbers of runs would be performed, projects would have additional risk of the unknown without the databases.

Slide 2:
Database generation time is measured in months, would like to get it down to be measured in weeks. Projects that this will be critical to future work.

Recently, over 250 simulations run, generated database in weeks

Slide 3:
Input data on the order of GB, none of them are large, there's just lots of them. Jobs are typically relatively small (1k-1.5k cores and 1-2 hour runs), moderately sized grids (20- to 200-million grid points).

What really matters is the throughput. If one job is sitting for days, but the rest are running, that's fine, but if most or all are waiting, it can really impact time of database generation.

Slide 4:
Future: Will hopefully see more scale resolving simulation using LAVA curvilinear flows, with a mix of steady-state and scale resolving, the latter of which can be 30 times more expensive than running a steady-state RANS simulation. Projects see that these simulations can be run and want to do more of them. More and more scale resolving simulations will creep into these CFD databases.

Pushing into the capability computing: 2,500 cores, takes a couple of days to run.

With increased computing capacity, we'll see more of the scale resolving simulations

Rehosting: looking to use GPUs and different types of accelerators. Relatively confident that scale resolving simulations will be ported over to accelerators and the next generation of hardware.

Q&A:
Michael Little: What is limiting the turnaround to months that needs to change to get to weeks?
Gaetan Kenway: A Couple of things. Biggest -- have been doing much better and spending lots of time optimizing older code to run more efficiently. Queue times – work discussed in slide did 250 cases in a week, but 5-6 days was spent setting in a queue. Times to get moderately sized jobs through can be increased by a factor 10x or more due to waits, limiting what can get done.

---

3. Use Case: Scale-Resolving Simulations with LAVA Cartesian AMR

Francois Cadieux
francois.cadieux@nasa.gov
ARC, TNA

Main points:
P1: Goal is to conduct simulations with a short enough turnaround to impact vehicle design, as well as augment or replace some of the physical testing being done.

P2: This work could not be done without HEC, partly because it's mostly ITAR protected.

P3: Frequent large amounts of output required by this work can slow the project down because writing to disk is one of the slowest operations you can do.

P4: It is difficult to do these types of jobs without a reservation, which is not seen as sustainable long term, as doing so presents a scheduling challenge.

P5: These simulations can be critical for NASA missions. At KSC, construction was halted for 5 weeks while CFD simulations were used to ensure safety of the final product due to discrepancy between design and actual construction.

Raw notes:
Slide 1:
Predicting ignition overpressure and acoustic vibration loads on launch vehicles (HEOMD).

Looking to be able to do these types of scale resolving simulations in a short enough turnaround time to help with the design of the vehicle.

Looking to augment or replace experimental tests, or to support and impact design of vehicles and environments. Looking for validation and having an impact on the project.

LAVA is developed from a specific capability need, no budget to work on features unneeded at the current time.

Convenient and responsive experience with HEC, looking for more capability computing resources.

Would not be able to do the scale resolving simulations without the use of supercomputing resources and storage -- much of the data is ITAR.

Would like to move forward to producing the scale resolving simulations more in databases. Close to 1 billion cells, running 3-6 weeks on 4-16K cores. In order to get these results, need long enough time to get convergence of statistics and acoustic spectra.

Slide 2:
KSC came to LAVA team to rerun simulations when construction was different than designs to ensure vehicle and personnel safety.

Generally, don't have a lot of input data, only data about boundary conditions or how components are moving. Lots of volume and surface data, comes out 200 TB per run. Writing to disk is a slow operation.

Aimed at capability computing, hard to run simulations of this scale without a reservation, which is unscalable and a scheduling challenge.

Slide 3:
Design reviews have biggest impact on load, but also publication cycles.

Trying to run more and more simulations to reduce risk; running against limit of capacity vs capability computing.

One factor impacting computing is algorithmic and numerical improvements.

Maturity: getting to a point where the physics modeling is there, but can we get enough resources to do that that is sustainable. For-sees complexity and length of job times will increase greatly. Simulations needing months to run, including an increase in number and complexity of simulations.

Q&A:
Michael Little: What kind of help would enable moving to new architecture?
Francios Cadieux: Could use more skilled developers to focus on new architecture and accelerators. There are questions about what programming model to adopt (OpenMP, OpenACC, etc.), don't have the experience or the time to explore all of them. More funding for current developers would also accelerate porting to new architectures. LAVA development is application-driven so our developers' time is generally tied up doing R&D and implementing new features required to perform simulations related to specific projects from which we are receiving funding. Funding earmarked specifically for porting and optimizing the code for new architectures would allow us to allocate more of our developers' time to that task – we currently receive very limited support from NESC.

Michael Little: Do you think, based on the algorithms and codes, that GPU would help you get to the timeframe you mentioned?
Francios Cadieux: We could possibly get 3-4x improvement without I/O given that our code is already deeply optimized for many-core CPUs, but don't suspect we could get more than 2x if we need frequent I/O. Currently constrained by how much data needs to be output. Can't get all the benefits from an asynchronous approach since we need to output surface data (<1GB) and ideally volume data (100GB) as frequently as possible for acoustics.

Brian Thomas: Mentioned data write being slow/constraining, what about improved data management? Skills from outside, or hardware or ?? would help? (e.g. what helps)?
Francios Cadieux: So far been pretty happy storing data on HECC systems. As far as storing data, anything that would reduce write times (100s of GB across many processors) would be great. Been a

hard task to find or develop technology that will allow us to output very frequently, efficiently, and that scales to 16K cores or more.

4. Use Case: Multidisciplinary Analysis and Optimization for Rotorcraft Aeromechanics

Li Wang
li.wang@nasa.gov
LaRC, D302/Computational AeroSciences Branch

Main points:
P1: HEC is critical for completion of mission. Loss of HEC resources would provide a real barrier. Potential solutions would be to buy a new system or write a proposal to access DOE resources.

P2: Access to resources would need to scale up significantly to reduce run times.

P3: Some coupled analysis/design computations were delayed due to long queue times. Greater job capacity is needed (100-500x in the next 5 years) to meet mission requirements.

P4: Use of the MDAO application being developed is expected to grow significantly in government and industry.

Raw notes:
Slide 1:
Goal: to produce powerful MDAO tool, combining/coupled tools from NASA and university partners

Envision technical challenges as work moves forward: Current CFD model uses RANS, which has been observed to not be as accurate as needed. Need to support several orders of magnitude of resolution to improve accuracy.

Due to capacity and run time limitations, been required to less than 10 million grid points.
Enabling capabilities needed requires more computational resources.

Slide 2:
Asynchronous backward in time at each timestep, 2 TB write to disk at each cycle.

Finds it almost impossible to use Tecplot directly at NAS, needs to pull down to local machines for visualization.

Mesh and time resolution; need to increase usage in order to increase optimization. Hardware expansion would need to be expanded by 100-500x.

Slide 3:
Run with much smaller geometries, low frequency noise reduction, experience extended long delay getting jobs to start (waiting 36 days) -- affects the total time we can get a solution.

The MDAO tool will be improved in both accuracy and efficiency in future. Will require long time integration, higher mesh.

Been ported to GPU and shown success with tests. Will do everything they can to adapt code to emerging hardware to reduce run time.

Q&A:
N/A

5. Use Case: Large Eddy Simulation for Propulsion Flows

Jim DeBonis
james.r.debonis@nasa.gov
GRC, LTN

Main points:
P0: Work is in scale-resolving turbulent simulations (large-eddy simulation (LES) and direct numerical simulation (DNS). At present, these types of simulations are the only clear way to improve turbulent flow predictions. These simulations are very computationally expensive. If NASA is to improve their flow simulation capability, a huge increase in computational capability is required.

P1: Work is extremely computationally intensive. Without HEC, the project would need to rent time from a commercial provider or DOE. Budget would be a major constraint.

P2: Use of reservations helps avoid queue wait times and is increasing, but this will merely shift the problem as more people begin to use them.

P3: Wait time depends on which architecture you choose to run on, and it's better to run a little slower if your job starts (and finishes) sooner. Some are better than others at taking advantage of this.

P4: If turnaround time were quicker, more complex problems could be run.

P5: Faster turnaround time would allow engineers to be more efficient and work serially. There is something lost when transitioning between tasks over a long period of time. "People aren't good parallel processors."

Raw notes:
Slide 1:
Scale resolving CFD propulsion tools for ARMD.

Fundamental work into LES, fluid flows -- doing a lot of work in turbulent heat transfer.

Extremely computationally intensive work -- without HEC would need to rent time on commercial or DOE computers, would reduce time/money availability and overall use.

There is the potential to use restricted applications, but currently using an open source code.

Nice to be able to offer time to industry or academic partners for collaboration.

Basically, running continuously throughout the year -- would have to scale back or spend lots of money on buying equivalent service.

Slide 2:
Needs multiple long-queue, five day runs, or the use of reservations. Helps avoid waiting in queues, but reservations are just going to shift the problem as more requests are made to avoid queues.

Shell-script driven bash runs.

Very useful to be able to update parameters while you're running, very frustrating to have code blow up while running after waiting a long time.

Wait time in the queue depends on which processors you pick; some researchers are better than others at guessing which processors will be available quicker -- would rather run slower but get data faster.

Run two different codes: finite difference code, flux reconstruction code.

Some of the work needs 30 or more time-averaged variables, files can get big.

Pull down to GRC workstations for post processing, looking into more NAS processing as simulation sizes increase.

Participated in GPU hackathon at Langley to port code over to GPUs, working to parallelize some of NASA's codes -- 99% Fortran, 5-15K cores per run, 10-200 GB

Slide 3:

Both the users and developers -- fairly even usage across the year.

As scale-resolving simulations become the more popular option, turnaround time will continue to increase.

Limitation: what can get through the system in a given time -- turnaround time determines how much work can be done. If we could get faster turnaround time, we could address more complex problems.

With minor exceptions, using mostly HECC resources.

As capability improves, will be limited by the capacity to do more complex geometries. Capacity of the machines will dictate how things progress.

Q&A:
Brian Thomas (across several comments): Can you estimate costs to the project for waiting? Do you have idle labor (at some level)? At what point does wait time slow the project (e.g. you run out of stuff to work on and need next CFD result from HEC in order to move forward...how tightly coupled is your workflow to getting computation done?) Another way to put this...what is the timescale of how long you wait in which it doesn't matter if the wait time is shorter? (you get as much work done even if the wait time is less)
James DeBonis: You submit a job and wait a week or more to get an answer, so there is undefined loss; “Why did I change that input?” Were good notes taken? -- delay messes with you. Would work more efficiently if there wasn't a delay, could work more serially. Lots of people aren't good at parallel work.
Gaetan Kenway: Mental task switching has a cost
Brian Thomas: "Context Switching"

---

6. Use Case: Aerodynamics & Aerothermal Aerosciences

Reynaldo J. Gomez
reynaldo.j.gomez@nasa.gov
JSC, D302

Main points:
P1: HEC is critical to the design and analysis of launch and entry vehicle operations that we cannot test in existing wind tunnel facilities. While HEC systems are expensive, the savings for eliminating a flight test can be on the order of $100 to $100s of million/flight test.

P2: Launch and entry vehicle database work is driven by design cycles and is constrained by the Program schedules and by existing computing hardware and system allocations. Ground winds airloads and engine ignition simulations are two relatively expensive data products where we are limited to 10s of solutions resulting in relatively large uncertainties in these databases. Progress has been made on the simulation of acoustics and buffet environments for launch vehicles, but these are very computationally expensive and we cannot afford to produce a database with the current tools and hardware.

P3: The Space Launch System (SLS) continues to analyze new configurations and plans to fly to new booster configurations that will generate continuing work for several more years. Plume chemistry and unsteady simulations are growing in importance and will drive up the costs of developing and refining aerodynamics and aerothermal databases for the MultiPurpose Crew Vehicle. Adding these physical modeling refinements will drive requirements up by 10x or more.

P4: We have been interested in GPUs for some time and many of our NASA developed codes have made attempts to use them but early hardware and code tools did not result in compelling improvements so they are not currently being used in production work. If GGPU systems are the way to get improved throughput we will consider participating or funding the efforts required to use them effectively.

P5: Would like to see a more aggressive growth in capacity.

Raw notes:
Slide 1:
HEC is critical to designing things that are challenging or impossible to test with current facilities. Considered essential to work. Could buy significant local assets or cloud computing for cost of flight testing.

ITAR, SBU, limited rights data

Reduction in current access would delay schedules, even impacting launch schedules. Some scheduling/priority changes based on administration focus.

Simulations are run daily by multiple users across many centers.

Relatively light weight data, but runs are archived in case they are needed in the future.

Slide 2:
Driven by design cycles: New geometries or trajectory can impact need for resources. Post flight analyses can be extensive, particularly for test flights, where as-flown conditions are analyzed and used for validation purposes.

Try to produce the best product that we can given the time and resources available. This sometimes/often means that the end product is good enough rather than a . More resources available quicker would improve end product.

Constrained by current capacity, would like to do more.

Would like to see a more aggressive approach, but more money is going into hardware and money for computational resources doesn't make the list.

Slide 3:
Lots of steady analysis, but current configurations are driving a need for more chemistry and unsteady analyses, driving costs up.

Would like to do more moving body cases, need to invest in the tools to do that.

Would like to do some things that can't afford the time/cost.

Need to be able to have quick turnaround -- data not as useful after the design has been built. Would like to substantially influence design whenever possible.

Always been interested in GPUs, OVERFLOW and FUN3D have looked at GPUs but not using the GPUs in production work.

Primary tools are coming from NASA or academic partners.

Q&A:
Tsengdar Lee: In the past you need remote visualization software and low latency network, are you still in need of those capabilities?
Reynaldo Gomez: There are groups that would like to have that. Many remote users use in-situ post processing to process as much at NAS as possible. We try to avoid transferring large results files whenever possible.

Robert Ciotti: Do you forecast code feature requirements looking forward, or you simply use what's available?
Reynaldo Gomez: Both. We use what is available but if we see a major need in the future we will work with developers to get there.

Robert Ciotti: How does NASA fund that?
Reynaldo Gomez: In the late 1980s the Space Shuttle Program identified a need to develop a capability to simulate Shuttle ascent flow fields and funded the early development of the OVERFLOW CFD code. We have a number of different funding methods if we want to support new development. Sometimes we will partner with the developers or fund them through our Programs to develop new capabilities. For smaller feature additions we might fund development work with our Branch, Division, or Center discretionary funds.

---

7. Use Case: TTT-RCA Technical Challenge of Accurate CLmax Prediction

Mujeeb Malik
m.r.malik@nasa.gov
LaRC, D302

Main points:
P1: Would save hundreds of millions of dollars if 50% of flight certifications can be done with CFD. [Note 50% of the high-lift (HL) portion of the flight certification test points. HL certification test points are 2/3rd of the overall test points].

P2: HEC is critical. DOE proposals are hit or miss and the project does not have the budget to go and buy time.

P3: Need 100 x the current capability in the next 5 years. Because we don't have that, we simply do not try to do some of the simulations we need.

P4: Codes are already being ported to GPU and will be ported to any new hardware for improved performance.

Raw notes:
Slide 1:
Scale resolving simulations; developing advanced CFD methods, high quality experiments and use that data to validate CFD tools.

Predict aircraft maximum lift -- would like to get certification by analysis. Would save hundreds of millions of dollars if 50% of flight certifications can be done with CFD.

Results have indicated that current RANS methods will not resolve the large scales of turbulence to predict at the accuracy that is required for certification by CFD

High lift area of both landing and take-off -- flow separation and unsteady flow, about 2/3 of flight certifications are done and RANS doesn't predict that.

1000x the RANS capability is not sufficient - why we don't do full aircraft DNS.

Ultimate goal to develop capability

Used NASA and commercial codes

HEC is critically important, cannot meet the mission without it, project doesn't have money to go out an buy time

Most of work is open sourced, don't want premature release of data

Slide 2:
Test case of DNS of simple configuration -- Current RANS not able to predict flow characteristics to the desired level of accuracy

Used GPUs running on DOE Summit and NAS

To do a simulation for 2x Reynolds number and the span for the simple configuration shown, need 16x resources for DNS compared to the lower Reynolds number DNS

Not storing all instance data, just not possible, on GPUs because it slows down the performance drastically.

We need 100x the capability in the next 5 years. Because we don't have it, we simply don't do simulations that we need because we can't.

Slide 3:
When we were granted access, we were able to run routinely 46K cores on DOE machines, not possible on NAS machines (because of large number of users). Have written DOE proposals in the past, but it's hit and miss. Project does not have money to just go out and buy time.

Experiments being done in parallel -- lots of users needing compute time in the coming years.

Anything we can do to increase capability computing (vs capacity).

Q&A:

William Thigpen: What is the number of GPUs that would work well for your case?
Mujeeb Malik: Right now, with 256 GPUs we could do that simple configuration. [should have said 256 nodes (1 node = 4 GPUs), but any increase in resource helps; we scale the simulation based on available compute resources but can't have useful outcome below a certain number]

---

8. Use Case: Modeling Parachute Dynamics for Entry Systems

Michael Barnhardt
michael.d.barnhardt@nasa.gov
ARC, TSA

Main points:
P1: Every entry mission in the past decade, if not further back, has faced problems with parachute qualification. Flight test driven qualification with very little additional analysis makes it very difficult to pinpoint cause of failure.

P2: Without access to HEC resources to perform this work, there would be greater cost and higher residual risk to agency missions dependent on old methods.

P3: Grid generation can be extremely challenging and post processing is also a limiting factor. Output can be terabytes of data.

P4: Work is driven by design cycle and will continue to increase with the growth of commercial space.

P5: Major constraint is capacity. Work is forced to reduce simulation fidelity to make analyses doable, which causes risk to programs who are relying on tests that aren't as representative as they need to be.

Raw notes:
Slide 1:
Every entry mission in the last decade has faced issues with parachute qualification; Mars2020 had several failures -- hard to understand what the root cause was. Expensive to have to build and fly to test. Flight-test driven qualification with very little additional analysis makes it very difficult to pinpoint cause of failure (such as, “What are the capsule dynamics?”).

Impacts both SMD and HEOMD, modeling project under STMD -- supporting NASA missions across the board.

Outside of the work being done by this project, there is no precedent of this level of analysis of parachutes, still depending on older approaches.

Multi-physics capability, couple high fidelity fluid dynamics; RANS does not accurately predict these systems, need hybrid RANS or LES; scale resolving, real-time, unsteady simulations in order to get the physics correct.

Can we predict failures? Bound the expected range of behaviors.

Current capabilities are limited. No way to run parametric analysis without HEC resources.

HEC resources are a critical path; would be depending on another agency to step up. Status quo would mean expensive long-term campaigns with greater risk.

Parachutes are a high-risk system, but the analysis is pretty focused.

Still almost totally in the realm of academic research.

Would want to get to a week cadence of runs but is driven greatly by design cycles. Small number of high-fidelity cases, per year.

Some potential for informing the design, but mostly risk analysis.

Slide 2:
Grid generation can be extremely challenging.

Post processing is also a limiting factor. TB levels of data that have to be sifted through afterward, how to do it quickly and efficiently.

User interaction is heavy -- quality assessment requires a lot of user oversight.

GB on average for the simulation data being read, but the node memory during run time.

Unsteady simulations and understand dynamics of what's going on, TB write periodically during runs.

May be re-assessing design every couple of years.

Slide 3:
Mission design cycles driven; seeing an increase of commercial space that could impact parachute test cases.

Major constraint is capacity. Forced to reduce simulation fidelity to make analyses doable, risk to programs, reliance on tests that aren't as representative as needs to be.

Q&A:
Brian Thomas: What tools do you utilize to shift through your output data? Would this post processing able to benefit from better tooling (data bases, services, etc.)?
Michael Barnhardt: Handful of different software; each approaches it separately. Usually individually written data extraction tools that go with the software you're using. Depend on commercial tools like Tecplot for visualization. Thinking ahead to high-order methods (Discontinuous Galerkin), since we're finding that available tools aren't great at being able to handle and make sense of that sort of information.

Robert Ciotti: Any plans to share results outside NASA, potentially years later? RE: commercial access to space?
Michael Barnhardt: We already do that.

---

9. Use Case: Transformational Tools and Technologies

Michael Rogers
michael.m.rogers@nasa.gov
ARC, TNA

Main points:
P1: TTT conducts foundational research and does not have funding to conduct these computations on other platforms.

P2: Program uses the largest amount of time within ARMD (55% YTD).

P3: Ability to add realistic geometry details and additional physics is limited and slows progress on simulating relevant problems. These requirements are expected to increase.

P4: An increased capacity of resources is needed to perform work (and to support future work).

P5: Codes would need to be refactored for efficient GPU utilization. May look at porting to new hardware that could potentially improve turnaround time.

Raw notes:
Slide 1:
Transformational Tools and Technologies Project, 55% of ARMD usage so far in FY20.

Anticipate growth in coming years across several projects.

Slide 2:
Supporting critical efforts for ARMD, level 1 milestones (Program Technical Challenges and agency Annual performance indicators) at a program level.

Narrow down design space and more rapidly find alternative tools.

What modeling improvements we need to do better in this area.

Slide 3:
How to make computations more affordable?

Doubling computational size in order to have multi-physics simulations as each new variable is added. Realism in details.

Foundational research, not generally sensitive or time critical.

Do not have funding to support these computations on other platforms.

Slide 4:
Pushing capacity on computing limits, need to look into efficiency with other hardware to increase abilities. OpenNCC currently runs using CPUs.

Increase of capacity needed for future work as complexity and additional physics will increase.

Re-hosting would be a significant effort.

Slide 5:
Anticipate more and more eddy type resolving flows, need additional capacity for increased physics

Need to be adding more functionality and more physics in order to more accurately predict, more resources will be needed.

Need to refactor the code to make use of future hardware and GPUs

Q&A:
Elizabeth Lee-Rausch: Can you give some understanding of why our DOE proposals for time were not successful?
Michael Rogers: Roll of the dice. Selection process was reaching out to NASA for priority in the past. There are so many folks trying to apply for that time and their research is broader than NASA flow physics. Important to show that our codes perform well on their hardware. If their hardware doesn't look like what is at NAS, then where can we go to run?

Mujeeb Malik: Two programs DOE provides time to outsiders. Your program has to be very relevant to their missions. Or they want you to do problems that use at least 20% of the machine to prove that you can scale (on Summit, it would be thousands of GPU nodes), hard pressed to use that kind of capability right now.

Brian Thomas: Just for clarification: as you scale up sounds like you just need more computational power/time, correct? No increased need for storage, network, etc.
Michael Rogers: As things gets bigger, all of those things get stressed. Would like more flops to do more complex computational physics.

Tsengdar Lee: **Can we have a chart** to describe the problem(s) we are experiencing to get resource from DOE?
Michael Rogers: Certainly, would be happy to provide that.

- NOTE: Charts and email explanation sent to Tsengdar Lee on 10 June 2020.

---

10. Use Case: Use Case: Computational Materials & Molecular Simulations

John Lawson
john.lawson@nasa.gov
ARC, IT

Main points:
P1: Goal: accelerate the development time of high-performance materials used on vehicles from decades to years through computational material modeling.

P2: NASA's computational power puts agency at a competitive advantage in computational materials science, and maturity of the hardware and software has allowed team to tackle very difficult problems.

P3: Can use molecular dynamics (MD) simulations and DFT computations separately or together in the same simulation. DFT-MD is the most challenging and computationally intensive.

P4: Load is fairly consistent throughout the year; the speed of throughput is most important

P5: Group has done some work with HECC to see if moving to GPUs could be beneficial. So far have not seen the improvements they'd expected, but may be worth revisiting.

Raw notes:
Slide 1:
Computational Materials Science: applying some of the tools of computational chemistry to materials. Molecular level computational tools. Supported mostly by ARMD, T3, CAS.

Accelerate development and gain insight for design. Re-entry and high-performance vehicles; materials that need to be developed for specific applications, not commercially available; accelerate the development time from decades to years. Metals, composites, batteries, etc.

Slide 2:
High performance computing has been critical to the realistic material modeling in the last 10 years. Have more computational muscle than academic or industrial researchers, puts NASA at a competitive advantage. Would have to use cloud computing options or purchase our own cluster without HEC systems.

Very computationally intensive applications; example: working on shape memory alloys, have material change shape on a vehicle during flight not through mechanical means.Shape change related to atomic scale phenomena. Subtle changes in the crystal structure. Understanding atomic motions enables tuning of the shape memory effects,e.g. controling at what temperature the shape changes.

Slide 3:
Bulk of computing time is in simulation run. Generating statistics, particle-based simulations.

Different computational types, including molecular dynamics, quantum simulations, with integrated calculations combining several methods being the most computationally intensive.

Using code from University of Vienna.

Slide 4:
Running all the time, fairly constant load.

Can always use more resources, the use case will go through several stages of more complex simulations and more refined codes.

Re-hosting could be interesting – will depend on some of the architecture.

Q&A:
Brian Thomas: What would be 'timely' in terms of getting cycles (can you quantify that)?
John Lawson: Some people mentioned that they have design reviews – we're not like that. For us, timely could be weeks, months. As long as we can get our results quickly (on the order of weeks) we can keep the project moving. It's more about the more you have to wait for things, the more you have to redirect work and time.

Michael Little: Regarding constraints statement - is there a knee on the curve for how much more capacity would really help?
John Lawson: A lot of the work is not infinitely scaled, we're probably more interested in speed than more nodes.

Michael Little: If large scale GPU processing became available what would you need to take advantage of it?
John Lawson: Tried GPU clusters available now, some molecular dynamics codes have been converted over. Worked with NAS staff to see if there was a benefit running on GPUs, did not see a benefit so far (kinda surprising), but maybe we just didn't try hard enough. May need to look more closely.

Piyush Mehrotra: Some MD codes can take advantage of GPUs.
John Lawson: The codes we have can run on GPUs, comparable to running on CPUs. Didn't see a benefit.

Group Discussion of Key Questions

No additional discussion on key questions from the group.

TL: These workshops have been very successful. Been in some discussion with ARMD program managers and the Institution Office and they have taken notice of this workshop. Only initial feedback.

Appendix D6 Workshop Session 6-9 Inputs

Christa Peters-Lidard, Sujay Kumar	GSFC 610	Land Information System (LIS)
Jacqueline LeMoigne	GSFC 407	AIST New Observing Strategies (NOS)
Kevin Jacobson	LaRC D308	Computational Aeroelasticity
Meelan Choudhari	LARC D302	Laminar-Turbulent Transition Physics & Modeling AAVP (HTP) and TACP (TTT)
Rahul Ramachandran	MSFC/IMPACT	Data Production (Forward processing and Reprocessing)
Manil Maskey	MSFC/IMPACT	Scaling pixel level smoke detection for high temporal geostationary satellite data
Susan Owen, Lei Pan, Hook Hua, David Bekaert, Gerald Manipon	JPL 329	Large-scale InSAR analysis
Shubha Ranjan	ARC TNC	AI/ML at Scale

National Aeronautics and Space Administration

Workshop Session 6-9

9 June 2020

National Aeronautics and Space Administration

Use Case: Land Information System (LIS)

Christa Peters-Lidard, Sujay Kumar
June 9, 2020
NASA Goddard Space Flight Center
ESD/ESTO AIST, THP, MAP, ASP/Water, HMA, WWAO, MC DO

Use Case: Mission Need

- Mission: The Land Information System (LIS) is a software framework for high performance terrestrial hydrology modeling and data assimilation developed with the goal of integrating satellite and ground-based observational data products and advanced modeling techniques to produce optimal fields of land surface states and fluxes.
  - LIS supports a breadth of current and planned NASA research interests/priorities, such as :
    - development of **OSSEs** in support for future missions such as Snow, MC, SBG, and NOS;
    - high resolution (1-km) **land reanalysis** over the Western U.S. (WWAO), HMA (HMA);
    - fully **coupled modeling and data assimilation systems** (MAP) integrating LIS, WRF, WRFhydro, and ParFlow;
    - **multi-variate data assimilation** for hydrologic science and applications (GRACE-FO, SWOT, NISAR); and
    - advanced **machine learning techniques** for data fusion, information extraction.
- Criticality: NASA's HEC is essential for this mission because it provides enough computational power (CPUs and GPUs) to handle large size integrations for substantially large model domains in high resolution and the massive data storage resources required.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - Purchase dedicated server: +no data egress cost, performance –sys admin/maintenance cost
    - Commercial cloud system: +no sys admin/maintenance; easier to share -

Use Case: Mission Need

- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - The code/data are **not confidential** and dissemination of data (modeled and ingested) is encouraged.
  - Availability of HEC is essential for running **routine LIS applications and analyzing the output**.
  - We work with many **external collaborators, NASA HEC** provides a shared environment for computational/data resources and for software development and scientific review of results.
- HEC Function: HEC supports this work by making sure to provide enough **allocation times (i.e., SBU)** so that our diverse numerical simulation experiments can be fully explored and executed to test our hypotheses thoroughly. At the same time, HEC provides **massive data storage** for model outputs and for post-processing that are run and required daily
- .
- Impact of HEC: NASA's HEC is essential in terms of providing adequate **computational power**. Also, the relatively fast allocations of the massive jobs submitted on the queue in the NASA's HEC are an important aspect of system requirements for the efficient progress of the mission.

Use Case: Mission Need

- Land Information System (LIS)

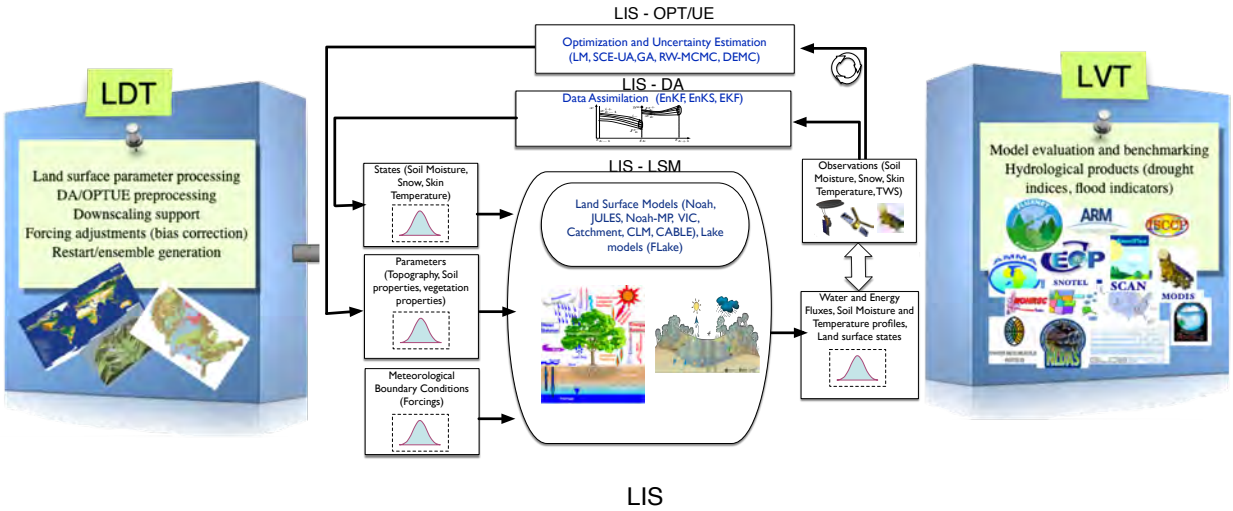

Use Case: Current HEC Application Characteristics

- Data:
  - In:
    - Surface meteorology (75 Tb, near-real time, various sources);
    - Remotely-sensed surface states (40 Tb, near-real time, SMAP/MODIS/GOES/etc.);
    - Land parameters (10 Tb, mostly static-e.g., soils, topography, various sources);
    - Reference/evaluation datasets (5 Tb, ~6-12 months latency for quality control, various sources)
  - Out:
    - Land surface (water, energy, carbon) states and fluxes.
    - Most output data does not require retention, only final/production data requires it.
    - Statistical analysis against reference/evaluation datasets (surface in situ point measurements, gridded reference products, etc.) is conducted using the LVT sub-system.
    - Some visualization monitoring on the LIS webpage.
  - Transfer:
    - Output goes to the GES DISC and/or the NCCS Data Portal for archiving as well as delivery of near-real time products to our users/customers.

Use Case: Current HEC Application Characteristics

- Computation:
  - Type:
    - Massively parallel (large number of computing cores networked together for MPI-based parallel processing) computing is required for the LIS use case.
    - These simulations require one or more of:
      - multiple ensemble members (for land data assimilation);
      - very fine land grid spacing (to resolve surface meteorology and soil/vegetation variability);
      - connection to the atmosphere (NU-WRF, for land-atmosphere feedbacks);
      - connection to sub-surface hydrology (WRF-Hydro; for streamflow and groundwater); and
      - forecasting (seasonal forecasts and climate scenarios, for drought/flood outlooks, etc.).
    - Access to a large volume of input data is also required. These require special HEC capabilities such as highly performant file systems and networking fabrics. Many shared libraries are required including NetCDF, hdf5, ESMF, gdal, geotiff, etc.
- Evaluation: The LIS team runs regression testing to ensure code/compiler/operating system stability. The LIS team also performs timing/speed tests to evaluate computational requirements/performance by exercising various processor configurations.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Large-scale **non-linear models, science data reprocessing, revision to major models.**
  - Within funding cycles another factor is **focal area of selected projects**. Projects that require multi-year, high-resolution, and large domains are more likely to demand more assets than short-term, regional prediction cases.
  - **Large data processing** is a key factor that affects computing performance. Therefore, parallel data I/O and processing using GPUs should be emphasized.
- Are there key functional gaps that exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: **Data storage capacity, turn around time, and external access** are the biggest limits that LIS experiences. In the future, as models and data move towards higher spatial/temporal resolution, more computational power will be needed. In turn, greater storage space (either local/virtual) will be required for both remote sensing data and model output. For example, the significant large volumes expected from future sensors (**NISAR**) will necessitate virtual data storage solutions and modifications to the code to support non-local data access.
  - Impact: These limitations affect LIS by negatively impacting schedules and imposing additional project management. Slow turn around and difficulty in getting external access makes it difficult to review results and to track mission success. **Limited data storage capacity** forces scientists to perform mundane tasks of shuffling data just to keep projects moving forward. Lack of computational power and storage space will limit the use of future higher resolution remote sensing data and novel developments in model physics.
  - External Sources: Our project does not run on a non-NASA computing facility.

HEC Needs Assessment Workshop Key Questions

- How will this Use Case evolve over the next decade?
  - Maturity: Several applications of LIS will move to high maturity level TRLs/ARLs and be operational with continued development, but new functionality will always be under development.
  - Growth: There will be increases in **time-step frequency and spatial resolution**, as well as other factors requiring both increased computational support and increased data demands, such as,
    - **data assimilation** (smoothing, variational methods, data fusion based on machine learning),
    - **ensembles** (multi model, multi parameter, seasonal forecasts, climate scenarios, DA),
    - **coupling** (with WRF or with hydrology models), and
    - **AI/ML and calibration** runs required for training.
  - Re-use: The LIS software is being used by a **wide and growing community of users** in both domestic and foreign institutions and organizations. We expect the user community to grow. LIS has a record of being included in several NASA solicitations, and the results of those influence the re-use and user growth rates.
  - Re-hosting: For typical HPC systems, LIS is easily portable. Code has a history of successfully being transitioned across systems (Discover, Pleiades), compilers, and operating systems (e.g., Sandy Bridge, Haswell/SLES11, SLES12).
    - For **GPU-based systems, much effort will be required to rework the software.**
    - For cloud-based systems, more work in **virtualization and containerization** will be required to run LISF effectively.

Use Case: Mission Need

- Mission Concept:

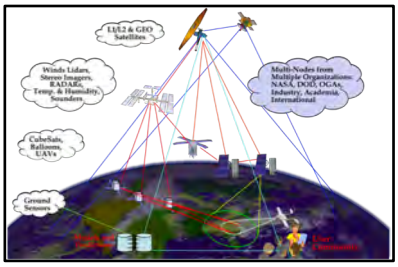

NOS = Intelligent and Generalized SensorWeb Concept, i.e., a distributed system of *sensing nodes* interconnected by a *communications fabric*, in which each node is either a single sensor, a group of sensors or a DSM.

Measurement Acquisition ("Mission" Design or Model-Driven)

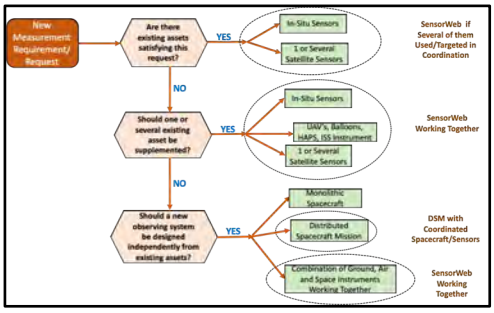

Observation Planning or Rapid Response to Event of Interest

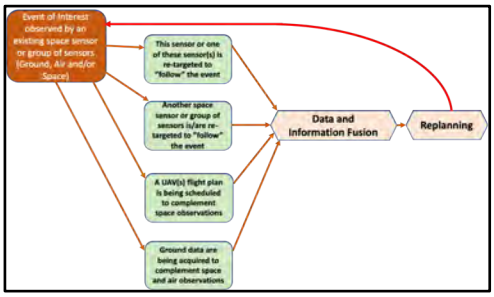

Use Case: Mission Need

- Mission: The goal of the **New Observing Strategies (NOS) concept** is to optimize measurement acquisition using many diverse observing capabilities, collaborating across multiple dimensions and creating a unified architecture:
  - Using the concepts of Distributed Spacecraft Missions (DSM) or SensorWebs at various vantage points
  - In response to Decadal Survey mission design needs, forecast or science model-driven, or application/event-driven
  - Using NASA- as well as non-NASA data sources or relevant services
  - HEC impact can be measured:
    - For NOS Design: by the number of observation trades/strategies/architectures that can be considered in near-real-time
    - For NOS Operations: by maximizing science return on investment/mission science value
- Criticality:
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - Would use commercial cloud computing
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Competition sensitive in some cases (competed missions)
  - Could impact some deadlines (e.g., proposal) or could miss detection of a transient event
  - Some of the NOS nodes could be non-NASA systems; data fusion might be done by NASA or non-NASA entities
- HEC Function:
  - How often application is run? How is the output used to achieve mission objective?
    - Design: A few iterations might be needed to optimize mission objectives and requirements
    - Data processing, fusion, analysis before launch: optimize algorithms/ML training to run onboard

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - NOS architecture selection: Problem dimensionality
  - Scalability: number of nodes; amount of data to be processed, analyzed and fused, science and telemetry data

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on NOS Use Case:
    - Communication speed S/C to ground and S/C to S/C
    - Time constraint at every step of NOS design and operations
  - Impact: Could make some of these NOS concepts impossible
  - External Sources: Does your project run on a non-NASA computing facility? If so, which and why? N/A yet

- How will this Use Case evolve over the next decade?
  - Maturity: Currently in development state
  - Growth: Will it increase resolution or time-steps to support evolving future mission needs? Yes, number of nodes will increase
  - Re-use: Will the number of people running this code/application grow? Yes
  - Re-hosting: What is needed to migrate code to future computing architectures? Probably

Future HEC Needs of NOS Use Case

- **On the Ground**
  - NOS and Constellation Design: Being able to evaluate quickly thousands of potential various architectures
  - Training Machine Learning systems on the ground for calibration, classification, event detection, etc.
  - Real-time sensor data/image processing, analysis and understanding for quick reaction to events of interest
  - Simultaneous modeling and visualization for quickly identifying missing/additional data/information of interest for better prediction or better accuracy of forecast and science models
  - Fast handling of multiple spacecraft/nodes, anomaly detection and mission management
  - Data archiving, management and mining of NOS and multi-mission concepts
- **Onboard Processing**
  - High Performance Spaceflight/Neuromorphic/Edge Computing for real-time onboard and in-situ:
    - Data processing and analysis, including Machine Learning and Transfer Learning
    - Data/information fusion and uncertainty assessment
    - Risk and mission value assessments
    - Autonomous planning, scheduling and decision making
  - Radiation Mitigation by Software for non-radiation hardened SOA hardware

National Aeronautics and Space Administration

Use Case: Computational Aeroelasticity

Kevin Jacobson
June 9, 2020
LaRC/D308 (Aeroelasticity Branch), FUN3D dev. team
ARMD - Advanced Air Vehicles Program (AAVP) &
Transformative Aeronautics Concepts Program (TACP)

Use Case: Mission Need

- Mission:
  - Vehicle design and analysis – X-planes, Launch/EDL vehicles, rotorcraft
    - Deformation under flight loading
    - Dynamic instabilities – vehicles will break themselves apart
  - Support of aeroelastic wind tunnel experiments
    - Provide experimentalists with expected stability bounds
    - Characterization of wind tunnel effects
  - Research: understanding aeroelastic physics, methodology and tool development
- Criticality: NAS is **absolutely essential** for our work
  - Cannot do wind tunnel testing in TDT to explore parameter space or every design iteration
  - Many critical aeroelastic behaviors are driven nonlinear physics
  - Without NAS, we would have to rely on local mid-range cluster – not nearly enough resources
- Mission Security Needs:
  - Security: The codes we use (primarily FUN3D) are EAR and some vehicles are ITAR
- HEC Function: ~25-50 jobs running or waiting in queues at any given time
- Impact of HEC: Export control restrictions mean we can't use commercial cloud computing
  - Possibly could get HPC hours for code R&D work from DoD, but unlikely to get necessary sustained resources access from DOE or DoD for production/mission work

X-59 QueSST

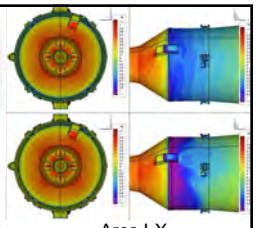

Ares I-X

Transonic Truss-Braced Wing Model

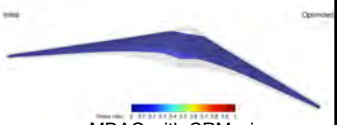

MDAO with CRM wing

KTH Generic Fighter in Transonic Dynamics Tunnel (TDT)

Use Case: Current HEC Application Characteristics

- Workflow: Shell or python scripts that submit/monitor PBS jobs, primarily running FUN3D
- Data:
  - In: Mesh and solution restart files (1 GB - 25 GB)
  - Out: Visualization of flow field and surface pressures, system identification of structural dynamics, retain outputs for later examination/comparison (25 GB – 1 TB)
  - Transfer: Data archived on Lou
- Computations:
  - 10-200 Skylake nodes, 72-216 wall hrs per simulation, typically <25 simulations per application
- Evaluation: Need more resources - shorter average queue times, more cores per job
  - Spending months on aeroelastic flutter analyses
    - With simplifying assumptions that have questionable validity – simplified geometry, smaller meshes, simplified physics models (RANS turbulence models in separated transonic flows, linear structures).
    - Must prioritize flight conditions based on intuition to get through queues -> potentially missing catastrophic flight conditions or relying on linear models which can be dangerous
    - Long queue times discourage us from submitting bigger jobs
  - **Given the resources, we (9 people) would fill Pleiades 24/7**
    - Have codes capable of doing the following, but…
      - we only sometimes can perform - mesh/timestep independence studies, eddy-resolving CFD
      - we never perform – parametric UQ, limit cycle oscillation analysis

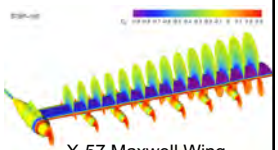

X-57 Maxwell Wing

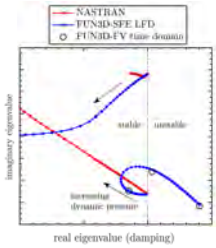

Failure of linear model to predict instability

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Nonlinear transonic and separated flow physics requiring small time steps coupled to low frequency structures requiring long simulations
  - Revisions to X-planes design, wind tunnel models, etc. require new batch of analyses
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: We are constrained by limited resources.
  - Impact: Complete study with higher fidelity + UQ will lead to more confidence in computations. Lay groundwork for Certification by Analysis. For X-57, there's no plan for an aeroelastic wind tunnel test, but we can't do higher fidelity CFD or UQ
  - External Sources: None
- How will this Use Case evolve over the next decade?
  - Maturity:
    - Analysis: presently rely on RANS + modal structural approximation. Expect eddy-resolving CFD + nonlinear structures to mature.
    - Expect more emphasis on parameter space exploration for optimization and UQ (Certification by Analysis)
    - Better integration with MDAO and MBE environments (OpenMDAO, ModelCenter)
  - Growth: Expect smaller and more time steps, bigger meshes, and more simulations
  - Re-use: Expect more users. Advances in CFD (adaptation methods, improved robustness, etc.) will decrease the required expertise. CFD in MDAO will become more common.
  - Re-hosting: Need access to emerging hardware to test and develop on. Starting to utilize GPUs. Expect migration of work to GPU accelerators as more of FUN3D is ported.

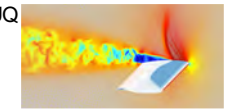

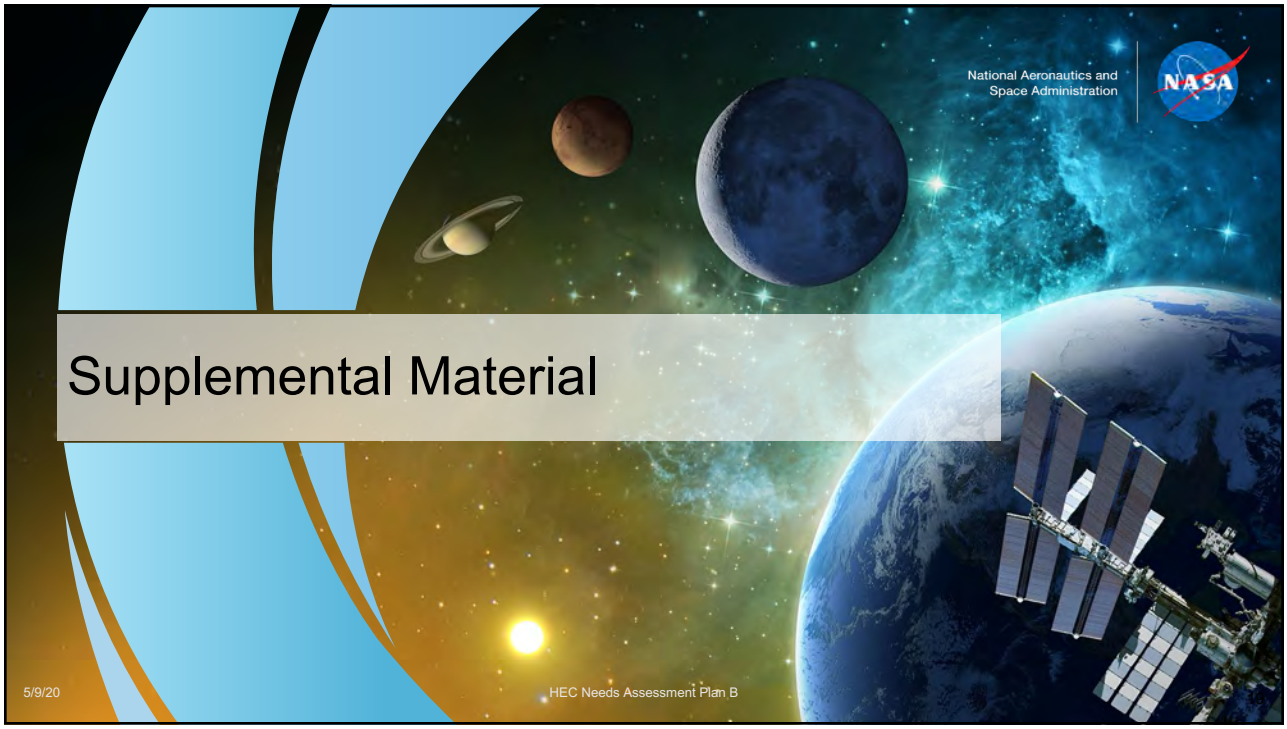

Current Use Case (Survey)

- Achieving the mission objective
  - Computational aeroelasticity – RANS CFD, CFD mesh deformation (linear elasticity on CFD mesh), Linear modal approximation for the structures
  - Code description – FUN3D finite volume and stabilized finite element solvers
    - ~90% Fortran, ~10% C++/CUDA, small amount of python to couple codes with OpenMDAO (written in python).
    - 1.2 million lines of code
    - Workflow mechanism – shell and/or python scripts
  - List the functional computational and storage requirements for the application
    - 10 – 200 Skylake nodes or equiv. 72 – 216 hrs walltime.
    - Starting early testing of aeroelastic simulations on V100s with CUDA
    - 5-6 GB per million grid points. 5 -100 million grid points total. 10,000 – 100,000 time steps
    - Domain decomposition with MPI, has some OpenMP (not utilized regularly)
    - Primary flow solver functions converted to CUDA C.
    - Run in batch jobs

5/9/20 HEC Needs Assessment Plan B 20

Current Use Case (Survey)

- Describe the data requirements for the application
  - Data loaded all at start. Front end has both serial and parallel input reading.
  - Visualization of surfaces, slices, and isosurfaces in flow field (25 GB – 1 TB). Some do via scripts on Pleiades, some done on local computers
  - Grids generated on local computers, then scp to Pleiades. Create other inputs (Fortran namelists) with scripts or manually edit the files.
  - Typically store all inputs, restart files, and sampling outputs to avoid redoing computations
  - Minimal data distribution
- Key computing obstacles to effective execution of the current application: long queue times, limited resources
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Have basic knowledge and experience of writing code for GPUs
  - Need expert help optimizing kernels

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - Describe how the application and workflow is expected to evolve over that time frame
    - Expect aeroelastic simulation to be more common in MDAO – coupling to new physics
    - Push towards integrating aeroelastic simulations into MBE environments
    - Parametric uncertainty quantification and optimization -> many more simulations per project
    - Increase in resolution and timesteps for higher fidelity physics in structural and fluid modeling
    - Expect continued refactoring for emerging platforms
  - Describe how the computational requirements will evolve over time
    - Increase in problem size summarized on slide 10
    - Expect migration to GPU accelerators and other architectures
    - Storage requirements will scale with mesh sizes and number of time steps on slide 10

A "Standard" Aeroelastic Simulation in the Next Decade

Today - **Resource bound**	0 – 5 Years	5 – 10 Years
CFD- RANS O(5-100 million) DOFs	CFD - Eddy resolving CFD O(50 million – 10 Billion) DOFs	CFD - Eddy resolving CFD + chemistry O(1 billion – 1 trillion) DOFs
Linear model reduction for structures O(10-100) DOFs	Linear finite element structures O(10,000 - 1 million) DOFs	Nonlinear finite element structures O(1 million – 10 millions)
10000s time steps for initial instability analysis (flutter)	Smaller time steps for CFD + longer periods of simulation (limit cycle oscillations) – O(1 million)	Smaller time steps to resolve finer CFD – O(10 million)
10-200 Skylake nodes 72 – 216 wall hrs	200-1000 GPUs 12 – 216 wall hrs	2000+ GPUs 6 – 216 wall hrs
Typically < 25 simulations	1000s-10,000s simulations for flight envelope exploration + optimization	10,000+ simulations for flight envelope + UQ

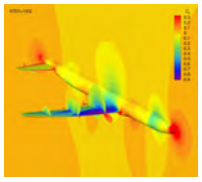

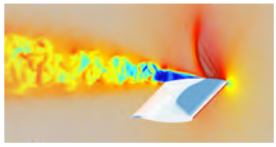

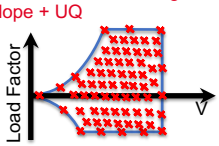

Unwanted aeroelastic behavior can occur anywhere in the flight envelope, not just at the edges

Application Drivers in the Next Decade

Certification by Analysis

- Reduce cost of certification augmenting/replacing flight tests with simulation
- Sometimes there are no other options but simulation
- Requirements:
  - High quality **large simulations**
  - **Many simulations** to perform uncertainty quantification (UQ) over flight envelope

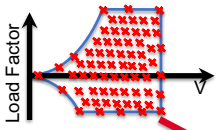

MDAO and MBE

- Design more efficient and safer vehicles to lower cost using multiphysics simulations and gradient-based optimization
- Requirements:
  - **1000s of simulations**
  - Modeling of more disciplines – every aspect of the aircraft/spacecraft. Better integration into MBSE environment

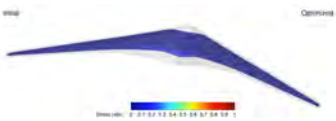

Transonic Flutter and Limit Cycle Oscillations

- Currently too difficult to do reliably with simulation due to RANS assumptions and questions about mesh/timestep independence
- Most aircraft fly in transonic regime
- Requirements:
  - Eddy-resolving CFD
  - **Longer simulations**
  - Nonlinear structures for more flexible aircraft

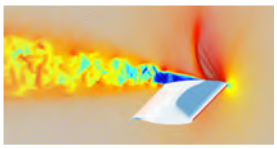

Unsteady Launch Vehicle Loads

- Potentially destructive unsteady loading anywhere from wind on the pad to transonic shock buffet
- Load prediction done by wind tunnel tests
  - Difficult to optimally design with wind tunnel
- Requirements:
  - High resolution to capture unsteady loading
  - **Long simulations** to simulate launch trajectory or gather statistics

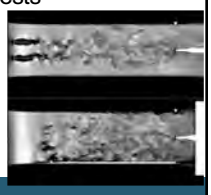

Other Feedback

- **Feedback from applications experts:**
  - Long queue times
    - We have codes capable of running bigger jobs, but don't due to resource availability
    - Most often want jobs to drop quickly regardless of CPU type
      - End up submitting same job to multiple queues and deleting jobs after the first one starts.
      - General queue for first available of any CPU type would be useful for our users
    - Users that are starting to explore GPUs spend long time in the queue
- **Feedback from developers:**
  - The devel queue not ideal for many code development activities
    - Submitting multiple jobs for performance/scalability testing
    - Majority of development compute cycles spent on continuous integration testing
      - Multiple jobs at the same time
      - Difficult to hook in continuous integration testing to NAS hardware
      - Would like to check that performance is maintained on various hardware
    - The intended purpose of devel queue gets abused - Gets used for production work on 2 hour restart cycles
      - Using the devel queue for production work is a sign that standard queue wait times are too long to get the work done in the necessary time frame for the engineering project
  - Availability of latest hardware for developers and production runs
    - At our local mid-range cluster, we had to freeze machine learning production work on V100s to have resources for GPU mini-app development during a Hackathon (Oct 2019). Shouldn't have to do that.

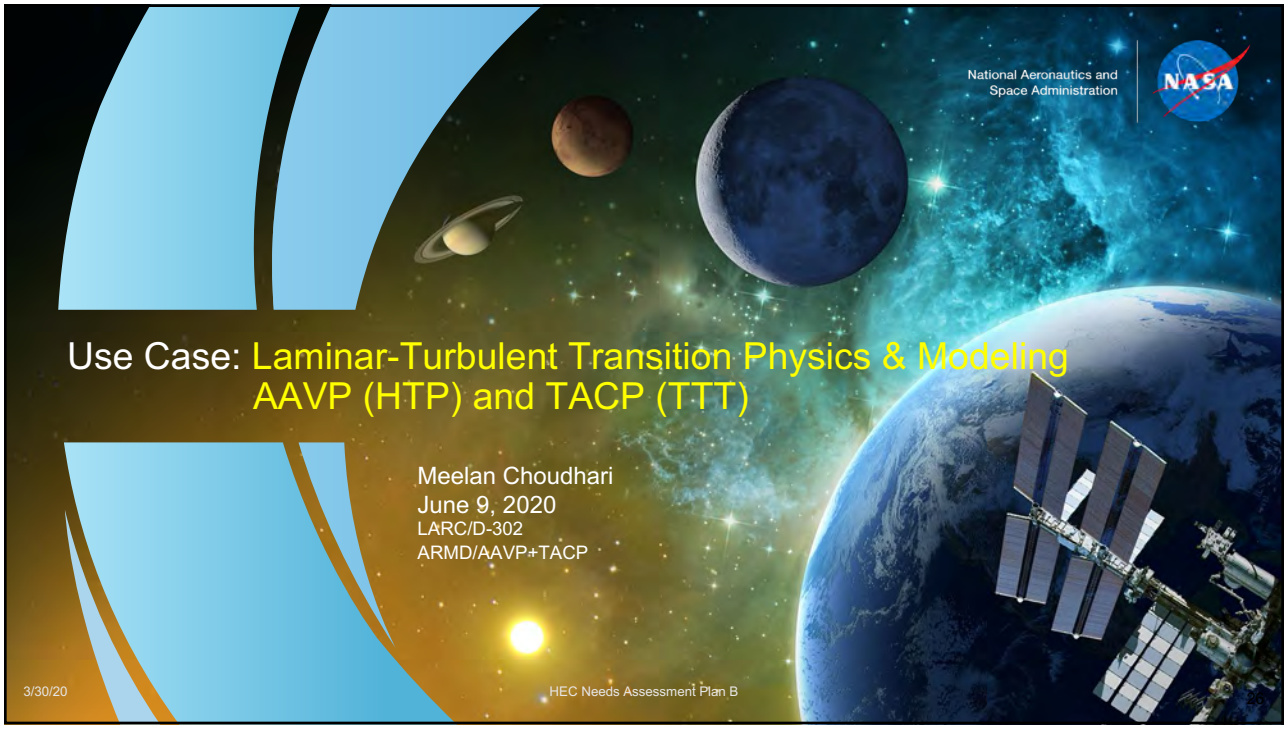

Use Case: Multi-fidelity modeling of boundary-layer transition (BLT)

- **Mission**: Accurate prediction of the state of the flow (laminar vs. turbulent) is a cross-cutting requirement for modeling the aerothermodynamic environment of aerospace vehicles across the speed regime.
  - The myriad paths to laminar-turbulent transition and their sensitivity to external disturbances requires a multi-fidelity prediction capability, with direct numerical simulation (DNS) assuming a critical role because of measurement (and facility) limitations, especially at high speeds. Routine DNS of natural transition to enable the development of accurate, multi-fidelity reduced-order-models (ROMs) would require order(s) of magnitude larger resources than the current NASA HECC.
- **Criticality**: The goal of physics based transition prediction to support CFD Vision 2030 cannot be accomplished without NASA HECC. The cost of not having this prediction capability would be substantially higher modeling uncertainty, enhanced risk for human re-entry missions, and inability to mature breakthrough concepts for reduced fuel burn. Inadequate substitutes for NASA-HECC include DoE/DoD machines and relia[illegible]on academic partners with resources provided by NSF/states, etc.
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - Depending on specific application, an inadvertent data release could violate export control requirements.
  - HEC interruptions amount to loss of productivity and compromised milestones
  - External collaborators without NASA credentials are a vital part of developing, enhancing, and applying the codes.
- **HEC Function:** HEC resources used continually for hierarchical simulations ranging from low fidelity stability models phenomenological transport equation based models, to DNS resolving all spatiotemporal scales in the flow. Typically, multiple hierarchical computations are ongoing at the same time. Currently feasible mesh sizes on NASA HEC (hundred million to a billion or somewhat larger) are barely adequate to serve the targeted goal. The highly competent and responsive user support staff cannot compensate for the limited hardware resources.

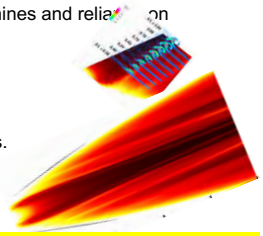

Hierarchical computations reveal the complexity of transition mechanisms on the upcoming BOLT flight experiment

- **Impact of HEC**: **NASA's HEC is absolutely critical to accomplishing this work**
  - Alternate options to perform this work would include: purchasing a system at Center or branch level, writing proposals seeking time on DoE and DoD facilities, renting time on a commercial facility, or to subscope the NASA effort to focus on reduced order models requiring moderate computing resources (with reliance on external partners to provide the data based on large scale numerical simulations

Use Case: Current HEC Application Characteristics

- **Work flow:** Overall HEC workflow includes a combination of high-order Navier-Stokes flow solvers, special purpose codes based on various levels of approximations to the N-S equations, and deep learning codes that encapsulate the knowledge base to build regression models.
  - The DNS simulations require substantial user level monitoring because of the compromises on mesh size to allow timely and reasonable quality output on the typically congested and oversubscribed NASA HECC. A lot of effort is spent on balancing turnaround time and the need for user intervention.
- **Data:**
  - **In**: Mesh sizes are fit to the available HECC rather than to the need. Currently permissible mesh sizes of a few hundred million to a billion or somewhat larger (approximately 3-30 GB)
  - **Out:** Time accurate data resolving high frequency fluctuations over long enough periods to allow reliable flow statistics. Again, data resolution and sampling frequency are dictated by practical considerations rather than by need (larger by 10x). Data is analyzed in multiple ways (flow-viz., statistical analysis, etc.) to distill insights into dominant physical mechanisms that are used to guide the development and validation of reduced order models. Alternately, millions of runs of low to intermediate fidelity tools are used to develop deep learning models that would help minimize the need for domain expertise during practical applications.
  - **Transfer:**
    - All data is archived on the HEC storage system; portions are brought down to local storage for graphics intensive postprocessing.
- **Computation:**
  - Primarily capacity computing on CPUs; special HEC capabilities such as GPUs needed for deep learning work
- **Evaluation:** An order of magnitude increase in HECC amounts to a mere doubling of mesh resolution. Thus, the ability to achieve targeted goals would easily require more than 100x increase relative to current levels. Data transfer speeds between local to HEC resources have worsened in recent years rather than improving..

HEC Needs Assessment Workshop Key Questions

- **What are the key factors today that affect the load on computing assets and cause it to vary?**
  - Load levels and variations thereof vary with geometric and physical complexity, modeling fidelity, project deadlines, and publication cycles.
- **Are there key functional gaps exist or are projected in the next 5 years or 10 years?**
  - Constraints on Use Case: Available resources are insufficient to meet mission needs, especially because the transition zone cannot be reliably modeled with lower order approximations such as LES and wall modeled LES. Current codes easily scale up to 10K-20K cores, but be rarely exercised with more than 5K cores because of excessive wait times. Extrapolations from the most resource intensive jobs we were able to run 5-10 years ago would paint a bleak picture in terms of meeting the widening gaps in the future.
  - Impact: Between increased restrictions on working with academic partners and the lack of increase in HEC resources, mission needs for physics based BLT models will not be met.
  - External Sources: DoE and DoD resources provide supplementary DNS data.
- **How will this Use Case evolve over the next decade?**
  - Maturity: Lower fidelity BLT models will be integrated into NASA CFD codes to allow increased benefit to mission needs.
  - Growth: Continued development of data driven modeling would entail increased emphasis on simulation databases for new classes of vehicles. The grid sizes will continue to increase to accommodate more complex configurations or to enable higher spatial resolution.
  - Re-use: High fidelity DNS modeling of BLT is likely to remain a specialist's domain, but increased criticality for high-speed vehicles may lead to a modest increase in the number of HEC users of this Use Case. A significantly larger number of users expected at the low to intermediate fidelity modeling as that end of the capability matures and, also, as data driven models for BLT get integrated into CFD codes.
  - Re-hosting: The DNS codes will need to be ported to emerging architectures including GPUs.

National Aeronautics and Space Administration

Use Case: Data Production (Forward processing and Reprocessing)

Rahul Ramachandran
NASA/MSFC IMPACT Project

HLS as a canonical example

- Merging Sentinel-2 and Landsat data streams can provide **2-3 day global coverage**
- Goal is "seamless" near-daily 30m surface reflectance record including atmospheric corrections, spectral and BRDF adjustments, regridding
- Project initiated as collaboration among NASA GSFC, UMD, NASA Ames
- Prototype for a multi-sensor Analysis Ready Data product

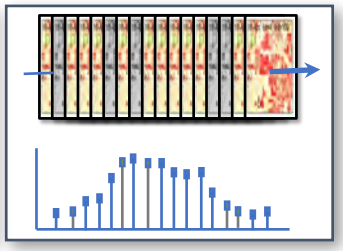

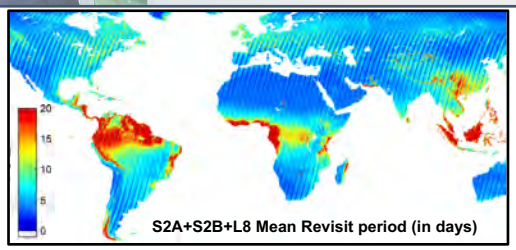

Potential Revisit using different Virtual Constellations

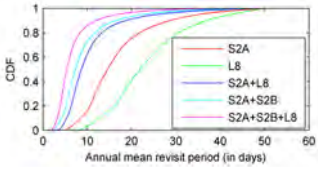

Mission Concept Example - HLS Production

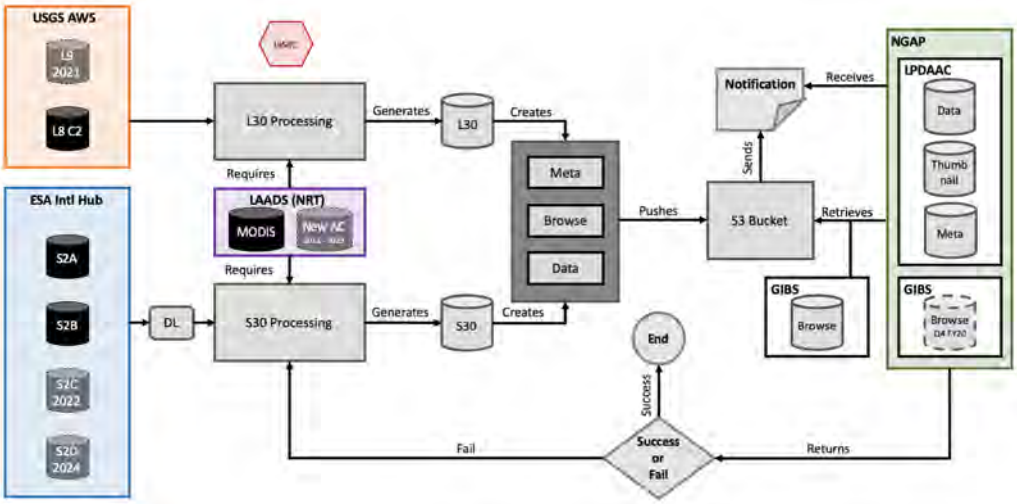

Use Case: Mission Need (HLS as an Example)

- Mission: [Briefly describe Agency-level need for this application or capability]
  - Data production (forward and reprocessing) are critical and can have different latency/cost requirements
- Criticality: [How essential is NASA's HEC to accomplishing this work]
  - Current paradigm leverages either commercial cloud or buying a system to manage locally
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Forward processing has latency requirements (data produced by certain time), reprocessing workflows do not and can be ideal candidates for running on an HEC
  - External collaborators will need access to systems to run workflows
- HEC Function: [Briefly describe how HEC supports performance of this work]
  - Access to computing to scale data processing
  - Need to run the application? Depends
    - Forward processing is event triggered
    - Reprocessing is as needed
  - [How is the output used to achieve mission objective?]
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - Need to find cost optimal solutions that could entail using a mix of Cloud and HEC

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - What triggers this workflow? SNS Notification (Event)
  - Is there a need for interaction by user, operator or other system in the computing workflow? If there are errors, jobs fail
  - Does this use case need quick reaction or near-real time computing? Explain briefly. Forward processing has latency requirements
- Data:
  - In: Daily Sentinel 2 data (9-10 TB), Landsat L2 data (1-2 TB)
  - Out: [Summarize data produced by computation and what happens to it]
    - Daily HLS L30 (2 TB), S30 (7-9 TB) are pushed to rolling storage
    - [How is the data analyzed to provide mission results and where? (statistical or AI analysis, visualization?)]
  - Transfer:
    - Notification is sent to the DAAC to initiate ingest
- Computation:
  - Type: [Summarize the type of computing used/needed and how HEC supports this today]
    - Cluster to scale batch processing
- Evaluation:

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.]
  - Data processing needs (sensors characteristics, type of processing, use of data from different instruments
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?]
    - Similar processing needs in the future leveraging data from multiple sensors
    - Be able to move workflow between cloud and HEC as easily as possible (without rewrites and only minor configuration change)
    - Quick access to resources without too much paperwork
  - Impact: [How does this affect fulfilling the mission need?]
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?]
    - Started on contractor cloud because of speed to get started
    - Migrating to GCC
- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?] Yes
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?] Yes
  - Re-use: [Will the number of people running this code/application grow?] Probably - different team
  - Re-hosting: [What is needed to migrate code to future computing architectures?] Same workflow needs to work on different platforms - need a common framework

Introduction

- Biomass burning smoke has numerous detrimental environmental and ecological impacts including
  - Respiratory and cardiovascular illnesses
  - Radiation budget
  - Nutrient availability
- Impacts realized both near source and potentially thousands of kilometers downwind depending on
  - Fire duration
  - Amount and type of biomass burned
  - Meteorological and fuel conditions
  - Vertical distribution in the atmosphere

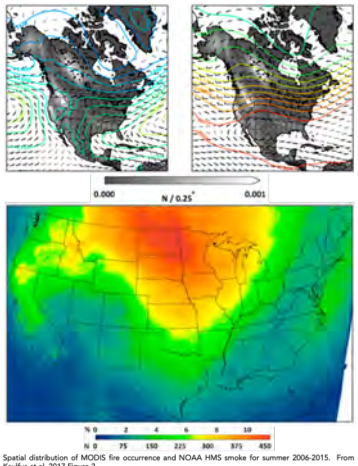

Spatial distribution of MODIS fire occurrence and NOAA HMS smoke for summer 2006-2015. From Kaulfus et al. 2017 Figure 2.

Objective

- Deploy a smoke detection model using machine learning on satellite remote sensing observations
  - Leverage observations from the new generation of geostationary satellite to overcome limitations in current smoke detection techniques
    - High spatial and temporal resolutions over large domains
  - Develop alternative to existing multispectral or subjective manual analysis methods
    - Automated single class classification
  - Leverage cloud computing resources
    - Scalable to large data volumes
    - Computationally efficient

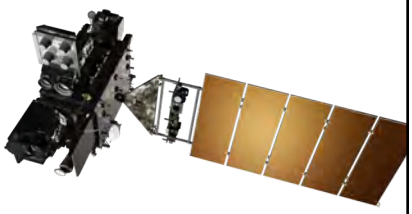

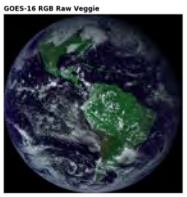

Approach

- NASA IMPACT team has developed an automated, deep learning-based detection model capable of identifying smoke plumes from shortwave reflectance for the Geostationary Operational Environmental Satellite (GOES) R series of satellites.
- Hand-labelled, past instances of smoke plumes from NOAA Hazard Mapping System, quality controlled for spatiotemporal accuracy by a subject expert, comprises the reference truth dataset.
- The detection pipeline comprises of pre-process, detection and post-process stages.
- A Custom Convolutional Neural Network (CNN), trained on a number of smoke events with varying optical thicknesses and viewing angles is used to predict the probability score for the pixel being a smoke pixel

Model Architecture

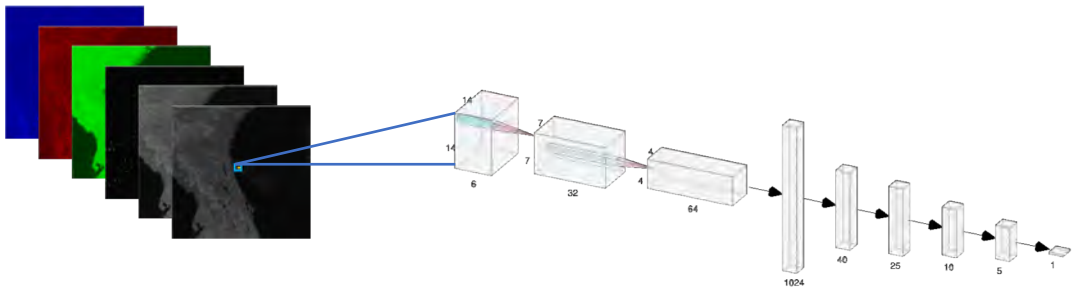

Apply a pixel-based Convolutional Neural Network (CNN)

- Input (N*2)*(N*2) neighborhood of reflectance values surrounding a center pixel (sample)
- 3 convolutional layers
- Each convolutional layer followed by max-pooling layer
- Convolutional outputs are flattened into vectors

Results - 2 May 2018 - Southern Florida

Smoke identified over both land and ocean

Probabilities resemble visually observed optical thickness

Predictions closer resemblance to quality controlled shapefiles

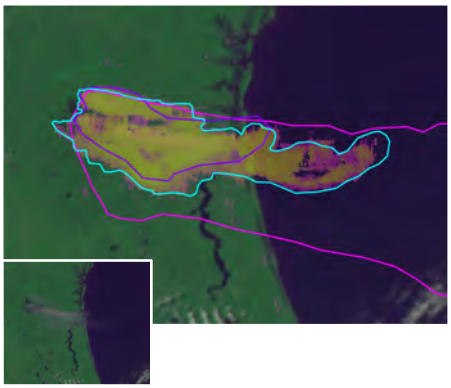

Results - 24 March 2018 - Southern Florida

Distinguishable from chlorophyll commonly found in coastal settings

Fair weather cumulus cloud discrimination

Spectral information for other classes not provided to the model

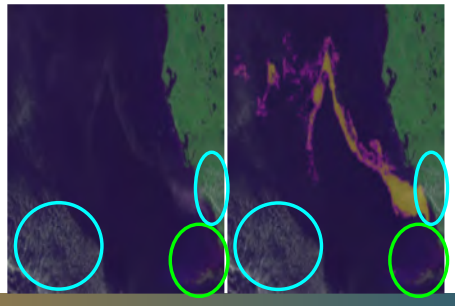

Results - 9 October 2017 - Central California

Large and small plumes

Identification over both land and ocean

Coastal stratus clouds

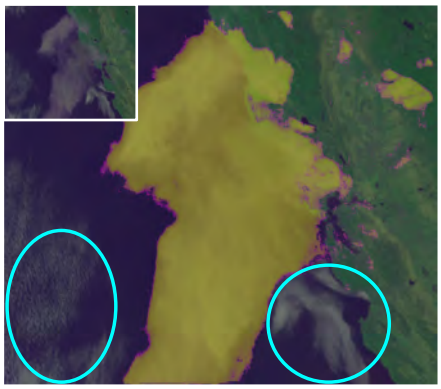

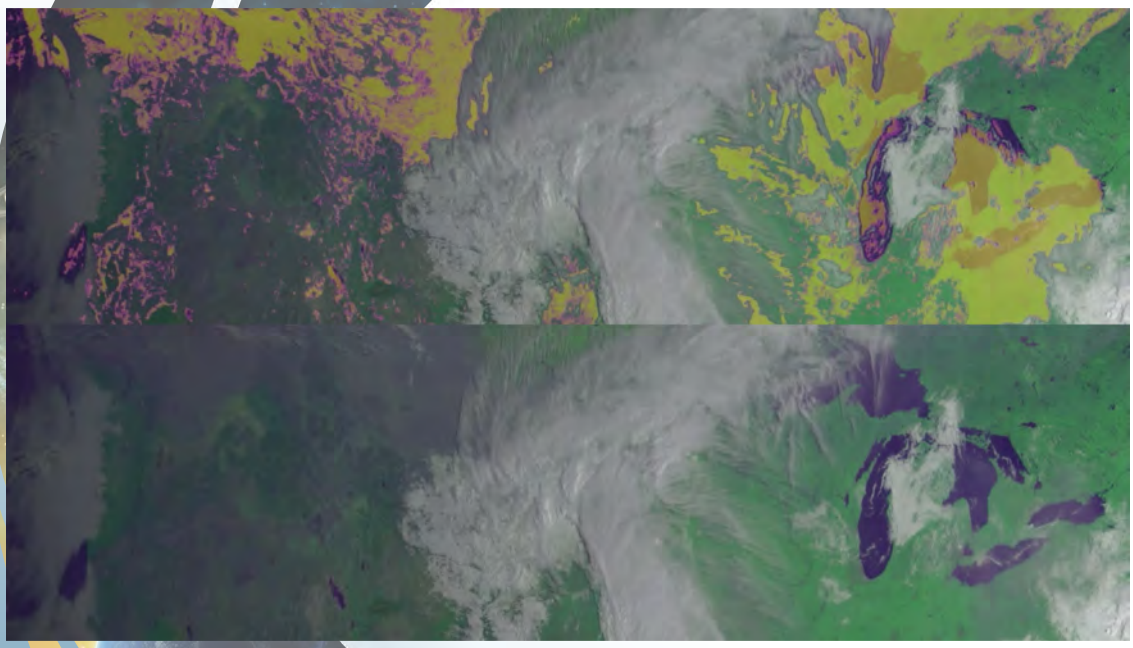

Opportunity to leverage HEC

- Pixel wise inference of full disc GOES-R is time consuming
- There is a potential to detect smoke pixels every 10 minutes (or hourly detection) with GOES-R
- The method can be extended to other geostationary satellites (i.e., Himawari)
- Approach is highly parallelizable
- Data is available freely at AWS (and may also available for in-time processing at HEC)
- Could serve as a pipeline for other phenomena detections
- Timely visualization within Worldview

National Aeronautics and Space Administration

Use Case: Large-scale InSAR analysis

Susan Owen, Lei Pan, Hook Hua, David Bekaert, Gerald Manipon
6/9/2020
JPL/Caltech
SMD/Earth Surface and Interior

Use Case: Mission Need

- Mission: Provide Capability to do Plate Boundary scale measurements of deformation with NISAR, other SAR data
  - We will measure the impact of the HEC by researcher demand, and impact of science from resulting data sets.
- Criticality/Impact of HEC: NASA's HEC should provide a low cost option for researchers
  - If NASA supercomputing was no longer available to meet your needs, we would use Commercial Cloud Computing, however for the size of the data processing, costs become prohibitive for analysis of large areas – e.g., all of the Tibetan Plateau. So researchers would most likely downscale their studies into small enough areas so as not to eat up all their ROSES funding in cloud computing costs (assuming ROSES funding amounts are not increased significantly)
- Mission Security Needs: Few Mission security needs - processed data is intended to be free & open
  - If processed data is accidentally 'released' before science review, there might be some confusion on the quality of the results.
  - Irregular interruptions in access to HEC would create delays in reaching objective to complete processing
  - As we would like this capability to be available to academic researchers as well as NASA researchers, ease of access to processed data is important. However, the automated system from JPL would do most of the processing.
- HEC Function: HEC provides the computational platform for generating L2 deformation maps from Sentinel 1 SAR L1 data.
  - We need to run analysis code on HEC to process the existing archive of Sentinel 1 data, then additional processing to keep up with new data.
  - The resulting L2 displacement maps are used in time series analysis to understand plate boundary deformation, a key science objective of ESI.

Use Case: Mission Need

- Mission Concept:

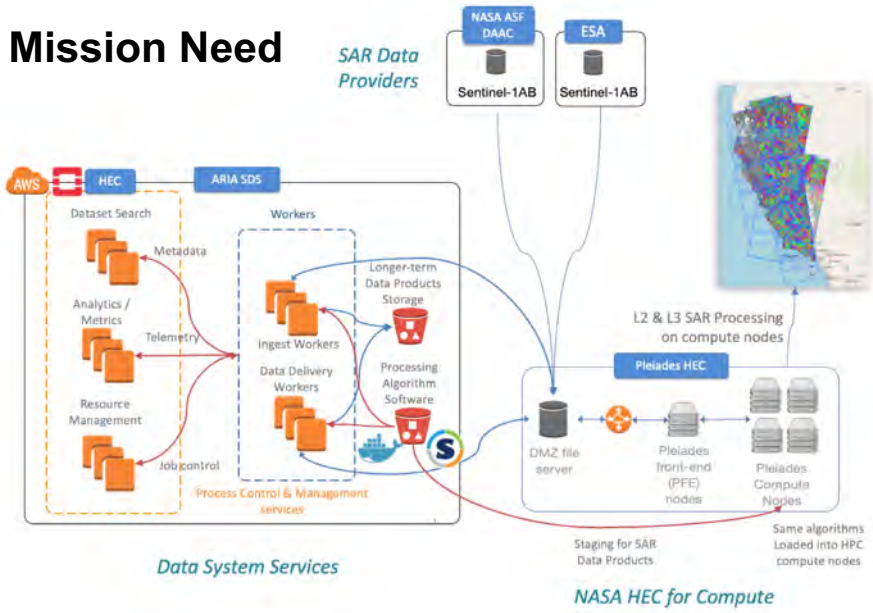

Use Case: Current HEC Application Characteristics

- Work flow:
  - HEC processing jobs are initiated from on-demand requests for processing over an area of interest (AOI) input from the ARIA science data system in AWS
  - Workflow processing dispatched across AWS and HEC, but large SAR compute jobs run on HEC Pleiades.
  - Ideal for bulk processing. Less ideal for lower latency real-time processing.
- Data:
  - In: ~5GB-20GB of input L1 SLCs per job
  - Out: L2 interferometric displacement maps per scene
    - Output data to be stored and distributed via NAS Data Portal
    - L2 displacement maps are used for generating L3 displacement time series for earth surface deformation analysis.
  - Transfer:
    - Baseline products stored and distributed via NAS Data Portal. Other standard products also delivered to ASF DAAC.
- Computation:
  - Type: standard 24-core Broadwell cluster on Pleiades used for L2 processing jobs
- Evaluation:
  - Pleiades can provide for cost-effective processing for R&A science projects.
  - SAR processing is both computationally and disk I/O intensive. Therefore, additional performance gains can be achieved with local on-board SSD disks on each Pleiades compute.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Need for science data reprocessing, proposal cycles (once capability is live and open to community)
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case:
    - Lack of on-board SSD disks on each Pleiades compute forces use of Lustre file system. SAR processing scratch disk usage may benefit more from SSD than Lustre.
    - Ability to run project core services on Pleaides will help to offload external infrastructure dependencies (e.g. AWS)
  - Impact: some additional performance optimization possible with on-board SSDs and collocating core services inside Pleiades.
  - External Sources: science data system core services currently running in AWS.

- How will this Use Case evolve over the next decade?
  - Maturity: Yes, it will be moving from a development state to an operational capability, and there will be with continued development (adding capability to process other SAR data sets, potentially additional workflows)
  - Growth: Future development may require different data volume needs (e.g., analysis of stacks for time series analysis)
  - Re-use: With launch of NISAR, the number of people running InSAR analysis should grow
  - Re-hosting: Algorithm processing codes are Containerized to run across HEC and cloud. Consolidate support of common software Container technologies would benefit. e.g. Docker vs Singularity.

National Aeronautics and Space Administration

Use Case: AI/ML at Scale

Shubha Ranjan
06/09/2020
AMES/Code TNC
HECC

Data Science-AI/ML at Scale

- HECC GPU resources and Data Science Platform
- Meet the ever-increasing demand of various scientific and engineering requirements.
- Focus on cloud computing as a complementary and seamless approach for the traditional HPC
- Access to resources that are not readily available at NAS, such as GPUs,
- Out bursting demand in computing
- In bursting for well-defined and mature workflows
- Support for multiple cloud vendors (ATOs in place)
- Common interface for multiple platforms
- Users can decide based of their requirements for specific functionality and features
- Pick based on allocations or funding
- Customizable computing environment using containers
- Successful establishment for above in the future
- User Training
- Pilot Projects

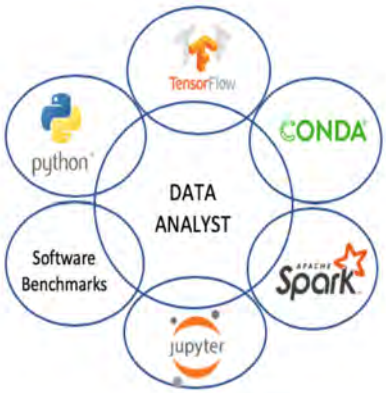

HEC Needs Assessment Workshop

Affinity Group Sessions

Monday 6/9/20 1pm – 4pm (AG6-9) (nominal affinity: Machine Learning) (10am – 1pm Pacific)
Facilitator: Chris Dwan <chris@dwan.org>
Scribes: Jill Dunbar <jill.a.dunbar@nasa.gov> and Michelle Moyer <michelle.c.moyer@nasa.gov>

Note: Most of the discussion was already captured on the slides; the notes do not repeat most of this key information. Also, we did not list chat questions that were not answered during the session, or were answered only in the chat.

Attendees

1. Mike Little, GSFC, convenor
2. Chris Dwan, facilitator
3. **Presenter #1: Christa Peters-Lidard, GSFC**
4. **Presenter #2: Jacqueline Le Moigne, GSFC**
5. **Presenter #3: Kevin Jacobson, LaRC**
6. **Presenter #4: Meelan Choudhari, LaRC**
7. **Presenter #5: Rahul Ramachandran, Presenter #6: Rahul Ramachandran (presenting for Manil Maskey) MSFC/IMPACT**
8. **Presenter #6: Susan Owen, JPL**
9. **Presenter #7: Shubha Ranjan, HECC**
10. Aaron Pina, NASA HQ
11. Arsi Vaziri, NAS
12. Brian Thomas, NASA OCIO
13. Daniel Duffy, NCCS
14. Elizabeth Hartman, HECC
15. Elizabeth Hook, GSFC
16. Ellen Salmon, NCCS
17. Henry Jin, HECC
18. Irina Kitiashvili, NAS
19. Jill Dunbar, HECC
20. Johnny Chang, HECC
21. Laura Carriere, NCCS
22. Leigh Ann Tanner, HECC
23. Mark Carroll, GSFC
24. Michelle Moyer, HECC
25. Peter Williams, Observer
26. Piyush Mehrotra, NAS
27. Robert Ciotti, HECC
28. Robert Ferraro, JPL
29. Robert Hood, HECC
30. Sherry Chang, HECC
31. Subhash Saini, NAS
32. William Thigpen, HECC
33. Ted Manning, HECC

KEY POINTS

Mission Use Cases covered during HEC Need Assessment

Land Information System (LIS), New Observing Strategies (NOS), Laminar-Turbulent Transition Physics, Data Production, Human Landing System (HLS), Pixel level smoke detection by geostationary satellite, Large-scale InSAR analysis, AI/ML at Scale

Impacts of HEC on Mission Needs: Range of Impacts and Significance

- Dependent: HEC provides enough computational power through CPUs and GPUs to handle the large size integrtations needed for large, high-resolution models and the massive data storage resources needed (Peters-Lidard)
- Dependent: Computational Fluid Dynamics 2030 vision depends on meeting goals related to physics-based transition prediction, which, in turn, require HEC.
- Foundational: HEC resources used all but continuously for hierarchical simulations from basic, low complexity components to integrated, multi-discipline models. (Choudhari)

Mission Security Needs relevant to HEC: Range and Significance

Security needs discussed confirm the five areas mentioned in the Use Case outline (bullets) and adds the importance of stability (system instability caused by security threats compromises mission integrity) and protection of intellectual property (of particular concern in public-private partnerships)

- Confidentiality
- Integrity
- Availability: Varies, with some missions dependent and others less so; downtime increases risk of missing transient events (rare and significant);
- Public Confusion/Controversy
- National Security: Export control requirements could be compromised by an inadvertent leak.

Criticality of HEC to Mission Needs

- Note: Similar to "Impacts of HEC to Mission" (above)
- Computational speed (FLOPS) and power (BYTES)
- Data storage capacity
- Allocation time (SBU)
- Financial consequences
- Mission Security
- Massively parallel and Embarrassingly parallel
- Risk reduction: manned missions (safety) and modeling uncertainty/instability
- Knowledge Growth: inability to mature/test breakthrough concepts, such as reduced fuel burn (Choudhari)

Workflow Issues and Needs

- This section and the subsequent one about "key factors affecting load" might be integrated

Data Input, Output, and Transfer: Issues and Needs

- Possibly keep this section focused on processing speed (FLOPS) and data volume (BYTES) as a specific look at these two elements of workflow issues or needs
- *Input, output, and transfer* seem generally associated with--if not dependent upon--computation and storage. Could combine this section with the following section or explain why these are distinctly important given the purposes of this needs assessment (The idea is that either lumping or splitting are fine as long as understandably explained based on meaningful considerations)
- *Computation and Storage:* A principle aspect of this challenge is that output from some work becomes input for other work, thus storage between runs becomes central to workflow. Storage on or as part of HEC system has costs, as does moving off/on (transfer) because the volume of data is significant.

Evaluation: HEC Effectiveness and Value Measures

- Might roll this category under the earlier "Impacts" section [or cross-walk impacts and evaluation to ensure each reinforces the other – evaluate based on critical impacts]
- Performance of computational efficiency, available disk storage, time in queue, work of the support team, code/compiler/system stability,
- Tools for the data analysis, visualization and expertise support

- Number of observation trades/strategies/architecture possible to evaluate in near-realtime (NOS-LeMoigne)
- Researcher demand and impact of resulting science (Owen)

Key Factors Affecting Current Load on Computing Assets

- Publication cycles and University calendars
- Revisions to major models
- Focal area of selected projects (temporal and geographic scales up to planet)
- Reprocessing of science data (preprocessing?)
- Scalability: Effect(s) of adding nodes, increasing amount of data to process
- Availability of adequate cores to complete job(s) in desired timeframes (Jacobson)
- Physics requiring small step times combined with low frequency structures requiring long simulations (tension?) (Jacobson)
- Practical considerations, like availability and queuing, drive decisions about data resolution and sampling frequency in models, as opposed to research needs or project needs (Chouhdari)
    - Data transfer speeds are increasingly a major constraint

Anticipated Evolution of Use Case and Mission Need

- Maturity: Eddy resolving Computational Fluid Dynamics and nonlinear structures will mature
- Growth: Increases in time-step frequency and spatial resolution, data assimilation, ensembles (multi-model, multi-parameter, seasonal forecasts, climate scenarios), model coupling, AI/ML and calibration needed for training, CFD expects smaller timesteps and more of those, NISAR launch likely to increase demand for InSAR analysis runs.
- Re-Use: Advances in CFD likely to attract more users and decrease the required expertise; CFD likely to become more common as part of Multi-Discipline Analysis (MDAO)
- Re-Hosting: Anticipate need for access to emerging hardware to test and develop application
- Other(s): Onboard processing, integration of mission management, anomaly detection, and multiple S/C and multiple nodes, fast integration,

Functional Gaps: Existing and Anticipated

- Constraints: big data analysis (including ML), disk space/data storage, computational allocation, queue, network, I/O, restrictions on working with non-NASA credentialed partners in academia (Choudhari)
    - Limited data storage capacity forces scientists to perform mundane tasks of shuffling data just to keep projects moving forward (Peters-Lidard)
    - "Extrapolations from the most resource intensive jobs we were able to run 5-10 years ago would paint a bleak picture in terms of meeting the widening gaps in the future." (Chouhdari)
- Communication speed between ground and S/C and S/C-S/C (spacecraft)
- Logarithmic relationship between HEC speed and mesh resolution for some modeling means an order of magnitude increase in HEC speed is needed to double resolution (Choudhari)
- Impact of Constraint or Gap: insufficient performance to obtain the required models and data analysis, mission failure, systems considered advanced today will be considered moderate in the near future as anticipated evolution of Use Cases and Mission Needs occur
- External Sources

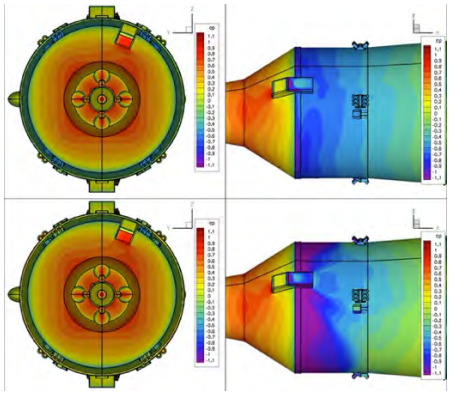

Figure 1. Aries I-X (Jacobson).

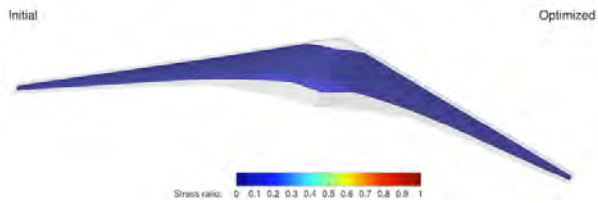

Figure 2. MDAO with CRM wing (Jacobson).

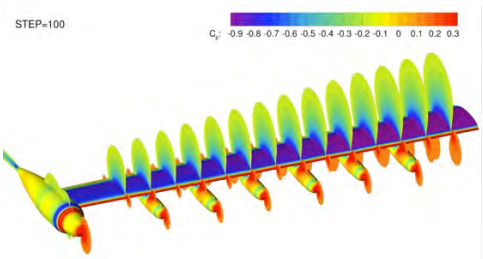

Figure 4. X-57 Maxwell Wing (Jacobson)

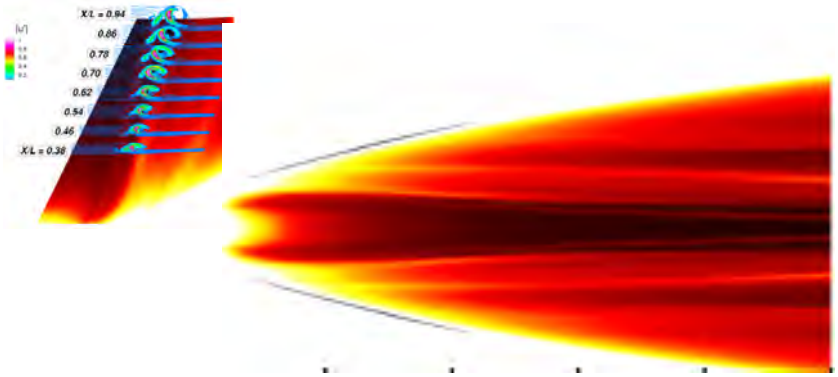

Figure 5. Hierarchical computations reveal the complexity of transition mechanisms on the upcoming BOLT flight experiment (Choudhari)

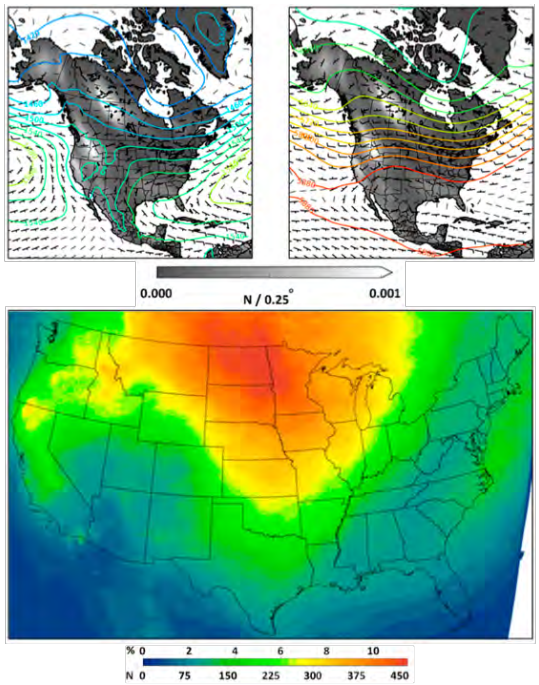

Spatial distribution of MODIS fire occurrence and NOAA HMS smoke for summer 2006-2015. From Kaulfus et al. 2017 Figure 2.

Figure 3. Large-scale InSAR Analysis (Owen)

1. Use Case: Land Information System (LIS)

Christa Peters-Lidard
christa.d.peters-lidard@nasa.gov
Goddard, Land Surface Modeling and Data Assimilation
ESD/ESTO AIST, THP, MAP, ASP/Water, HMA, WWAO, MC DO
Mission Directorate: SMD

Main points:
P1: Overall purpose: being able to provide data products to customers and collaborators in a timely matter. Need seamless access for many users at other government organizations.

P2: HEC is essential for computational power, mass data storage, and data sharing; turnaround time is an issue now for Land Information System (LIS) work. We are using hundreds of cores (~256-512)

P3: In the next 5-10 years, data storage capacity will be an issue. We spend too much time moving data around. The HEC resources should be able to directly access the data repositories to use as input for processing.

P4: If using cloud, recommend a standard NASA suite of libraries and compilers to eliminate sysadmin costs for average user.

Slide 2 Use Case: Mission Need

- LIS framework supports many NASA research/interest priorities. There is a great need for HPC including data storage.
- Without access to HEC, there are two options: purchase dedicated servers, or use commercial cloud. Both involve costs (for server, system administrator staff and maintenance; for cloud, major egress costs).
- Currently provide data to the public using NASA's Goddard Earth Sciences Data and Information Services Center (GES DISC),

Slide 3: Use Case: Mission Need

- Earth science data is open and available on GitHub.
- External collaborators expect products, for example to support drought assessment. If HEC systems aren't available, we can't provide the data to external collaborators. A shared environment is important.
- Need to provide timely runs and timely outputs.
- Identity issues are getting complicated, so getting access to external collaborators is harder.

Slide 4: Use Case: Mission Need

- All Land Surface Data Toolkit (LDT) and Land Surface Verification Toolkit (LVT) all run on HEC resources
- Run with input such as weather data; models predict temperature, evaporation, and so on; data assimilation and uncertainty estimation all require iteration back with model.

Slide 5: Use Case: Current HEC Application Characteristics

- Big challenge: connecting HPC with data program.
- Recently using NCCS data portal for archiving.
- How can data and visualization be better connected?
- Providing graphics output for users via a web page.

Slide 6: Use Case: Current HEC Application Characteristics

- Massively parallel; lots of data, lots of shared libraries, science and computational evaluations, regression testing, timing tests.
- Ensembles; very fine land grid spacing (1km or finer).
- Connection to hydrology – NOAA; forecasting – NCAR.

Slide 7: Key Questions

- Major models are changing and require a whole suite of revision work.
- Changing focal areas, proposal requirements; new instruments (SMAP, GRACE -> SWOT; NASA-ISRO SAR Mission [NISAR]).
- Need seamless access for many users at other government organizations.
- Issue: Having to compress data due to limited disk space.

Slide 8: Key Questions

- Changing data simulation methods, adding ensembles, more complexity, huge community of users, solicitation to use LIS, which is very portable.
- Coupling causes more complexity and increased compute resources.
- Moving to GPUs requires a lot of work.
- Containerization for cloud becoming more important, also requires lots of work.
- Doing a lot of work with LIS; lots of trial and error.
- Need to make LIS accessible to general users.

Q&A:

from Chris Dwan (Ext) to Everyone: 10:21 AM

First question: Are these data sizes per collection? If so, what is the overall data volume per year? How does this grow over time?

from Piyush Mehrotra (Int) to Everyone: 10:22 AM

How come you would have no system administration cost on commercial clouds?

A: My statement was not quite accurate; we handled the sysadmin internally. We have some very savvy computer scientists in our group who did library installs themselves; with cloud you have to install everything yourself - maybe there could be a standard NASA suite of libraries and compilers to eliminate sysadmin costs for average user

from Brian Thomas (Int) to Everyone: 10:22 AM

How much of this data is shared with public (rough size in Xbtyes)? Any idea of how much data are downloaded (volume)?

(Please answer in the survey that will be sent out to all presenters)

from Piyush Mehrotra (Int) to Everyone: 10:22 AM

Typical core counts for a single job?

A: Typically, hundreds of cores, ranging from 256-512, some related to I/O; should be able to scale better soon.

2. Use Case: AIST New Observing Strategies (NOS)

Jacqueline LeMoigne
Jacqueline.LeMoigne@nasa.gov
Earth Science Technology Office (ESTO)
Mission Directorate: SMD

Main points:
P1: If gaps in compute and communications resources are not filled, this could impact the NOS concept.

P2: Need to be able to react to events, and to model and visualize quickly; Need to analyze and process very large amount of data in real time.

Slide 2: Chart

- Not presenting one use case, but a collection of cases.
- NOS is a concept for the future, where all nodes interact and collaborate with each other, as a web or fabric; observation systems that take advantage of space, air, ground sensors.
- This would include many paths for collaboration.
- Coordination can be done on the ground or on board; handling telemetry data; going to science data management, and data mining.

Slide 3: Mission Need

- Using NASA and non-NASA sources.
- From design point of view, this concept could potentially be thousands of different architectures.
- Without HEC, would use commercial cloud.
- Operations: time-sensitive for some deadlines or transient events
- Data communication needs to be secure; also, non-NASA collaborators might need to access the data.

Slide 4: HECC Application Characteristics
Slide not used.

Slide 5: Key questions

- With much more HEC resources, could look at much larger observation trades, which are limited by computing availability.
- With a larger amount of trades, the amount of data will explode in future; HEC needs to take into account scalability.
- Constraints: communication speed – need to take into account response to transient events in real time, and speed of reconfiguration of different assets.
- Looking at the NOS concept over next decade: mega-constellations of nodes (web/fabric) that could communicate with other NASA constellations (and other agency constellations, for example NOAA); integration of multi-source data is going to grow. Data explosion: much more data will need to be processed. If those scalability gaps are not filled, it could make some of the NOS concepts impossible.
- Possibly could run on non-NASA resources, but much of this is still undefined.
- We hope reuse will be possible.

Slide 6: Future HEC Needs of NOS Use Case

- Brainstorming ideas to figure out what could be affected by different concepts
- Looking at all aspects of HEC: cloud, quantum, etc.
- On ground: Fast handling: data management in real time
- Onboard: Fusion in real time of individual data sets and of fusion data. Radiation mitigation by software might require HEC onboard processing (in order to use state of the art hardware that is not radiation-hardened).

Q&A:
None

---

3. Use Case: Computational Aeroelasticity

Kevin Jacobson
kevin.e.jacobson@nasa.gov
LaRC Aeroelasticity Branch, FUN3D dev. team
Advanced Air Vehicles Program (AAVP) &
Transformative Aeronautics Concepts Program (TACP)
Mission Directorate: ARMD

Main points:
P1: NAS is "absolutely essential" for this work, and the work is essential for identifying catastrophic flight conditions; but more resources are needed. Can't use commercial cloud resources for security/confidentiality reasons, export control.

P2: In order to accurately study aeroelasticity, NASA uses experiments (transonic win tunnel) and modeling and simulation. The wind tunnel in heavy demand, so we rely heavily on computation. The computational results are used to guide the wind tunnel tests; among other things, to avoid breaking models inadvertently.

P3: We have the codes; with unlimited resources we could fill the machine 24/7 with bigger and more computations.

P4: Bottom line: is the vehicle safe to fly over the entire flight envelope? Proper assessment of aeroelastic behavior is required or there could be dangerous consequences. A hard question that requires a lot more analyses than we are currently doing. Codes are available, but we're unable to do detailed, larger computations due to resource limitations. We simplify the model to "squeeze through the NAS keyhole" and this reduces the fidelity of the answer.

Slide # 2 Mission Need

- Fun3D simulations for mission work is a big user at NAS.
- Aeroelasticity simulations coupled to low-frequency structures, interactions can rip the plane apart; can be slow (fatigue) or instantaneous (flutter); can happen any time to any vehicle (planes, launch/EDL vehicles, rotorcraft).
- NASA invests a lot of time and money in test models; pretest analysis is important. Aeroelasticity analysis supports wind tunnel experiments before testing, to estimate where instabilities will happen. Without pre-tunnel analysis, researcher doesn't know when the model will break.
- Instabilities can happen anywhere in flight envelope, so analysis must be very thorough. Without proper aeroelasticity assessment, can be life threatening situation.
- Computations are very important because can't get enough time in tunnel.

- Can't use cloud computing for security/confidentiality/export control reasons; and DOD is interested in their own vehicles—not a practical solution for getting resources. Local LaRC cluster is not big enough.
- Limited by resources: need more cores per job for aeroelastic analysis.

Slide 3: Current HEC Application Characteristics

- When limited by resources, making assumptions with questionable validity.
- Failure of linear model example: X-57. The CFD model showed instabilities missed by linear model. So without CFD model we won't know these potential issues. No planned aeroelastic wind tunnel testing—it relies completely on computations.
- Assumptions affect answers—we have to simplify models (for example, use wing but no fuselage) to get jobs through the queues at NAS.
- Tasked with question: is the vehicle safe to fly? Much harder question that requires a lot more analyses than we are currently doing which is analyzing a handful of flight conditions. Codes are available, but we're unable to do detailed, larger computations due to resource limitations.
- In the next decade: Will run more complex simulations, Currently can do 10-100 times bigger on Summit; integrating into different frameworks to be accessible to more users.

Slide 4: Key Questions

- Running FUN3D primary flow solver on Summit (DOE); starting to use GPUs more often for basic work. Expect more migration to GPUs in future as more aeroelasticity in Fun3D gets ported; also, integrating into other frameworks, w/ mesh adaptation.
- Right now, approaching push-button, can be used by more and more users.
- Now tasked with determining: is this vehicle safe to fly? And we need more resources to answer that questions.

Q&A:

from Chris Dwan (Ext) to Everyone: 10:55 AM

Question: Have you done the cost comparison between the cost to deploy more computational resources as opposed to building or waiting for a new physical experiment site (wind tunnel)?

A: Don't know the answer, but NASA is struggling to support wind tunnel so there is not likely a plan to build another. The facility manager says the wind tunnel is valued at $225 million. Building a new tunnel would cost roughly double that. The current tunnel's operating cost is about $150k per day, not including model construction costs and hours spent on test preparation and data analysis.

from William Thigpen (Int) to Everyone: 11:00 AM

When you look at GPUs, how many do you need for your work per run?

A: Currently using 8-12 nodes on Pleiades; on Summit, with hundreds of GPUs (tens of billions degrees of freedom).

---

4. Use Case: Laminar-Turbulent Transition Physics & Modeling AAVP (HTP) and TACP (TTT)

Meelan Choudhari
m.m.choudhari@nasa.gov

AAVP+TACP
Mission Directorate: ARMD

Main points:
P1: HEC is critical—95% of computations done on NAS systems.

P2: NAS staff have been outstanding over the years to make applications feasible, but the hardware resources are barely adequate.

P3: Due to hardware limitations, exploring getting cycles on DOE and DOD machines, and through academic partners at the National Science Foundation; but it's not a practical solution.

P4: NAS experts could play a much more important role in porting codes to GPUS than in the past.

Slide 2 Use Case: Multi-fidelity modeling of boundary-layer transition (BLT)

- Niche area of flow modeling: how ordinary laminar flow becomes chaotic or turbulent. Transcends all speed regimes and many types of vehicles. Drag and surface heating goes up order of magnitude during transition.
- Use variable fidelity modeling of different types of waves that lead to transition to capture the whole process. Solving for all of the frequencies simultaneously against evolving them in time.
- Transitions at all scales are involved, in order to resolve region full Navier-Stokes equations. This is most of the NAS usage.
- It takes lots of resources to design experiments … goal is to develop reduced order models that can be integrated into NASA CFD solvers, also to help design experiments, and interpret experiments.
- Computations are very important because of the sensitivity in minute flow details; can't model at full-scale Reynolds numbers in wind tunnel experiments, especially at higher speeds—get limited results.
- Direct numerical simulations give more details.
- Dependent on NASA HEC - about 95% of computations take place at NAS. Due to hardware limitations, now also exploring getting cycles on DOE and DOD machines; also w/ academic partners thru NSF or individual states.
- Configurations are basic and fundamental, mostly, but some are SBU; w/ supersonic, there is increasing sensitivity.
- HEC is used continually for different types of simulations from low to high fidelity. Current mesh sizes range from 100 million to 2-3 billion; even those are barely adequate to serve the goals. Even though NAS staff are outstanding throughout the years, HW limitations are increasingly constraining for this type of work.
- Using other systems to fill the computational gap is not realistic; this is not a satisfactory substitute for HEC. HEC is critical for ARMD projects for this type of work. No satisfactory substitute; increasing use of DOE and DOD resources not realistic option in my opinion.

Slide 3: Current HEC Application Characteristics

- Recently using deep learning to build more robust models; so far, machine learning work has been limited to building regression models based on reduced order models; but in the future, as we ramp up DNS effort, time-dependent data will be used increasingly.
- Could use 10x capacity now and in the future. Post-processing includes full visualization and statistical analysis; need millions of low-fidelity models to support machine learning work; most is archived on HEC storage.

- Doubling the resolution of 3D/4D models in every direction for higher fidelity is an order-of-magnitude increase; increasing HW by 10x gets only a 2x increase in resolution.

Slide 4: Key Questions

- Constrained by available capacity rather than types of problems; codes can easily scale up to 50 – 100 K or more, but we have never been able to get more than 5K cores in order to optimize.
- Doesn't paint an encouraging picture to bridge the gap.
- Expect user base will increase substantially.
- With the resurgence of interest in high-speed vehicles, we anticipate much higher usage of high-fidelity computations using HEC—a couple of orders of magnitude higher.
- Using GPUs is a work in progress: haven't made significant breakthroughs in porting codes. NAS staff would need to help.

Q&A:
Few questions; moved on.

---

5. Use Case: Data Production (Forward processing and Reprocessing)

Rahul Ramachandran
rahul.ramachandran@nasa.gov
MSFC/IMPACT
Earth Science Data Systems Program
Mission Directorate: SMD

Main Points

- Key gaps: infrastructure middleware seamlessly between cloud and HEC without having to rewrite code
- Collaborative science platforms for researchers to use that can utilize both cloud and HEC.

Slide 2: Harmonized Landsat and Sentinel 2 (HLS) project as canonical example

- Team took an existing science algorithm for harmonizing Landsat, Sentinel 2 for specific sites; built a workflow using the algorithm to generate global product.

Slide 3:

- Building processing pipeline.
- Working within cloud have to focus on cost optimization. Landsat data already in cloud.
- Get data from three different sources, including models to initialize processing.
- Goal to migrate processing to GCC, data to NGAP cloud

Slide 4: HLS Production

- More and more migrating to commercial cloud. Most doesn't need security.
- Need a mix of cloud and HEC to optimize cost.
- Many external collaborators require access to computer resources.
- Forward processing: Must be produced in timely manner to be distributed to different users (not true for backward processing).

Slide 5:

- Operator manages errors.
- Fairly high volume datasets.

Slide 6: Key factors

- Data processing needs are changing; use from different instruments is increasing; may become norm to use different sets of data from the archive. HHLS is a good example: multiple data streams from different sensors.
- Need quick access and to move between HEC and cloud without too much effort – can take a lot of time. (Need access to resources on quick timelines, permissions for large teams.)
- Moving to GCC cloud.
- Same workflow needed for data from different instruments; need a common framework; shouldn't be recreating worflows optimized for processing.

Slide 7:

- Conceptual vision: Cloud-HEC integration: need ways to access data and seamlessly move processing between cloud and HEC; collaborative science platforms for researchers to use that can utilize both cloud and HEC.

Q&A:

from Chris Dwan (Ext) to Everyone: 11:29 AM

I'm interested in the effort required (if any) to move from a file-based data store to an object based (cloud buckets). In other domains, I've seen that come with a fairly high premium in terms of refactoring.

A: Worried about cost, would have to use large instances, lots of memory, would be very costly.

from Brian Thomas (Int) to Everyone: 11:32 AM

for moving computing between HEC and cloud, what do you expect to use? Containers?

A: Yes, but try to break them down so they can leverage lambda function.

---

6. Use Case: Scaling pixel level smoke detection for high temporal geostationary satellite data

Manil Maskey (Original presenter)
manil.maskey@nasa.gov
Rahul Ramachandran (Actual Presenter)
rahul.ramachandran@nasa.gov
NASA/MSFC IMPACT
Earth Science Data Systems Program
Mission Directorate: SMD

(Note: Rahul Ramachandran spoke in Manil's place and went through slides quickly so no discussion)

Slide 2: Introduction

- Large volume of temporal data.
- Smoke impact dependent on various factors.

Slide 3: Objective

- Using GOES data to detect smoke in real time.

Slide 4: Approach

- Hand-labelled not very accurate, had to be recalibrated.

Slides 5-9: Model architecture

- Uses specific region.
- Algorithm detected well.

Slide 7: Opportunity to leverage HEC

- Opportunity to use HEC in near real time, cloud is expensive (especially pixels).
- Near real time (every 10 minutes from GOES data).
- Specific example of deep learning that can be used not only for training but also for inference in near real time.

No questions, will follow up later with summary and survey.

---

7. Use Case: Large-scale InSAR analysis

Susan Owen
Susan.Owen@jpl.nasa.gov
Earth Surface and Interior
Mission Directorate: SMD

Main points:
P1: SAR processing is always under development; will continue for next decade.

P2: Preparing for launch of NISAR, so users are increasing and will continue to grow.

P3: Have had to figure out two container technologies, Docker and Singularity, in order to user both HEC and AWS; we need a common container technology.

Raw notes: About halfway through developing capability.

Slide 2: Mission Need

- Traditionally, analysis is done over small areas; with the upcoming launch of NISAR in 2022, will be able to look at larger scale; this creates data processing challenges. Difficult to do on the cloud, as cloud resources are prohibitively expensive; this would eat up all the ROSES funding.
- Being able to run on Pleiades would be a large advantage. Developing tools to allow researchers outside of (?) to use Pleiades. Interruptions/delays not a huge problem.
- Want academic researchers as well as NASA to have access to data. Will measure impact by demand and by science results.
- Security: Open and free data; Sentinel data stored in DAAC; generally scientists want to look at results before they get stored but w/ large amt of processing, there will be an automated quality check.

- Function: Running SAR analysis code on Pleiades to process Sentinel 1 data; backlog of data exists from this mission; scientists want to analyze data with more than one pair of data (?) resulting time series needed to address questions and analyze data, understand plate boundary.

Slide 3: Mission Need
(Workflow picture)

- ARIA SDS developed to analyze SAR (Synthetic Aperture Radar) data for quick response (for example, earthquakes or flood maps from hurricanes), but also needed for large amount of data for scientists. Need the ability to respond quickly to an event.
- Need Pleiades for very large spatial scales for analyzing data in standard ways.
- Had an existing local system and AWS; processing done using Docker containers. Docker won't work on Pleiades. Test case for Singularity container, working closely with NAS team to work out issues.
- Currently, use automated software packaging and Singularity instead of Docker in order to run on Pleiades.

Slide 4: Current HEC Application Characteristics

- Large-scale SAR jobs run on Pleiades; rapid response run on AWS.
- Currently working out challenges of launching at scale to run across 100-200 nodes on Pleiades. Close to running end-to-end but not completed yet. Still being developed.
- Alaska Satellite Facility – Users will be able to request what they want.
- Data: 5-20 GB of input for SAR per job; output distributed by data portal; time-series processing done on local workstation.
- ARIA has standard product stored at AWS; user can make requests and that data will be available via NAS data portal.

Slide 5: Key Questions

- Dependent now on AWS for running core services jobs; would like to run more of the core services on Pleiades.
- More storage closer to compute = additional performance optimization.
- SAR processing always under development; will continue for next decade; hope to have it operational because lots of data products can be generated with SAR.
- Getting ready for launch of NISAR in 2022; SAR users are increasing and will continue to grow.
- Have had to figure out two container technologies, Docker and Singularity, in order to user both HEC and AWS; we need a common container technology.

Q&A (Extended chat excerpt regarding Docker and security)
06/09/2020 15:04:53 PM from Michael Little (Int) to Everyone: Sue, have you been able to figure out how hard it is to convert docker to singularity and back again.
A: No answer captured

06/09/2020 15:05:39 PM from Robert Ciotti (Int) to Everyone: Security needs arise when connecting their system (HySDS) to pleiades. there are a number of security issues in interfacing the cloud based system to 'reverse burst; to pleiades. containers were also an internal development activity that took a couple years to secure and included testing different container frameworks.

06/09/2020 15:06:12 PM from Brian Thomas (Int) to Everyone: I'm a newbie, so you all probably know the answer to this, but why can't you use Docker on Pleiades? Security concern(s)?

06/09/2020 15:07:16 PM from Mark Carroll (Int) to Everyone: Correct Brian. Security concerns with Docker

06/09/2020 15:08:37 PM from Brian Thomas (Int) to Everyone: From an agency perspective, we would save work/cost if we can address security for Docker. In other projects I have worked in, we have developed hardened (docker) containers which we could require in a security plan for processing containers on Pleiades. The Air Force is using a similar approach.

06/09/2020 15:08:45 PM from Robert Ciotti (Int) to Everyone: Many critical security vulnerabilities. many exploits. The fundamental design required root access that we can't give to users. addressing that issue was extremely time consuming and actually required us to work directly with singularity development to address various problems to remove the need for root. One can convert docker containers to singularity. That's what HySDS does

06/09/2020 15:11:04 PM from Mark Carroll (Int) to Everyone: You can also run Docker inside Singularity which limits exposure to vulnerabilities

06/09/2020 15:11:34 PM from Robert Ciotti (Int) to Everyone: That's not always the case.

06/09/2020 15:11:49 PM from Robert Ciotti (Int) to Everyone: Unfortunately - we have had trouble with that for various reasons.

06/09/2020 15:12:14 PM from Chris Dwan (Ext) to Everyone: We have substantial experience with securing Docker in clinical contexts. The underlying vulnerability is the challenge in maintaining user identity in accessing external resources like file stores and external APIs.

---

8. Use Case: AI/ML at Scale

Shubha Ranjan
shubha.ranjan@nasa.gov
HECC, Code TNC
SMD

Main points:
P1: In the future, ML needs to become more seamless and scalable; containers and/or common platform are key. Lots of work needed to make it more user friendly.

Slide 2: Data Science-AI/ML at Scale

- ML/AI at HECC; we are interested in use cases at HECC; HECC survey coming out soon to gauge future usage/requirements for Data Science in ML/AI.
- Current HECC functionality: experimental clusters of GPUs (V100 and K40) to gauge demand. Nowhere near close to demand coming up in 3-5 years. Could increase with in-house or commercial.
- Current software includes "all the usual" software (see chart) provided in convenient conda environments that work on GPUs and CPUs (HECC has more CPUs).
- Held a recent training session for ML; deep learning training session coming up.
- Team has had successful pilot projects helping teams move into ML; can help others do the same.

- Future: Cloud can have big role in scaling, or for providing features we don't have (like science or ML packages). For example, Google AI has a toolset to gain insight in modeling while building neural network. Their Explainability Initiative
- Containers are key for seamless use between from cloud to HEC. Should be scalable, more user friendly, as it currently needs a lot of intervention. Need a common platform.
- Outbursting to meet demand or to provide capability. Pilot project – astronaut health model – expanding for use by doctors; patient retina scans could be hosted on public cloud for research to be used by doctors worldwide.
- In-bursting; train model on subset of the datasets on AWS, then bring back in-house to run on larger datasets save costs.
- 3-5 years out. Challenges: (1) data format, how to optimize and distribute; (2) ability to add individual clouds to common platform, each cloud vendors' services require explicit compliance certification; (3) integration with existing systems for accounting and account management.
- Long queue wait times. Cloud could mitigate some, but very expensive. Need to find a balance.

from Brian Thomas (Int) to Everyone: 12:22 PM

Another newbie question : Is the software from Data Science platform at ARC also available for GSFC HEC resources? If not, what are the blockers to having a shared software platform for HEC across NASA?

A: What we have installed at Ames are open source software, can be installed at Goddard. We can provide assistance. There is comprehensive, publicly available information on the HECC Knowledge Base: https://www.nas.nasa.gov/hecc/support/kb/173/

Christa Peters-lidard: RE: Re-Hosting: Beyond containers, we had challenges in SMCE trying to use kubernetes, then we moved to ECS to make it work. What is the best practice converting MPI code to cloud-ready parallel code? Is AWS Parallel Cluster standard now in AMCE? How can we make this transition easier? In addition to egress costs, It is a roadblock for us to move from NCCS to cl oud

Rahul: [Related to Shubha's comments about outbursting} Cloud works for long tail science, but we need the capability to inburst if needed to scale out research cost effectively

---

Group Discussion of Key Questions

Brian Thomas: Clarification: it takes a lot of effort to maintain secure, compatible software environment. Why are software platforms not being developed across agency? Shubha is proposing this.

Bob Ciotti: Are you referring to the system that Susan was talking about?

Brian Thomas: Broader: a common platform solves a number of problems. Puts all under security plan/review; continual maintenance for usability downstream; shared KB that users can have access to and build experience for future use.

Bob Ciotti: What do you envision the platform would be? We are trying to provide the lowest underlying features (below the view of what the user sees) that the pipelines can be built on top of and moved around. There are issues with container problems (authoring), requires root access, anything fast moving—development like Docker or Singularity—we don't want to take that risk. HySDS is a basic

platform used to run generic data science data pipeline workloads on Amazon or Pleiades to be used for SAR; numerous different pipelines. The SDS platform requires underlying support, long-term activities, may require specific systems, have to use particular interface system. Various issues on the systems side: we provide a minimally needed set of tools to enable the pipelines.

Chris Dwan: Question of what a platform means on a cloud environment is industry wide. What are key factors today that drive usage and load on assets: for example: seasonality, interns, etc.

Christa Peters-Lidard: Key driver is spatial and temporal resolution of the model, solving differential eqations across time/space scales.

Meelan Choudhari: In response to the question regarding the seasonality of loads, I would like to mention that the ability of large scale simulation to impact high cost (flight) experiments is limited by what can be achieved in the relatively condensed timeframe for the planning and execution of the flight experiment. In other words, the project deadlines are an important factor in driving the load.

Chris Dwan: What are you not doing due to resource limitations?

Christa Peters-Lidard: Ensembles. Can be used for different reasons - physics, parameters, data assimilation; multiplier of model. Used for uncertainty; can't do full analysis because can't fit it all.

Chris Dwan: What's the impact of that?

Christa Peters-Lidard: Increasing as a challenge; smaller-scale ensembles. Physics gets more complicated as you drill down.

Kevin Jacobson: Problem size driven by UQ, we don't really touch it. If we aren't doing it, we could lose a vehicle if analysis is not good enough (esp. for Mars missions).

Rahul Ramachandran: Similar use cases that will require production with multiple data streams; demand is constant not episodic; for ML example of phenomena detection, need computational power will be episodicepisodic

Chris Dwan: In terms of shifting to other architectures, do you have a re-architecting on the horizon? If so, what resources would accelerate the move to new architecture?

Kevin Jacobson: GPU migration in progress. Lack of resources has been an issue; local cluster had competition for resources; participated in DOE GPU hackathon but had to freeze development work to do production work. Need dedicated resources for development.

Chris Dwan: Any other non-nasa facilities in play?

No response.

Chris Dwan: Forward looking, what is changing as capabilities mature and move to operations? Going from RND to operational?

Christa Peters-Lidard: A two-category system is better for us. Have a development version that's constantly pushing the boundaries of data and physics; need for higher resolution, etc.; with an operational version behind. Constant transition between the two. And, there is always a pipeline of new development coming up.

Rahul Ramachandran: In my use case, we don't have capability of migrating seamlessly to HEC.

Chris Dwan: What's holding you back?

Rahul Ramachandran: To do it correctly requires resources, agreements, common framework, not one-off implementation that we can do often.

Chris Dwan: Best practices around implementation?

Rahul Ramachandran: Need more than best practices; not just technical problem. Need common platform, well defined processes

Shubha Ranjan: Across HEC and cloud for in- and out- bursting.

Jacqueline LeMoigne: R&D is not as predictable as operational.

Chris Dwan: Bursting from HEC resources to cloud-anyone doing it without undue effort?
That doesn't rise to a project-level effort.

Christa Peters-Lidard: Need a whole project to get it done.

Rahul Ramachandran: When you have dependencies on other groups, you can have issues.

Chris Dwan: Speaks to needs for common standards and practices. Growth of user community: expanding? Teams or collaborators?

Christa Peters-Lidard: We would expect growth of people wanting to do SAR analysis as NISAR launches; greater demand on Pleiades for SAR analysis.

Rahul Ramachandran: We have done some surveys on adopting ML; there is steady growth. Based on that, there will be a need for more usage of ML approaches to conduct science.

Mark Carroll: A lot more interest in ML at GSFC and requests for help have dramatically increased over the last 2 years.

Chris Dwan: Lightning round: number one ask for HEC?

Christa Peters-Lidard: Seamless connection between HEC and DAAC or other data distribution. Having to move lots of data now, I don't want to move anything: want it to be instantly available. Two other asks: 1. GPU code conversion; 2. standard NASA cloud setup with compilers/libraries and assistance with porting MPI code to cloud.

Jacqueline LeMoigne: Benchmarking for ML algorithms on different platforms.

Kevin Jacobson: Ability to get bigger and more jobs thorugh the queues. More capacity.

Rahul Ramachandran: Tighter integration between future data systems and HEC. Science analysis on cloud, integration to scale out analysis through HEC if they have an option. Improved 1) access, 2) procurement and replacement, 3) allow long-tail science users the ability to scale

Susan Owen: Having onboard solid state disks (on the compute nodes); also, common container technology.

Shubha Ranjan: Seamless platforms for HEC, other NASA data centers and clouds.

Brian Thomas: Should be possible to create architecture where we enforce containers which are run non-root. Basically require all base images for containers must be those created by HEC staff (HEC container registry) which are built so that they change the user. Then, to prevent reversion, insert step to scan for changes back to ROOT before allowing registration of user created images on the system

Appendix D7 Workshop Session 6-19 Inputs

William Putman	GSFC 610.1	Global simulations of increasing resolution and complexity
Nikki Privé, Will McCarty, Yanqiu Zhu	GSFC 610.1	Meteorological OSSEs
Ricardo Todling	GSFC 610.1	Forefront Numerical Weather Analysis & Prediction
R. Gelaro, E. Hackert, B. Putman, R. Todling, A. El Akkraoui	GSFC 610.1	Coupled Analysis/Reanalysis: MERRA-3 and Beyond
Andrea Molod, Eric Hackert, Yury Vikhliaev, Bin Zhao, Donifan Barahona, Anna Borovikov, Robin Kovach, Siegfried Schubert, Zhao Li, Young-Kwon Lim, Lauren Andrews, Richard Cullather, Randal Koster, James Carton, Kazumi Nakada	GSFC 610.1	S2S2D Prediction and Predictability: Retrospective Forecast Suite
Steven Pawson, Anton Darmenov, Christoph Keller	GSFC 610.1	GEOS-CF
Peter Colarco	GSFC 610.1	Chemistry-Climate Modeling

National Aeronautics and Space Administration

Workshop Session 6-19

19 June 2020

National Aeronautics and Space Administration

Use Case: Global simulations of increasing resolution and complexity

Bill Putman
June 19, 2020
GSFC Global Modeling and Assimilation Office, Code 610.1
Science Mission Directorate

Use Case: Mission Need

- Mission: Forefront simulations pushing the resolution and complexity of the represented Earth system form the basis for future application development, observing system design and adaptability to emerging HEC trends.
  - This application of GEOS pushes the limits of the current HEC capability computing and exposes weaknesses in algorithm, infrastructure and workflow design for GEOS systems that will be production applications in 5-10 years.
  - These simulations provide:
    - Representative 'Nature Runs' for OSSE development
    - High fidelity output for process studies
    - Large amounts of training data for parameterization and artificial intelligence development
- Criticality: High-end computing (HEC) at NASA provides an essential resource for completion of this work.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - This work would rely on proposals to external leadership class computing at the DOE and elsewhere
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Given the research nature of this work and a lack of direct real-time mission support, this work is more easily adaptable to fluctuations in compute availability.
  - Vetted data can be shared with the community through open dataportals and external users are not currently required to execute jobs
- HEC Function:
  - This work is often performed as 'simulations of opportunity' on new and emerging platforms as they become available at NCCS and NAS

Use Case: Mission Concept

- Mission Concept
  - Project DYAMOND: Simulating August 2016 (40-days)
    - Dynamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains
    - 3-km global convection permitting simulations executed as an ensemble of simulations with global collaborators to understand the successes & limitations of current modeling systems to represent the fidelity & complexity of the Earth system at the resolution of our observing system
    - 1 instance requires 40,000 cores at NCCS for ~2 weeks
    - Producing ~600TB of data

GEOS 3-km DYAMOND Simulation : August 2016

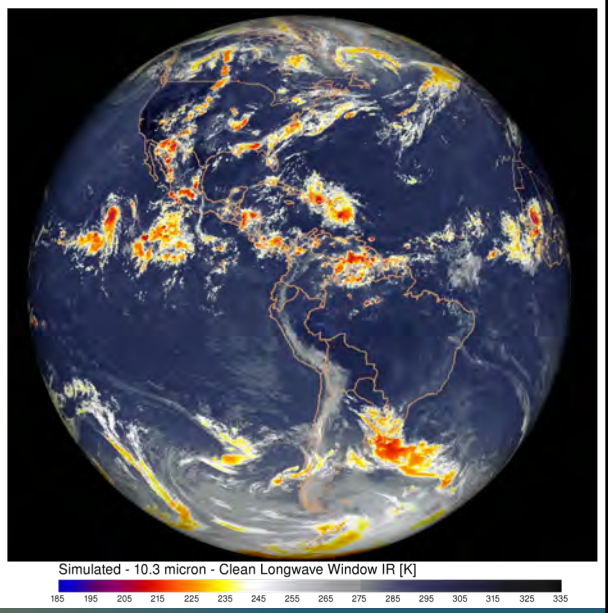

Use Case: Mission Concept

- Mission Concept
  - Project DYAMOND: Simulating August 2016 (40-days)
    - Dynamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains
    - 3-km global convection permitting simulations executed as an ensemble of simulations with global collaborators to understand the successes & limitations of current modeling systems to represent the fidelity & complexity of the Earth system at the resolution of our observing system
    - 1 instance requires 40,000 cores at NCCS for ~2 weeks
    - Producing ~600TB of data
  - In reality, this requires many attempts (and more cycles) to:
    - Resolve performance, infrastructure and HPC issues
    - Improve science capability
    - Evaluate the role and limitations of current algorithm/parameterization design
    - Explore the impacts of increasing complexity

GEOS 3-km DYAMOND Simulation : August 2016

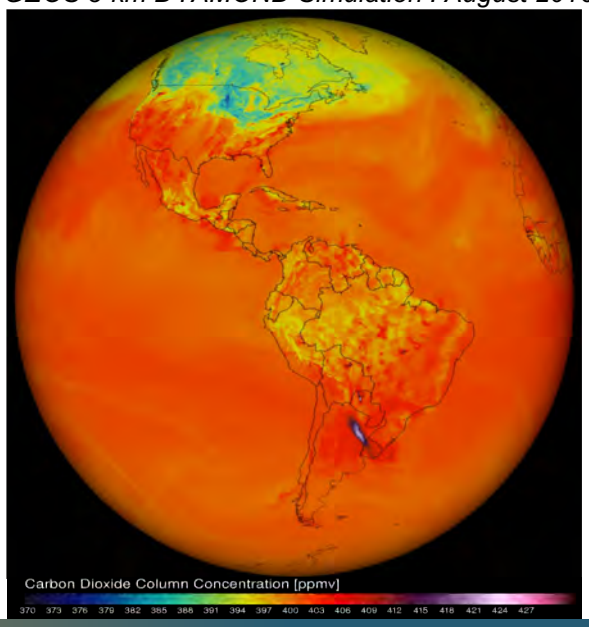

Use Case: Mission Concept

- Mission Concept
  - Project DYAMOND: Simulating August 2016 (40-days)
    - Dynamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains
    - 3-km global convection permitting simulations executed as an ensemble of simulations with global collaborators to understand the successes & limitations of current modeling systems to represent the fidelity & complexity of the Earth system at the resolution of our observing system
    - 1 instance requires 40,000 cores at NCCS for ~2 weeks
    - Producing ~600TB of data
  - In reality, this requires many attempts (and more cycles) to:
    - Resolve performance, infrastructure and HPC issues
    - Improve science capability
    - Evaluate the role and limitations of current algorithm/parameterization design
    - Explore the impacts of increasing complexity
  - Limited by availability of HEC resources and the need to balance GMAO production, pre-production testing and research simulations among the other applications

Estimated Footprint of NCCS Compute Capability in 2021

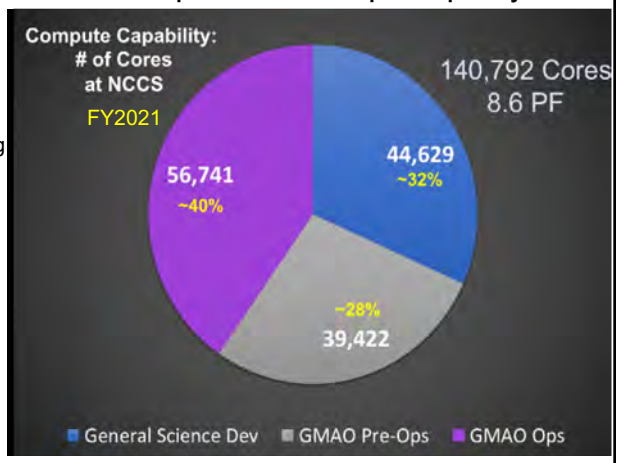

We are utilizing capability at NAS as well to complete coupled ATM/OCN simulations in 2020, next slide…

Use Case: Mission Concept Evolution (Coupled Atm/Ocn)

2020 Capability		by 2025		by 2030	
Resolution	*Cores/Data*	*Resolution*	*Cores/Data*	*Resolution*	*Cores/Data*
Atmosphere 6 km *137 Levels*	**O(10,000)** **500 TB** per simulation	**Atmosphere** 3 km *181 Levels*	**O(100,000)** **>5 PB** per simulation	**Atmosphere** 1.5 km *>200 Levels*	**O(1 million)** **>100 PB** per simulation
Ocean *4 km* *90 Levels*		***Ocean*** *2 km* *100 Levels*		***Ocean*** *1 km* *150 Levels*	

- High resolution means lots of work for HEC, ideal for accelerator technologies but application needs adaptation
  - **Domain Specific Languages** – application portability and efficiency for the dynamics and transport
  - **Machine Learning** – development for physics and chemistry (observation driven and performance oriented)
- Use of these applications for OSSE Nature Runs requires a new workflow for partners *(See Arlindo's Use Case)*
  - Current design tries to capture everything for all OSSEs for the entire simulation at once (*this is unsustainable*)
  - Envisioning one long compute (1-year) with minimal output to evaluate and catalogue simulated processes
  - Restart capability to reproduce segments of the Nature Run with OSSE required output and frequency
  - On-the-fly observation operators to eliminate complex post-processing workflow
  - Containerizing the system for OSSE partners to execute their own segments and produce their own data

Use Case: Current HEC Application Characteristics

- Work flow: Basic submit, post-process, evaluate, repeat…
  - User submission
  - Is there a need for interaction by user, operator or other system in the computing workflow? **NO**
  - Does this use case need quick reaction or near-real time computing? **NO**
- Data:
  - In: High resolution boundary data for atm/ocn and emissions [minimal]
  - Out: See previous slide…O(10)PB
    - Enough data to evaluate, catalogue and reproduce the run
    - The data are analyzed statistically against Observations & Reanalysis and through visualization
    - Create a catalogue of events [such as Tropical and Extra-Tropical storms, Mesoscale Convective Clusters, extreme drought/flood/temperature/wind/aerosol/chem events]
  - Transfer:
    - Data remains in centralized storage on-site, shared publicly via dataportal
- Computation:
  - Type: We can utilize peak computing capability at NCCS/NAS O(10-100) Thousand Cores.
    - Would benefit from accelerator adaptation
- Evaluation: HEC is effective if we can improve our representation the Earth system via resolution/complexity

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Major revisions to models and availability of new HEC hardware dictate the frequency of runs with high resolution/complexity
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Capacity to execute large jobs and improve our models
  - Impact: These research runs would benefit from acceleration platforms to improve turnaround and minimize impact on production and other research across the GMAO and SMD
  - External Sources: Could leverage emerging Exascale platforms at leadership computing facilities
- How will this Use Case evolve over the next decade?
  - Maturity: Continued development: higher resolution, increased complexity, and improved algorithms
  - Growth: Roughly a doubling of resolution every 5-years, including time-step fidelity
  - Re-use: Potentially, with design for containerization of Nature Runs for the OSSE community
  - Re-hosting: Code rewrite via DSL, and use of AI in physics and analysis techniques

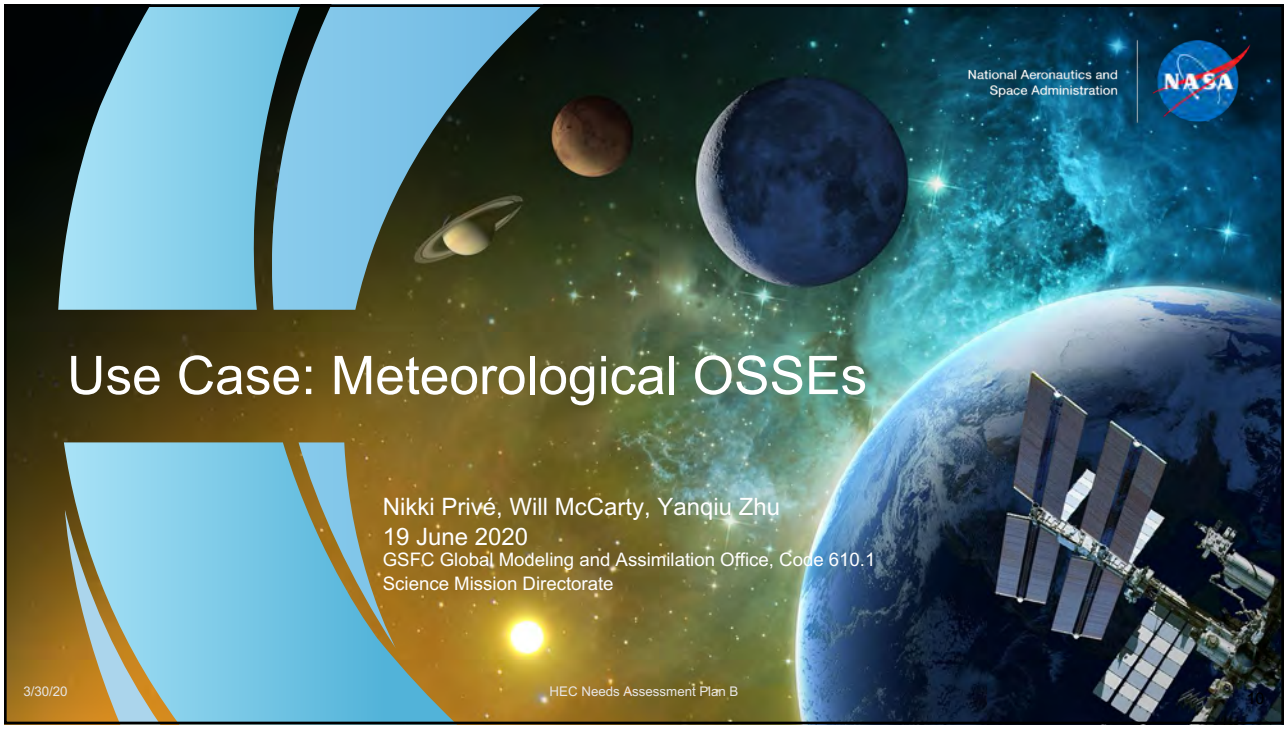

Meteorological OSSEs: Now and Future

Mission Need (1 of 2)

- **Mission Objective**:
  - Simulation of proposed new satellite observations provides guidance on instrument design
  - Utilize simulated framework to accelerate the implementation of real data into operational and research production systems
- **Use Case Example:** ➤

OSSE Control

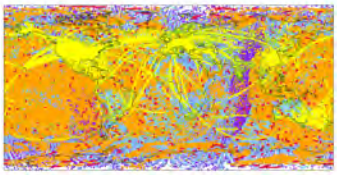

A simulated baseline that reflects reality against which we can test new observation concept

OSSE Experiment

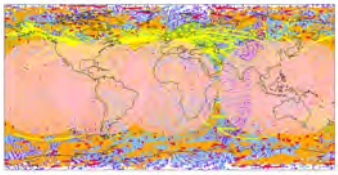

By adding a new observing system to the control, we can now quantify its (simulated) impact

The current global observing system (top-left) only measures hyperspectral infrared radiance spectra from polar orbit (e.g. AIRS, IASI, CrIS, cyan)

The GOES Program has funded the GMAO to use our OSSE system to help establish a trade space in quantifying the impact of future hyperspectral IR sounders (pink, top-right) in geostationary orbit

Meteorological OSSEs: Now and Future

Mission Need (2 of 2)

- **Criticality**: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - High quality OSSEs would not be possible without HEC resources
  - OSSEs are useful in developing a trade space, which requires multiple sets of experimentation
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - Accidental release of data has potential risk; the exposure of proprietary information associated with instrument development is possible
    - This can range from competition-sensitive information to bordering on ITAR
  - Continuous availability of HEC resources is not required. Sufficient resources for intensive calculations and high resolution modeling are needed on an intermittent basis, perhaps surging during competed proposal cycles
  - Simulated observations and model outputs need to be accessed by external collaborators (universities, other agencies). Collaborators may also need to perform high resolution modeling runs to test observation simulators.
- **HEC Function**:
  - A global high-resolution model with 4D ensemble variation data assimilation system will be run for multiple 2-3 month periods for each proposed instrument. The frequency of experiments depends on demand from instrument teams.
- **Impact of HEC**: See “Criticality” above.

- **Current Workflow**:
  - Linear and segmented
  - Nature run is run, written, and validated
    - Significant IO, computationally expensive, large footprint
  - Nature run output is used to generate observations
    - Significant IO, parallel/scalable, moderate footprint
  - Data assimilation experimentation is performed using the observation
    - IO & computation consistent w/ other standard DA experiments
- Current System is an atmosphere-only NR and data assimilation system

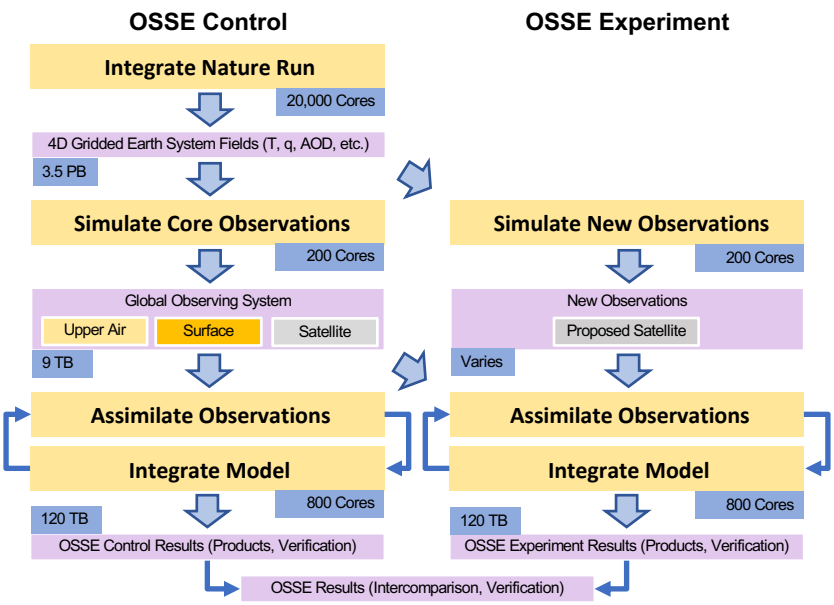

- **Future Workflow**
  - Limit disk footprint & IO at the expense of more computations
  - Observations are generated within the Nature Run
    - Requires re-integration of NR for new observing systems, which requires bit-wise reproducibility of the NR/model over its usable lifespan
  - Experimentation similar to current workflow
- Future workflow should be designed to leverage GMAO Earth System modeling capabilities and be extendible beyond atmosphere-only applications

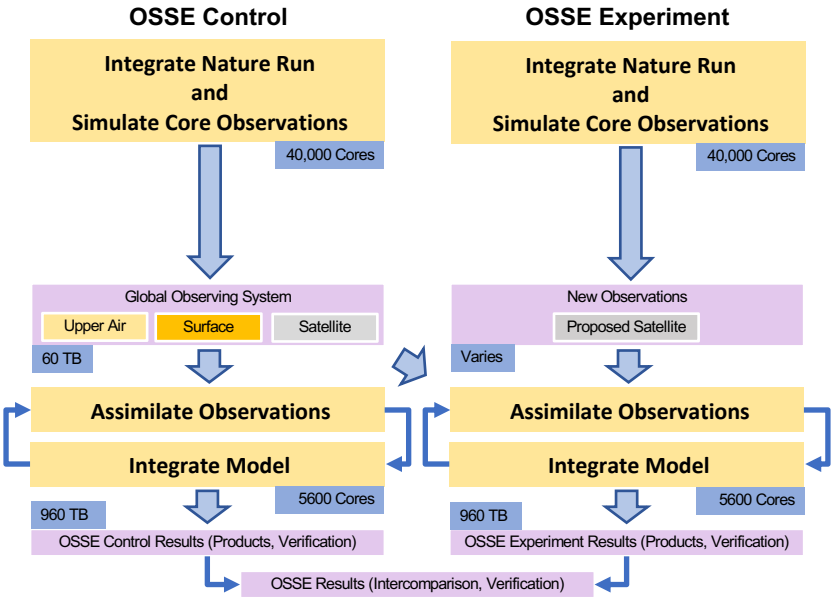

Current/Future HEC Application Characteristics

	Current	Future
Data	Nature Run: 7 km horizontal resolution, 72 level model Experiment: 3 Months, 25 km horizontal resolution, 72 level model **In:** • Nature run: 3.5 PB (resolution ~ 4x finer than experimentation, integration length of 2 years) **Out:** • Observations: 9 TB per season • Model and DA Outputs: 120 TB per perturbation (note, each OSSE contains numerous perturbations) **Transfer:** • The G5NR is distributed via OPeNDAP • OSSE experimentation stays at NCCS, with negligible transfer across the research community	Nature Run: 6 km/181 L Atm, 4 km/90 L Ocean Experiment: 3 Months, 12 km horizontal resolution, 132 level model **In:** • Nature Run: 500 TB reduced output **Out:** • Observations: 60 TB per season • Model and DA Outputs: 960 TB per perturbation **Transfer:** • Strategy to deliver observations versus Nature run
Computation	• GEOS-5 Nature Run (G5NR) is no-cost beyond IO since it has already been run • Observation generation computational expense is primarily IO limited • 800 cores per perturbation	• To simulate observations, new nature runs will require 40000 cores for NR re-integration and in-line observation simulation • In-model simulation will require the development of a modular and easily-used infrastructure for external users to integrate their instrument simulators • Each perturbation estimated at 5600 cores based on current FP
Eval	Nature Run must be: • Capable of producing the complexities needed for observation simulation • Validated and verifiable – re-integration requires bit-wise reproducibility to ensure that only one-validation-per-nature-run is required	

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Production of new Nature Runs
  - Updates of baseline OSSE system with concomitant calibration and ensemble controls
  - Needs of instrument teams, proposal cycles
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Future Nature Runs require large number of computing cores to simulate observations
  - Impact: Given current queuing structure it may have a large impact on time-to-solution. A public cloud solution may be applicable
  - External Sources: None
- How will this Use Case evolve over the next decade?
  - Maturity: Current capability is mature but requires frequent updating. Future transition to new Nature Run will require substantial modification of observation simulation infrastructure.
  - Growth: Both spatial and temporal resolutions will increase in tandem
  - Re-use: This capability should extensible and highly re-usable, hopefully engaging new space missions. Much of the workflow is not limited explicitly to “Meteorological OSSEs”
  - Re-hosting: This will be addressed within GMAO’s general strategy for evolving the GEOS modeling system.

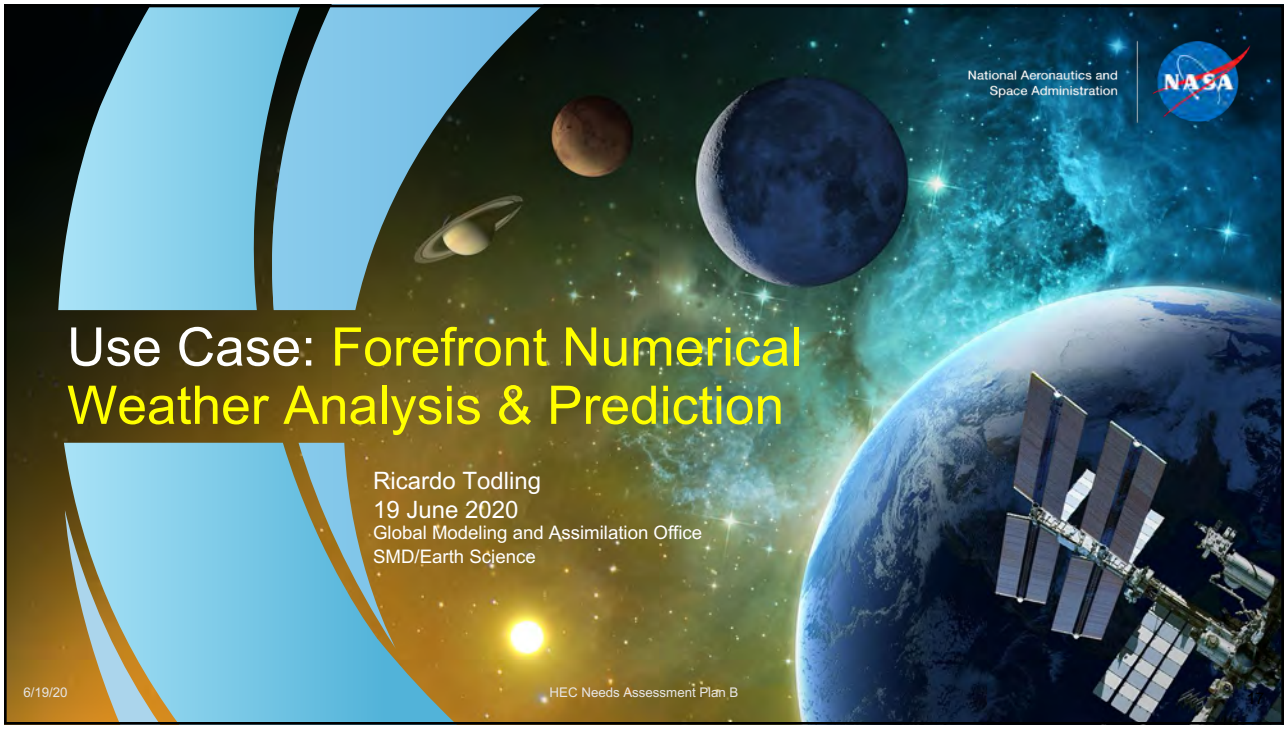

GEOS Forward Processing: Mission Need

- Mission: Cycling a high-resolution near-real-time Advanced Data Assimilation Weather Analysis & Prediction System that supports NASA Instrument Teams, Field Campaigns, and Scientists to plan, develop, and assess the contribution from new and under-utilized Earth Observing Systems for weather and climate.
  - The impact of GEOS-FP can be assessed by the number of costumers it supports. Data is typically freely available and is accessed not only by NASA, but also by others, e.g., NOAA, JCSDA, University Researchers, etc (over 120 subscribers).
- Criticality and Impact of HEC: GEOS-FP is a highly intense computer application requiring thousands of cores on HEC, and heading toward needing hundreds of thousands of cores as we enter the era of Exascale Computing.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Turning over to partners is unrealistic as most have no access to HEC capabilities; buying own Computing system is out of question; renting adds to uncertainties related to cost fluctuation; use Advanced Cloud Computing *might* be a possibility in the future.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - GEOS-FP products are made available in near-real-time, rough Quality Assurance checks are performed before delivery, pre-operational tests validate GEOS-FP and typically ensuring the overall Integrity and Reliability of its products.
  - The stability of HEC is essential for generating products that costumers can rely on a regular basis; glitches and interruptions in HEC cause delays in product delivery, add to customer frustration and possible disinterest.
  - Data dissemination is done through a *reliable* Data Portal and OpenDAP; where these sit and ownership are unimportant.
- HEC Function: GEOS-FP runs four-times daily, launching multiply jobs, some taking a few minutes, a few taking 1-4 hours.
  - Output from GEOS-FP supports multiple NASA Instrument Teams and real-time Missions; output available as generated.

GEOS Forward Processing: Mission Concept

Present Mission: Hybrid 4DEnVar

- Broad Workflow View of FP.
- FP runs two concurrent DA systems:
  - Hi-Res Deterministic Hybrid DA;
  - Low-Res Ensemble DA.
- Boxes are main separate executables.
- Arrow represent file data exchange.
- Boxes show number of cores used.
- Boxes show resolution (all 72 levels).
- Arrow texts show data size.
- All numbers are per cycle of DA.

Near Future (<1.5yr) Mission:

- Approximately double vertical levels.
- Deterministic DAS max cores: 7500.
- Ensemble DAS max cores: 4500.
- 430 nodes on Haswell; 300 on Skylake.
- Increase in assimilated (all-sky) radiances.
- BKG and IAU double in size.
- Size of ICs depend on aerosol configuration; at least a factor of 1.5 increase.
- Size of BCs and Products remain largely unchanged.

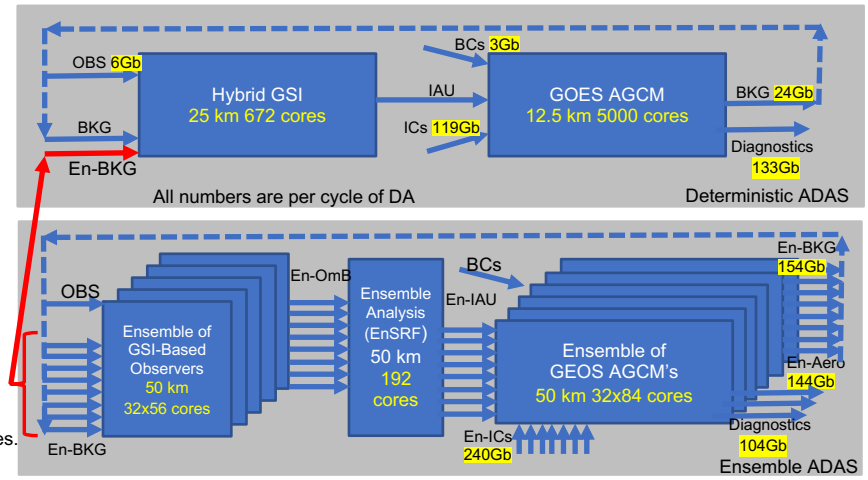

BKG: backgrounds
ICs: model restarts
OBS: observations
IAU: assimilation restart
BCs: boundary conditions

GEOS-FP: Current HEC Characteristics

- Work flow:
  - Near-real-time arrival of Earth System Observations and required Boundary Conditions; these trigger preprocessing and quality control steps each time (and while) GEOS-FP runs.
  - Each cycle of GEOS-FP entails a high-resolution, 12.5 km, Hybrid 4DEnVar application that runs concurrent with a supporting 50 km 32-member Ensemble DA System, both with 72 vertical levels.
  - Typically all jobs are automated and controlled by a scheduler (D-BOSS); glitch on HEC and input data require intervention by operators; operators work round-the-clock to ensure timely delivery of GEOS-FP products.
- Data:
  - In: Sources GES-DISC, EOSDIS, MODAPS, AERONET, NOAA, and JPL. Amounting to 200 Gb daily, with 10-hour latency.
  - Out: GEOS-FP produces roughly 1.0Tb of data per cycle (4 cycles daily)
    - About 10-15% of data generated is release as product; the other are data needed for diagnostics and cycling.
    - Up to 30 days of output is retained on live disc for diagnostic evaluation.
    - Product evaluation is made available through GMAO FLUID Application, HTTPS, and OpenDAP.
  - Transfer:
    - HEC permanent storage capability; NCCS Data Portal; Goddard DAAC.
- Computation:
  - Heaviest reliance is on MPI Fortran/C++ programs requiring 5000 Cores & roughly 1 Tb of scratch area per cycle (4 cycles daily).
  - Further relies on Python, Perl, Shell Scripts, IDL, Grads, and to some extent Matlab.
  - Presently GEOS-FP runs under special agreement with NCCS and relies on reserved resources.
- Evaluation: NCCS HEC has been extremely helpful supporting GEOS-FP (and other GMAO operational applications)

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Shared disc access is a major hog to operational systems such as GEOS-FP; a dedicated disc env is a must.
- Are there key functional gaps existing or projected for the next 5 years to 10 years?
  - Present:
    - Constraints on GEOS-FP: Availability of disc space has shied GEOS-FP from providing costumers with full resolution output (GEOS-FP runs at 12.5 km, but output is written at 24 km). Sporadic HEC hiccups cause near-real-time to fall behind at times.
    - Impact: Customers prevented access to high fidelity output; some delay to product assessment and delivery.
  - Future:
    - Constraints on possible upgrades to GEOS-FP: Resources (cores and disc) to fit expected increase in resolution and ensemble size.
    - Impact: Quality of product might not be as good as can be if resources are not available.
  - External Sources: the bulk of the Input Data comes from outside sources (NOAA, JPL, others).
- How will this GEOS-FP evolve over the next decade?
  - Maturity: GEOS-FP is already (near-real-time) operational and is to continuously evolve; expected to become JEDI-based.
  - Growth: Increases are expected in horizontal and vertical resolutions in both Hybrid and Ensemble components, including size of ensemble.
  - Re-use: GEOS-FP is accompanied by a parallel (pre-implementation system, FPP) running in delay mode; *two such systems run at all times*.
  - Re-hosting: Complexity of code migration to future architecture depends on how underlying technology evolves. Containers might provide temporary help as they allow keeping status quo, but they do not facilitate exploitation of possibly new features made available by new tech. JEDI is relying on Atlas to shield code from architectural changes.

GEOS Forward Processing: Future (5+ yrs) Mission Concept

- Full reliance of JEDI.
- In-core DA.
- Foresees horizontal res increase.
- To run Ensemble of Data Assimilations.
- Assumes increased number of members
- Forecasts might use Coupled Model.
- Workflow reliance on Cycl.
- Adaptability to arch changes thru Atlas.
- Some ability to run in the Cloud.
- Two such systems to run at all times.

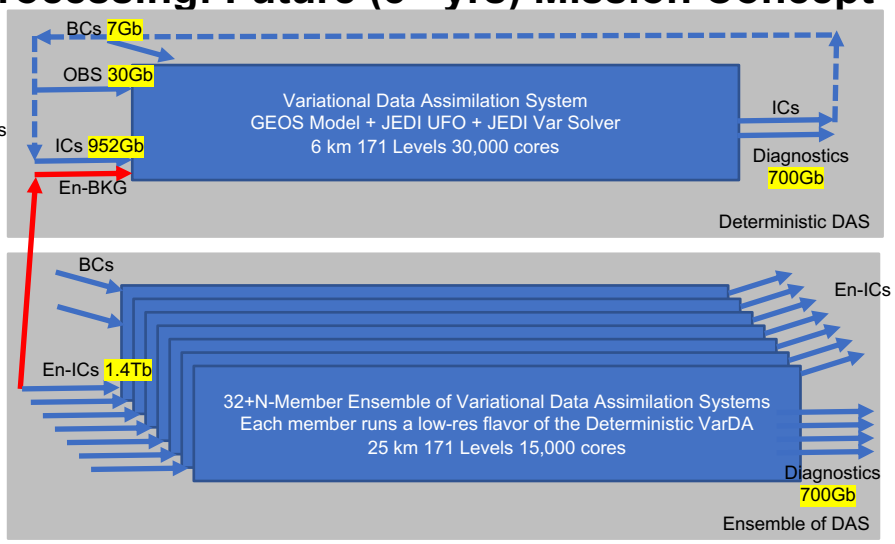

	Present	Future
Inputs	368 Gb	~ 2.4 Tb
Outputs	559 Gb	~ 1.4 Tb
Products	102 Gb	~ 1 Tb
Cores	7,688	~ 45,000
Throughput	2d/d	~ 2d/d*

* Operational constraints are such that the Future system must have throughput at least similar present.

National Aeronautics and Space Administration

Use Case: Coupled Analysis/Reanalysis: MERRA-3 and Beyond

R. Gelaro, E. Hackert, B. Putman, R. Todling, A. El Akkraoui

19 June 2020

GSFC Global Modeling and Assimilation Office, Code 610.1

Science Mission Directorate

Coupled Analysis/Reanalysis: MERRA-3 and beyond

Mission Need (1of 2)

- **Mission Objective**:
  - Develop a coupled Earth system data assimilation capability to explore the combined value of NASA observations in air, land, ocean, ice
  - Produce comprehensive Earth system reanalyses to place these observations in a climate context, enabling a broad range of research and applications
- **Use Case Concept:** ➤

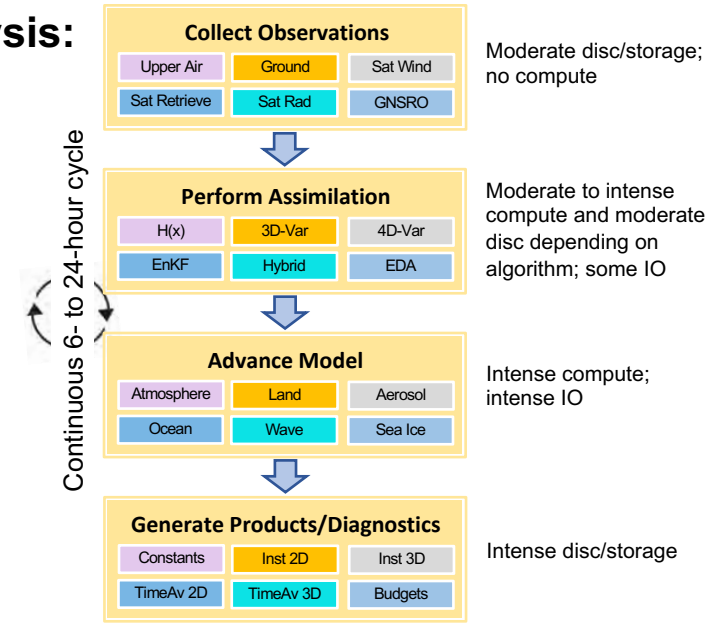

Coupled Analysis/Reanalysis: MERRA-3 and beyond Mission Need (2 of 2)

- **Criticality**: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - HEC resources are essential for reanalysis production and prerequisite testing. Without these capabilities it would be impossible to meet our deliverable of a realistic GEOS-based coupled reanalysis. Commercial cloud resources as a supplement, but cost would likely be prohibitive for long-duration production.
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - The inadvertent release of science or testing data does not pose any major risk. However, some level of confidentially is warranted while the new coupled data assimilation system is still in development.
  - Continuous availability of HEC resources is important during development and critical during production, as continuous delivery of reanalysis products is expected.
  - Output products are intended to be accessed by a broad range of internal and external collaborators, as well as users world-wide. Their feedback provides essential quality assurance and directly informs our efforts to meet users' needs.
  - **HEC Function**:
  - A reanalysis runs in round-the-clock, continuous data assimilation cycles, and typically in multiple streams simultaneously. Because these streams run behind real-time, there is no need (and considerable detriment) to stop the cycle.
  - The output represents a synthesis of NASA modeling and observations required to demonstrate and explore the value of these assets for understanding the Earth-system and advance national capabilities in Earth-system analysis and prediction.
- **Impact of HEC**: See "Criticality" above.

Current/Future HEC Application Characteristics

- Current Work-flow (uncoupled, out-of-core DA)
  - Observations for an individual component (usually atmosphere) are assembled and stored in separate formats, with non-generic data-handling and QC codes
  - Forecast and assimilation steps for single-component model are performed by separate executables, with significant IO and moderate-intense compute depending on DA solver
  - Product generation by a single-component model requires significant data transfer and storage
- Future Work-flow (coupled, in-core DA)
  - Observations for multiple components (atmosphere, ocean, ...) are assembled and stored in a common format, with generic data handling and QC
  - Forecast and assimilation steps for a multi-component model are performed in-core by a single executable, minimizing IO but will require advanced memory utilization.
  - Product generation by a multi-component model will significantly increase data transfer and storage needs

Current/Future HEC Application Characteristics

	Current	Future
Data	**GEOS 21RC Atmospheric Reanalysis** 20-yr C360/L72 Atmosphere; Constituent Replay **In:** • Restarts/Backgrounds: **190 TB/yr** (10-d latency) • Observations: **10 TB/yr** (10-d latency) **Out:** • Products: **90 TB/yr** • Total: **340 TB/yr** • Statistical analysis and visualization of products (dowmload) **Transfer:** • Products served at NASA GES DISC • Development output at NCCS; **tiered data management plan**	**MERRA-3 Coupled Reanalysis** 30-yr C720/L91 Atmosphere; 12km/L72 Ocean **In:** • Restarts/Backgrounds: **830 TB/yr** (10-d latency) • Observations: **20 TB/yr** (10-d latency) **Out:** • Products: **360 TB/yr** • Total: **1500 TB/yr** • Statistical analysis and visualization of products (in place) **Transfer:** • Products served at NASA GES DISC • Development output at NCCS; **tiered data management plan**
Compute	• Job Cores: **1,000** • Throughput: **1 analyzed decade / wall-time year** • Total: **1000 core for 2 years**	• Cores: **10,000** • Throughput: **1 analyzed decade / wall-time year** • Total: **10,000 cores for 3 years**
Eval	HEC Effectiveness: Near-continuous availability of required HEC resources is paramount for reanalysis, to ensure that (1) long-duration historical streams are completed within projected time frame, and (2) near-real time stream meets monthly product delivery targets	

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Reanalysis production is given high priority. However, reanalysis development is impacted by other production jobs (GEOS-FP, -CF), mission support, and GEOS development cycles.
  - AGU, AMS meetings, intermittent climate assessment exercises, can have intermittent impact.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Coupled DA will be one of the most intensive compute applications of GEOS systems. In-core DA will require advanced memory utilization compared with current applications.
  - Impact: Time-to-solution, while not as critical as for real-time applications, will significantly impact scope.
  - External Sources: None.
- How will this Use Case evolve over the next decade?
  - Maturity: New reanalysis development and production occurs continually, with a 5- to 10-yr replacement cycle.
  - Growth: Resolution will increase but lag that of uncoupled systems (GEOS-FP); this is not the main cost driver.
  - Re-use: Coupled DA will gain wider use for other applications (S2S, NWP), but long reanalyses will remain limited in number at any given time.
  - Re-hosting: This will be addressed within GMAO's general strategy for evolving the GEOS modeling system.

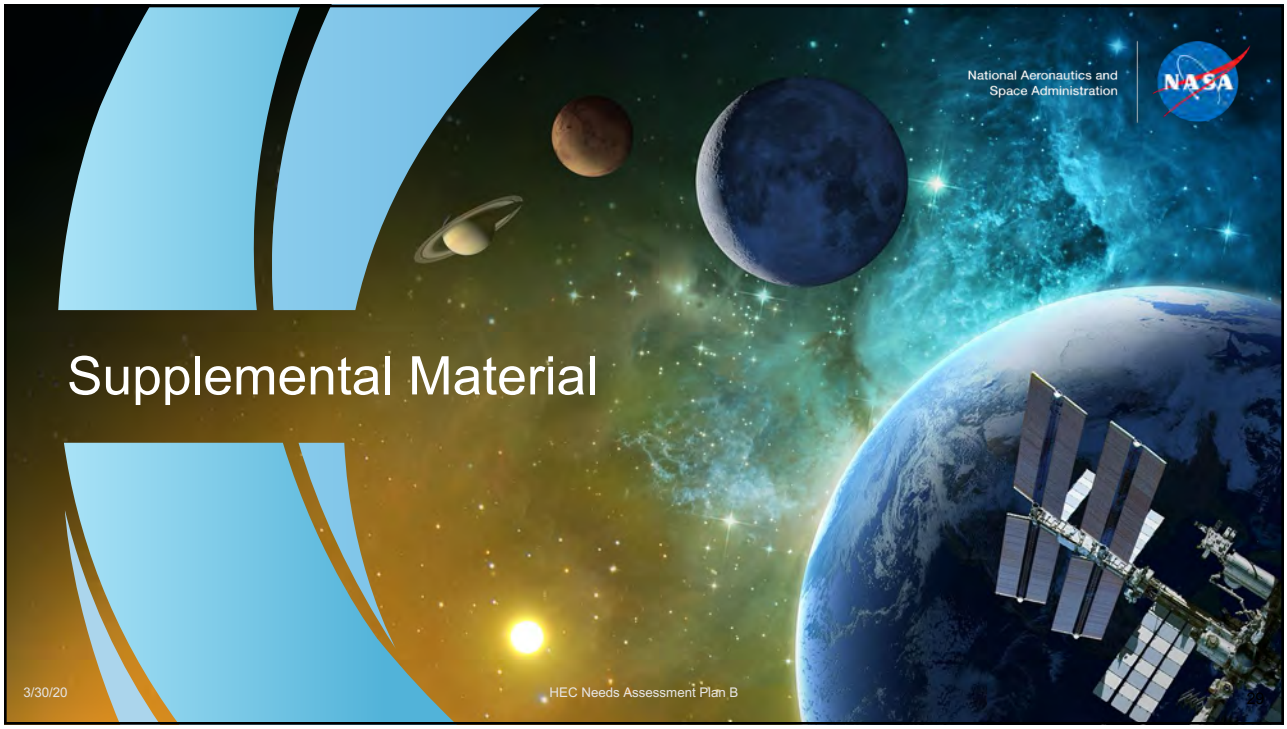

GEOS R21C Product Data Table

Data volume listed in GigaBytes(GB) unless otherwise noted

GEOS-R21C Central c360 72L	Data Volume per 00z Cycle	Data Volume per 06z Cycle	Data Volume per 12z Cycle	Data Volume per 18z Cycle	Data Volume per Day	Data Volume per Month	Data Volume per Year	Production (TB/Month) (< 1 month)		Tier-1 Storage (TB/6-Months) (1-6 months)		Tier-2 Storage (TB/Year) (> 6 months)	
								Destination	Fraction	Destination	Fraction	Destination	Fraction
Central Restarts	35	35	35	35	141	4230	50760	Tier-1	100%	Tier-2	25%	Tier-3	8%
Model Diagnostics	5	5	5	5	20	600	7200	Tier-1	100%	DAAC	0%	DAAC	0%
Analysis	18	18	18	18	70	2100	25200	Tier-1	100%	Tier-2	25%	Tier-3	8%
Aero-Chem	3	3	3	3	12	360	4320	Tier-1	100%	DAAC	0%	DAAC	0%
Observations	7	7	7	7	28	840	10080	Tier-1	100%	DAAC	0%	DAAC	0%
Forecast Products	136	19	74	19	248	7425	89100	Tier-1	100%	Tier-2	47%	Tier-3	2%
TOTAL (TB)	0.2	0.1	0.1	0.1	0.5	15.2	182.3		15.2		29.7		7.8
Analysis period [2000-2024]	Total Product Storage [2000-2024] (PB): 4.5								Reduced storage requirements [2000-2024] (PB): 0.107				

Data volume listed in GigaBytes(GB) unless otherwise noted

GEOS-R21C Ensembles (32 members) c90 72L	Data Volume per 00z Cycle	Data Volume per 06z Cycle	Data Volume per 12z Cycle	Data Volume per 18z Cycle	Data Volume per Day	Data Volume per Month	Data Volume per Year	Production (TB/Month) (< 1 month)		Tier-1 Storage (TB/6-Months) (1-6 months)		Tier-2 Storage (TB/Year) (> 6 months)	
								Destination	Fraction	Destination	Fraction	Destination	Fraction
Restarts	50	50	50	50	201	6030	72360	Tier-1	100%	Tier-2	25%	Tier-3	8%
Met-Background	35	35	35	35	141	4230	50760	Tier-1	100%	Tier-2	50%	Tier-3	8%
Aero-Background	10	10	10	10	40	1200	14400	Tier-1	100%	Tier-2	50%	Tier-3	8%
Diagnostics	4	4	4	4	16	480	5760	Tier-1	100%	Tier-2	50%	Tier-3	8%
Analysis	3	3	3	3	10	300	3600	Tier-1	100%	Tier-2	100%	Tier-3	8%
Ensemble Stats	5	5	5	5	20	600	7200	Tier-1	100%	Tier-2	100%	Tier-3	100%
Observation Stats	5	5	5	5	20	600	7200	Tier-1	100%	Tier-2	100%	Tier-3	100%
TOTAL (TB)	0.1	0.1	0.1	0.1	0.4	13.1	157.5		13.1		34.9		25.5
Analysis period [2000-2024]	Total Product Storage [2000-2024] (PB): 3.8								Reduced storage requirements [2000-2024] (PB): 0.224				

3/30/20 HEC Needs Assessment Plan B 30

MERRA-3 Product Data Table

Data volume listed in GigaBytes(GB) unless otherwise noted

MERRA-3 Central c720 91L	Data Volume per 00z Cycle	Data Volume per 06z Cycle	Data Volume per 12z Cycle	Data Volume per 18z Cycle	Data Volume per Day	Data Volume per Month	Data Volume per Year	Production (TB/Month) (< 1 month)		Tier-1 Storage (TB/6-Months) (1-6 months)		Tier-2 Storage (TB/Year) (> 6 months)	
								Destination	Fraction	Destination	Fraction	Destination	Fraction
Central Restarts	178	178	178	178	713	21385	256620	Tier-1	100%	Tier-2	25%	Tier-3	8%
Model Diagnostics	80	80	80	80	320	9600	115200	Tier-1	100%	DAAC	0%	DAAC	0%
Analysis	70	70	70	70	280	8400	100800	Tier-1	100%	Tier-2	25%	Tier-3	8%
Aero-Chem	12	12	12	12	48	1440	17280	Tier-1	100%	DAAC	0%	DAAC	0%
Observations	7	7	7	7	28	840	10080	Tier-1	100%	DAAC	0%	DAAC	0%
Forecast Products	545	74	297	74	990	29700	356400	Tier-1	100%	Tier-2	47%	Tier-3	2%
Ocean Obs+Diagnostics	64	64	64	64	256	7680	92160	Tier-1	100%	Tier-2	25%	Tier-3	8%
TOTAL (TB)	0.9	0.5	0.7	0.5	2.6	77.2	926.3		69.7		125.4		35.7
Analysis period [1980-2024]					Total Product Storage [1980-2024] (PB):		39.8				Reduced storage requirements [1980-2024] (PB):		0.488

Data volume listed in GigaBytes(GB) unless otherwise noted

MERRA-3 Ensembles (32 members) c180 91L	Data Volume per 00z Cycle	Data Volume per 06z Cycle	Data Volume per 12z Cycle	Data Volume per 18z Cycle	Data Volume per Day	Data Volume per Month	Data Volume per Year	Production (TB/Month) (< 1 month)		Tier-1 Storage (TB/6-Months) (1-6 months)		Tier-2 Storage (TB/Year) (> 6 months)	
								Destination	Fraction	Destination	Fraction	Destination	Fraction
Restarts	201	201	201	201	804	24120	289440	Tier-1	100%	Tier-2	25%	Tier-3	8%
Met-Background	141	141	141	141	564	16920	203040	Tier-1	100%	Tier-2	50%	Tier-3	8%
Aero-Background	40	40	40	40	160	4800	57600	Tier-1	100%	Tier-2	50%	Tier-3	8%
Diagnostics	16	16	16	16	64	1920	23040	Tier-1	100%	Tier-2	50%	Tier-3	8%
Analysis	10	10	10	10	40	1200	14400	Tier-1	100%	Tier-2	100%	Tier-3	8%
Ensemble Stats	5	5	5	5	20	600	7200	Tier-1	100%	Tier-2	100%	Tier-3	100%
Observation Stats	5	5	5	5	20	600	7200	Tier-1	100%	Tier-2	100%	Tier-3	100%
TOTAL (TB)	0.4	0.4	0.4	0.4	1.6	49.0	587.8		49.0		118.7		60.0
Analysis period [1980-2024]					Total Product Storage [1980-2024] (PB):		25.3				Reduced storage requirements [1980-2024] (PB):		0.527

National Aeronautics and Space Administration

Use Case: S2S2D Prediction and Predictability: Retrospective Forecast Suite

Andrea Molod, Eric Hackert, Yury Vikhliaev, Bin Zhao, Donifan Barahona, Anna Borovikov, Robin Kovach, Siegfried Schubert, Zhao Li, Young-Kwon Lim, Lauren Andrews, Richard Cullather, Randal Koster, James Carton, Kazumi Nakada

19 June 2020

GSFC Global Modeling and Assimilation Office, Code 610.1
Science Mission Directorate

Use Case: Mission Need (1 of 2)

- Mission:
  GMAO's development of sub/seasonal prediction systems is designed to demonstrate the beneficial impact on forecast predictability and prediction skill of the addition of **additional interactive earth system components** in the model, **increased resolution and forecast ensemble size**, and forecast **initialization from a modern assimilation** that incorporates new ocean data types. This activity relies heavily on the availability of computing resources.
- Criticality: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  HEC resources at NCCS and NAS are essential for timely integration of the coupled reanalysis and the suite of retrospective forecasts. Without these capabilities we would need to resort to **a limited number of ensemble members, rely on low-resolution forecasts, or lose near-real time capabilities**. Each of these choices would substantially affect the success of S2S activities by making estimates of predictability difficult, degrading prediction skill, and delaying system upgrades.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  Output files need to be accessed by external collaborators and other users, such as participants in multi-model activities to which GMAO contributes. Because of the massive dataset size, there will be a need for partners to execute algorithms on the same nodes where data resides. However, some confidentially is warranted while the system is still in development.
- HEC Function:
  HEC resources will be used to conduct a retrospective suite of sub/seasonal forecasts at high resolution, spanning the period 1981-present, consisting of an ensemble of 40 sub/seasonal forecasts per month. The forecasts will also continue into near-real time. A global weakly coupled reanalysis will also be run and used to initialize the forecasts.
- Impact of HEC: See "Criticality" above.

Use Case: Mission Need (2 of 2)

- Mission Concept: The new version of GMAO's seasonal forecast system includes higher ocean resolution and a larger ensemble size. The advantages of increased ocean resolution include more realistic bathymetry and surface currents. Our new ensemble strategy takes maximum advantage of computing resources and provides more accurate predictability estimates and increased forecast skill.

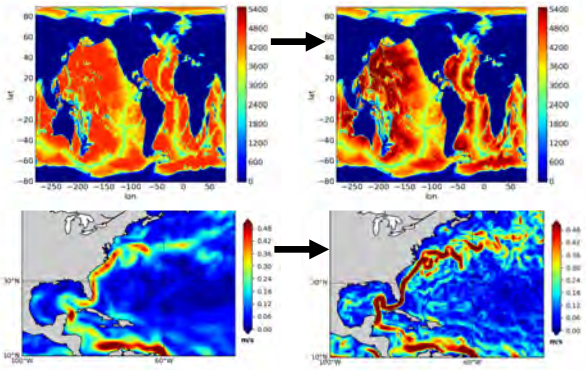

Top panels show ocean bathymetry (m) from GEOS-S2S-2 on left and GEOS-S2S-3 on right. High resolution brings a deeper ocean and more precise throughflows. Bottom panels show the improved surface currents near Brazil and the Gulf Stream.

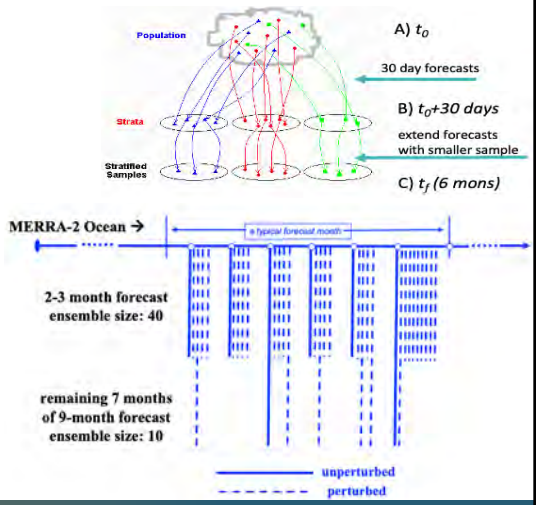

Use Case: Current/Future HEC Application Characteristics

- **Workflow**: Conduct both retrospective and near-real-time phases of the activity using resources at NCCS and NAS.
  - **Retrospective Phase**: Workflow begins with the coupled assimilation, as all input data are available. Initial conditions for forecasts may need to be transferred between NAS and NCCS, ensemble perturbations are calculated, and forecasts are run.
  - **Near-Real-Time Phase**: Input data for ocean analysis are downloaded, pre-processing of input data occurs, then assimilation is run. Ensemble perturbations are then calculated, and forecasts run. **Latency is up to a few days behind real time, as reanalysis depends on arrival of data and on multiple near-real-time procedures.**
  - Although scripts automate much of the process in both phases, there is frequently need for interaction by staff.
- **Data**: **50 km atmosphere, 25 km ocean horizontal resolution, 72 level atmosphere, 50 level ocean model**
  - **In**: See workflow description for input data flow during retrospective and near real time phases.
  - **Out**: Assimilation output totals ~0.2 PB, Forecast output totals ~1.5 PB for a 40-year retrospective record.
    - Data from assimilation will be retained on disk; data produced during near real-time forecasts will be retained at NAS.
  - **Transfer**: Data should be made available on data portals.
- **Computation**: Split workload between NCCS and NAS
  - **Retrospective**: 180 hours of wall clock time on 122 Skylake nodes for a single month of simulation, with the additional requirement that 8-10 forecasts are in the machine concurrently. **This phase requires continuous access to nodes for one year.** Data assimilation needs are ~100 nodes and 18 hours for each month.
  - **Near-Real-Time**: **Dedicated nodes are needed to assure access during forecast runs**, specifically the forecasts need 8-10 sets of 122 nodes for total of 72 hours each month, plus an additional ~100 nodes for 24 hours each month for post processing.
- **Evaluation**: Current computing is adequate for evaluation process.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - GEOS development cycles: atmosphere, coupled, chemistry reanalysis production, preparation for AGU, AMS and other highly attended meetings, if resources are not dedicated.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - **Constraints on Use Case**: Retrospective suite requires regular access to large number of cores to complete in a timely manner, near-real-time forecasts require time-sensitive access to large number of cores and to a stable and reliable file system for short periods as I/O requirements are substantial.
    - **Impact**: The current queuing structure could have a large impact on time-to-solution, particularly during the development phase, and on ability to complete the retrospective suite and near-real-time production. The current file systems at some centers lack necessary stability.
  - **External Sources:** Not applicable.
- How will this Use Case evolve over the next decade?
  - **Maturity:** The GEOS S2S Version 3 software is relatively mature, but **GEOS-S2S is under continuous development**. Each future version over the next decade will necessitate new retrospective suites.
  - **Growth: Increased horizontal and vertical resolution** in the atmosphere and ocean is expected. In addition, the expansion of the model is expected to include **more interactive components of the earth system** such as interactive vegetation and ice sheet models. In addition, increased use of new data types, such as sea ice thickness, during assimilation is expected.
  - **Re-use:** This capability should be extensible and highly re-usable.
  - **Re-hosting:** This will be addressed within GMAO's general strategy for GEOS modeling system evolution.

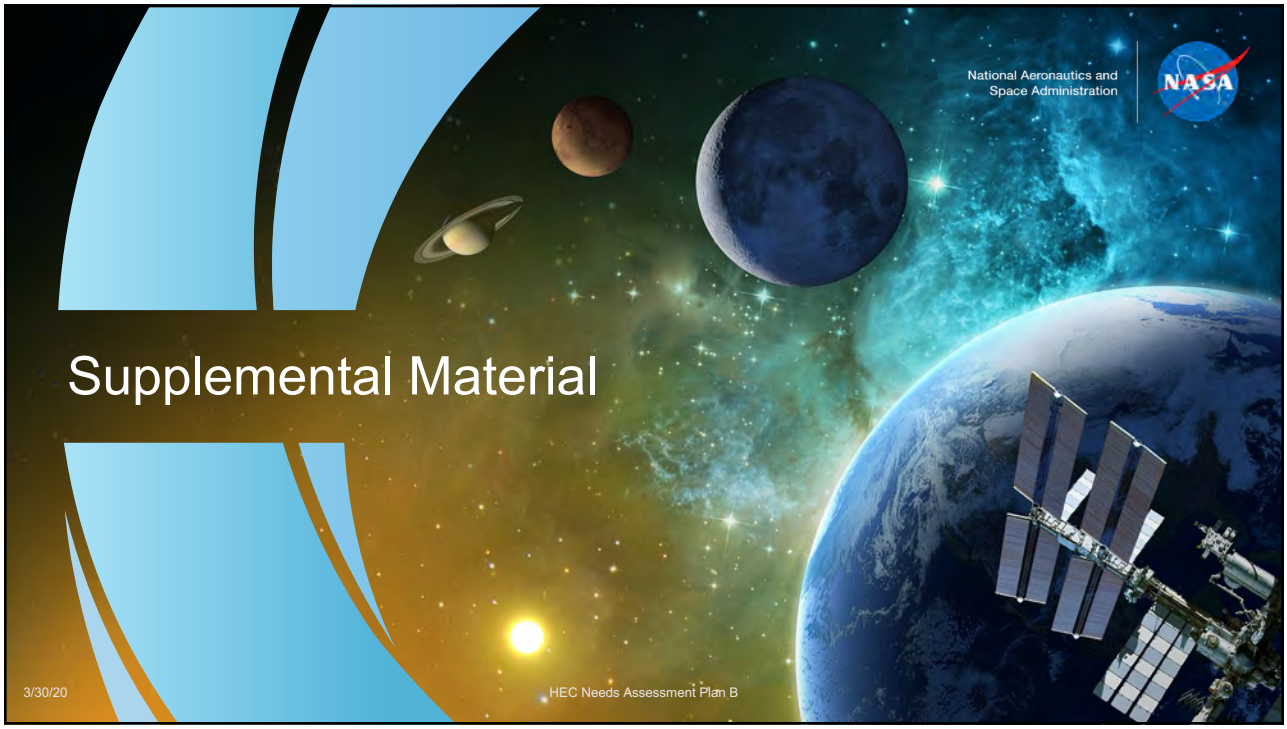

Information about the Survey

The Use Case described in these slides builds on the evolution of GMAO modeling capabilities, as as such needs to be planned in that context. Therefore, survey material will be provided separately by Steven Pawson together with other GMAO's HEC Needs.

GEOS-CF: Mission Need

- Mission: Produce global analyses and forecasts of atmospheric composition – focus is the use of NASA's space observations on constituent analysis/forecasting. Applications include global health/air quality assessments and support of NASA's in-situ field missions and basic research in to transport, emissions and composition
  - Without HEC these products would not be available to support NASA's mission
- Criticality:

  If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?

  This work is not viable without NASA HEC, or NASA acquisition of HEC resources (interagency or commercial cloud)

  Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Data are freely released to the community
  - Such applications depend on timely generation and release of the constituent analyses and forecasts
  - All systems are run by users with NASA credentials; end users with no credentials access data from public domain
- HEC Function:
  - Production system is run on a daily cycle, presently four analyses and one forecast per day (will increase)
  - Output are: monitored daily; used constantly by end users; used with prior days for scientific R&D studies
- Impact of HEC: NASA supercomputing is essential for this work. Agency mandated cloud may be suitable.

GEOS-CF: Mission Concept

- Present configuration:
  - C360L72 model
  - ~250 advected reactive species, including aerosols
  - Meteorology/transport: replay (to GEOS-FPIT)
  - Specified/computed emissions
  - Analysis: transport only
  - Four analyses and one forecast per day

- Future objectives:
  - Increase, in line with GEOS-FP
  - More complex aerosols; simpler reactive chemistry
  - Transport with GEOS-FP; feed back composition to FP
  - Local adjustment of emissions based on DA/obs
  - Implement multivariate constituent DA
  - Ensembles of analyses plus two forecasts per day

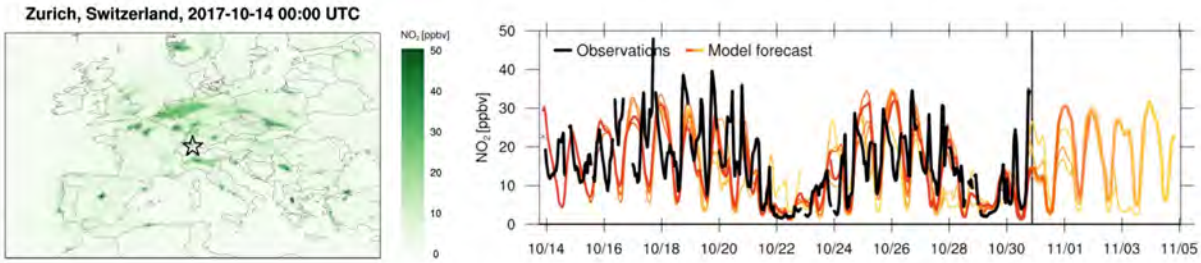

Nitrous oxide: CF map over Europe on 14 Oct, 2017 (left); time series of observations, CF analyses and forecasts over Zurich (right)

Use Case: Current HEC Application Characteristics

- Work flow:
  - Triggered by time of day (regular schedule)
  - Oversight by operators and scripted progress: human intervention if jobs fail.
  - Scenario includes NRT computing for timely delivery within window
- Data:
  - In: Meteorological analyses and emissions information [+ satellite data to be assimilated]
  - Out:
    - Output data need retention for science investigation and applications
    - Data analysis: pointwise comparisons, statistical analysis, emerging AI applications, visuals
  - Transfer:
    - Local HEC storage – a realistic archiving plan is needed for the longer term
    - Transfer to public-facing sites (temporary/weeks)
- Computation:
  - HEC allocation through assignments at NCCS – aligned to GMAO's production streams
    - Data analysis is using ML techniques on GPU – expect more of this
- Evaluation: the project has evolved in close cooperation with NCCS, without whom it would not have scceeded

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Large job footprint: GEOS model and GEOS-Chem chemistry are in themselves large, complex models
  - Regular production runs (daily) is steady while R&D runs require sporadic (but multi-day/week) additions
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: present HEC resources & (increasingly) data volumes bound the spatial resolution
  - Need to develop methods and tools for data analysis, including more ML based capacity (info content)
  - Impact: resolution and complexity will grow as resources grow
  - External Sources: This application runs on NASA HEC; R&D could run elsewhere (external partners)
- How will this Use Case evolve over the next decade?
  - Maturity: This is a production system; it will likely become increasingly coupled to GEOS-FP; it will be used for reanalysis applications in the (near) future
  - Growth: Resolution, complexity, and level of feedback to meteorological; systems will increase with time
  - Re-use: To accelerate development of capabilities, it is imperative that more users adopt this system
  - Re-hosting: DSL techniques, as in core GEOS modeling; additional ML on chemistry code and outputs

National Aeronautics and Space Administration

Use Case: Chemistry-Climate Modeling

Peter Colarco
June 19, 2020
NASA GSFC, Atmospheric Chemistry and Dynamics Laboratory, Code 614
SMD/Earth Sciences – MAP
Contributions from: K. E.Knowland, A. Molod, E. Sherman (GMAO), Q. Liang, M. Manyin, L. Oman (ACDL/614)

Use Case: Chemistry-Climate Modeling (1/3)

- **Mission Objective**
  - The CCM project develops and integrates atmospheric chemistry and aerosol components within the GEOS model that
    - inform international assessments (e.g., Montreal Protocol)
    - support the utility of NASA satellite measurements (e.g., scientific analysis, *a priori*)
    - advance research on coupled Earth system response to atmospheric chemistry on seasonal-to-decadal timescales
  - The CCM project recognizes a need to move toward higher resolution full stratospheric and tropospheric chemistry simulations to support recent (e.g., TROPOMI, GEMS) and upcoming (e.g., NASA's TEMPO) satellite missions
- **Use Case Concept**

NASA Data Informs Model Improvements

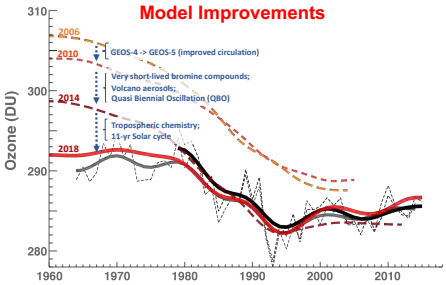

High Quality Modeling of Past Informs Projections

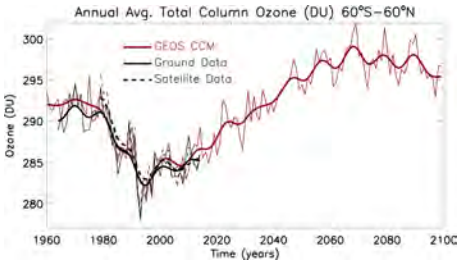

The Need for High Resolution to Properly Simulate Stratospheric O_3 Intrusions

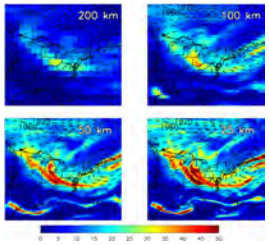

Use Case: Chemistry-Climate Modeling (2/3)

- **Mission Objective**
  - The CCM project develops and integrates atmospheric chemistry and aerosol components within the GEOS model that
    - inform international assessments (e.g., Montreal Protocol)
    - support the utility of NASA satellite measurements (e.g., scientific analysis, *a priori*)
    - advance research on coupled Earth system response to atmospheric chemistry on seasonal-to-decadal timescales
  - The CCM project recognizes a need to move toward higher resolution full stratospheric and tropospheric chemistry simulations to support recent (e.g., TROPOMI, GEMS) and upcoming (e.g., NASA's TEMPO) satellite missions
- **Use Case Concept**

Understanding the spatial and temporal variability of tropospheric ozone and its key precursors, CO and NO2

Data

CCM

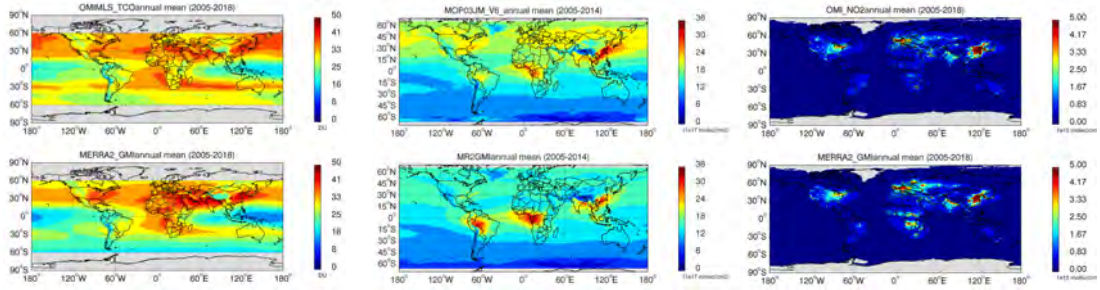

Use Case: Chemistry-Climate Modeling (3/3)

- **Criticality:** If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - HEC resources are essential for realistic chemistry simulations. Without these capabilities we would only be able to perform a limited number of cases and conditions, or to run only at low resolution, degrading the effectiveness of the activity.
- **Mission Security Needs:** (Confidentiality, Integrity, Availability)
  - CCM science and algorithms are not confidential as development and analysis is shared with the wider scientific community. That said, some level of confidentially is warranted while a system is still in development.
  - System availability is necessary to meet deadlines for assessments and scientific meetings, especially in cases where longer integrations are required, but our use case does not fundamentally impose real-time requirements.
  - Output needs to be accessible by external collaborators. Because of the massive size of datasets, there are benefits to allowing partners to perform some analysis on the same nodes where data resides. Options to share output by web services and OPeNDAP are necessary.
- **HEC Function:**
  - Contributions to assessments require multiple ensemble members of ~150 simulated year duration; more HEC resources would allow sufficient throughput for higher spatial resolution simulations.
  - Providing satellite *a priori* datasets and targeted scientific analyses which require multiple integrations of typically year- to decade-long simulations at high horizontal resolution with full stratospheric and tropospheric chemistry.
- **Impact of HEC**: See "Criticality" above.

Use Case: Current HEC Application Characteristics

- **Work flow**: Compilation, configuration, and execution of simulations carried out with customized GEOS model
  - R&D and science runs are "hands on" (i.e., expert user interaction needed for configuration and setup, testing parameterizations); Benchmarking and assessment runs are centralized to team POC tasked with execution
- **Data**:
  - In: Minimal input data requirements for boundary conditions or emission source functions, prescribed sea surface temperatures and sea ice; these are small datasets that vary typically at monthly to annual timescales
  - Out: Data is analyzed in comparison to observations and by visualization
    - Longer integrations (e.g., assessments) require saving and archiving monthly global gridded output; data volume can be high depending on number variables saved, prescribed by the assessment protocol
    - Shorter integrations will typically involve higher frequency, high resolution model output targeted toward evaluation data sets (e.g., satellite chemical constituents, airborne measurements) and need not necessarily be saved globally
  - Transfer: Most data is saved and analyzed at HEC center; some is retrieved to local servers (e.g., 614 "cluster")
- **Computation**: Current: C90L72 = 1 year/2 days O(1000); Future: C180L132 = 1 year/2 days O(5000)
  - Use case requires typical HPC resources (e.g., Fortran, MPI) and available compute capacity for multi-core calculations; ensemble members can be run in parallel; scalability is a consideration or else reduced resolution
- **Evaluation**: Stability of the compute environment (filesystems and uptime) available nodes to meet needed throughput are the prime drivers

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Generally steady cadence of development activities, but expect load to ramp up toward annual conferences (e.g., AGU, AMS); ramped up need to assessment activities (usually every four years, with 1-2 year lead up)
- Are there key functional gaps that exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Main constraint on use case advancement is scalability of compute environment; higher resolution and higher complexity simulations are the future need, but require ever more compute nodes and archive volume
  - Impact: Inability to get timely access to the required number of nodes drives toward a need to sacrifice model resolution and degrade the quality of the modeling products; lack of storage space for data archive (current and historical) weakens scientific utility of CCM output
  - External Sources: N/A
- How will this Use Case evolve over the next decade?
  - Maturity: New capabilities developed under competed ROSES proposals and integration into use case through core team support; couplings of system components (e.g., atmosphere-land; atmosphere-ocean) bring additional cost and complexity
  - Growth: There is a need to perform simulations at ever higher spatial resolutions, but see limitation above
  - Re-use: There has been a stable user base over the years
  - Re-hosting: Fundamental changes to compute architecture will require migration of code

National Aeronautics and Space Administration

Use Case: Observing System Simulation for Earth Science Space Missions

Arlindo da Silva
1 June 2020
GSFC Global Modeling and Assimilation Office, Code 610.1
Science Mission Directorate

NOTE: Included with all other GMAO Presentations

Observing System Simulations: Mission Need (1 of 2)

- **Mission Objective**:
  - Detailed simulation of global observing systems is critical por pre-launch algorithm development and assessment, and for conducting quantitative science trade studies during mission formulation.
  - Measuring impact of HEC: The maturity and readiness of algorithms at launch, prediction of performance after launch

- **Use Case Concept**

Nature Run: Earth System Model

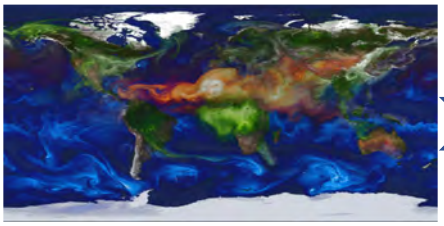

Sampled at Instrument Footprint

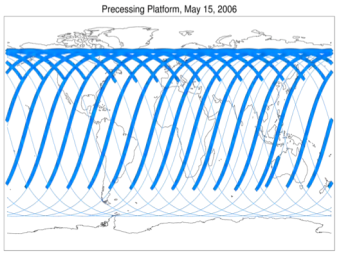

Detailed Instrument Simulator

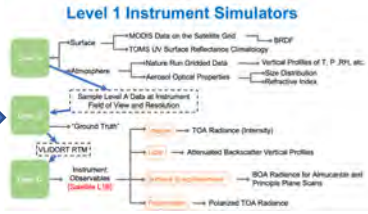

Observing System Simulations: Mission Need (1 of 2)

- **Criticality**: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  - HEC resources are essential for realistic observing system simulations. Without these capabilities one would need to resort to a limited number of cases and conditions, or to rely on low-resolution, low-fidelity simulations, all affecting the effectiveness of the activity.
- **Mission Security Needs**: (Confidentiality, Integrity, Availability)
  - While algorithm development and science trade studies are generally the domain of the science team, the inadvertent release of science data does not pose any major risk. However, some level of confidentially is warranted while a system is still in development.
  - Continuous availability of HEC resources is only marginally important as this is not a RT applications. However, capacity will be needed to perform high resolution simulations on demand
  - Output files will need to be accessed by external collaborators (assuming a diverse Science Team including US Universities and international participants). Because of massive size of datasets, there will be a need for partners to execute algorithms on the same nodes where data resides.
- **HEC Function**:
  - A global, very resolution Nature Run will be performed for a period of at least 1 year, with checkpoints at discrete intervals(very little I/O otherwise) – earth system model is the data compression algorithm.
  - For specific missions and applications (e.g., clouds/aerosols for ACCP), model will restart at key periods and produce carrying out in-line, as the model reruns, mission specific calculations
- **Impact of HEC**: See "Criticality" above.

Current/Future HEC Application Characteristics

- Current Work flow:
  - A very high-resolution Nature Run with detailed microphysics is performed for a period of at least one year, saving to disk full output for all anticipated uses – compromises are made, such as output frequency and variables
  - For a given Space Mission application, files are sampled at instrument footprint, saved to disk, and detailed instrument simulators are performed off-line
  - Nature Run Generation is a highly parallel application meant for batch processing
  - Instrument simulators are embarrassingly parallel applications, often I/O bound.
- Future Work flow:
  - Phase I: A very high-resolution Nature Run with detailed microphysics is performed for a period of at least one year, with frequent checkpoints, and limited output (mostly for browsing/case selection).
  - Phase II: For a given Space Mission application, "forecasts" are run from saved restarts, sampling model parameters at instrument footprint, performing detailed instrument simulation on-line, in parallel.
  - User provides instrument simulator plug-ins, prepares configuration files, submit jobs.
  - Both phases are highly parallel applications meant for batch processing.

Current/Future HEC Application Characteristics

	Current	Future
Data	7 km horizontal resolution, 72 level model **In:** • Nature Run: minimum input requirements O(Gb) • Observing System Simulations: ~ 500 Tb/year for instrument/orbit **Out:** • Nature Run: ~ 800 Tb/year • Observing System Simulations: e.g. PACE: 5 Tb/year for OCI L1 • Statistical analysis usually performed where the data resides. **Transfer:** • Output stays at NCCS data portal where it can be accessed via https:// or OPeNDAP	3 km horizontal resolution, 132 level model **In:** • Nature Run: minimum input requirements O(Gb) • Observing System Simulations: ~ 13 Tb/year **Out:** • Nature Run: ~ 13 Tb/year (weekly restarts) • Observing System Simulations: Depends on Observing System • Statistical analysis usually performed where the data resides. **Transfer:** • Output stays at NCCS data portal where it can be accessed via https:// or OPeNDAP
Computation	• GEOS-5 Nature Run (G5NR) was performed at NCCS discover • Several Observing System Simulations (e.g., PACE/ACCP) are being performed at NCCS because of "data locality"	• Simulations can be performed at any NASA HEC Center or in the pub;lic Cloud, provided some form of "data portal" is provided. • Both NR and Observing System Simulation need to be performed at a HEC center, but not necessary NASA's (input is small).
Eval	Nature Run must have resolution and microphysics complexity to simulate the Observing System at hand. HEC capability must be able to support simulations long enough to fully document impact of the new observing platform.	

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - GEOS development cycles: reanalysis & NR production, development
  - AGU, AMS and other highly attended meetings if resources are not dedicated
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Under the new paradigm, observing system simulations require regular access to large number of cores for production
  - Impact: Given current queuing structure it may have a large impact on time-to-solution. A public cloud solution may be applicable.
  - External Sources: N/A
- How will this Use Case evolve over the next decade?
  - Maturity: NR capability relatively mature, while in a state of continuously development. New paradigm for in-line instrument simulator needs definition of protocols and user training (plug-ins, etc.)
  - Growth: increase in resolution will lead to smaller time steps; with sampling being done on-line there is dependence on output frequency.
  - Re-use: this capability should extensible and highly re-usable, hopefully engaging new space missions
  - Re-hosting: This will be address within GMAO's general strategy for evolving the GEOS modeling system.

National Aeronautics and Space Administration

Use Case: GMAO Support to JCSDA

Ron Gelaro
GMAO
JCSDA Executive Team

17th JCSDA Technical Review Meeting and Science Workshop
May 29-31, 2019

GMAO Systems

System	Product Description	Distribution	Availability
GEOS FP	Real-time weather analysis and prediction using latest validated GEOS system, includes FSOI	NCCS Portals	Public
GEOS FP-IT	Real-time stable analysis product for NASA Instrument Teams, spanning 2002 onwards	GES-DISC	Subscription
New! GEOS CF	Once daily near real-time composition forecast, includes full chemistry	NCCS Portals	Public
GEOS S2S	Nine-month coupled AOLI forecasts with ocean DA and including SubX	Push from NCCS	To NMME By request
MERRA-2	Atmospheric reanalysis 1980-present, including meteorology, aerosols, ozone, sea ice	GES-DISC (with DOIs)	Public
M2-AMIP	10-member GEOS model ensemble using MERRA-2 version of GEOS GCM	NCCS	Public
7km-G5NR	7-km GEOS-5 Nature Run (produced in 2014) including aerosols and CO_2	NCCS Portals	Public
SMAP L4	Soil moisture, carbon fluxes as a part of the SMAP Mission (Level-4 products)	NSIDC (DAAC)	Public

National Aeronautics and Space Administration

GEOS CF 5-day composition forecast (from GMAO web page)

NASA GMAO – GEOS CF Forecast Initialized on 12z 05/06/2019
Surface O_3

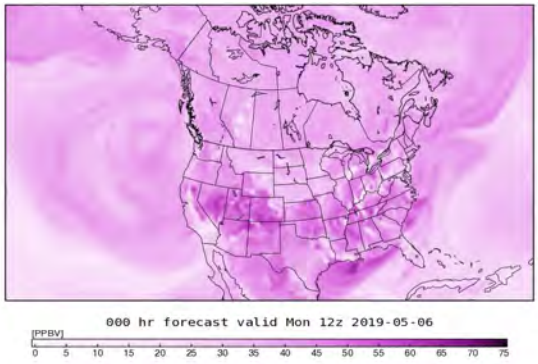

NASA GMAO – GEOS CF Forecast Initialized on 12z 05/06/2019
Surface NO_2

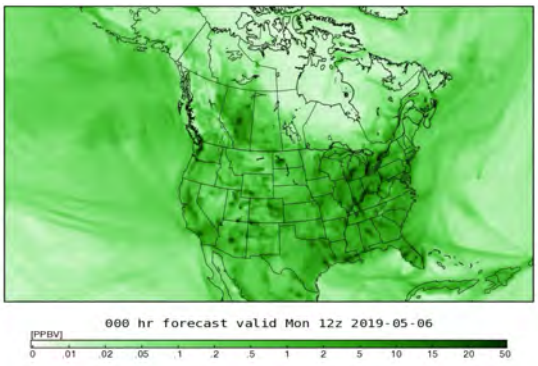

GEOS FP-IT meteorology + GEOS-Chem full chemistry ⇨ GEOS CF
- 1-day replay (DA to come)
- 5-day forecast
- 25-km resolution (c360)
- 250 reactive gases

National Aeronautics and Space Administration

GEOS FP current status

System configuration

- FV3 dycore + GEOS physics suite
- GSI hybrid 4D-Envar, 6-h cycle
  - GCM: C720 L72 (13 km)
  - GSI: 1152 × 721 L72 (25 km)
  - EnKF: C180 L72 × 32 (50 km)
- Incremental analysis update
- AO interface + skin SST analysis
- Aerosol assimilation
- FSOI at 00UTC

Key updates since last JCSDA workshop

- VLab Version-0 FV3 dycore
- RRTMG LW
- Retuned cloud ice/liquid effective radii
- Retuned topography and GWD

- GPM GMI all-sky radiances ✓
- NOAA-20 CrIS and ATMS
- SNPP OMPS total ozone
- Channel-correlated obs error for AIRS, IASI
- Variational scan-angle bias correction
- Retuned obs error for polar AMVs

National Aeronautics and Space Administration

FSOI for satellite radiances in GEOS FP

24-h Global Moist Energy Aug 2018 – Apr 2019

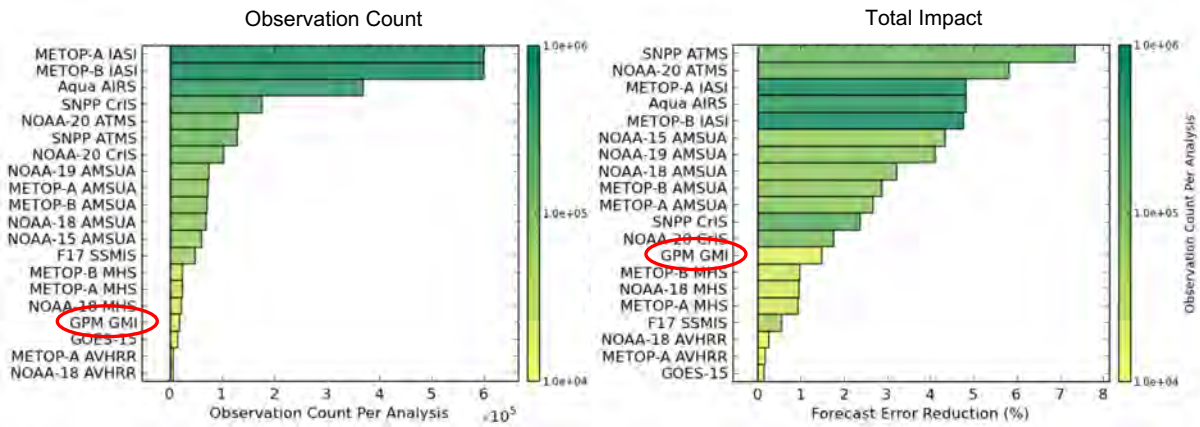

National Aeronautics and Space Administration

FSOI for satellite radiances in GEOS FP

24-h Global Moist Energy Aug 2018 – Apr 2019

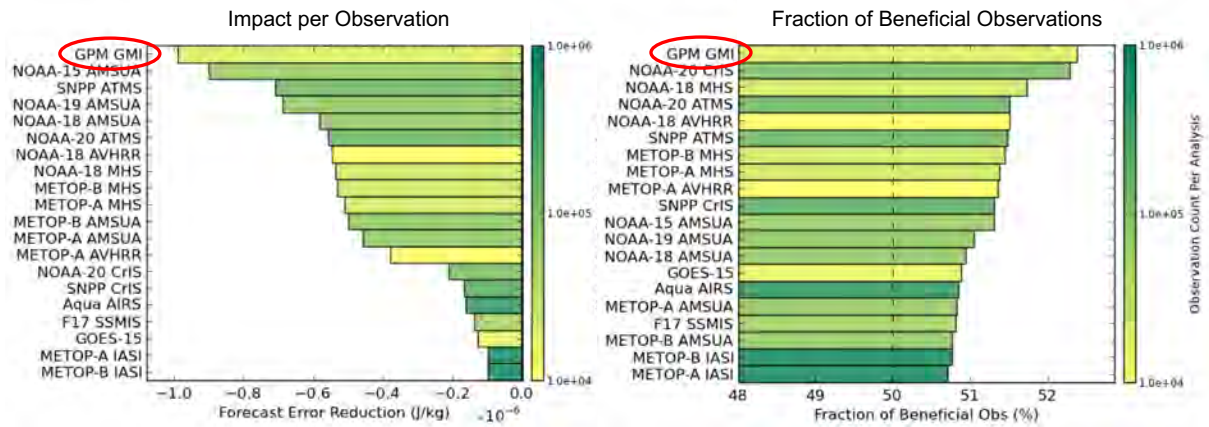

National Aeronautics and Space Administration

GEOS FP implementation plans for 2019/2020

2019 Q3

- VLab Version-1 FV3 dycore (as in GFS)
- RRTMG SW
- Grell-Freitas scale-aware deep convection
- UW shallow convection
- Retuned microphysics, PBL, turbulence
- Updated land model
- Coupled AO model-capable

- Stochastic physics in EnKF (SPPT)

2020 Q1

- All-sky MHS, AMSU-A
- Metop-C AMUS-A, MHS, ...
- AIRS AMVs
- Aeolus LOS wind (quality permitting)
- Channel-correlated obs error for CrIS FSR
- Revised thinning of GEO AMVs
- Aircraft bias correction

2020 Q3

- 132 vertical levels
- Unified PBL - EDMF or SHOC
- Unified GWD
- RRTMGp (high-performance)

National Aeronautics and Space Administration

Extending all-sky MW radiance assimilation to MHS and AMSU-A

***GMI, MHS** and **AMSU-A** observed brightness temperatures near Typhoon Maria on 2018-07-06 00UTC*

GMI Ch.10 (166 GHz)

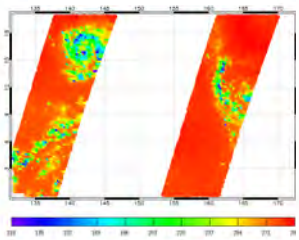

MHS Ch.2 (150 GHz)

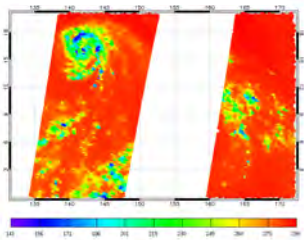

AMSU-A Ch.4 (53 GHz)

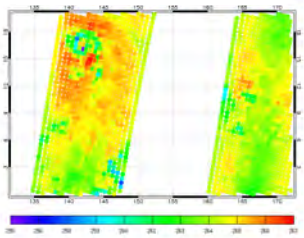

Like GMI, **MHS** has channels near the 183 GHz water vapor absorption band, but flies on more platforms and has a wider swath width (increases refresh rate and coverage).

AMSU-A has channels near the oxygen absorption band and is sensitive to temperature.

Min-Jeong Kim, GMAO

National Aeronautics and Space Administration

Impact of all-sky MHS and AMSU-A on forecast skill

Early cycling experiments

GEOS Hybrid 4D-Envar, 6h cycle

GCM: 25-km
GSI: 50-km
EnKF: 100-km

Control: full observing system (includes GMI)
MHS All-sky: Control + MHS all-sky
AMUSA All-sky: Control + AMSUA all-sky

Mostly neutral impact so far, but significant positive impact of all-sky MHS on low-level tropical humidity (95%)

Tuning and experimentation continues...

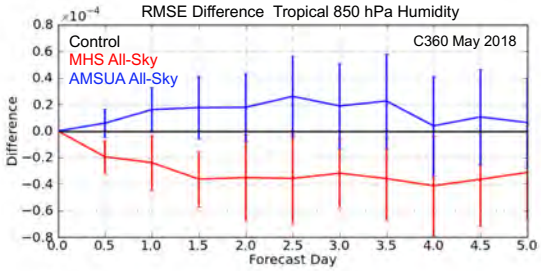

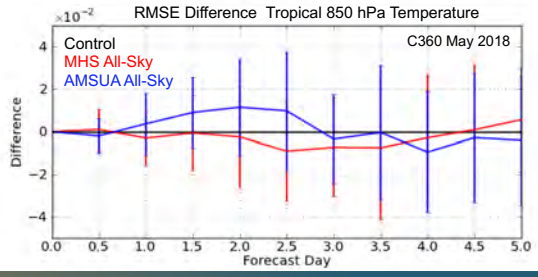

National Aeronautics and Space Administration

Moving toward a JEDI-based GEOS system for coupled DA

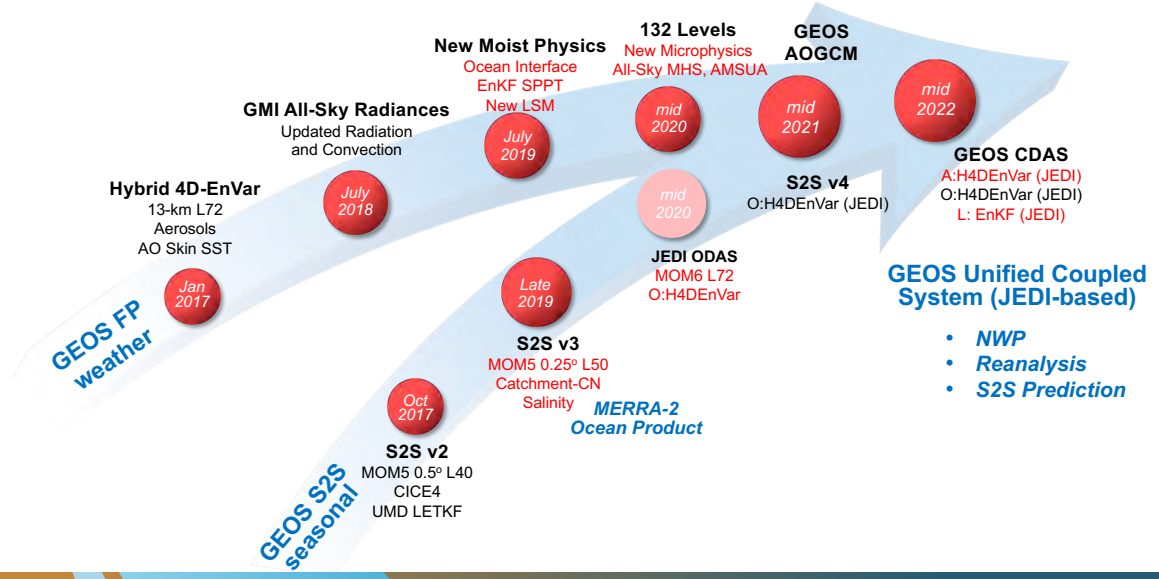

GEOS Unified Coupled System (JEDI-based)

- *NWP*
- *Reanalysis*
- *S2S Prediction*

National Aeronautics and Space Administration

NASA support of JCSDA for 2019

3.00 FTE Core Team
1.75 FTE In-kind ...

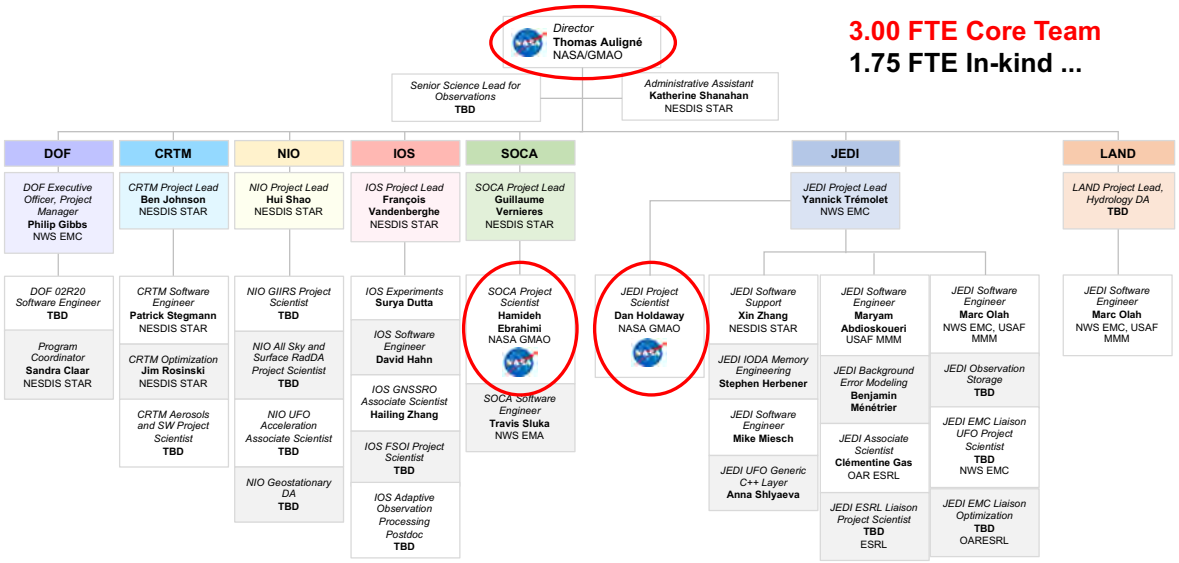

National Aeronautics and Space Administration

GMAO in-kind contributions to JCSDA for 2019

	DOF	CRTM	NIO	IOS	SOCA	JEDI	Total	Work Description
Akella, Santha					0.1		**0.1**	Unified Forward Operator (UFO) for marine DA
Gelaro, Ron	0.05						**0.05**	Executive Team
Karpowicz, Bryan						0.2	**0.2**	UFO support, QC filters
Kim, Min-Jeong			0.1				**0.1**	GPM/GMI in UFO
Mahajan, Rahul					0.3	0.3	**0.6**	Work flow for marine DA; SST UFO; coupled UFO
McCarty, Will		0.1	0.1				**0.2**	SW IR in CRTM; Aeolus in UFO
Moradi, Isaac		0.2					**0.2**	General CRTM development; coefficient generation
Todling, Ricardo						0.3	**0.3**	Adjoint DA solver
Total	**0.05**	**0.3**	**0.2**	**0**	**0.4**	**0.8**	**1.75**	

DOF-Director's Office; **CRTM**-Community Radiative Transfer Model; **NIO**-New and Improved Observations; **IOS**-Impacts of Observing Systems; **SOCA**-Sea-ice, Ocean, Coupled Assimilation; **JEDI**-Joint Effort for Data Assimilation Integration

National Aeronautics and Space Administration

One day of observations in JEDI Ocean 3D-Var

SST (IR)
NOAA-19, Metop-A AVHRR, SNPP VIIRS

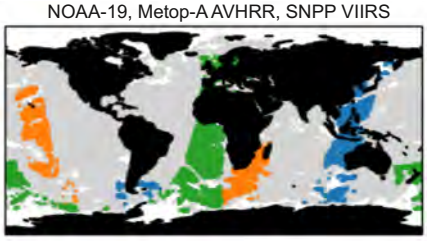

SST (MW)
GMI, AMSR2, WindSat

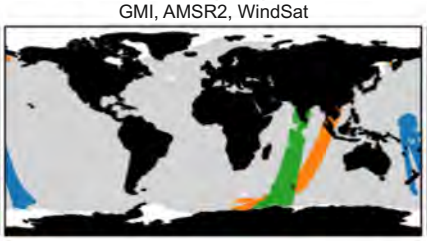

SSS
SMAP

Altimetry
Jason-2,-3, Sentinel-3a, Cryosat-2, SARAL

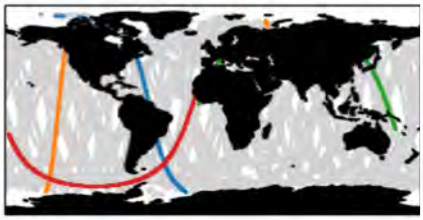

Insitu T/S
TAO, PIRATA, XBT, CTD, Ships, Drifters

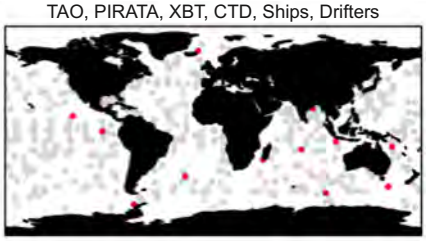

24-hr assimilation window on 15 April 2018

National Aeronautics and Space Administration

JEDI-based GEOS atmospheric 4D-EnVar tests

850 Temp Increment 2018-04-14 21UTC

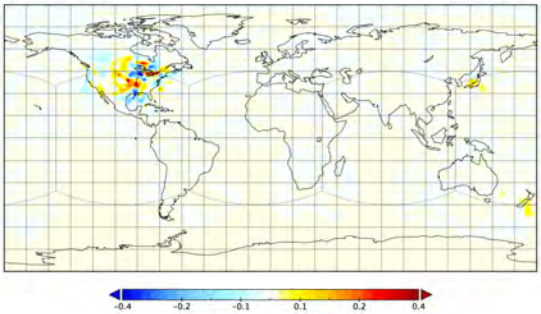

250 U-wind Increment 2018-04-14 21UTC

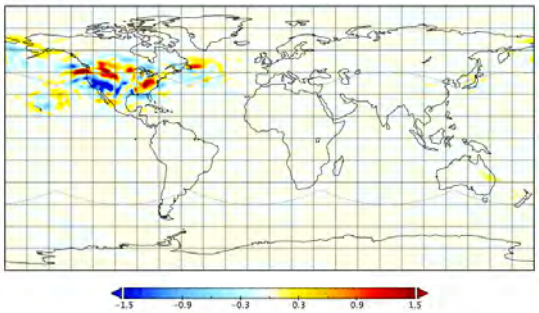

- 6-hr window centered 2018-04-15 00UTC
- GEOS C180 (50-km) 10-member ensemble
- Aircraft and radiosondes (~400K observations)
- 2 sub-windows of 3 hours
- BUMP localization, prescribed length scales
- Background check for QC

Dan Holdaway, JCSDA

HEC Needs Assessment Workshop

Affinity Group Sessions

Friday 6/19/20 1:00 p.m. – 4:00 p.m. (AG_GMAO) (nominal affinity: GMAO) **(10am – 1pm Pacific)**
Facilitator: Peter Williams <peter.williams6850@gmail.com>
Scribes: Jarrett Cohen <jarrett.s.cohen@nasa.gov> and Sean Keefe <rosemary.e.keefe@nasa.gov>

Attendees

1. Mike Little, Convener, NASA GSFC, ESTO and HEC Program
2. Peter Williams, Facilitator
3. Presenter #1: William Putman, GMAO
4. Presenter #2: Nikki Privé, GMAO
5. Presenter #3: Ricardo Todling, GMAO
6. Presenter #4: Ronald Gelaro, GMAO
7. Presenter #5: Andrea Molod, GMAO
8. Presenter #6: Steven Pawson, GMAO
9. Presenter #7: Peter Colarco, NASA GSFC Atmospheric Chemistry and Dynamics Laboratory
10. Aaron Pina, NASA HQ
11. Anton Darmenov, GMAO
12. Christoph Keller, GMAO
13. Dan Duffy, NCCS
14. David Considine, NASA HQ
15. Donifan Barahona, GMAO
16. Elizabeth Hook, NASA HQ
17. Ellen Salmon, NCCS
18. Eric Hackert, GMAO
19. Henry Jin, NAS
20. Jarrett Cohen, NASA GSFC, HEC Program
21. K. Emma Knowland, GMAO
22. Laura Carriere, NCCS
23. Lauren Andrews, GMAO
24. Lesley Ott, GMAO
25. Piyush Mehrotra, NAS
26. Sean Keefe, NCCS
27. Thomas Clune, GMAO
28. Tsengdar Lee, NASA HQ, HEC Program
29. William Thigpen, NAS

NOTE: Arlindo daSilva's Presentation and notes from discussion in Session 6_1 are duplicated here to put all GMAO-related materials in a single location.. A Presentation by Ron Gelaro to JCSDA in May, 2019 is also included in the slides to highlight GMAO strategic direction without discussion notes.

1. Use Case: Global Earth System Modeling with High Complexity and Resolution

Bill Putman, NASA Goddard
william.m.putman@nasa.gov
GSFC/Code 610
SMD

Raw Notes

Slide 1:
Allowing systems to adapt to trends. Push what we can do from c_. Systems tend to be 5-10 years ahead of system applications used. OSSE is the main format via Nature Runs. Arlindo presented on that in the last meeting. Can produce large amounts of data, used for AI training data. Criticality: HEC is a huge compute intensity problem major benefit to work with support staff closely. If you can't demo your models, it's hard to get through that process.

Confidentiality: no direct real-time support. Data generally shared to community through data portal to external users to access data we provide. Describe runs as simulations of opportunity.
Huge problem takes up a lot of compute in a short time, so when new platforms come in, these jobs can take up machine use and help identify any issues.

Slide 2:
Most recent application: project diamond, part of an ensemble atmosphere only simulation, no chemistry, just hydro and clouds, take 2 weeks to complete a simulation, producing 600 TB of data from one simulation.

Mission: These simulations tend to form the basis for a number of things, including future application development, informing observing system design through OSSEs, and conforming to HEC trends.
We push what we can do with current hardware (algorithm, workflow, etc.). These systems tend to be 5 to 10 years in advance of what our modeling and assimilation systems can do.
Criticality: Getting on DOE leadership systems is difficult; you have to demonstrate performance on them.
HEC Function: Run them during the break-in period.

Slide 3:
DYAMOND was a collaboration with other centers around the world.
Takes many more tries and attempts to get model running efficiently, testing new algorithms and more than __ can increase run complexity.

Slide 4:
NCCS footprint: 40k cores for 2 weeks: in 2021, estimate 32% consumption of NCCS for our runs. We tend to have issues when we start on the new hardware and try to push new resolutions and complexity. We like to test increasing science capabilities, new parameterizations, etc. Those can all add on the compute that these runs take up.

Slide 5:
[No additional notes: see slide content]

Slide 6:
Evolution of concept design: working on a coupled ocean-atmosphere model, produces 500 TB of simulation. Expands by doubling resolution every [5?] years. Extremely compute-intensive potential problem there. We rely on NCCS, would benefit from accelerator technologies to increase portability and efficiency. Observational-based ML for chemistry and physics, oriented on improving performance, physics accuracy. Future: data volume is too disruptive, New approach: containerize the OSSE to reproduce [?]. A run like this would take up all of the blue section of the pie chart.

Slide 7:
Workflow is a basic, user-submitted process. Data: small input vs output. Future: be able to evaluate, catalog, reproduce hi-resolution Nature Runs. Transfer: 10-110k cores today

Slide 8:
Key factors: Driving force to make new runs or model revisions or new hardware. Gaps: generally, capacity to execute large jobs and compute models.

The driving forces in when we make big runs are pushing resolution or increasing complexity.
We consume as much as we can and push those boundaries.

Q&A [with Bill Putman]:

Peter Williams: Any questions about security with your work? Some other presenters have mentioned concern about what has been called public confusion if work gets our early. Do you have any concern?

Bill Putman: No, it's not shared broadly. They are retrospective runs that tend to be shared with a small user group.

from Tsengdar Lee (Int) to Everyone: 1:36 PM
Is there a scalability issue in the current code?

Bill Putman: At this point, we don't think so. On traditional hardware, we want to take advantage of an OpenMP-Hybrid MPI approach; the dynamics would scale quite well. Accelerator hardware would help a lot.

Mike Little: Bill, can you talk about any additional help you might need to take advantage of with acceleration, code conversion, or does GMAO have that?

Bill Putman: The heart of this is being done with the dynamics, in collaboration with NOAA and Vulcan (externally funded). The challenge is finding application developers onsite who can help with that. NCCS hired a new staff member to start in July. They will dedicate themselves to doing that work. finding application developers is always a challenge. We need more people who can dedicate themselves to this type of work.

---

2. Use Case: Meteorological OSSEs, especially for decadal survey

Nikki Privé, TBD
nikki.prive@nasa.gov

Slide 1:
I will reference a lot of what Bill and Arlindo said. Team: Nikki Privé, Will McCarty, and Yanqiu Zhu.

Slide 2:
OSSEs: Mostly what we do is try to simulate new, proposed, satellite instrument data types in a simulation environment where we can do testing. Simulations based on Nature Runs, simulate a new proposed instrument type, add to mix, do experiments. Work funded by GOES in geostationary satellite IR satellites, test in control. Benefits: compare different, proposed instrument setups, evaluate how quickly we can use that data as it becomes available, ingest and see how it performs operationally.

We made a control simulation based on our Nature Run and simulated existing observation types. Then we simulated a proposed instrument type. It is helpful to compare different proposed instrument setups. We can ingest data and see how it performs in the operational and research systems.

Slide 3:
Criticality: extremely reliant on computing resources, without HEC, we would be stuck using highly simplified models, no realistic results. We do have security issues, because we work with a proprietary instrument information on new instrument proposals, bordering on ITAR, enough info. to be an issue if it was insecure.

We are not real-time, don't need continually-available resources, but we do need sufficient hi-resolution modeling resources, intermittently. Varying demand depends on funding cycles.

Most work done in house, some work with external collaborators, some university, NOAA, or other agencies. In future, collaborators may need to make some hi-resolution models themselves to test model simulators. In addition to Nature Runs, we need to run our experiments, 2-4 months of ensemble models. Sometimes we have a lot of demand to run a lot of runs at the same time (e.g., proposal periods).

Slide 4:
Current workflow: Test simulated observations with new observations, needs will be similar to forward processing model runs: always do control case with full 40-ensemble— experiments may not require a full ensemble. The system is updated regularly – the Nature Run is updated approximately every 10 years, requiring a complete overhaul. The experiment framework and simulated observations are updated approximately every 2-3 yearsyears.

Slide 5:
We will not be able to continue doing it as we have in the past. We will need to do smaller segments of the Nature Run on 40,000 cores to simulate the observations. We might have to repeat these segments several times. We need to maintain the ability to re-run the Nature Run and have bit-wise reproducibility. In the future, we might couple in an ocean.

We have had the system for more than 15 years, so it is fairly mature. We have to update it. As we transition to the new Nature Run regime, it will require substantial change to get that code set up in the new Nature Run. The update cycle frequency may increase once the simulated observations are integrated into the Nature Run, with more frequent updates to the Nature Run.

Slide 6:
In future, Nature Runs will be so large, I/O so intensive, that Nature Runs and observation simulations will be done differently. Actual observations datasets are not large, and calibration requires rerunning data over and over again, so we need to maintain the ability to run Nature Runs in a reproducible way for shorter times. Experimentation in future: simulations of workflow, coupled oceans.

Slide 7:
Current/Future characteristics. Most of our expense is IO. Future: only 500 TB for Nature Run outputs. 6 km or less, experiments still 3 months in length at twice the current resolution. Observation datasets may be larger if we use more data, larger hyperspectral datasets. Factor of 8 increase in future for higher resolution model runs. Require 40k cores available to do repeated observation simulationsobservation simulations.

Factors: Production of Nature Runs, overhaul our entire system now, same in future. More frequent updates toto baseline system to keep not-too-far behind real-time system, recalculate full ensemble. Key gaps: May be interested in a public cloud solution. Evolve: 15 year system is fairly mature system, frequently update to stay in line, as we transition to new NR regime, require substantial manpower to set up code in new Nature Run. Growth: Hoping we can keep a flexible, reusable capability. Current constraints are storage available for experiment output, and sometimes queue time for running when there is heavy demand on the system.

Q&A [with Nikki Privé]:

from Michael Little (Int) to Everyone: 1:44 PM
Nikki, can your runs be interrupted, or do they require continuous processing?

Nikki Privé: We are cycling over a period of a few months. At any point it can be halted if there is a problem. We don't need to have the resources carved out to run an entire experiment all at once.

from Tsengdar Lee (Int) to Everyone: 1:55 PM
It can be interrupted and restarted but it cannot be multi-threaded such as in reanalysis.

Nikki Prive [review addition]: To elaborate on Tsengdar's response, we cannot multi-thread an experiment, but we do run multiple experiments at the same time. Also, with the future Nature Run efforts, if we have multiple restarts, the observation simulation could be multi-threaded.

from Daniel Duffy (Int) to Everyone: 1:50 PM
Do the storage estimates include compression? If not, how much could be saved using compression?

from William Putman (Ext) to Everyone: 1:52 PM
The model numbers (including what is Nature Run here) do account for compression.

from Daniel Duffy (Int) to Everyone: 1:55 PM
How broadly do you want to distribute future Nature Runs? Only enable use within the GMAO or distribute the data externally?

Nikki Privé: There are politics involved in that. Could be distributed if NOAA, NRL, or other government agency wished to perform an OSSE and asked, "how we can use the Nature Run forfor our own OSSEs, are we able to get their own copy of the NR and run it?" Bill has suggested containerized cloud computing, and I don't see thisthis as realistic, otherwise. We are not set up to produce simulations for other groups. StevenSteven or Bill, would you like to speak to that?

Nikki Privé: On the HPC side, I am not entirely sure. Planning for the next Nature Run is at an early stage. Going forward, there is the JEDI framework that we would want to be staying in line with. I am not going to say that we can't distribute the NR, but I can't think of any other solutions off the top of my head.

from RICARDO TODLING (Int) to Everyone: 1:58 PM
With all moving to JEDI, I believe this could actually be done.

Peter Williams: Let's come back to that in general discussion.

---

3. Use Case: Forefront Numerical Weather Analysis and Prediction (GEOS-FP)

Ricardo Todling,
ricardo.todling@nasa.gov

Slide 1:
[No additional notes: see slide content]

Slide 2:

One of the ways to evaluate Impact is in how many costumers our products serve and how satisfied with our products these costumers are. Criticality: HEC is fundamental to support what we presently run; our systems can typically stress HEC availability to an extreme and we look forward to the days of Exa-computing..

Mission: It helps the development of new technologies. A version of it gets frozen at some point for reanalysis.

Criticality: We expect to use more and more cores. We hope to benefit from advanced technologies. I can hardly believe that NASA would rely on outside sources for providing those services (e.g., cloud computing).

Slide 3:
In terms of FP, runs multiple jobs, 4 cycles a day, 6 hours of observations, immediately released to the public. Within the boxes, I highlight the highest physical resolution and the largest number of cores. The ensembles, since they are "copies," run at lower resolution. I also indicate the amount of input and output (lines, arrows, and yellow-highlighted numbers).

Slide 4:
Mission concept: Total of 700 cores used, see slides for input, output, products. Near future: new data types, incr. # of levels (doubling), estimates of cores in slide.
The data that come into the model are the boundary conditions and the observations.

Slide 5:
Current HEC characteristics: 10-15% of data generated is released to DAACs. Diagnostics available via FLUID. Computation: FP system relies on NCCS, essential for performance and reliability for all we do, cloud computing can be tricky. Evaluation: excellent working with NCCS teams and systems.

Slide 6:
It will still run GEOS as the model but embedded within the JEDI infrastructure.
Use Cycl for scheduling; more modern.
Atlas developed by ECMWF; helps with adapting to architecture changes.

GEOSGEOS Forward Slide. We envision a system that relies on JEDI, staffed by the Joint Center for Satellite Data Assimilation (JCSDA). Embedded in a new structure, resolution increase, larger # of cores. Ensembles will change several levels of assimilation. Rely on advanced technologies, change scheduling using cycle or something more modern. JEDI is relying on atlas, facilitates adaptability to changes in architecture, make codes more easily adaptable to changes. JEDI is able to run in the cloud. (See table for increase in resources.)

Q&A [with Ricardo Todling]:

from Daniel Duffy (Int) to Everyone: 2:10 PM
We have talked about this in the past, but I am not sure I recall the answer. Which is more important to the GMAO's requirements: an increase in the number of ensemble members or higher resolution ensemble members?

Ricardo Todling: Tricky question, both would be great. We could benefit from resolution change presently, ideally both, if we reach a certain point with members, resolution has limits.

Steven Pawson: This is a good example of the type of question we need to answer by doing some research based simulations to get a quantitative answer for that.

from Daniel Duffy (Int) to Everyone: 2:10 PM
Also, what is the target overall throughput in simulation days per day?

Ricardo Todling: See the table: presently we are at the point where we run close to what allows us to have a near-real-time system, anything slower will be too late for anyone to want it. Whatever increase in resources, resolution, we must come up with a system that fits that purpose. We are at the limit, which is not ideal, but ok.

from Daniel Duffy (Int) to Everyone: 2:17 PM
Thanks! Would love to talk about how to support those experiments.

from Daniel Duffy (Int) to Everyone: 2:18 PM
Thank for the responses Ricardo. I am interested in a discussion about how we could increase the throughput of the cycled DAS to much greater than 2 days/day.

Bill Putman: It is not necessary to use more processors in the system to speed it up, rather find ways to improve the workflow of the system to improve performance, on the edge of real-time. The system can be faster.

Ricardo Todling: Ultimately, what will make a difference is the JEDI approach of in-core DA. A lot of what we spend time on now is reading and writing files. With everything being done in-core you avoid a lot of this headache.

from Daniel Duffy (Int) to Everyone: 2:18 PM
Thank for the responses Ricardo. I am interested in a discussion about how we could increase the throughput of the cycled DAS to much greater than 2 days/day.

---

4. Use Case: Coupled Analysis/Reanalysis (MERRA-3 and Beyond)

Ron Gelaro, NASA Goddard, GMAO
ronald.gelaro-1@nasa.gov

Slide 1:
Reanalysis in the MERRA-3 era and beyond, 5+ years from now. Input came from R. Gelaro, E. Hackert, B. Putman, R. Todling, and A. El Akkraoui.

Slide 2:
Reanalysis using a fixed data assimilation system to do a long reprocessing of the observational record to give a consistent, long-time, high-temporal picture of the atmosphere over 20 to 30 years. It is the most public-facing product that the GMAO produces. Each year, hundreds of terabytes get distributed to people around the world.

Not a forefront system with highest resolution, typically done with a tried and trusted friend, whose strengths and weaknesses you know well. We are headed towards doing coupled analysis/reanalysis, multiple components of Earth systems state (ocean-earth). A coupled ocean and atmosphere will be our highest priority for MERRA-3. We want to demonstrate the value of NASA observations in all these Earth system components.

Atmosphere: 6-hour cycle of assimilation.
Ocean: likely be longer cycles of assimilation, 24 hrs to a few days.

Not just an atmospheric model, generation of products. See graphic for high level resource used. Future: demo NASA observations across NASA. 6 to 24-hour basis for assimilation.

Slide 3:
Internally it is used constantly, not just for science and research, but for bias correction, validation, etc.

Criticality: essential for reanalysis production and testing. Impossible to meet deliverable without HEC resources. Commercial cloud: productions are 20-30 years of simulated Earth states, run on 2-3 years, then catch up to real time on commercial cloud would be prohibitive.

No critical security needs, test processes should not be released, some level of confidentiality required. Continuous availability is the most important aspect during development, critical for production/public-facing product, multiple streams (though for future coupled reanalyses, fewer streams must be used to ensure continuity), around the clock runs to near real-time, continuous availability critical. Products are meant to be used worldwide, internally used constantly for validation and verification, calibration, bias calculation, most trusted record of assimilated state of earth system. The production runs behind real-time, at least 1-2 weeks behind real-time.

Slide 4: Major change in workflow:
Current: atmospheric reanalysis to start next year [2021], single-component data assimilation system with a lot of observations gathering of information from various places, resources, mostly atmospheric model. Computing is moderately intense, not the highest resolution, but there is a lot of file swapping between components, substantial product generation.

DA=data assimilation
You gather observations from disparate sources. It involves a lot of IO. Computing is moderately intense, but there is a lot of file swapping between components.

Future: Involves staging of observations from multiple Earth system components, and the use of ocean-atmosphere coupled models. Assimilation will be moving toward an in-core system, minimizing I/O, but once we step a coupled ocean-atmosphere model, memory will be a huge issue, thus advanced memory use will be needed.

Slide 5:
Current/Future apps: Try to keep same throughput to get 20-30 runs done. GEOS 21RC Atmospheric Reanalysis will begin in 2021, using O(1000) cores. They will need 10 times the cores they use now for future coupled reanalysis, and continuous usage will be critical.

Slide 6:
Key factors: Once we are in production, fighting it out with development, other operational runs, GEOS development cycles impact machine availability. Coupled data will be more utilized.
Moving to coupled DA will impact the strategy for evolving the GEOS modeling system.

This effort is not trying to keep up with the latest, greatest version of GEOS.
We need to be able to finish these things in a couple of years, or they are not valuable anymore. The transition to MERRA-3 probably represents the largest inflection point for growth.

Q&A [with Ron Gelaro]:

from Michael Little (Int) to Everyone: 2:43 PM
A question for the discussion at the end: any resources you would need to advance your code to run on new platforms that aren't already identified in your current plans.

from Michael Little (Int) to Everyone: 2:54 PM
Another question for the end discussion: How do you evaluate whether or not HEC is supporting your Use Case satisfactorily (i.e., no wait time vs. longer wait times; embedded system support in your organization to no user support)?

---

5. Use Case: Use Case: S2S-to-Decadal Prediction and Predictability (GEOS-S2S)

Andrea Molod
andrea.m.molod@nasa.gov

Slide 1:
[No additional notes: see slide content]

[No additional notes: see slide content]

Slide 2:
Lots of contributors for this seasonal forecast: Andrea Molod, Eric Hackert, Yury Vikhliaev, Bin Zhao, Donifan Barahona, Anna Borovikov, Robin Kovach, Siegfried Schubert, Zhao Li, Young-Kwon Lim, Lauren Andrews, Richard Cullather, Randal Koster, James Carton, and Kazumi Nakada. This use case is a little bit different than what else has been described because of nature of seasonal predictions: we are taking about a huge number of small jobs running at the same time. Timing of a storm 3 weeks out is not predictable.

Part of the important element here is to be able to figure out what is predictable and what is not. The timing of a storm 3 weeks out is not predictable. Many ensemble members in forecast makes your system more reliable. Don't forecast temperature, but a temperature anomaly next season. Requires a lot of HEC resources to do this. Fewer ensemble members lower res, without HEC. Recent system w/more resolution: The skill of the multi-model ensemble has gone up because of our contribution.

Because of the massive dataset size, this will require a user or partner who wants to use this for their research to have some access to do averaging or processing. We need the retrospective suite to figure out the skill and to have a climatology from which we can make forecasts.

New system we are testing, up to 40 sub/seasonal forecasts per month. Weakly-coupled data assimilation system, focus on forecast suite. If NASA supercomputing weren't available, we would have to run fewer ensemble members. This would hit us badly.

Slide 3:
Extra resolution and ensemble size add to computational burden. Images show what high resolution buys you. You see beautiful structure on the right, with the Gulf Stream further from the coast, where it needs to be. We figured out a new ensemble strategy that takes advantage of the computing. You need at least 40 ensemble members. There is no evidence beyond 3 months that a large ensemble size gives you a huge advantage.

Every 5 days, we kick off perturbed and unperturbed forecasts, run them for a few months, and then run a subset for future months.

Slide 4:
We are convinced that we will need to make use of resources at both NAS and NCCS.

We only do data assimilation for the ocean, but we replay to a pre-computed, atmospheric simulation. Everything is in-house, so essentially we just have to run the forecasts.

The previous system generated 2 petabytes of output at lower resolution. Now it is 1.5 petabytes; we have to be selective.

A single month uses 122 nodes over 9 or 10 days. For retrospective, we need to stay on the machine for a year. We are testing at NAS and are very happy with the throughput.

Slide 5:
We have had success in getting into interstitial spaces for the retrospective runs.

Turnaround time is 3 or 4 years, when we upgrade the system.

Results show that increased horizontal resolution in the ocean can improve storm tracks.

Q&A [with Andrea Molod]:

from Laura Carriere (Int) to Everyone: 3:04 PM
When you move to the next version, do you need to keep the output from the previous version on disk?

Andrea Molod: In the beginning, yes; eventually, no.

from William Thigpen (Int) to everyone: 3:05 pm
Do you have an idea for reliability you need of the file systems?

Andrea Molod:
I don't know what type of reliability. The systems at NAS are certainly less reliable than at NCCS. We can do the retrospective at NAS but need to do the near-real-time at NCCS. Reliability has to come from NAS: their file systems keep jumping up and down.

Peter Williams: File system stability/file reliability a constraint? Are there any difference/nuances of these terms worth pulling out, or are they synonymous?

Andrea Molod: They are more or less synonymous.

from RICARDO TODLING (Int) to Everyone: 3:07 PM
File system behavior can be very dependent on user activity. Ideally, that would not be the case.

Andrea Molod: From her review
I would add a speculation that the irregular behavior of the NAS file system is not related to system load. We have made tests running one simulation at a time and 10 at a time, and there seems to be little correlation. We continue to experience sporadic failures to read or write a file during a run after long wall-clock time wait periods, and when we resubmit all is well.

---

6. Use Case: Atmospheric Composition – Air Quality Analysis and Forecasting (GEOS-CF)

Steven Pawson
steven.pawson-1@nasa.gov

Slide 1:
[No additional notes: see slide content]

Slide 2:
Produce global analyses. Things we need to do that nobody has ever done before:

Went into production a little over a year ago. We are at the forefront of international capabilities here. That is one of the challenges in projecting it into the future.

We would not be able to run this system without NASA HEC resources. We run in real-time on the NCCS systems.

We have to get them out on the [same] day, so timeliness is a constraint for us.

We need to archive our analyses in the long term.

GEOS-CF: Footprint 130 nodes (3640 CPUs); job runs 8.5 hours wall clock - we have to reduce this by about 2.5 to synchronize with GEOS-FP. Public data: 8 Tb/year analysis and ~40 Tb/year for five-day forecasts. On NCCS we archive 80 Tb/year for analyses.

Slide 3:
GEOS-FPIT is stable and reliable in real time.
We are not doing an assimilation of constituents yet, but that is coming.

Imagery: We can actually compute pretty well the weekly cycle of nitrous oxide in Zurich. We also get the diurnal cycle pretty well.

Down the road, we will increase resolution, to get it in line with GEOS-FP, but it likely won't be at the same resolution. With 256 species, that would be a prohibitive goal. We would like to use GEOS-FP for transport.
To meet our goals, we would need to speed up GEOS-CF 2½ times.

All of the uncertainties make it difficult to give the footprint of the job in 5 to 10 years.

Slide 4:
There is a lot of potential with emerging techniques such a ML and visualizations done by SVS. Draft in Laura Carriere's hands for an archive plan. HEC allocation at NCCS: we use ML and GPUs and we will be expanding that in future, we work closely with NCCS, which is critical for success.

Slide 5:

RD runs require sporadic, additional resource requirements on multiple days, multiple weeks that rely on trace features. Next 5-10 years: resolution and complexity limited by HEC resources. Increasing problem: data volume we can save to do reanalysis. Need to develop more tools for data analysis: ML will grow in this arena. Resolution and complexity of system, feedback to chemistry on meteorology and vegetation dynamics will grow as resources grow over time. R&D can be done elsewhere: ex Harvard on cloud, not high-resolution. Imperative that more users use this and making it available to more users. Speeding up core model will affect the performance of this system.

Putting GEOS and GEOS-Chem together makes an even larger model.
The resolution and complexity is limited by the HEC resources. One of the biggest challenges we will face are the observations as well as [missed what he said].

Q&A [with Seven Pawson]:

from Michael Little (Int) to Everyone: 3:25 PM
Steven: do you expect to include the changes to the Planetary Boundary Layer measurement under the decadal survey?

Steven Pawson: Yes, in fact, that's going to be a focus of the GMAO down the road. We just hired a new civil servant to focus on that. The PBL is essential for meteorology and the trace gas problem, a stronger focus down the road.

from Michael Little (Int) to Everyone: 3:44 PM
How do you envision the evolution of your application to need to be supported by new capabilities by the HEC Program?

---

7. Use Case: Chemistry-Climate Model (GEOS-CCM)

Pete Colarco, NASA Goddard, Atmospheric Chemistry and Dynamics Laboratory
peter.r.colarco@nasa.gov

Slide 1:
[No additional notes: see slide content]

Slide 2:
We are not a couple climate modeling effort *per se*, GEOS system is yet to be vetted in that form, we are doing analysis now on coupled system, as Andrea mentioned, for system response on multi-year timescales. Need for higher resolution: impact of ozone impact/intrusions is improved by high-resolution sensors, will drive models at higher and higher resolutions. See: Use of NASA data to inform model improvements [slide graphics].

This is currently a research model.
Have added complexity to the model to serve these international assessments, satellite retrieval teams (*a priori*), etc.
Images:

1. Total ozone column in the tropics. Have converged to the brown and gray lines, the satellite data.

2. Projection aspect: response of ozone layer as reduce ozone-depleting substances in atmosphere. Can we use the coupled system?
3. Need for higher resolution: shows the impact of stratospheric ozone intrusions at different resolutions. We will be doing simulations at higher and higher resolutions.

Slide 3:
Chemistry-climate modeling: Significant role of NASA satellite data improving model examples.

Slide 4:
Criticality: We do most of our work at NCCS, with increasing amounts at NAS on Pleiades.
These resources are essential for doing these simulations, else we would run at low resolutions or fewer use cases. This is a research activity, so we don't have the confidentiality and ITAR concerns of other applications, and algorithms are shared. Data analysis is done with wider scientific community, more concern for security during proposals.

We don't have day to day production activities, but we have a 4-year cadence of ozone assessments which require 1 year advanced – 150 years of simulated time. Massive datasets: advantages to NCCS/NAS systems available to do analysis where data resides. We don't have real-time requirements with this system. We want to get the model to the fidelity of the satellite observations.

Slide 5:
This project is somewhat of a research activity, using GEOS and coupled system components. We need to work across components that haven't been coupled before, scientists experts user work is needed for parameterizations.

Slide 6:
Main constraints: scalability of how many resources are available for higher spatial resolutions and complexities.

We have a cadence of every 4 years.
The frontier is at the coupling of new components.
As we go to different computing environments, we will need assistance in migrating to these new architectures.

Q&A:
No questions.

Group Discussion of Key Questions

Peter Williams: Now we will ask general questions about all GMAO use cases and needs:

> **from Michael Little (Int) to Everyone: 2:43 PM**
> A question for the discussion at the end: Any resources you would need to advance your code to run on new platforms that aren't already identified in your current plans? Answer: GMAO would need help adapting its codes to run on advanced architectures with new processor types (e.g., GPUs), as well as restructuring to exploit advanced memory utilization strategies.
>
> **from Michael Little (Int) to Everyone: 2:54 PM**
> Another question for the end discussion: How do you evaluate whether or not HEC is supporting your Use Case satisfactorily (i.e., no wait time vs. longer wait times; embedded system support in your organization to no user support)? As mentioned in my presentation, for reanalysis, continual availability is a significant metric.

from Michael Little (Int) to Everyone: 3:44 PM
How do you envision the evolution of your application to need to be supported by new capabilities by the HEC Program? NAS or NCCS?

Peter Williams: As the application evolves, are there additional, particular needs that you will want from HEC Program?

Bill Putman: Two major aspects where we need support: 1. adaptation for new architectures and accessibility to them for new development, as well as human resources to make development happen. 2. Access to advanced memory management, techniques, and hardware devices to avoid large data volumes going to and from disk rapidly.

Steven Pawson: One obvious aspect, as our systems get to higher resolution and with more components, our data volumes get larger. We will need help on the data management side as well, and not just management but the ability to access data after the fact.

Tsengdar Lee: I know some of the data will go to DAACs, but what about Portals? I have heard so much about people wanting a portal, which would mean significant data management at the centers but also curation of the data, like at a DAAC at computing center. Are we moving in that direction? Is there a need for that?

Steven Pawson: If we had a coherent strategy across the agency, or at least Earth science, that problem could be solved. With a coherent strategy across the agency, discussions among users and DAACs would be good, so we don't all try to solve the problem individually.

Ron Gelaro: Dan Duffy has said many times that we don't want the data center to be the distribution center. [See Dan Duffy's comments later.]

Piyush Mehrotra: Do you want us store data in DAAC? If you store all the data in the DAAC and have to do reanalysis…

Steven Pawson: Two aspects: data that goes into our systems, data produced with Systems–we are talking about the latter at this point. The data computed at the systems, we have to decide what to do with it.

Ron Gelaro: The idea of people downloading these datasets is prohibitive. The OSSE people talked about regeneration of files. For reanalyses, you will need to bring analytics to the data source rather than download huge data sets. Downloading will not be practical anymore.

from ELLEN SALMON (Int) to Everyone: 3:49 PM
As I understand it, Dan Duffy's point is that the HEC systems are not FUNDED to be long-term archives.

Piyush Mehrotra: Depends on how you look at that. We haven't deleted anything since NAS systems came online.

Tsengdar Lee: This is a requirements workshop. The whole purpose of this workshop is to figure out future requirements, so if there is a need, we go back and argue for those resources. We don't want to limit that by what we are currently funded to do.

Laura Carriere: So, we are not a final archive for products. But someone's output is someone else's input, and that is why we have Centralized Storage, so you don't have to download it again. There are some conflicting requirements about how we go about doing this. Christa said it's hard to move data over, how do we get around those things? If it's unique, it gets attention. If it isn't, we don't. What are your requirements? What is going to make your life easier?

Peter Williams: Any thoughts on storage and archiving question?

Dan Duffy: Ellen put it well, it's something we need to consider: what is the HEC Program's requirements for archiving, and how does that fit into the puzzle? We need to understand if it fits within what we have, or is it something we need to do? Some of the DOIs that are kept on this data, so the question is, can we fit all that?

from Michael Little (Int) to Everyone: 3:54 PM
If you move to a new processing architecture (i.e., GPU) to get a large-scale increase in cores, what would resources would you need to do this?

Peter Williams: As you change processing architecture on cores, are there part resources you need to do this? Any thoughts on Mike's question?

Ricardo Todling: I am not sure what resources mean. It is both hardware and software. What I can see is, this question is associated with the question of how easy is it to transition to a new architecture? I think we need new approaches such as how new codes relates to new architecture. The Atlas software (from ECMWF) will hide all that from the users, so software doesn't need to change. It will make transitions much easier.

We need a guide to some type of framework of that sort. From the data assimilation component, the part that handles the observations. JEDI will incorporate Atlas. For example, the GOES model will be linked to JEDI but not the framework, so all of it may not benefit from Atlas. So, we at GMAO need to pay attention to that development.

Bill Putman: That is the DSL (domain-specific language). We are looking at Cocos or grid tools, specifically, the dynamics or the physics which has a broader [community] development. The intent of DSL is hardware-agnostic coding. Get hardware-agnostic coding done for doing work on CPUs or GPUS. That is not easy. The challenge is getting qualified developers.

Ricardo Todling: People need training, Like PKI(?)

Peter Williams: Any other thoughts?

from Michael Little (Int) to Everyone: 4:00 PM
What metrics do you use to evaluate HEC?

Mike Little: The evaluation—How do you evaluate whether the HEC Program is supporting you? What metrics do you use to evaluate HEC Program as it relates to use case?

from Peter Williams (Ext) to Everyone: 4:00 PM
Mike's Q: how do you evaluate whether or not HEC is supporting your Use Case satisfactorily. (i.e., no wait time vs. longer wait times; embedded system support in your organization to no user support)

Nikki Preve: For the OSSEs, one metric is whether or not we can provide support for instrument teams on the timeframe needed for their proposal or decision-making process.

Steven Pawson: It's a bit qualitative. The individual user sees turnaround of their jobs as a key metric. Are they getting answers as fast as they want? I tend to see it a little differently. Is the group as a whole making progress and reaching their goals as a group in a reasonable timeframe—that's not measurable, quantitatively.

Nikki Prive [review addition]: For the OSSEs, one metric is whether or not we can provide support for instrument teams on the timeframe needed for their proposal or decision-making process.

Bill Putman: When we begin to feel as though we are holding back science because compute can't keep up is when we're not being supported sufficiently. There is a chance that we push that when we think about complexity of systems. There is some complexity we could benefit from that we can't do right now.

Ricardo Todling: You made a good point, Bill, which goes to what Steven said. Is HEC favoring a handful of users or the community under that environment? I am one of those favored users.
I would say that, as long as there are strong limitations in the system, it's probably overall not satisfying. People have to be able to do science. Computing should not be their main problem. It should be whatever science they are doing.

Peter Williams [Mike Little?]: Any final big-picture questions or responses to Steve's question:

Mike Little: Thanks Peter, great job facilitating this meeting, especially presenters. Extremely illuminating. What we will do is collect the notes and slides, go into a cycle where scribes and Peter with coherency in the next few weeks, make further notes about them later. Any notes will go out to presenters with specific questions and clarifications to get everything clear, collect the feedback on notes, use that as input into Users Needs Assessment Report. There will be a survey sent out such that it will not duplicate what we have discussed here. We are hoping to have a draft to HEC Need Assessment by the end of July, so you should see it in the course of the next 6 weeks.

Ricardo Todling: from RICARDO TODLING (Int) to Everyone: 4:06 PM
Will you need the supplemental material? if so when?

Mike Little: We will ask those questions in the survey with time to respond. Thanks very much all!

---

1. Use Case: Observing System Simulation for Earth Science Space Missions

Arlindo da Silva
arlindo.m.dasilva@nasa.gov
GSFC Global Modeling and Assimilation Office, Code 610.1
SMD

Main points:
P1: Simulating future NASA space-based Observing Systems based high-fidelity earth system model and instrument simulators. Earth system model simulations must span a full seasonal cycle at global scale.

Criticality: this must be done an a HEC system for this type of model.

P2: Earth system model simulations can be shared with the public; early observing system simulations and trade studies may not be appropriate for public release, but there are no critical security issues.

P3: Future flow: frequent checkpoints with limited output. I/O speed becomes a bottleneck with higher resolutions if one continues to produce a full set of output. ML can be used to speedup model parameterizations and instrument simulator. Minimize I/O, use model as a compression algorithm: simulations can be restarted from any point during the year, with calculations performed online with full access to model state.

P4: Future place: HEC Center or public Cloud. Trading disk space to processing ability.

P5: Requires regular access to large number of cores. Instrument simulator plug-in protocol and user training required.

Raw notes:
Slide 2: Observing System Simulations: Mission Need (1 of 2)
Specific topic chosen for formulation of Earth science missions.

Mission concept: want to simulate that observing system with a high degree of fidelity. The flow is high-resolution, high-fidelity global physics model –> sample output of model –> apply instrument simulator, create a data stream similar to the instrument.

The goal is global coverage and fidelity.

Slide 3: Observing System Simulations: Mission Need (1 of 2)
It must be done on a HEC system. There is no way to do it on a desktop.

They are sharing data among the science team, not so much the public.

Slide 4: Current/Future HEC Application Characteristics
With increasing resolution and complexity, how does the workflow need to evolve?
It becomes a game where you have to read data from your files. I/O becomes more and more of a bottleneck.

You have to make the high-resolution run sequential. You do so with relatively limited I/O.

The user would provide instrument plug-ins that you would dynamically load libraries into as you go along.

Use model as the data compression algorithm. The idea is that instead of writing huge amount of 3D data every few minutes one would dump the “model state” much less frequently (say, weekly), and restart the model from anywhere along this trajectory. Once one restarts the model and integrates it forward in time you have full access to the model state, every time step, and by means of plug-ins can then deploy instrument simulators on-line.
This calls for running the earth system model, sequentially for a year, with weekly state dumps (checkpoints). A mission designer may want to simulate their observing system for select days/weeks during the year. So, she develops a plug-in for their instruments and restarts the model for the periods of interest, running it for a relatively short amount of time.

Slide 6
Massively parallel, but I/O gets in the way.

Q&A:

Q from Daniel Duffy (Int) to Everyone: 1:38 PM

Can the Nature Run simulation be broken into multiple streams, or does it require a single stream to process the entire time period?

A: Ideally we would like to do that in a single stream, because in reanalyses we have the observational data to bring it back to same point. With a freely running model, you will not have that; it will be [asynchronous, starting at different points?].

Q from William Thigpen (Int) to Everyone: 1:40 PM

Do you have a number of cores you require?

A: See slide 14 with Table 2.1. It is 40,000 cores for 6 km, perhaps less with optimization.

Appendix D8 Non-workshop Inputs

NOTE: These are provided since they are direct communications and typically not available online.

ID	Mission Directorate	Division/ Program	Presenter	HEC Project Title
D8.1	SMD	Astro	Kenyon	Simulation of Planetary Formation
D8.2	SMD	Astro	Powell	TESS Full Frame Light Curves
D8.3	SMD	Astro	Jiang	A Global 3-D Radiation Magneto-hydrodynamic simulation of super-Eddington Accretion Disks
D8.4	SMD	Astro	Brooks	Simulation Galaxy Formation and Evolution
D8.5	SMD	Astro	Cen	Galaxy Formation and Evolution
D8.6	SMD	ESD	Su	JPL OSSE Support
D8.7	SMD	ESD	Chu	OCO-2 and OCO-3 Science Data Processing
D8.8	SMD	Helio	DeVore	Solar Corona and Space Weather Models
D8.9	SMD	Helio	Chen	MMS/RDSP/THEMIS/Psyche
D8.10	ARMD		Malik	Need for Modern HPC Hardware to Solve NASA Challenge Problems (Draft)
D8.12	ARMD		Capps	Industry Inputs into Future NASA Project Planning (9/2/19)
D8.13		DT	McLarney	DT Artificial Intelligence and Machine Learning APMC 6/11/20
D8.14	ARMD		Malik	NASA CFD Vision 2030 Study – A path to Revolutionary Computational Aerosciences (9/30/14)
D8-15	ARMD	AOSP	Oza	Machine Learning for Aviation Safety
D8-16	ARMD	TTT, AATT	Gregory/ Schneck	Supporting Safe, Robust Aerospace Structures Through HPC Enhanced Computational NDE
D8.17	SMD	HECC	Thigpen	HECC Overview, March 2019
D8.18	ARMD		Capps	Industry Inputs into Future NASA Project Planning, March 6, 2019
D8.19	OCIO	DT	McLarney	Digital Transformation: AI/ML Driefing to NASA Administrators Program Management Council
D8.20	ARMD		Malik	NASA CFD 2030 Study September, 2014
D8.21	OCE	HPCI	Gregory	Supporting Safe, Robust Aerospace Structures Through HPC Enhanced Computational Non-destructive Testing (CNDE)
D8.22	SMD	ESD	Fenty	ECCO-ACCESS: a Cloud-Native Storage and Data Analysis System for Ocean Climate Research
D8.23	SMD	NAS	Mehrotra	NASA's Drive Towards Exascale Computing, July 19, 2020

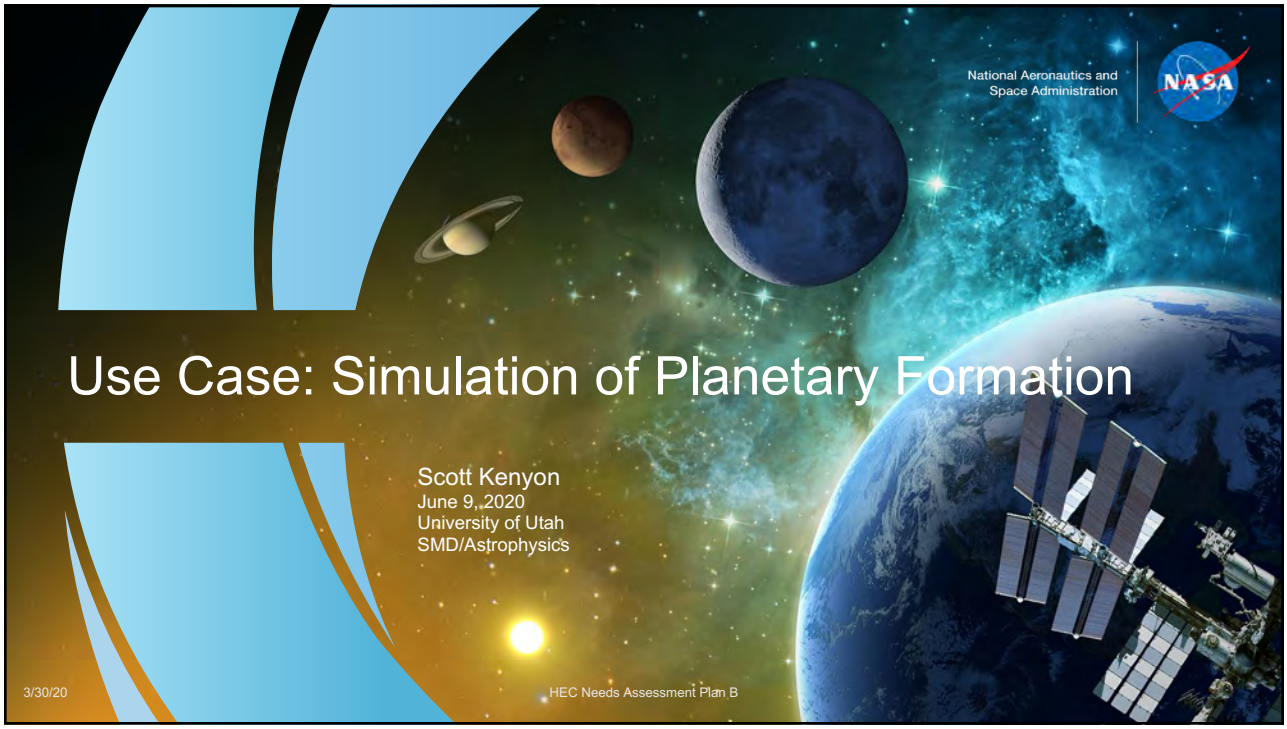

Use Case: Description

- Ben Bromley and I would continue our program of planet formation calculations. Our overall goal is to simulate the formation of a planetary system from a circumstellar disk of gas and dust. We have built a C++ code that follows the evolution of a gaseous disk and a swarm of solid material ranging in size from 1 micron up to planet-scale objects. Currently, we focus on the formation of the Pluto-Charon satellite system and the evolution of Kuiper belt objects.
- In the future, we can complete projects that follow the growth of many planets within the disk. Some test cases that we completed in the last week or so suggest that following the formation of the solar system and other planetary systems is possible on the Discover system.
- We use 56 to 212 processors for each of our calculations.
  - A full calculation requires a few to several tens of twelve hour runs.
  - The suite of calculations needed to understand how the architecture of planetary systems depends on initial conditions may take ten to several tens of full calculations. Thus, a typical project needs 212 * 12 * 25 *25 cpu hours for completion.
  - We can usually complete a few projects a year, so that gives you an idea of what we need to finish papers expanding our understanding of the formation of planetary systems.

National Aeronautics and Space Administration

Use Case: TESS Full Frame Image Light Curves

Brian P. Powell
HEASARC, Code 660.1
NASA/GSFC

Use Case: Mission Need

- Mission: Produce and analyze light curves from TESS full frame images
- Criticality: [How essential is NASA's HEC to accomplishing this work]
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
    - This work is not possible without HEC and would require the purchase of time and storage space on commercial computing systems.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - [What is impact of inadvertent release of data? (i.e., public confusion, professional embarrassment, National Security, Privacy, etc?)]
    - Raw data is already public, inadvertent release consequences would only be the exposure of data for which there is unpublished research.
  - [If HEC availability is interrupted regularly, how does that affect your ability to achieve objectives?]
    - This work would not be possible without consistently available HEC.
  - [What is mission need for external collaborators (without NASA credentials) to access system or to run jobs?]
    - None.
- HEC Function:
  - [Briefly describe how HEC supports performance of this work]
    - The production of light curves requires, on average, 3 CPU-years per TESS sector. Prior to the use of NCCS HEC, these light curves containing valuable scientific information would not be available to the scientific community, as had been the case for nearly two years prior to the start of this work.
  - [How often do you need to run the application to support mission objectives]
    - Build light curves (Discover): Once per month for approximately 3 CPU-years.
    - Analyze light curves (ADAPT): Several times per week.
  - [How is the output used to achieve mission objective?]
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - [If NASA supercomputing was no longer available to meet your needs, how would you perform this work? (i.e., buy a system, rent time on a commercial facility, use Commercial Cloud Computing?)]
  - This work is not possible without HEC and would require the purchase of time and storage space on commercial computing systems. ******NOTE**** this is the same question as "criticality" above**
- Mission Concept: [Provide a graphic if possible, depicting this application or use case]
  - Graphic on next slide

TESS Light Curve Mission Concept

Raw Data: MAST

Light Curve Construction: **NCCS Discover**

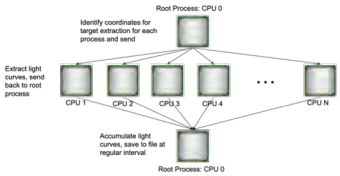

Unsupervised AI/ML: **NCCS ADAPT**

Supervised AI/ML: **NCCS ADAPT**

Embedding space representations

Classification

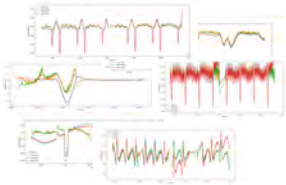

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - What triggers this workflow?
    - New sector of TESS data released from MAST (avg 1/month)
  - Is there a need for interaction by user, operator or other system in the computing workflow?
    - Yes, it is a complex multi-step process. NCCS is currently providing expertise to help streamline.
  - Does this use case need quick reaction or near-real time computing? Explain briefly.
    - No, although faster processing means more science output in a shorter time period.
- Data:
  - Discover:
  - In: [Summarize Data used to feed into the code/application in terms of volume, latency, source]
    - TESS full frame images, approximately 2 TB per sector
  - Out: [Summarize data produced by computation and what happens to it]
    - Processed light curves, approximately 4 million files, 100 GB per sector
    - [Does the output data require retention? How is that selected?]
      - Retention provides on centralized storage (/css)
    - [How is the data analyzed to provide mission results and where? (statistical or AI analysis, visualization?)]
      - AI/ML analysis conducted on ADAPT
  - ADAPT:
  - In: [Summarize Data used to feed into the code/application in terms of volume, latency, source]
    - Discover output, Processed light curves, approximately 4 million files, 100 GB per sector
  - Out: [Summarize data produced by computation and what happens to it]
    - [Does the output data require retention? How is that selected?]
      - No, outputs are extracted by the user.
    - [How is the data analyzed to provide mission results and where? (statistical or AI analysis, visualization?)]
      - ADAPT is used to create data representations and classifications,
  - Transfer:
    - [Where does the output go after it is generated? (i.e., local to HEC storage system or to another location)]
    - Processed light curves are stored in CSS, AI/ML output extracted by the user.
- Computation:
  - Type: [Summarize the type of computing used/needed and how HEC supports this today]
    - [Is this capacity computing or are special HEC capabilities needed for this computation?]
      - Both HEC CPUs and GPUs are needed for this project
- Evaluation: [Summarize how you evaluate the effectiveness of HEC in supporting this work]
  - Evaluation is subjective at this point since the work was not possible without HEC. As the process improves in efficiency, we can check run times and memory usage.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.]
    - Data release schedules, AI/ML expertise availability
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?]
    - None, the HEC is actually removing the limitations that previously existed.
  - Impact: [How does this affect fulfilling the mission need?]
    - N/A
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?]
    - No
- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?]
    - CISTO is providing resources to streamline the process and neural networks are being refined. Eventually, this process will be entirely automated.
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?]
    - The telescope will be increasing the frequency of measurements in the near future and Discover/ADAPT will be able to accommodate this change.
  - Re-use: [Will the number of people running this code/application grow?]
    - It is likely that more researchers will use the Discover light curve output for AI/ML analysis, hence the central storage.
  - Re-hosting: [What is needed to migrate code to future computing architectures?]
    - Code can be adapted as needed to future architectures.

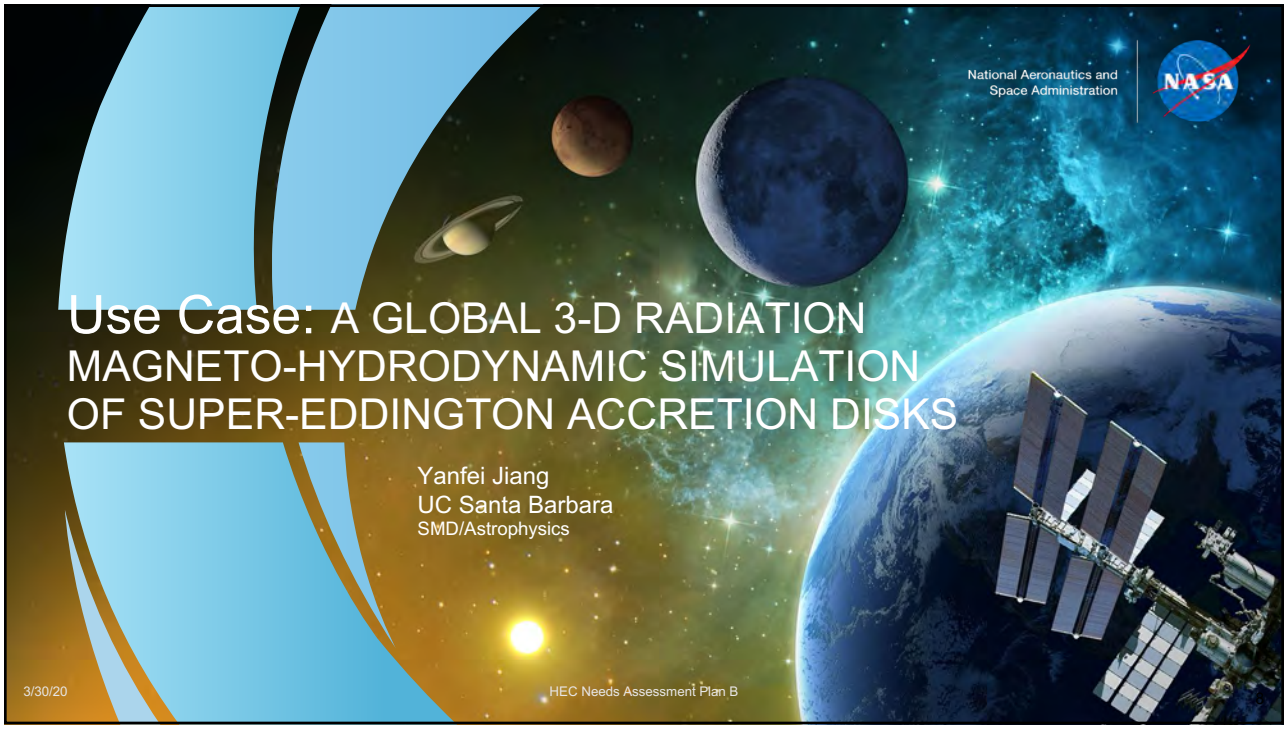

Use Case: Mission Need (Supplied by User)

- Mission:
  - We typically need on the order of 100M cpu hours per year, which is far beyond what we can get from local computing resource.
- Criticality: [How essential is NASA's HEC to accomplishing this work]
  - If NASA supercomputing was no longer available to meet your needs, there is no way we can perform the research projects we anticipated. We will have to scale down the problem size so that we can do it in the local computing resource.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - The impact of inadvertent release of data will be minimal for our project as most people will not understand the scientific data we will generate.
  - If HEC availability is interrupted regularly, it will significantly slow down the progress as we need to access the data frequently.
  - We do not need external collaborators to access system or to run jobs.
- HEC Function: [Briefly describe how HEC supports performance of this work]
  - We need to run each application continuously for one to two months. Then analyze the data for another month.
  - We typically analyze the data and publish the results in scientific journals.
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - We will not be able to do any of the simulation without NASA's HEC.
- Mission Concept: Figure 3 of this paper: https://iopscience.iop.org/article/10.1088/0004-637X/796/2/106/pdf

Use Case: (Synthesized from Paper)

- Mission:
  - Super-Eddington accretion flows onto black holes studied using global 3-d radiation magneto-hydrodynamical simulation
- Criticality:
  - This was a competitively selected project through ROSES solicitation
- Mission Concept: Figure 3 of this paper: https://iopscience.iop.org/article/10.1088/0004-637X/796/2/106/pdf

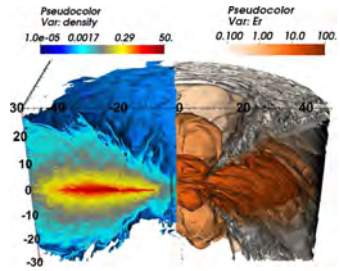

Figure 3. Snapshot of disk structures for density (left) and radiation energy density (right) at time $1.13 \times 10^4 t_s$. Units for ρ and E_r are ρ_0 and $a_r T_0^4$ respectively.

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - The applications can start with very simple input files and we take the output for analysis.
  - There is no need for interaction by user, operator or other system in the computing workflow.
  - We do not need quick reaction or near-real time computing. We can just take the output after the application is done.
- Data:
  - In: We just need to load a plain text file to start the application. The file size is smaller than 1MB.
  - Out: The application generates a series of hdf5 data files with the total size varying from 1TB to 10TB.
    - We need to keep the data for analysis. The saved files are based on the time resolution we need.
    - We produce visualization of the data as well as statistical analysis based on the scientific questions we need to answer.
  - Transfer:
    - Most of the output is saved in local HEC storage system.
- Computation:
  - Type:
  - Normal HEC resource is sufficient for our project.
- Evaluation: Support from HEC is very efficient and can help with any issues I have.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - The key factor for us is the use of large-scale, non-linear models, which requires a long cycle to compute.

- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: The key limit is the turn around time.
  - Impact: This significantly increases the time to finish our project.
  - External Sources: We typically use local resource to run small scale problems prepare the large scale runs.

- How will this Use Case evolve over the next decade?
  - Maturity: It will be an operational capability with continued development.
  - Growth: More physics will be added to support evolving future mission needs.
  - Re-use: The number of people running this code/application will grow.
  - Re-hosting: Convert the current code to a GPU version is one key goal.

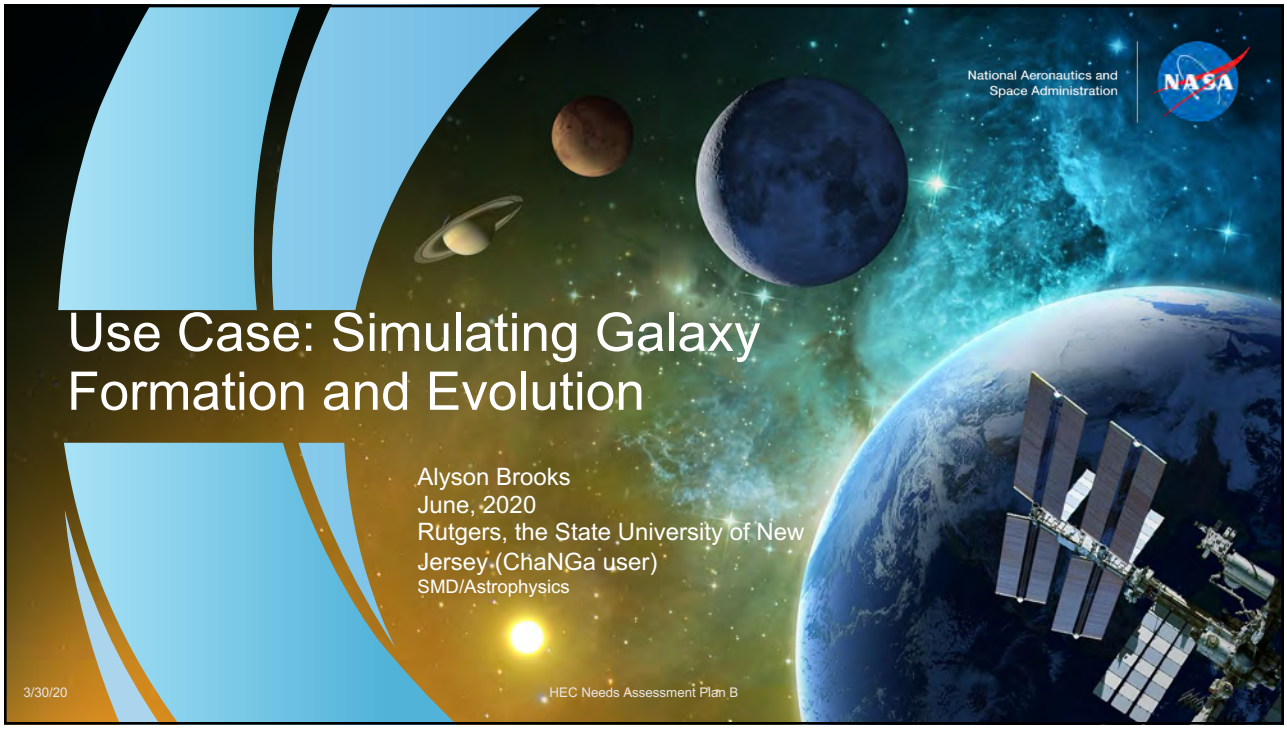

Use Case: Mission Need

- Mission: Discover the origin, structure, evolution, and destiny of the universe. We do this through fully cosmological, high-resolution, simulations of galaxy formation to interpret observations and make predictions. Impact can be most directly measured by publications, in which we acknowledge HEC resources.

- Criticality: Absolutely essential. There are few resources available to us that provide the hundreds of nodes needed to run our code.
  - The only current alternative option is NSF XSEDE resources, and these have not been as productive as NASA HEC for our work. Buying computing time is not feasible given the number of nodes needed and the hefty prices.

  - Mission Security Needs (Confidentiality, Integrity, Availability):
  - This work is not private/secure. No harmful impact if data is inadvertently released.
  - If HEC availability is interrupted regularly, this delays analysis/publication of findings.
  - All collaborators on this project are external to NASA.
  - HEC Function: Need supercomputers (hundreds of nodes at a time) to run the simulations.
  - We have used tens of millions of core hours in recent years. We run nearly continuously throughout a year.
  - We analyze the simulations to “discover the origin, structure, evolution, and destiny of the Universe.”

Use Case: Current HEC Application Characteristics

- Work flow: Simulations are run using ChaNGa (Menon et al. 2015), a smoothed particle hydrodynamics (SPH) + N-body code. ChaNGa uses the charm++ (Kale & Krishnan1993) runtime system for dynamic load balancing and communication to allow scalability up to thousands of cores.
  - Simulations are run from a few million years after the Big Bang to the present day. Astrophysical systems are dominated by self-gravity, leading to a large range in densities (large dynamic range).
  - Little to no user interaction, other than job submission. Simulations are checkpointed. Most jobs take weeks to months of wallclock time to complete.
- Data:
  - In: Initial conditions consist of one tipsy binary format file, not more than a few GB in size, a few ascii parameter files.
  - Out: Outputs in binary format. One main file, a dozen auxiliary files with additional particle information. Each output timestep is a few Gb.
    - Outputs are written based on how frequently evolution of galaxy needs to be analyzed. Typically 50-100 output files per simulation.
    - Some post processing is done using NASA HEC resources, e.g., generating galaxy catalogues.
    - Data is analyzed by plotting of particle data to answer specific questions. Done off site, usually at universities.

Use Case: Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - High resolution, non-uniform distribution with large dynamic range, short timesteps required.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: Current NASA resources have led to the most efficient use of ChaNGa compared to any other resource. However, code could be sped up with new processors. Also, access is tied to grant funding, which is variable.
  - External Sources: XSEDE (Stampede2, Bridges, Frontera)
- How will this Use Case evolve over the next decade?
  - Maturity: Continued development (more advanced physics modules, higher resolution)
  - Growth: More realistic physics modules, speed up, dual GPU/CPU usage
  - Re-use: number of users should stay roughly constant, though there is a publicly available version of ChaNGa which could increase user numbers
  - Re-hosting: ChaNGa is adaptable to many architectures

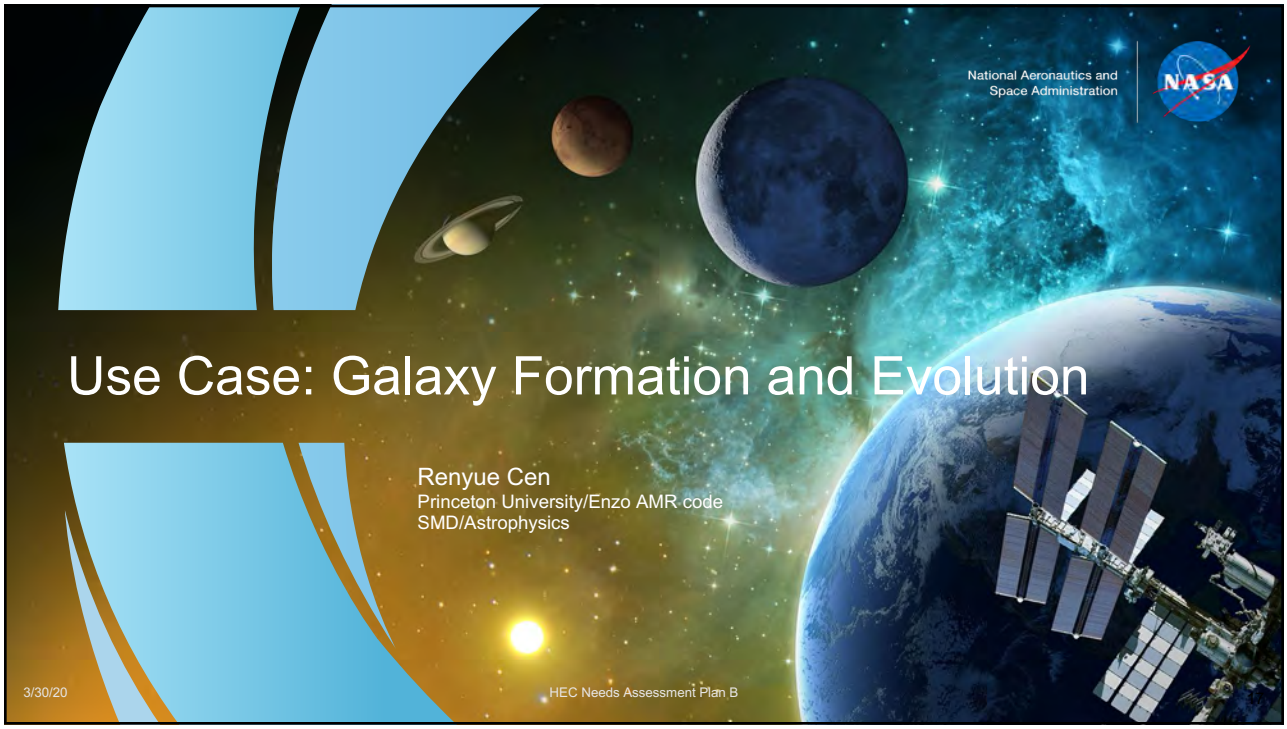

Use Case: Mission Need

- Mission: My general research area is on galaxy formation and evolution to drive NASA goals.
  - NASA HEC facilities enable state-of-the-art cosmological hydrodynamic simulations that are difficult to do otherwise.
- Criticality: NASA computing facilities are critical for my research
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? I would not be able to carry out my cosmological hydrodynamic simulations on scales and resolution enabled by NASA HEC platforms, particularly Pleiades.
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Not critical
  - Somewhat, as long as distruptions are not long and frequent
  - Ability to access data on NASA HEC platform by external users is important for greater impact
- HEC Function: General support from NASA HEC is good for general computing issues
  - We run applications nearly constantly.
  - We analyze outputs to make products to interpret and guide NASA missions.
- Impact of HEC: NASA's HEC is essential to accomplishing this work.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? I would not be able to carry out my cosmological hydrodynamic simulations on scales and resolution enabled by NASA HEC platforms, particularly Pleiades.
- Mission Concept: https://www.astro.princeton.edu/~cen/ShortAMRMovie2.mp4

Use Case: Current HEC Application Characteristics

- Work flow: 1. Run run hydrodynamic simulations. 2. Analyze simulation outputs. 3. Produce science.
  - Normal procedure
  - Yes. We often need assistance from HEC staff to enable simulation executions and data handling.
  - Yes. For example, often I need HEC staff to extend a job's execution time in real time to enable checkpointing.
- Data:
  - In: input or restart data sizes are in the 10-100GB range
  - Out: HDF data files are produced and some portion remains on disk for analysis
    - Simulation data need permanent storage space (e.g. tape systems)
    - We produce statistical results and visualizations, and AI applications are relevant as well.
  - Transfer:
    - All data generated that need to be stored are stored on lou, and on disks.
- Computation:
  - Adaptive Mesh Refinement hydrodynamic simlations.
    - HEC capabilities, especially Pleiades, are essential.
- Evaluation: Excellent

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - large-scale nonlinear models
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: The HEC capabilities are good in terms of enabling running large-scale simulations, however, more allocations times are desirable for running more simulations.
  - Impact: This limits the scope of explorations and statistical sample size of modeling.
  - External Sources: None.

- How will this Use Case evolve over the next decade?
  - Maturity: Our work will evolve in terms of simulation size and sophistication with more resources.
  - Growth: It will increase resolution or time-steps to support evolving future mission needs.
  - Re-use: The number of people running this code/application will likely grow.
  - Re-hosting: For future new computing platforms, it is essential to provide migration support.

Use Case: Mission Need

- Mission: Perform Observing System Simulation Experiments (OSSEs) to assist future mission formulations
  - The quantitative assessment of future mission impacts from OSSEs enabled by HEC is critical to optimizing mission design and assist decision-making. Improvements in mission design represent the HEC impact.
- Criticality: HEC computing resources are indispensable to our OSSE work.
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Use JPL internal super-computing facilities or Commercial Cloud Computing (not yet explored)
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Mission design is competition sensitive and proprietary in nature. OSSE results can be published and open to the public.
  - Interruption to HEC would delay mission formulation progress, which can have severe negative consequence.
  - [What is mission need for external collaborators (without NASA credentials) to access system or to run jobs?] N/A
- HEC Function: [Briefly describe how HEC supports performance of this work]
  - We run OSSE simulations on HEC throughout the year. More intense when specific requests arise.
  - All outputs are archived on storage.
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - HEC is indispensable to our OSSE work. If HEC is not available, we have to explore JPL internal super-computers or resort to commercial cloud computing.
- Mission Concept: [Provide a graphic if possible, depicting this application or use case]

JPL OSSE Support Example of OSSEs

Forecast Impact

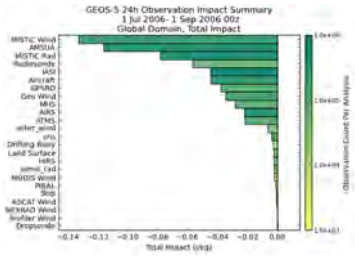

GMAO

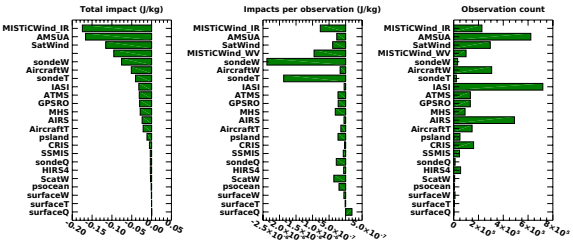

JPL

- We have successfully reproduced results from Will McCarty's MISTiC (IR sounder) AMV OSSE, updated with improved data thinning and new version of the GEOS model
- We have built an AMV error function that assigns AMV uncertainties in the OSSE system based on the analysis results
- AMV OSSEs are being conducted using JPL error function

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - Needs for future mission formulations come up regularly.
  - Is there a need for interaction by user, operator or other system in the computing workflow? Yes
  - Does this use case need quick reaction or near-real time computing? Explain briefly. No, our work is not real-time operation. Prompt technical support from HEC is appreciated.
- Data:
  - In: We use GEOS5 Nature Run and satellite observations. The data volume is huge. Latency is not an issue.
  - Out: [Summarize data produced by computation and what happens to it] OSSE data generated by GEOS5 are outputted every 6 hours and saved on /archive. The data volume is on the order of hundreds of TBs.
    - [Does the output data require retention? How is that selected?]: OSSE results are re-analyzed regularly.
    - Visualization and statistical analysis are performed at HEC.
  - Transfer:
    - [Where does the output go after it is generated?] Outputs are stored locally on HEC storage system.
- Computation:
  - Type: SBU and data storage are needed. HEC allocations are fair and reasonable.
    - [Is this capacity computing or are special HEC capabilities needed for this computation?] No
- Evaluation: The support from HEC is excellent. We appreciate the effectiveness of HEC.

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.] We use large-scale, non-linear models with revision to major models.
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?] OSSEs are time-consuming. It takes several weeks to run a one-month simulation.
  - Impact: [How does this affect fulfilling the mission need?] It may limit our capability to understand the seasonal cycle of the mission (such as CO2) in a rather short time period.
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?] No. Large amount data need be transferred. Compilation of the model system is machine-dependent.
- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?] Maybe
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?] Yes
  - Re-use: [Will the number of people running this code/application grow?] Yes
  - Re-hosting: [What is needed to migrate code to future computing architectures?] Recompiling code is not easy.

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs: General circulation model (GCM) simulation
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones: discover
    - Programming languages used (If multiple languages used, provide estimate percentage of the code in each languages.): Fortran (80%) C (10%) python (10%)
    - Number of lines of code
    - Workflow mechanism used to control code execution (perl, jupyter, shell script, etc.): shell script
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores): 16 nodes
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other: MPI
    - Storage is covered below
    - Run in batch, interactive, or near-real time: batch

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other): parallel
  - How much input and output are generated per run: 100 GBs for input and 10 TBs for output
  - Summary of the total amount of input and output data per year: 100 TBs
  - How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena?: statistics, visualization
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other: object storage
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used.
  - How much data needs to be curated (archived)? 100 TBs
  - How much data needs to be distributed and, generally, to whom? 100GBs to other colleagues
- What are the key computing obstacles to effective execution of the current application?
- Are there any analytic functions performed on the output of the application?
  - Timing, computing load, workflow
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed.

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.) : Yes
    - Increase in resolution and/or timesteps: Yes
    - Rewrite to improve performance
    - Refactor for emerging platforms
    - Rehost onto cloud
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs (nodes or cores)
    - Accelerator usage (type and memory requirement)
    - Amount of memory required (overall or per process or per core)
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other
    - Storage requirements
    - Run in batch, interactive, or near-real time

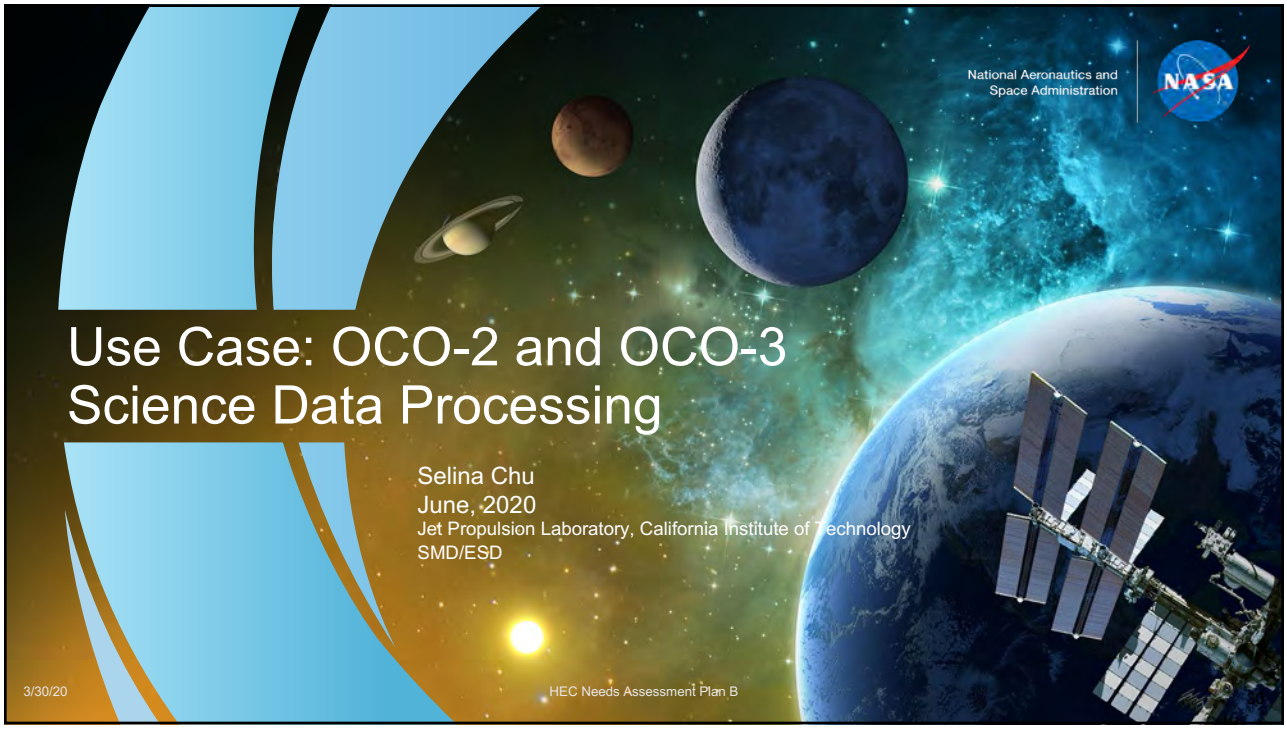

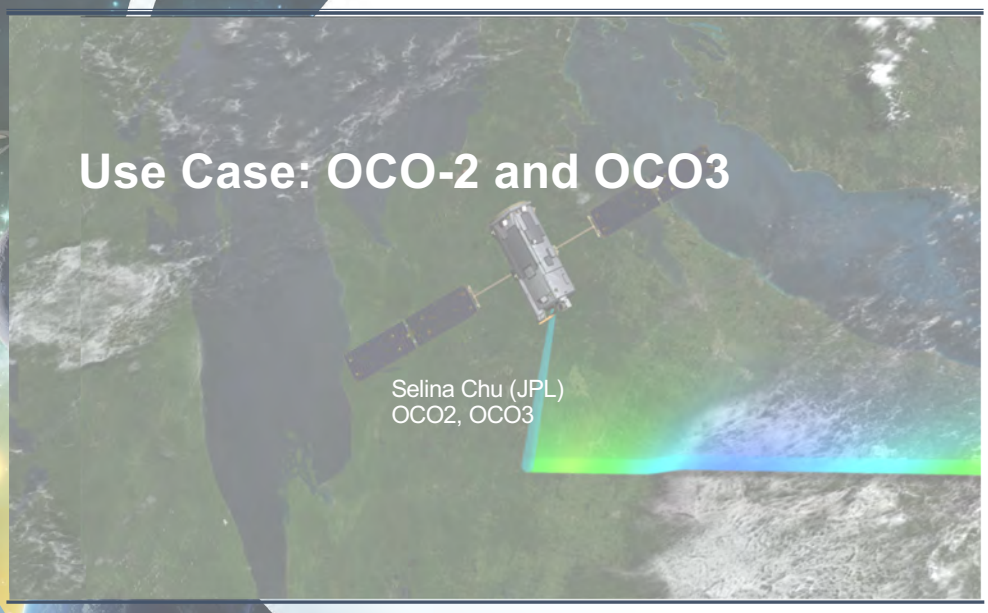

Use Case: OCO-2 and OCO3

Selina Chu (JPL)
OCO2, OCO3

Use Case: Mission Need

Mission

The Orbiting Carbon Observatory-2 (OCO-2) is a satellite with a three-channel imaging grating spectrometer that returns high-resolution spectra of reflected sunlight in molecular oxygen (O_2) A-Band and carbon dioxide (CO_2) bands.

- It is being used to study carbon dioxide distribution in the atmosphere and to detect emission hotspots and volcanoes since its nominal operation in 2014.

Orbiting Carbon Observatory-3 (OCO-3), the immediate successor to OCO-2, was installed on the International Space Station in May of 2019. It has been in Nominal Ops phase since August 2019.

- Both (OCO-2/OCO-3) are both in nominal ops phase, and the data product are being used by scientists from all over the world to study carbon dioxide distribution to a better understanding of carbon dioxide and the global climate.

Our data science system utilizes an optimal estimation retrieval code to extract CO_2 values from the OCO-2/OCO-3 observations

- It estimates the column-averaged CO_2 dry air mole fraction (XCO_2) with single sounding with random errors near 0.5 ppm (0.125%) over most of the globe. Another product derived is the solar induced chlorophyll fluorescence, which provides additional insight into the CO_2 uptake by the land biosphere.
- The main difference between OCO-2 and OCO-3, aside from being installed on the ISS, is that it provides a more detailed measurement and mapping of carbon dioxide. It has a pointing mechanism that allows for measurements of carbon dioxide emissions from almost any 50 mile by 50 mile region of interest.

All products from OCO-2/OCO-3 are delivered to the Goddard Earth Sciences Data and Information Services Center (GES DISC) for archival and public distribution

Use Case: Mission Need (2)

Criticality

- NASA's HEC provides both missions the capability to perform high-volume processing that results in higher quality products
- The optimal estimation retrieval code (or the full-physics algorithm) to extract CO_2 values from the OCO-2/OCO-3 observations has a very high computational cost
  - Individually, each mission collects approximately 1 million soundings per day.
  - If they are processed locally at JPL, it takes about 5 minutes to process just a single sounding.
    - Due to the local resources at JPL, it limits us to process only about 6% of the data that results in Level 2 products
- NASA's HEC enables our operational data systems to reprocess all cloud-free scenes on a monthly basis with the most current calibration coefficients. This results in a much larger amount of Level 2 products and with higher resolution
  - OCO-2 and OCO-3 are able to complete monthly reprocessing in a timely manner
- In addition, OCO-2 have been able to perform a reprocessing campaign on mission data, collected from the beginning 2014 to 2020
- In the last year, Pleiades supercomputer provided OCO-2 with over 20 million processor-hours
- If NASA supercomputing was no longer available, alternate options include:
  - Use commercial cloud system like AWS
  - Purchase dedicated server, which require the cost of have additional system administrators and maintenance cost
- With the aforementioned options (and our limited budgets), we might not be able to provide quality data, which comes from the monthly reprocessing using up-to-date calibration coefficients, that is necessary for scientists to perform state-of-the-art studies

Use Case: Mission Need (3)

Mission Security Needs

- The data and code are not confidential, and both, daily nominal and reprocessed, data are available to the public
- However, if data are corrupted or hacked due to security issue, the reprocessed products might not be accurate
  - Affect the accuracy of scientific studies that could result in misinformation to the public
  - Since the reprocessed products contain the most information, scientists mainly use the reprocessed data for their studies
- Availability of HEC is necessary to process mission data and provide accurate and most recent results for the public and scientists to perform studies

HEC Function

- HEC supports this work by providing us with the displace and usage of their supercomputers (with enough SBU) for our projects to perform reprocessing of the data, which is running the full-physics algorithm (or the optimal estimation retrieval code) to extract CO_2 values from the OCO-2/OCO-3 observations, at a very high computational cost
- OCO2 has been using this HEC for our monthly reprocessing and our reprocessing campaign, which are vital in providing scientists with high quality results for their research

Impact of HEC

- NASA's HEC is essential in providing the high performance computing power allowing us to processing massive jobs in an efficient manner that is significant in the progress of the mission

Use Case: Mission Need (4)

Mission Concept:

Collect spectra of CO_2 & O_2 absorption in reflected sunlight over the globe

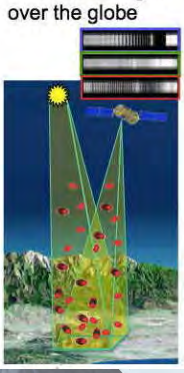

Retrieve variations in the *column averaged CO_2 dry air mole fraction, X_{CO2}* over sunlit hemisphere

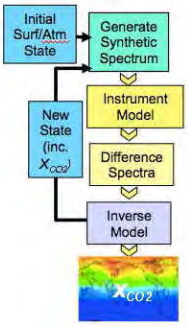

Validate measurements to ensure X_{CO2} precision of 1 - 2 ppm (0.3 - 0.5%)

Use Case: Current HEC Application Characteristics

Work flow

We have a processing system set up on Pleiades and on our local machines at JPL. Our workflow operates by monitoring a dedicated directory.

We transfer input and signal files from our local machines to Pleiades; once the files arrives, it gets picked up from our system to start processing, by submitting the jobs to the PBS

Once the job is done, we perform some post processing, and then the results gets transfer backed to JPL

Data

In: OCO-2/OCO-3 observations are high resolution of spectrometer readings of reflected sunlight in molecular oxygen (O_2) A-Band and carbon dioxide (CO_2) bands

Out: Optimal estimation retrieval code (or the full-physics algorithm) is used to extract CO_2 values from the OCO-2/OCO-3 observations, and L2 products are created from these results

Transfer

Output are delivered to the Goddard Earth Sciences Data and Information Services Center (GES DISC) for archival and public distribution on a monthly bases

Use Case: Current HEC Application Characteristics

Computation

- Parallel processing and computing using a large number of networked computing cores, facilitate by the PBS
- Large volume of disk space to store input data
- High performance file system to store and access data

Evaluation

- We perform testing and checks to ensure the correctness of our jobs
- We also perform timing tests to evaluate the computational performance

National Aeronautics and Space Administration

Use Case: Solar Corona & Space Weather Models

C Richard DeVore
June 2020
NASA GSFC Code 674
SMD-Heliophysics

Use Case: Mission Need

- Mission: Physics-based understanding of solar coronal dynamics and the origin of space weather events
  - Advances in understanding are enabled by HEC resources used to test and improve theories via high-fidelity simulations
- Criticality: NASA (or alternative) HEC resources are absolutely critical to accomplishing this work
  - If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work? Acquire resources at alternative centers (NSF/DOD/DOE), if possible, or transition work and personnel to other agencies/institutions
- Mission Security Needs: (Confidentiality, Integrity, Availability)
  - Inadvertent release of data could surrender science leadership to competitors and/or adversaries
  - Interruptions to HEC availability delay research progress
  - External collaborators are critical; they are issued NASA credentials sufficient to access systems and to run jobs
- HEC Function: HEC supports performance of this work by providing the resources to run massively parallel applications
  - Applications are run essentially continuously to support mission objectives
  - Output data from simulations are analyzed to develop physical insight and understanding to further science objectives
- Mission Concept: Three-dimensional numerical simulation of a solar coronal jet (still frames and animation):

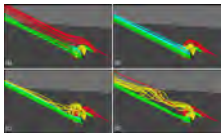

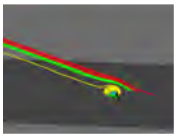

Use Case: Current HEC Application Characteristics

- Work flow: Design numerical experiment; generate input data as needed; generate and analyze output data
  - Physics-based understanding requires carefully designed and controlled experiments for subsequent analysis
  - There is a need for interaction by the user to monitor the computing workflow through multiple runs
  - This use case does not require quick-reaction or near-real-time computing
- Data:
  - In: Input data fed to the code/application are local in most cases; restart capability is used prodigiously
  - Out: Output data produced by computation are compiled and analyzed graphically in post-processing steps
    - Output data are retained in archival storage for both immediate and subsequent retrieval and analysis
    - Data are analyzed and visualized to provide mission results for publication in science journals of record
  - Transfer:
    - Key output data are generated locally and archived permanently on HEC storage systems; data transfer to user desktop platforms enable post-processing analysis and visualization
- Computation:
  - Type: Massively parallel multiple-instruction multiple-data computations use Message Passing Interface
    - FORTRAN/C-based MIMD MPI simulations do not require special HEC capabilities
- Evaluation: NASA HEC quite effectively supports this work; however, NCCS queue limits are highly constraining

HEC Needs Assessment Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - Varying loads on computing assets result principally from proposal/presentation/publication cycles, secondarily from revisions to input data and to major physics models within the application
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: A major constraint is the queue limits, especially on run time, imposed by NCCS current policy; managing this constraint requires frequent user intervention in cycling successive runs
  - Impact: This constraint inevitably delays completing simulations and fulfilling mission needs
  - External Sources: DOD, FR, and UK collaborators also run the application on non-NASA computing facilities
- How will this Use Case evolve over the next decade?
  - Maturity: Transition from research to operational capability is possible with continued development
  - Growth: Increased spatial resolution and temporal duration support both current and future mission needs
  - Re-use: Increased functionality and fidelity of the application will promote growth in code utilization and value
  - Re-hosting: Complete rewrites are needed to migrate code to alternative computing architectures (e.g., GPUs)

Current Use Case (Survey)

- Achieving the mission objective
  - Computation type: Multiple-instruction multiple-data (MIMD)
  - Describe the application
    - Antecedents: Application has evolved over two decades on numerous platforms (Cray, SGI, etc.)
    - Programming languages used: Mainly FORTRAN 77/90; some C
    - Number of lines of code: Thousands
    - Workflow mechanism used to control code execution: Shell script
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs: 140 to 2280 cores
    - Accelerator usage: None
    - Amount of memory required: 2 GB per process
    - Describe the parallelism of the application, internodal and intranodal communication requirements: MPI
    - Run type: Batch

Current Use Case (Survey)

- Describe the data requirements for the application
  - Data flow: Application runs autonomously from input to output, generating output data at prescribed intervals
  - Data loading: Application loads data in parallel, all at the start of the run
  - Data generated per run: Input <= 100 MB per core; output <= 2 GB per core
  - Data generated per year: Input <> 5 TB total; output <> 100 TB total
  - Output data analysis: Quantitative and graphical analysis and visualization are all used to obtain insight
  - Application storage needed: High-speed shared file systems for runs; archival file systems for preservation
  - Input data: Some runs use observational data that must be cleaned up and regridded prior to use
  - Output data curating (archiving): 10 to 20 TB per year
  - Output data distribution: 5 to 10 TB per year, to all application users
- What are the key computing obstacles to effective execution of the current application?
  - Limited availability of sufficient numbers of processor cores and sufficiently long run times in queues
- Are there any analytic functions performed on the output of the application?
  - Yes; they make no substantial contribution to computation time
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - No, not at this time; skilled programmers would be needed to transition to GPUs, for example

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals?
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity: New, compute-intensive physics components are in final testing now
    - Increase in resolution and/or duration: Resolving the lower solar atmosphere is numerically highly demanding and requires many more time steps to execute the runs
    - Rewrite to improve performance: Paradigm-changing speedups (100X) are needed for real-time applications such as space weather forecasting
  - If possible, describe how the computational requirements will evolve over time
    - Size of runs: Simulations using thousands of cores should become routine
    - Accelerator usage: Requires major application rewrites
    - Amount of memory required: Application is processor-limited, not memory-limited
    - Describe the parallelism of the application, internodal and intranodal communication requirements: MPI
    - Storage requirements: Tens of TB per year should become routine
    - Run mode: Batch for science investigations; near-real time for future forecasting applications

Use Case: Mission Need

- Mission: **NASA missions need simulations to guide interpretations and analyses of data to achieve physical understanding, and to develop predictive capabilities such as space weather forecasts. My team performs simulations ranging from global simulations for the Earth, Moon, and asteroid Psyche, to support the MMS (primary mission support), RBSP (radiation belt), THEMIS (nightside magnetosphere dynamics), ARTEMIS (Moon), and Psyche (Discovery mission to the metallic asteroid 16 Psyche) missions, for example. Our lunar simulations supports the Human Exploration Directorate, in addition to SMD.**
- Criticality: If NASA supercomputing was no longer available to meet your needs, how would your mission perform this work?
  We would be forced to obtain supercomputing time from NSF and DOE and compete with other fields, and hence productivity and efficiency of mission support will be significantly reduced.

. Mission Security Needs: (Confidentiality, Integrity, Availability)
  - [What is impact of inadvertent release of data? (i.e., public confusion, professional embarrassment, National Security, Privacy, etc?)] **Low impact**
  - [If HEC availability is interrupted regularly, how does that affect your ability to achieve objectives?] **Major disruption of workflow and scheduling of projects**
  - [What is mission need for external collaborators (without NASA credentials) to access system or to run jobs?] **easy mechanism for sharing simulation output would be highly useful**
- HEC Function: [Briefly describe how HEC supports performance of this work]
  - [How often do you need to run the application to support mission objectives] **Weekly**
  - [How is the output used to achieve mission objective?] **The output enables interpretations and analyses of mission data to achieve physical understanding, and to develop predictive capabilities such as space weather forecasts.**
- Impact of HEC: [How essential is NASA's HEC to accomplishing this work]
  - [If NASA supercomputing was no longer available to meet your needs, how would you perform this work? (i.e., buy a system, rent time on a commercial facility, use Commercial Cloud Computing?)] **HEC is essential to accomplishing our mission support and discoveries of new physics. Our large-scale computation requires supercomputers. Only major agencies have these facilities (NASA, NSF, DOE).**
- Mission Concept: [Provide a graphic if possible, depicting this application or use case] **See next page**

Global and Local Simulations to Support the MMS Mission

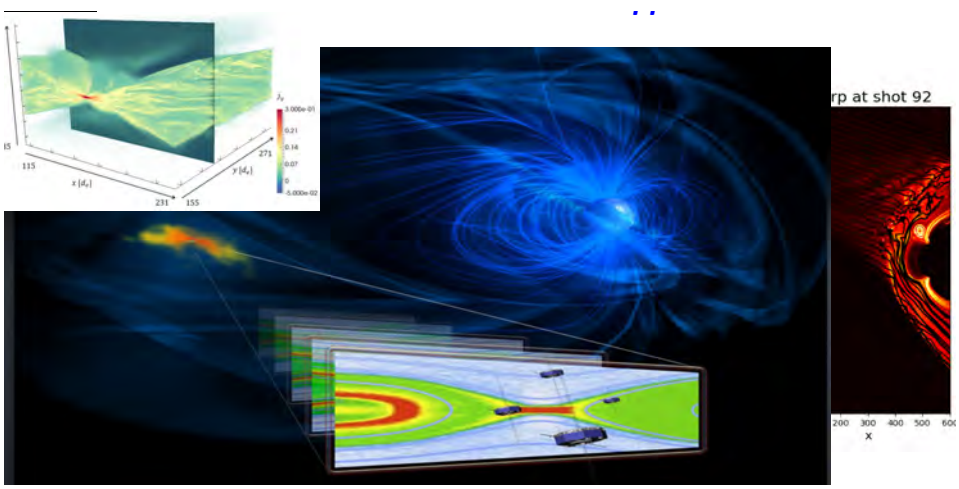

Use Case: Current HEC Application Characteristics

- Work flow: [Describe high level steps in the work flow for accomplishing the use case]
  - What triggers this workflow? **Mission discoveries, preparation for new science phases or mission to be launched.**
  - Is there a need for interaction by user, operator or other system in the computing workflow? Yes.
  - Does this use case need quick reaction or near-real time computing? Explain briefly. No.
- Data:
  - In: [Summarize Data used to feed into the code/application in terms of volume, latency, source] Simulation parameters, source codes, etc.; very low volume
  - Out: [Summarize data produced by computation and what happens to it] **simulation output in various forms, up to 10s of TB per run: time evolution of various quantities in the simulations**
    - [Does the output data require retention? How is that selected?] **YES: retention is based on expectation of what is important in the run and how it will be used. It is not uncommon to come back later (even a few years later) and to process the simulation in a different way or to even look for new physics. As such, we try to save as much a possible.**
      - [How is the data analyzed to provide mission results and where? (statistical or AI analysis, visualization?)] **Visualization, statistical analysis, experimenting with AI.**
  - Transfer:
    - [Where does the output go after it is generated? (i.e., local to HEC storage system or to another location)] **HPSS (e.g. Lou) + selected pieces go to local workstations easy access and post-porcessing**
- Computation:
  - Type: [Summarize the type of computing used/needed and how HEC supports this today] **General computing with almost the entire supercomputer (most powerful ones at HEC (Skylake, Aitken). Our codes are highly parallelizable, so we need large-scale computation.**
    - [Is this capacity computing or are special HEC capabilities needed for this computation?]
- Evaluation: [Summarize how you evaluate the effectiveness of HEC in supporting this work] **very effective for small-to-mid scale jobs. Ineffective for large-scale jobs (10's of thousands of cores).**

HEC Needs Workshop Key Questions

- What are the **key factors today** that affect the load on computing assets and cause it to vary?
  - [Such as use of large-scale, non-linear models, science data reprocessing, field campaigns, publication cycles, revision to major models, etc.] **Large-scale non-linear models, mission proposal cycles, publication cycles, revision of major models, for example.**
- Are there key functional gaps exist or are projected in the next 5 years or 10 years?
  - Constraints on Use Case: [What, if any, limits are imposed by the current assets, including capacity, turn around time, external access, processor types?] **Remote visualization is currently possible, but difficult to use. This will only become more important as simulation sizes continue to grow. As the mission needs grow, HEC will need to grow at the same pace to provide the necessary amount of CPU-hrs in 5-10 years. We expect a growth factor at least 100. current Capacity is limited to a few tens of thousand cores.**
  - Impact: [How does this affect fulfilling the mission need?]**. This limits the type of science that could be done at HEC, excluding critical problems such as self-consistently model the Earth's magnetosphere including the radiation belt.**
  - External Sources: [Does your project run on a non-NASA computing facility? If so, which and why?] **Yes, NERSC (DOE), Summit (DOE), and Frontera (NSF). These are large-scale computational facilities. Summit and Frontera are larger than HEC by orders of magnitude. When the mission needs support from large-scale computation runs, we have to apply allocations at NSF and DOE. Is this what the Agency wants? NASA missions have to rely on NSF/DOE supercomputers.**
- How will this Use Case evolve over the next decade?
  - Maturity: [Will it move from a development state to an operational capability with continued development?] **The codes are operational and continually being further developed for higher/broader capacities.**
  - Growth: [Will it increase resolution or time-steps to support evolving future mission needs?] **Yes!**
  - Re-use: [Will the number of people running this code/application grow?] **Yes!**
  - Re-hosting: [What is needed to migrate code to future computing architectures?] **The codes are portable. We are developing GPU version for one of them (VPIC).**

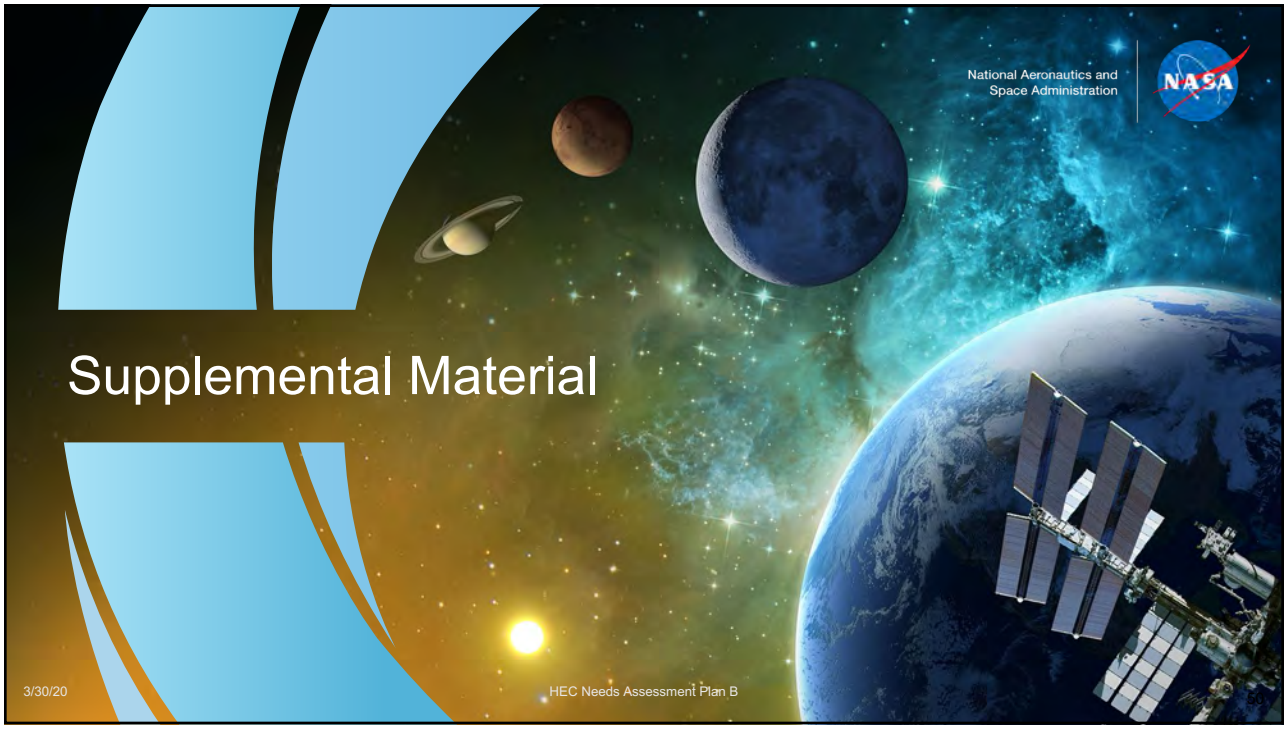

Current Use Case (Survey)

- Achieving the mission objective
  - Describe the type of computation it performs **several codes performing kinetic plasma simulations. particle-in-cell (VPIC), and hybrid (ion PIC, electron fluid) HYPERS is a unique asynchronous scalable compile-time configurable 2D/3D hybrid code that performs particle, fluid and EM field calculations on MPI-enabled parallel Linux systems.**
  - Describe the application
    - Antecedents: does this application evolve from work on existing platforms and which ones **The code VPIC has been run on many other systems (eg TACC Frontera, NERSC Cori, LANL Jaguar, Grizzly).**
    - Programming languages used (If multiple languages used, provide estimate percentage of the code in each languages **VPIC : 100% C/C++. HYPERS: C++ (50%), Fortran 77 (45%), perl (5%)**
    - Number of lines of code **VPIC: >100000**
    - Workflow mechanism used to control code execution (perl, jupyter, shell script, etc.) **Shell scripts, perl**
  - List the functional computational and storage requirements for the application (consider these for operational or science runs, not debugging or testing)
    - Size of runs (nodes or cores) **VPIC and HYPERS 3D runs > 50k-100k cores, 2D runs 2k-10k cores.**
    - Accelerator usage (type and memory requirement) **Currently uses v8 and v16 vectorization. GPU version is in development**
    - Amount of memory required (overall or per process) **Roughly 2-4 GB per core**
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other **VPIC: SIMD + OpenMP(or pthreads) + MPI .**
    - Storage is covered below
    - Run in batch, interactive, or near-real time. **VPIC & HYPERS: typically batch. Postprocessing runs in batch/interactive modes**

Current Use Case (Survey)

- Describe the data requirements for the application
  - Describe the data flow of the application from input to output **Simulation parameters -> raw output (N-to-N) -> postprocessing to combine files (N to 1) -> analysis and visualization.**
  - How are the data loaded in the application (serial, parallel, all at start, piecemeal for different elements, other) **the only loading happens at restart, This is N-to-N , start only, partially serialized read (a few 100's of processes read at a time)**
  - How much input and output are generated per run **varies. From 10s of GB to 10s and 100s of TB**
  - Summary of the total amount of input and output data per year **This year, output was about 900 TB (VPIC +HYPERS), input was negligible.**
- How does the output get analyzed (machine learning, statistics, visualization, other) to produce insight into the phenomena? **Output is mainly analysed using visualization and spectral analysis in python, matlab or IDL. 3D simulations require Paraview run on multiple nodes in the interactive queue.**
  - Describe the type of storage that this application needs, such as high-speed shared file systems, local file systems, object storage, other **We use storage on nobackup and lou. High-speed i/o is necessary for the large simulations (eg for the 3D run, it took an hour to write the restart files, and there were sometimes crashes during testing on electra/aitken).**
  - How is input data brought together for use? Does it require clean-up, regridding, etc. before being used. **NA**
  - How much data needs to be curated (archived)? **Order of hundreds of TB for 3D simulations (we perform single digits of these per year)**
  - How much data needs to be distributed and, generally, to whom? **Within team members within HEC.**
- What are the key computing obstacles to effective execution of the current application? **mainly computational time and available space in the file system**
- Are there any analytic functions performed on the output of the application?
  - Timing, computing load, workflow **Timing diagnostics are written during the run and can be compared to the ideal performance of the particle pusher (roughly 15-20 million pushes/cpu/sec)**
- Do you have available the workforce necessary to re-host your code on any new platform necessary?
  - Describe any skills or capacity needed. **PARTIALLY. VPIC is supported by LANL. Other codes may require re-working, depending on the type of the new system.**

Future of the Current Use Case (Survey)

- Describe the computing needs of the application as it evolves over the next 5-10 years
  - What are the key computing obstacles that could effect these goals? **computational resources required to perform simulations with realistic parameters, and the difficulty in analysis of the large data sets produced.** Other: sporadic node failures in big runs (when using ~the entire machine) and slow/unreliable parallel MPI input/output.
  - Describe how the application and workflow is expected to evolve over that time frame
    - Increase in complexity through the introduction of new components (i.e., physics, chemistry, etc.) **New electron physics (fast scales) and chemistry (charge exchange).**
    - Increase in resolution and/or timesteps **New mesh refinement options (stretched and block-adaptive mesh).**
    - Rewrite to improve performance **New programming paradigm: hybrid MPI and OpenMP execution**
    - Refactor for emerging platforms **GPU optimization**
    - Rehost onto cloud **NA**
    - Introduction of Artificial Intelligence (AI), Machine Learning (ML), and/or Deep Learning (DL) components **Currently experimenting using these for post-processing and analysis. HYPERS uses AI. We hope to use AI/ML for developing predictive capabilities using the results of our simulations (for instance providing parameters for global MHD codes).**
    - If possible, describe how the computational requirements will evolve over time **MOSTLY INCREASE IN RESOLUTION and SIZE of the SIMULATION. EXISTING TECHNICAL LIMITS FAR EXCEED PRACTICAL LIMITS at HEC. Hope HEC can grow by at least 100 times to bring simulations closer to predictive capabilities.**
    - Size of runs (nodes or cores) **Runs are still a few orders of magnitude away from realistic parameters, hence we will use the maximum resources available.**
    - Accelerator usage (type and memory requirement) **NA**
    - Amount of memory required (overall or per process or per core) **Memory per core remains the same**
    - Describe the parallelism of the application, internodal and intranodal communication requirements, i.e., does the application use MPI, Shared Memory, other **continue to use MPI and threads with shared memory**
    - Storage requirements **same**
    - Run in batch, interactive, or near-real time **Batch**

Need for Modern HPC Hardware to Solve NASA Challenge Problems

DRAFT

Mujeeb R. Malik, Senior Aerodynamicist (ST)
Eric Nielsen, Research Scientist (HPC SME)
David M. Schuster, NESC Technical Fellow for Aerosciences (ST)
Dennis M. Bushnell, Chief Scientist (ST)
NASA Langley Research Center

LAVA OVERFLOW

55

Outline

- The Need:
  - Critical NASA Mission aerosciences challenges can't be computed/solved in a timely fashion due to lack of sufficient machine capability
- The Solution
  - Modest investment to acquire a GPU-based machine
  - Periodic assessment/updates, as HPC hardware continues to evolve
- HPC Landscape and NASA's Standing
  - *NASA compute capability lags by 2 orders compared to national and international peers*
- Summary

56

Top NASA Aerosciences Challenges Require Emerging Exascale Compute Capability

- Four Aerosciences areas overwhelmingly challenge the prediction of environments and performance for flight vehicles

1. **Aero-plume Interaction**
   - Airframe/propulsion interaction for aircraft efficiency, emissions, and acoustics.
   - Capsule Backshell Reaction Control Systems in separated flow.
   - Plume-forward systems: Orion Launch Abort System, propulsive entry/descent.
   - Plume-surface interaction for landing.
2. **Unsteady Separated Flows**
   - Aeroacoustics and buffet for both aviation and space applications.
   - Aircraft high lift systems.
   - Space vehicle protuberances, base flows, vehicle backshell flows, parachutes, and decelerators.
   - CFD state-of-the-art does not yet address this problem accurately, reliably, or efficiently.
3. **Aerothermal Predictions**
   - Coupled aero-material ablator performance.
   - Boundary layer transition for sustained hypersonic flight.
   - CFD for aero-heating still sensitive to algorithm and geometry meshing.
4. **Decelerator System Performance**
   - Multi-physics predictive capability for parachutes, inflatable, and propulsive decelerators.

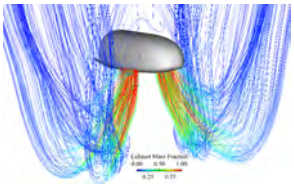

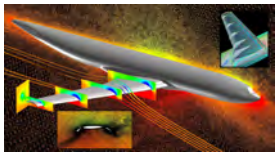

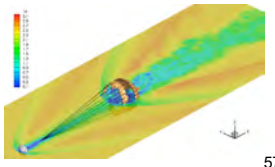

- Requires weeks/months for critical mission computations; some leading edge computations simply not tried because of lack of compute resources
- Critical need for investment in new computing hardware

57

CFD Technical Challenge

- **TTT TACP06:** "Develop and demonstrate computationally efficient, eddy-resolving modeling tools that predict maximum lift coefficient (CL_{max}) for transport aircraft with the same accuracy as certification flight tests" (Level 1 milestone due 9/30/2025).
  - High-lift certification flight tests account for roughly 2/3 of flight certification test points
  - 50% of flight tests replaced by computations would save about $300 million from each aircraft development program
- Promise to expand capability of CFD to predict unsteady phenomena such as stall, buffet and flutter
- Eddy-resolving methods require an order of magnitude more computational resources than RANS

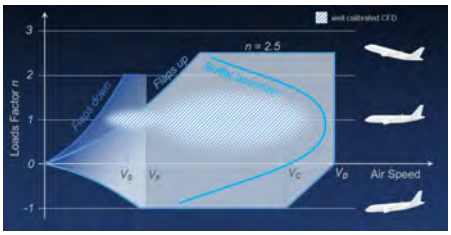

- CFD has been calibrated only in **relatively small regions of the operating envelope** where the external flow is well modeled by current RANS methods
  - High-speed cruise (aero design)
  - Low-speed at nominal attitude with moderate flap settings

Euler Equations (inviscid) → *Reynolds-Averaged Navier-Stokes: "RANS" (turbulence modeled)* → *Large-Eddy Simulations: "LES" (large scales resolved)* → *Direct Numerical Simulations: "DNS" (all scales resolved)*

Increasing physics, increasing cost →

58

TTT/RCA Developed Technology Leveraged for Human Space Flight

- **Building Booster Separation Aerodynamic Databases for Artemis II**
  - "CFD simulations made using FUN3D flow solver to formulate the booster separation databases for Artemis II. The CFD calculations include simulating the aerodynamic effects of 22 different plumes firing during the separation, counting the four core stage main engines, the two boosters, all 16 booster separations motors"
  - "These databases will be used by the Guidance, Navigation, and Control (GNC) group at NASA's Marshall Space Flight Center to model the booster separation event for Artemis II. The GNC team will incorporate the aerodynamic data into their dynamic simulations, which will be used to ensure that the boosters can separate successfully without recontacting the core under all possible flight conditions. Ultimately, the results of those simulations will be critical in reducing the risk to the crew of Artemis II during booster separation."

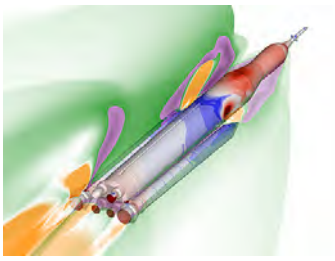

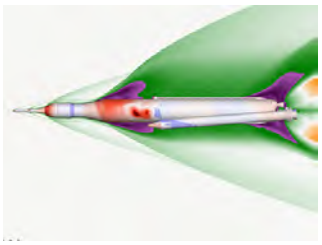

Large database generation task could be reduced from months to days with new computing hardware

https://www.nas.nasa.gov/SC19/demos/demo27.html

HPC in the CFD Vision 2030 Roadmap

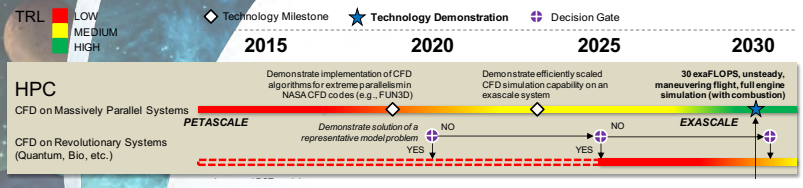

- *"Current simulation software must be ported to evolving and emerging HPC architectures with a view toward efficiency and software maintainability."*
  - *FUN3D and other research DNS/LES codes ported to GPUs*
    - *Ready to be used for aero database generation, and large scale simulations to unravel flow physics*
- *"Investments must be made in the development of new algorithms, discretizations, and solvers that are well suited for .... anticipated in future HPC architectures."*
  - *Developments being sponsored by TTT/RCA*
- *"Increased access to the latest large-scale computer hardware must be provided and maintained, not only for production runs, but also for algorithmic research and software development projects…"*
  - *This is a current, urgent, unfulfilled need*

60

Need for Modern HPC Hardware to Solve NASA Challenge Problems

- Some NASA missions/codes are ready to exploit the power of GPUs
  - FUN3D has shown significant performance gains, now ready for production (database) runs (SLS, Mars entry, ARMD applications) with much faster turnaround than currently possible
  - Inhouse research DNS codes, with 30x speed up, ready for scale resolving simulations of turbulence for smooth body flow separation
  - GPU version of wall-modeled LES code Charles (Cascade Technologies) yields an order of magnitude faster turnaround time
- Proposal for new investment for both capacity and capability computing

61

Options to get access to state of the art capabilities for NASA

- **Purchase NVIDIA GPU-based hardware (similar to DOE's Summit)**
  - Cost for 512 nodes (4 GPUs/node) machine: $48 million [Note: Summit has 4608 nodes with 6 GPUs/node]
  - To be housed at NAS (ARC)
- **Compute landscape is changing fast**
  - Will need a refresh after 3 years; assess at that time which new hardware will be best to acquire
    - Informed by decisions/experience of DOE with newer hardware
    - And, a 3 year effort led by LaRC to evaluate new emerging hardware and porting selected NASA codes to the same (budget request of $3 million/year for 3 years)
  - The above machine will continue to be used for production runs for another 3-4 years

62

Top 10 Supercomputers

From Nov 2019 Top 500 list

HPL Rank	System Name	System Architecture	Site	Cores (M)	HPL R_{max} (PF)	R_{peak} (PF)	HPL R_{max} / R_{peak}	Power (MW)	GF/W	HPCG R_{max} (PF)	HPCG R_{max} / R_{peak}	HPCG Rank
1	Summit	IBM AC922: IBM P9 (22 cores) + NVIDIA Volta + dual-rail IB-4X EDR interconnect	USA (DOE/ORNL)	2.4	148.6	200.8	74%	10	14.7	2.93	1.5%	1
2	Sierra	IBM S922LC: IBM P9 (22 cores) + NVIDIA Volta + dual-rail IB-4X EDR interconnect	USA (DOE/LLNL)	1.6	94.6	125.7	75%	7	12.7	1.80	1.4%	2
3	TaihuLight	Sunway (260 cores) + Chinese interconnect	China (Wuxi)	10.6	93.0	125.4	74%	15	6.1	0.48	0.4%	8
4	Tianhe-2A	Intel IvyBridge (12 cores) + NUDT Matrix-2000 + Chinese interconnect	China (Guangzhou)	5.0	61.4	100.7	61%	18	3.3	0.58	0.6%	4
5	Frontera	Dell: Intel Cascade Lake (28 cores) + IB-4X HDR	USA (NSF/TACC)	0.4	23.5	38.7	61%	6	3.9	0.31	0.8%	
6	Piz Daint	Cray XC50: Intel Haswell (12 cores) + NVIDIA Pascal + Cray Aries interconnect	Switzerland	0.4	21.2	27.2	78%	2	8.9	0.50	1.8%	7
7	Trinity	Cray XC40: Intel Haswell (16 cores) + Intel KNL (68 cores) + Cray Aries interconnect	USA (LANL/SNL)	1.0	20.2	41.5	49%	8	2.7	0.55	1.3%	5
8	AI Bridging Cloud	Fujitsu: Intel Skylake (20 cores) + NVIDIA Volta + IB-4X EDR interconnect	Japan	0.4	19.9	32.6	61%	2	12.1	0.51	1.6%	6
9	SuperMUC-NG	Lenovo: Intel Skylake (24 cores) + Intel Omni-Path interconnect	Germany	0.3	19.5	26.9	72%	4	5.2	0.21	0.8%	
10	Lassen	IBM S922LC: IBM P9 (22 cores) + NVIDIA Volta + dual-rail IB-4X EDR interconnect	USA (DOE/LLNL)	0.3	18.2	23.0	79%	1	15.1	0.35	1.5%	

NASA's position continues to slip in the HPC landscape

#32 4.4	Pleiades	SGI ICE X/Intel Xeon	NAS	0.24	5.95	7.1
#39 1.7	Electra	SGI ICE X/Xeon	NAS	0.12	5.44	8.3

63

DOE Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

- DOE national labs have a large number of machines to support current computing needs of the DOE mission
- In addition, a 7 years $4 Billion Exascale Computing Project (ECP) launched in 2016 to acquire exascale computing systems (includes $1.8 Billion for R&D)
  - 3 focus areas: Hardware & Integration, Software Technology, Application Development
  - Supported by 80+ R&D teams, 1000 researchers
  - Co-design centers/teams help to ensure that applications would effectively utilize exascale systems

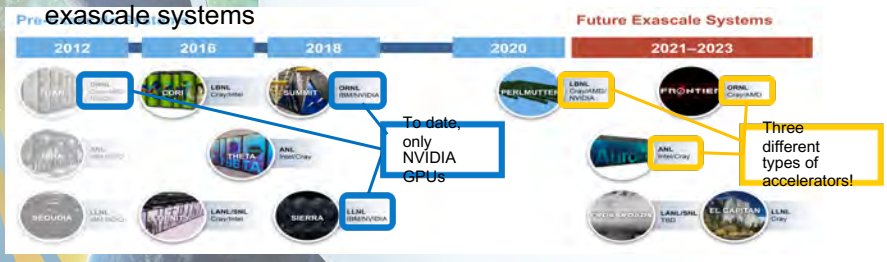

Doug Kothe, ECP Director, at HPC Forum 360 organized by AIAA CFD Vision 2030 Integration Committee, SciTech 2020

Potential Partnership with Other Government Agencies

Can we get compute resources from DOE?

Department of Energy is the only Government Agency that provides access to others

It has 3 standard means by which access is provided:

- Director's discretionary allocation
  - Relatively small allocation (1-4 million CPU hours) for code porting/get scaling data for proposal preparation for the two competitive programs below
- ALCC Program
  - Annual proposal based allocation, designed to facilitate partnerships that benefit DOE mission https://science.osti.gov/ascr/Facilities/Accessing-ASCR-Facilities/ALCC
    - Was awarded about 66 million and 84 million cpu hours in two successive years
    - However, proposals were not selected in past two years even though 'strongly recommended' by the reviewers. The reason given was that selection was very competitive (only 37 out of 75 proposals were provided allocation)
    - It is also difficult to find relevance of our project interests to the DOE mission areas
- INCITE Program
  - Proposal-based allocation program, which is the largest publicly available resource in the world, but with very stiff competition. https://www.alcf.anl.gov/news/us-department-energy-s-incite-program-seeks-proposals-2020
    - Requires that proposers use at least 20% of the machine. For Summit (the current #1 on the Top 500 supercomputers list), this means 920 GPU nodes (equivalent to 1 million cores). So, this is for one of a kind simulation, not for large number of small simulations
    - Two proposals were funded in 2019, one using FUN3D and another using a high order method for transition simulation

DOE also allocates compute time under the CAAR program (https://www.olcf.ornl.gov/caar/) to ready software applications for their next-generation supercomputers

- An application using FUN3D, in support of RCA's high-lift challenge, was rejected because it proposed multiple

DOE resource allocation is very competitive, a hit or miss process, and these resources cannot be in any project's critical path

65

NASA is No More the Leader in HPC

NASA capability lags 2 orders of magnitude, compared to national and international peers

Clearly, "we are being left in the dust" with major mission cost, schedule and capability impacts

- The huge workload on NASA machines necessitates weeks/months delays; in some cases reached "not worth trying for" status

On a path to quickly lose in-house leading edge capability in modeling and simulation

NAS High End Computing (HEC) is an extremely well-managed system, but they don't have sufficient resources to explore/provide the emerging HPC hardware

- HEC annual budget is $45 M, which allows only $15 M for hardware upgrade

A major push toward HPC needed at the Agency level (as already

66

The Impacts

Progress toward the following goals could be significantly accelerated with additional machine capabilities:

- Certification by computation, required for mission success going forward, by cost savings and reducing development cycle
- Computation of Turbulence, required for effective advanced concept evaluation and optimization to achieve reduced mission costs, much improved capabilities
- Incorporation of AI/Big Data/ deep learning/ autonomy, increasingly required for cost and capability mission metrics
- Efficient and safe Mars (human) landing
- Efficient sustained hypersonic flight
- Development of the enabling autonomous ATC system for millions of EVTOL/AAS/ UAS/UAM/ODM, PAV aircraft → A new $ 1T/yr aero market

67

Summary

- High Performance Computing capability has a major impact on NASA missions
- Propose that agency budget $57 million over 3 years ($51M, $3M, $3M) to acquire a new GPU-based machine and to continue assessment of new emerging hardware
- Recommend internal periodic review of the HPC needs and frequent leading-edge machine updates going forward
- Conduct an NRC study to assess NASA HPC needs and make recommendations, for the long run

68

National Aeronautics and Space Administration

Use Case: Machine Learning for Aviation Safety

Nikunj Oza
July, 2020
NASA Ames Research Center (Code TI)
ARMD/AOSP

3/30/20 HEC Needs Assessment Plan B

AI/ML Success Example: Machine Learning for Aeronautics

- Industry mostly uses rules over 2-3 variables to find safety issues
- Can only find known safety issues. Want to find unknown problems
- Developing data-driven methods for
  - Anomaly detection
  - Precursor identification: Finding precursors to particular safety issues
- Using data from aircraft operations (Flight Operations Quality Assurance—-FOQA), trajectory data on takeoffs and landings at airports and metroplexes

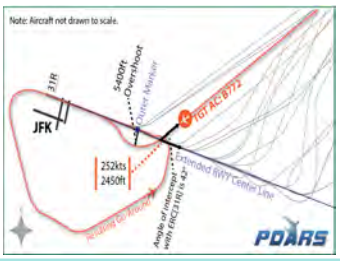

Technical Challenges (why ML was needed):
- Substantial time series data (150+ variables, many thousands of flights daily)
- Very few operationally significant anomalies---low false alarm rate difficult to achieve, but necessary

Approach (how ML was employed):
- Developed new algorithms for anomaly detection and precursor identification over heterogeneous data
- Developed active learning algorithm to learn from domain expert feedback to reduce false alarm rate
- Working with MITRE to deploy algorithms for use with 50+ airlines as part of Aviation Safety Information Analysis and Sharing (ASIAS) partnership

Principal Investigator: Nikunj C. Oza, PhD, ARC

Quantifiable Benefits:
- Have identified several anomalies deemed safety relevant and not identified by current rule-based methods. Used to modify rule-based method at one airline.
- Working with four airlines and MITRE to deploy algorithms for operational use

Lessons Learned, Best Practices:
- Allow substantial time for simple data analysis and cleaning before applying algorithms
- Existing commercial and open-source tools are almost never sufficient---machine learning algorithm R&D is required for application
- Machine Learning expertise and domain expertise are needed for both data preprocessing and results interpretation

https://ntrs.nasa.gov/search.jsp?R=20190032265 2020-07-31T16:34:55+00:00Z

High-End Computing Overview

William W Thigpen
Manager, High-End Computing Capability Project
william.w.thigpen@nasa.gov

NASA's HEC Requirements: Capacity

HEOMD (engineering-related work) requires HEC resources that can handle large numbers of relatively-low CPU-count jobs with quick turnaround times.

Over 1500 simulations utilized ~2 million processor hours to study launch abort systems on the next-generation crew transport vehicle.

The formation of vortex filaments and their roll-up into a single, prominent vortex at each tip on a Gulfstream aircraft.

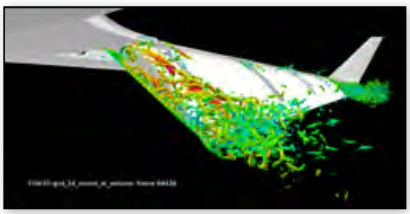

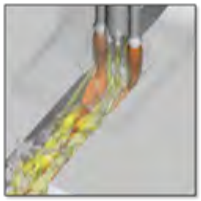

Over 4 million hours were used over a four-month project to evaluate future designed of the next generation launch complex at Kennedy Space Center.

NASA's HEC Requirements: Capability

ARMD and SMD (aeronautics and science related work) require HEC resources that can handle high fidelity relatively-large CPU-count jobs with minimal time-to-solution. Capability enables work that wasn't possible on previous architectures.

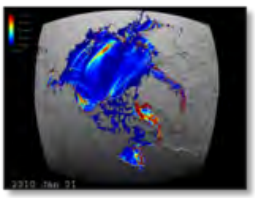

NASA is looking at the oceans, running 100's of jobs on Pleiades using up to 10,000 processors. Looking at the role of the oceans in the global carbon cycle is enabled by access to large processing and storage assets.

For the first time, the Figure-of-Merit has been predicted within experimental error for the V22 Osprey and Black Hawk helicopter rotors in hover, over a wide range of flow conditions.

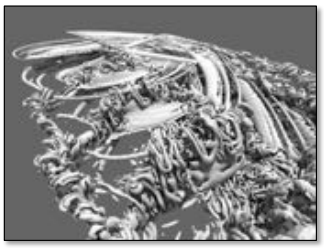

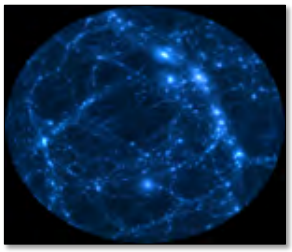

To complete the Bolshoi simulation, which traces how the largest galaxies and galaxy structures in the universe were formed billions of years ago, astrophysicists ran their code for 18 straight days, consuming millions of hours of computer time, and generating massive amounts of data

NASA's HEC Requirements: Time Critical

NASA also needs HEC resources that can handle time-sensitive, mission-critical applications on demand (maintain readiness).

HECC enables enormous planetary transit searches to be completed in less than a day, as opposed to more than a month on the Kepler SOC systems, with significantly improved accuracy and effectiveness of the software pipeline.

REENTRY

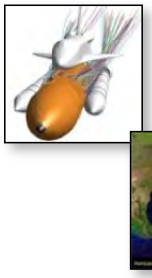

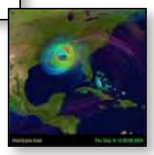

STORM PREDICTION

UAVSAR produces polarimetric (PolSAR) and interferometric (repeat-pass InSAR) data that highlight different features and show changes in the Earth over time.

HECC Mission Statement

Ensure the ability of NASA to meet its computing, computational, and data analytic requirements for science and engineering by identifying and preparing for emerging information technologies and by providing access to high-end computing systems inside and outside of the Agency together with services to maximize productivity. Through internal and external partnerships, advance and develop fundamentally new approaches in high-end computing; this includes R&D in hardware and hardware subsystems, software, architectures, system performance, computational algorithms, data analytics, development tools, and software methods for extreme data- and compute-intensive workloads.

Our mission is guided by the vision that:

NASA's HECC resources are relied on as an essential and pervasive partner by the breadth of Agency science, engineering, and technology activities, enabling rapid advances in insight and dramatically enhancing mission achievements.

HECC Goals

- Identify, research, and develop well-suited emerging computing and computational technologies that will enable NASA science and engineering.
- Establish a roadmap for including the new technologies in HECC's research and production environments.
- Work across NASA mission directorates, academia, and industry to research and develop algorithms that provide better insight and/or more rapid time to solution by exploiting emerging computational and data analytics technologies.
- Provide access to effective pathfinding testbeds employing emerging computing technologies that enable teams to implement computational algorithms for NASA critical science and engineering requirements.
- Infuse HEC into NASA's scientific and engineering communities.
- Assure preparedness to meet NASA's future modeling, simulation, and analysis needs.
- Ensure that NASA HECC resources and activities are well-managed and wisely used.

HECC Objectives

- **Objective 1, Service Excellence**: Provide excellent High-performance computing (HPC) services to NASA customers and users, enabling pervasive, timely, and successful significant mission outcomes through continuous improvement.
  - Provide the agreed annual HPC allocation to NASA Mission Directorates.
  - Provide demonstrable improvements in products and services to NASA Mission Directorates.
  - Gather and analyze customer requirements on a regular basis (via surveys, workshops, etc.).
  - Support reviews of the HECC capital refresh process.
- **Objective 2, User Development**: Infuse effective HPC usage practices and knowledge into NASA's scientific and engineering community.
  - Maintain a robust, multi-faceted interface to help users increase productivity.
  - Help the users on mapping their problems to the latest computing architecture.

HECC Objectives (cont.)

- **Objective 3, Future Viability:** Ensure that scientists and engineers are able to exploit emerging information technologies to solve NASA's most challenging problems.
  - Identify and research emerging computing, computational, and data analytic technologies and establish a roadmap for including them in HECC's research and production environments.
  - Develop algorithms that exploit emerging computing, computational, and data analytic technologies.
  - Provide access to emerging information technology testbeds.
- **Objective 4, Prudent Management:** Manage HECC in a responsible, effective, equitable, and cost-conscious manner.
  - Comply with federal and NASA requirements for Portfolio management.
  - Report routinely on portfolio status and performance.

HECC Enables or Directly Contributes to NASA's Four Strategic Goals

NASA Strategic Goal 1: Expand Human Knowledge Through New Scientific Discoveries

- Modeling the Earth and assimilating observational data to shorten the time from observations to answers for important, leading-edge science questions.
- Analyzing massive amounts of Earth-observing satellite data at the global scale, extracting information and turning it into knowledge and wisdom.
- Modeling the solar environment to better understand the causes of space weather as it affects the Earth and other planets in the solar system.
- Modeling physical regimes ranging from the solar system to the universe, to help understand observations.
- Developing and refining theories of the evolution of the universe using computational modeling.
- Performing simulations supporting understanding of the space environment's effects on astronaut health.
- Analyzing, distributing, and enabling exploration of data and data products captured or produced from solar and space satellites.

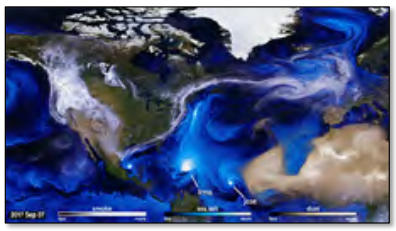

Sea salt particles get caught up in the swirling winds of hurricanes Irma and Jose on September 7, 2017. NASA captured the interaction of hurricanes and aerosols during the 2017 Atlantic hurricane season by combining satellite data with sophisticated models that describe the underlying physical processes. *William Putman, Anton S. Darmenov, NASA/Goddard; Matthew R. Radcliff, USRA; Aaron E. Lepsch, Ellen T. Gray, ADNET Systems, Inc.*

NASA Strategic Goal 2:
Extend Human Presence Deeper in Space and to the Moon for Sustainable, Long-term Exploration and Utilization

- Simulating NASA and commercial spacecraft systems and subsystems supporting design and operational scenario evaluation.
- Simulating complete life-sustaining environments to accelerate their planning, design, and evaluation for long-term human presence in space.
- Modeling and simulation of instrument behavior, performance, and impact.

Space Launch System Block I B booster flow field simulated using NAseparationSA's FUN3D code. The crewed vehicle features a different, more powerful second stage. *Jamie Meeroff, Henry Lee, NASA/Ames.*

NASA Strategic Goal 3:
Address National Challenges and Catalyze Economic Growth

- Conducting and fostering research in emerging information technologies that enable future discoveries.
- Establishing partnerships across NASA and with academia and technology companies to develop critical algorithms that exploit emerging information technologies.
- Engaging industry partners in the use of NASA modeling and simulation systems architected and operated for aerospace applications.
- Offering advanced HEC education programs and internships and reaching out to Science, Technology, Engineering, and Math (STEM) education and professional organizations through conference participation, supercomputing center tours, and presentations.

Visualization of the flow of NASA's modified design of a complete DJI Phantom 3 quadcopter configuration in hover. Simulations revealed the complex motions of air due to interactions between the rotors and the airframe. *Patricia Ventura Diaz, Tim Sandstrom, NASA/Ames*

Strategic Goal 4: Optimize Capabilities and Operations

- Providing users with computing architectures well suited to their scientific and engineering workloads and delivering a full-service HEC offering.
- Providing users with access to pathfinding testbeds that enable development of algorithms for emerging architectures.
- Striking a balance between upgrading HEC technologies and minimizing impact on users.
- Maintaining a high level of system availability and providing uninterrupted access to computational resources and user data.
- Determining appropriate shares of the resources for each NASA Mission Directorate and closely tracking usage to ensure maximum productivity.

The Pleiades supercomputer's rack-based architecture allows NASA to continually increase the system's computing capability through hardware upgrades, without needing to expand its physical footprint. The current configuration is nearly 15 times more powerful than it was when the system was originally installed in 2008.

Meeting Today's Requirements

HECC Hardware Assets

5 Compute Clusters

- Pleiades 158 ½ Racks / 11,215 nodes / 7.85 PF / 8,222 SBU/hr
- Electra 24 Racks / 3,422 nodes / 8.32 PF / 4,815 SBU/hr
- Aitken 4 E-Cells / 1,152 nodes / 3.69 PF / 2,290 SBU/hr
- Merope 56 ½ Racks / 1,792 nodes / 252 TF / 520 SBU/hr
- Endeavour 3 Racks / 2 nodes / 32 TF / 44 SBU/hr

1 Visualization Cluster 245 million pixel display / 128 node / 703 TF

10 Lustre File Systems 46.0 PB

6 NFS File Systems 1.5 PB

Archive System 1,000 PB

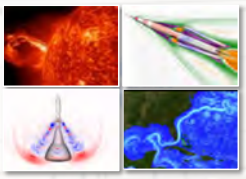

D-Wave 2000Q quantum system

- Whistler processor with 2031 qubits

NAS Facility Extension

- A one-acre site with 30 MW power to house HPC systems in modules.

HECC Services

HECC provides a suite of complimentary services to the user community to enhance the scientific and engineering results obtained from the hardware assets.

- Systems: Customized solutions including compute and storage solutions to meet specific project or mission requirements. Cloud access for immediate or non-standard computing.
- Application Performance and Productivity: Software solutions provided to research/engineering teams to better exploit installed systems.
- Visualization and Data Analysis: Custom visualization during traditional post-processing or concurrent during simulation to understand complex interactions of data.
- Networks: End-to-end network performance enhancements for user communities throughout the world.
- Data Analytics: Exploitation of data sets through neural nets and emerging new techniques.
- Machine Learning: Custom environments to enable learning through advanced data techniques.
- Custom Data Gateways: Custom data portals to support diverse programs and projects.

Preparing for the Future

- New algorithms to match hardware realities.
- Mixed precision and machine learning/AI in computation.
- New methods in visualization and data analytics.
- Application enhancements to exploit I/O advancements.
- Fault resilience.
- Performance optimization.

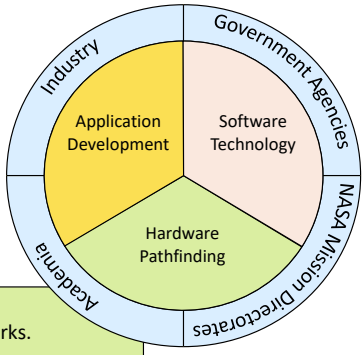

- Research and develop programming approaches, for accelerated and non-accelerated architectures.
- Optimize and port techniques/tools.
- Optimize and port techniques/tools for visualization, data analytics, AI, and machine learning.
- Research and develop efficient math libraries.
- New file system software technologies.
- Research and develop machine learning techniques for problem solving and system error tracking/prediction.
- Evaluate software packaging technology, such as containers.

- Examine requirements.
- Develop necessary benchmarks.
- Assess I/O approaches.
- Develop models to predict performance.
- Make recommendations for pathfinding systems.

Validating the Path

- **Metrics**
  - Research
  - Operational
  - Performance
  - Environmental Impact
- **Reviews**
  - Bi-Weekly Major IT CPIC Call
  - Monthly Internal HECC Review
  - Quarterly OMB Data Center Inventory Update
  - Annual Internal Budget Review
  - Annual PPBE Review
  - Annual Report
  - Annual ARC HECC Management Review
  - HECC Annual User's Meeting
  - Triennial Customer Meeting
  - Triennial External Review
- **Customer Feedback**
  - Annual User Survey
  - Triennial HECC Requirements Workshop

Quarterly Utilization Over Last 14+ Years

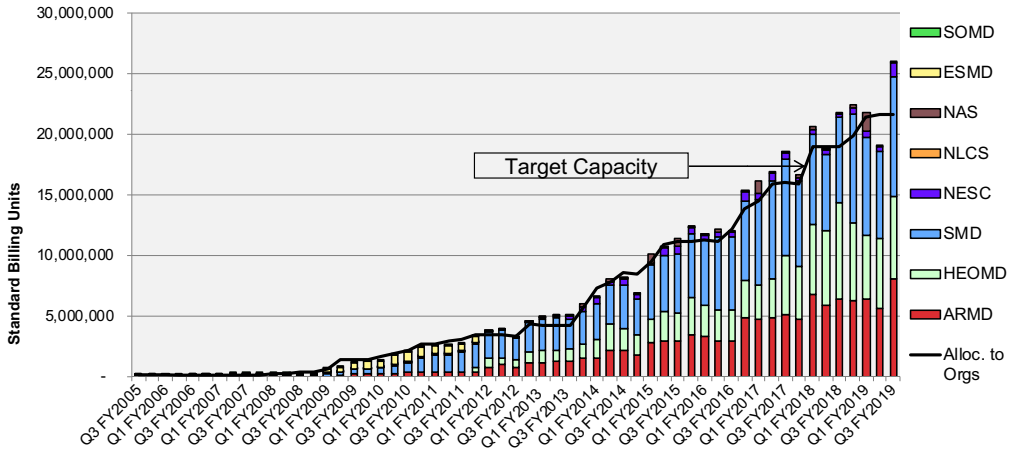

* 1 SBU (Standard Billing Unit) represents the work that can be performed on a dual-socket Broadwell node in one hour.

Aeronautics Support (25,490,311 SBUs)

Advanced Air Vehicles

of projects: 71
# of SBUs used*: 6,597,602

- HECC is used to develop concepts and technologies for dramatic improvements in the noise, emissions, and performance of transport aircraft.
- HECC is used to develop concepts and technologies to increase rotorcraft speed, range and payload, and decrease noise, vibration and emissions.
- HECC is used to develop advanced computer-based prediction methods for supersonic aircraft shape and performance and to develop technologies that will help eliminate today's technical barriers (such as sonic booms) to practical, commercial supersonic flight.
- HECC is used to develop computer-based tools and models and scientific knowledge that will lead to significant advances in our ability to understand and predict flight performance for a wide variety of air vehicles.

Transformative Aeronautics Concepts

of projects: 51
# of SBUs used*: 14,313,302

- HECC is used to develop and utilize Reynolds-averaged Navier-Stokes (RANS) and Large Eddy Simulation (LES) methods, and hybrid RANS-LES techniques to improve calculation methods for propulsion flows dominated by turbulent boundary layers and mixing.
- HECC is used to assess natural laminar flow concepts, to elucidate the physics and control of boundary layer transition in swept wing flows and drag reduction concepts for compressible boundary layers.
- HECC is used to validate chemistry, chemistry-turbulence and spray models being developed under the National Jet Fuels Combustion program.

Airspace Operations and Safety

of projects: 2
# of SBUs used*: 53,695

- HECC is used for developing reliable computational tools for predicting and analyzing stability & control characteristics of aircraft prior to or while encountering loss-of-control flight conditions characterized by abnormal flight , abnormal vehicle conditions, external upsets, and icing.
- HECC is used to develop methods for computing aerodynamic performance degradation associated with ice accretions on swept wing geometries.

Integrated Aviation Systems

of projects: 15
# of SBUs used*: 4,474.068

- HECC is used for accurate prediction of airframe noise from a full scale aircraft and evaluation of flap and landing gear noise reduction concepts in flight environments.
- HECC is used to develop technology for compact, high-power-density electric motors to power an all-electric general-aviation aircraft or helicopter, a hybrid turbine-electric regional airliner or a large transport with many small engines around the aircraft.

Aeronautics Evaluation and Test Capabilities

of projects: 3
# of SBUs used*: 51,644

- HECC is used to develop simulations of the contraction, test section with plenum & model support system, and high speed diffuser of the National Transonic Facility.
- HECC performs computational fluid dynamic (CFD) analyzes in support of the GRC 9'x15' Low Speed Wind Tunnel (LSWT) rake calibration effort.

*October 1, 2017 to September 30, 2018

Human Exploration and Operations & Safety Support (26,537,011 SBUs)

Multi-Purpose Crew Vehicle
of projects: 6
# of SBUs used*: 1,038,826

- HECC is used to support the creation of hundreds of computational solutions that model the flow field around the Crew Module and Launch Abort System for all flight regimes to be used as input for the aerodynamic databases.
- HECC is used to run computational fluid dynamics simulations to study the aerodynamic and aerothermal environments for the Multi-Purpose Crew Vehicle.
- HECC is used to develop and deploy a prototype system for rapid aerodynamic performance database generation and to use it on real-world problems faced by the Human Exploration and Operations mission directorate.

Space Launch Systems
of projects: 15
# of SBUs used*: 21,428,285

- HECC is used for computational fluid dynamics simulations of Space Launch Systems ascent to assess aerodynamic performance, protuberances, stage separation, and plume effects (such as plume-induced flow separation) for evolving vehicle designs.
- HECC is used for computational fluid dynamics analysis of Advanced Booster development efforts in the combustion stability areas.
- HECC is used for prediction of the launch induced environment for the Space Launch System including liftoff acoustics, ignition over-pressure, separation environments, debris, Launch Pad Abort Environments and hydrogen entrapment.
- HECC is used to simulate tanks and main propulsion system components (including feedlines, valves, manifolds, ducts, and pogo accumulators) for evaluation of criteria such as flow uniformity and component pressure drop.

HEOMD - Space Flight Operation & General
of projects: 41
# of SBUs used*: 4,069,899

- HECC is used to simulate the effect of larger solid rocket boosters and new propulsion systems on the launch facility at Kennedy Space Center, such as investigating whether ignition overpressure waves generated during liftoff are fully suppressed by the existing water suppression system.
- HECC is used to evaluate visiting-vehicle induced loads on the International Space Station (ISS) during mated and rendezvous operations and to evaluate crew Extra-vehicular Activity/Intra-vehicular Activity and attitude control loads on ISS.
- HECC is used in developing a combustion response model to investigate combustion instability in hydrocarbon-fueled rocket engines.
- HECC is used for technology development for entry, descent and landing systems.

NASA Engineering & Safety Center
of projects: 5
# of SBUs used*: 1,720,752

- HECC is used for simulations to provide guidance to the Space Launch System advanced booster designers by providing aerodynamic loading implications for various potential advanced booster geometric configurations.
- HECC is used to improve the capability to predict combustion stability in liquid rocket engines to increase NASA engineers' capability to more confidently and efficiently identify and mitigate combustion stability issues in engine development programs.
- HECC is used to used for studies of large eddy simulations of oblique-shock / supersonic hot jet interaction, aimed at prediction of plume-induced vibroacoustics.

*October 1, 2017 to September 30, 2018

Science Support (29,879,752 SBUs)

Astrophysics
of projects: 87
# of SBUs used*: 9,608,529

- HECC is used by the Kepler mission to find Earth-sized planets around other stars and to fully analyze the Kepler data to find any undiscovered planets still "hiding" in the data.
- HECC is used to understand the physics of high redshift galaxy formation and make detailed predictions that can be used to guide NASA observations of the first galaxies.
- HECC is used for quantifying the redistribution of matter in galaxies when supernova energy is deposited; exploring the growth of black holes and the impact of active galactic nuclei on galaxy evolution; and determining whether the ultraviolet light from stars in galaxies can "escape" to re-ionize the universe.

Earth Science
of projects: 86
# of SBUs used*: 8,953,439

- HECC is used to combine observational data with numerical simulations of the global ocean circulation to provide vital information for understanding climate change and its impact on land and sea ice, ocean ecology, and the global carbon cycle.
- HECC is used for high-resolution cloud resolving model simulations to provide unique and detailed insights into the processes that form tropical clouds and cloud systems, which account for approximately two-thirds of global rainfall.
- HECC is used to explore the feedback mechanisms between polar ice sheets and atmosphere circulation in order to determine how global temperature changes translate into increased sea level rise.
- HECC is used to improve the understanding of the current balance of carbon in the Arctic and to provide a framework for early detection of future carbon destabilization.

Heliophysics
of projects: 93
# of SBUs used*: 7,697,471

- HECC is used for modeling solar magneto-convection in order to understand how magnetic fields emerge through the sun's surface, heat the sun's outer atmosphere, and produce sunspots, spicules, and flares.
- HECC is used for realistic multi-scale simulations to understand the complicated physics of the turbulent convection zone and atmosphere of the sun and for analyzing and interpreting observations from the NASA space missions.
- HECC is used to simulate small-scale magnetic fields generated by turbulent dynamo action just beneath the solar surface in order to accurately predict space weather events that impact the Earth environment.

Planetary Science
of projects: 73
# of SBUs used*: 3,620,313

- HECC is used to decipher the structure of the lunar interior to understand the origin and thermal evolution of the moon and to extend this knowledge to other bodies in the inner solar system.
- HECC is used to model the origin and evolution of Kuiper belt objects to determine how their properties constrain our current models of planet formation.
- HECC is used perform modeling and simulation of asteroid entry, breakup, airburst, blast propagation, and tsunamis to assess the risks that potentially hazardous asteroids could pose to populations and infrastructure in the event of an Earth strike.

*October 1, 2017 to September 30, 2018

Return on Investment

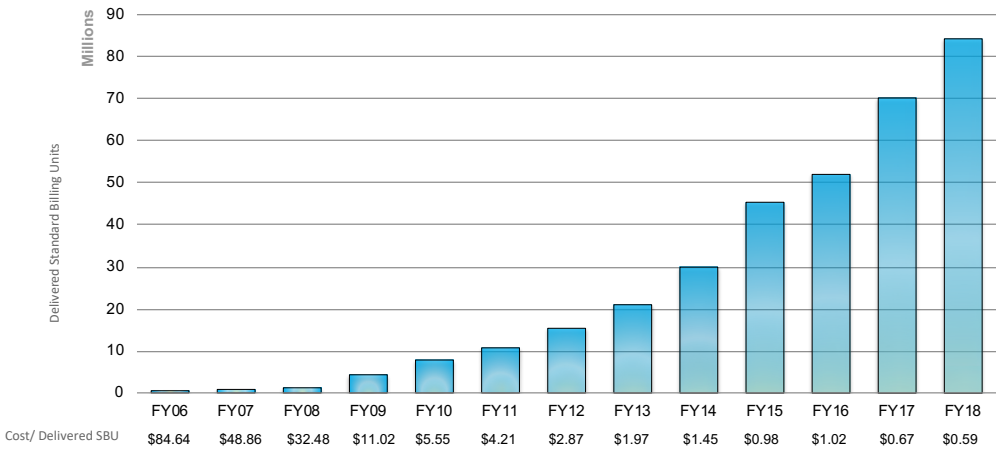

Why are We Doing This

- **The calculation used to be very simple…**
  - When the cost of maintaining a group of nodes for three years exceeded the cost to replace those nodes with fewer nodes that did the same work, we replaced them.
- **Now, not so much…**
  - We look at the total computing our users get and procure new nodes within our budget and remove enough nodes to power and cool the new nodes.
  - This means that we are not able to actually realize all of the expansion we are paying for.

But That's Not All

- **Our computer floor is limited by power and cooling.**
- **Our current cooling system:**
  - Open-air cooling tower with 4 50HP pumps.
  - Four 450-ton Chillers.
  - Seven pumps for outbound chilled water.
  - Four pumps for inbound warm water.
- **Our Electrical System:**
  - Nominally, the facility is limited to 6MW.
  - 25% is used for cooling (1.33 PUE).
  - 4MW – 5MW for computing.

DCU-20

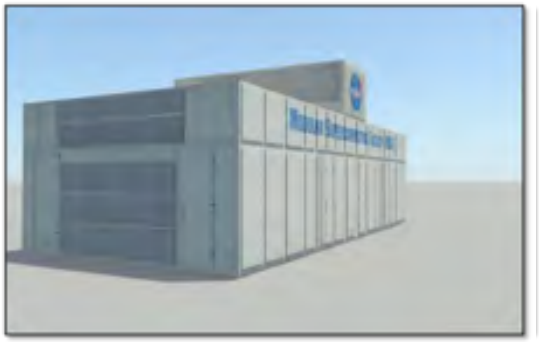

Module 2 Assembly

Module 2 Assembly

Annual Energy Impact

Electra System (2,304 Nodes)	N258 Facility	R&D 088 Facility	% Savings
Water Utilization Per Year	2,920,000 gal*	128,000 gal**	96%
Electricity per year	9,554,000 kwh° 2,424,150 kwh°	7,350,892 kwh°° 221,026 kwh°°	23.0%(overall) 90.9%(cooling)

* Assumes 2,304 nodes represent 20% of N258 facility load
** Year 1 (Oct 2017-Sept 2018 actuals)
° 1.34 PUE (Oct 2017-Sept 2018 actuals)
° ° 1.031 PUE (Oct 2017-Sept 2018 actuals)

NAS Facility Expansion (NFE)

NAS Facility Expansion Accomplishments

- **Deployed energy-efficient and cost-effective infrastructure capable of housing a 1,000,000-core.**
  - One-acre site can hold 12 compute modules and 3 data modules.
  - Site is 14 feet above mean sea level.
  - A new 115kV breaker & 30MW 115kV–25kV transformer was installed in the N225B substation.
  - Conduit to hold the electrical cables connecting N225B to R&D-099 and the data cables from N258 to R&D-099 were installed and all of the power cables and 1 of the data cables were run and terminated.
  - 25kV switchgear was placed on the site to distribute up to 30 MW to the modules.
  - A fire road and fencing were built.
  - The concrete pad is complete.
  - The first module was shipped on three specially equipped truck and installed via crane.
  - The 2.5MW 25kV-480V transformer, two adiabatic cooling units, pipes, pumps and electrical cabling were installed and tested.

Coming Online Soon

- **The facility was commissioned on July 24, 2019 with a 35-item punch list that will be completed by August 16, 2019.**
- **The first four E-Cells are installed and being tested by engineers and systems experts prior to release to the user community.**
  - The E-Cells have 1,152 Cascade-Lake nodes providing a theoretical peak of 3.69 PF.
  - The system, Aitken, was named after the lunar crater, Aitken, part of the very large South Pole-Aitken basin. The crater was named after the astronomer Robert Aitken (12/31/1864 – 10/29/1951).
  - Aitken will connect to N258 through four 288-strand Single Mode fiber cables.

Supercomputing Systems Team (SST)

Greg Matthews
Systems Admin
gregory.matthews@nasa.gov

Supercomputing Systems

Our Mission

"Evaluate, Test, Design, Purchase, Prototype, Build and Maintain High End Computing Hardware and System Software components in support of NASA missions."

Our Customers

- Individual Users with Unique Requirements
- Projects/PI's
- Programs
- NASA Mission Directorates

System Growth

- **We are continuously updating system hardware and software, including:**
  - Adding/removing racks of compute nodes.
  - Adding/removing filesystems.
  - Patching/updating kernel and other system software.
  - Updating device firmware or other software operating outside the user level.

[redacted]

Pleaides Upgrades 2008 – 2016

[redacted]

System Updates

- **Activity that Impacts System Availability**
  - Rolling Updates
    - » OS updated at job termination.
  - Suspend/Resume
    - » Live jobs receive SIGSTOP/SIGSTART.
  - Dedicated Time
    - » Can be filesystem specific (e.g. /nobackupp2).
    - » Full System Outage (we try to keep these to a minimum).

Supercomputing Systems

- **Systems Components**
  - On-going integration and life-cycle management of hardware and software.

- **Hardware**
  - Pleiades/Electra/Merope (16,000 nodes)
  - Hyperwall 3 (130 nodes)
  - V100 GPU Cluster (16 nodes)
  - K40 GPU Cluster (64 nodes)
  - Data Analysis (6 nodes)
  - Shared Memory Nodes (2 systems)
  - Data Archive and Tape management
  - Support Systems Infrastructure (150 nodes)
  - RAID Devices (> 70)
  - Interactive Front Ends

- **Software**
  - Operating Systems Builds (SLES, CentOS)
  - Job Scheduling (PBS)
  - Lustre Filesystems
  - NFS Filesystems
  - BeeGFS Filesystem
  - CXFS/DMF Filesystems
  - Modules Environment
  - Secure Unattended Proxy
  - Shift Transfer Tool (shiftc)
  - NFS Re-Exporter
  - Containers (early release)

What We Can Help You With

- **System Level Problem Resolution**
  - We track and plan our work through submitted tickets. Please report problems to Support so that we can begin to troubleshoot the issue.
    - » System issues can be difficult to track down, so ***we often rely on users to identify and report problems.***
- **Systems Feature Development**
  - Software Features
    - » Custom software requirements or components
    - » Proprietary software license servers
  - Dedicated hardware testing/procurement
    - » Lustre or NFS dedicated to a user or project
    - » Dedicated Front End or Small Cluster

Custom Systems Requirements

- **Dedicated 'Front End' Systems**
  - Systems directly attached to HECC supercomputing environment, which could include:
    - » Customized system with specialized hardware (e.g., RAID).
    - » Restricted access for project personnel.
    - » More control over system updates/testing.
    - » Temporary prototyping with no operational commitments. If we have an older/unused system this prototyping could have no cost and happen in less than a couple months.
  - Requirements for dedicated systems:
    - » HECC purchases equipment with some requirements on configuration and maintenance.
    - » Typically covers 3 years maintenance.
    - » May have administrative cost for tailored support.
    - » Minimize software support issues (required to use our operating system builds).

Recent/Near Term Capabilities

- **/swbuild Filesystem**
- **Reservable Front Ends**
- **Reservations**
- **NFS Re-Exporter for External Publishing**
- **Shiftc/SUP Enhancements**
- **Node Type Independent Scheduling (soon)**
- **Baby Lustre Servers**
- **Archive Upgrades (capacity and performance)**

Some Success Stories

- **A [JPL containers?]**
  - Details
  - Details
- **Accomplishment #1**
  - Details
  - Details
- **Accomplishment #1**
  - Details
  - Details

Caption

We Listened: Your Feedback Created Changes

- **sky_wide Queue**
  - Allows large (600 - 1,100 node) Skylake jobs a chance to run each week.
  - Job requests can be tuned to support recovery/resubmission in case of unexpected failure.
- **PBS Reservation Requests**
  - Users may request PBS reservations via support@nas.nasa.gov.
- **Development of User-initiated PBS Reservations**
  - Users will be able to treat a job like a reservation.
  - If the job dies early you'll be able to resubmit to the same nodes.
- **Development of Job Start Time Windows**
  - Altair developing a feature to specify start window (e.g., job start between 8a-5p M-F).

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time
Greg Matthews
gregory.matthews@nasa.gov
1-650-604-1321

HECC Cloud Offering

Robert Hood
Technical Manager, APP & Vis
robert.hood@nasa.gov

2018 Trade Study:
When Does It Make Sense for HECC to Use Clouds?

- **Main Finding:**
  - Commercial clouds do not offer a viable, cost-effective approach for replacing in-house HPC resources at NASA.
- **Additional Finding:**
  - Commercial clouds provide a variety of resources not available at HECC. Certain use cases may be cost-effective to run there.
- **Example Conditions That May Warrant Cloud Usage:**
  - When utilization would be low—such as when using new resource types
    - use cloud-based resources until demand rises to the point that it's more cost effective to acquire on-premises HECC resources to meet that need.
  - When there are other costs to consider, such as opportunity costs associated with high utilization (longer queue waits).
  - When there are real-time requirements, such as web services.

Follow-up Work Identified by Trade Study

- **Gain a better understanding of potential benefits and costs of having a portion of the HECC workload in the cloud:**
  - Understand performance characteristics of jobs that might run there.
  - Define a comprehensive model that allows accurate comparisons of cost of running jobs depending on resources used.
- **Prepare for a broadening of HECC services to include a portion of its workload running on commercial cloud resources:**
  - For HECC users.
  - For non-HECC users.

Pilot Project: Enabling Cloud Usage

- **Move Jobs from HECC Resources to Amazon Web Services (AWS)**
  - HECC user logs into an AWS cloud front end to build executable.
  - User annotates batch scripts, indicating files that need to be staged to/from cloud.
  - User submits batch jobs to “cloud” queue.
  - PBS server moves jobs to server running at AWS; stages input files to AWS.
  - Server in cloud allocates resources, runs job; PBS server in HECC stages output files back.
  - Accounting is done manually; HECC pays.
  - Limited to non-export controlled codes and data (i.e. “low” security plan).
- **User-Defined Software Stacks**
  - Container technologies (Charliecloud, Singularity) are under evaluation.
- **User testing of the “cloud bursting” started in September 2019.**

Modeling of Spiral Wave Instabilities (SWI)

- Researchers increased resolution of spiral waves in protoplanetary disks by 4x and captured cascade of SWI-driven turbulence at smaller scales.
  - Previous resolution was limited to 512*128*128 due to computational costs.
  - Using NVIDIA V100 nodes at AWS, the resolution increased to 2048*512*512.
- Higher-resolution computation took 100 hours on single AWS compute node with eight V100s.
- The researchers plan to test the SWI case with more realistic models and higher resolutions.

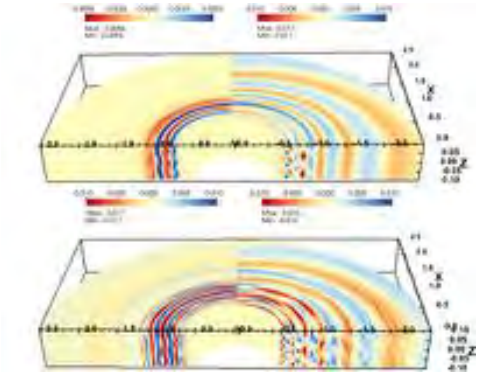

Top: Results with standard resolution (512x128x128). Bottom: 4x higher resolution (2048x512x512). While the spiral wave instability (SWI) develops with both resolutions, the SWI-induced turbulence is better resolved with a larger number of grid cells. *Jaehan Bae, Carnegie Institute of Washington.*

3D Aerospace Manufacturing Simulations

- Material scientists narrowed the gap between experiment and simulation of polymer-polymer interfaces by more than two orders of magnitude.
  - Previously, coarse-grained approximations used the LAMMPS molecular dynamics (MD) code running without GPUs (red squares).
  - New fine-grained simulations used GROMACS MD package running on GPUs at AWS (blue squares).
- The researchers will extend these results to predict other important aerospace manufacturing parameters, paving the way for exploration of novel frontiers of nanomanufacturing.

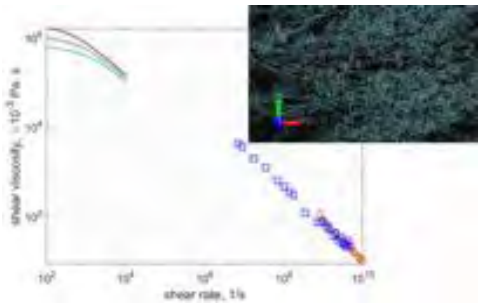

This graph shows atomistically resolved molecular dynamics (MD) predictions of the shear viscosity for flat polymer-polymer interfaces (squares at lower right). These can be compared with experimental data on bulk samples (lines at upper left). The inset shows MD predictions obtained on the Pleiades supercomputer. *Dmitry Luchinsky, SGT, Inc., Intelligent Systems Division, NASA Ames Research Center.*

Current Work

- **We Have Extended the Pilot AWS Cloud Project to Include:**
  - Moderate security plan.
  - Full accounting, with account limits and automated tracking of consumption,
    - Users will bring their own funding,
    - Accounts will be charged for:
      - Job resource usage.
      - Ongoing storage usage.
      - Data transfer bandwidth.
      - Overhead.
- **The HECC AWS Cloud Offering is Now Operational (Starting August 7th):**
  - Full accounting, with account limits and automated tracking of consumption.
  - Processes for:
    - Account setup requests.
    - Transfers from NASA WBS to HECC cloud account.
    - Monthly statements.

Future Work

- **Extend Services as Required by HECC Users.**
  - New resource types at Amazon.
  - Other cloud providers (Google? Azure?)
- **Extend Services to Include Non-HECC Users from NASA.**
  - They also need to bring their own funding.
  - Provide web-based user interface for defining and running jobs, moving data, etc.
  - HECC would add cost-recovery fee to make this self-sustaining.
- **Devise Cost Methodology,**
  - Must be able to do meaningful cost comparisons between on-premises and in-cloud resources in order to determine which would be most cost effective.
    - Include opportunity costs.
    - Establish benchmark suite.
  - Adjust processes for acquisition and phase out of resources to include commercial cloud.

Getting Started with HECC AWS Cloud

- **Principal Investigator (PI) of Group:**
  - Consult with HECC about workflow requirements, potential costs, advantages, etc.
  - Determine whether an ITAR/EAR99 environment is needed:
    - No → Proceed with Public Cloud.
    - Yes → Wait till Gov Cloud is approved for HECC use.
  - Send NASA funding information to HECC to establish an ARC WBS billing account.
  - Provide a list of users allowed to use their HECC Cloud account.
  - Provide an initial desired cloud configuration (can be adjusted later).
- **HECC:**
  - Set up the account and configure the environment.
  - Track usage of all resources:
    - Front end, PBS server, compute, filesystems, storage, network, support, etc.
  - Charge expenses against PI's funding and provide usage report.

Acknowledgements and Additional Information

- **The Trade Study Team:**
  - S. Chang, R. Hood, H. Jin, S. Heistand, J. Chang, S. Cheung, J. Djomehri, G. Jost, D. Kokron
- **The Trade Study: NAS Technical Report NAS-2018-001**
  - Posted on our website at:
    https://www.nas.nasa.gov/assets/pdf/papers/NAS_Technical_Report_NAS-2018-01.pdf
- **The Cloud Team:**
  - S. Chang, S. Heistand, H. Yeung, M-Y. Koo, J. Chang, S. Cheung, J. Djomehri, R. Hood, H. Jin
- **Knowledge Base:**
  - See the sixteen new articles published under Cloud Computing > AWS Cloud:
    https://www.nas.nasa.gov/hecc/support/kb/175/

We're listening!

Ideas
Suggestions
Questions

Contact me any time
Robert Hood
robert.hood@nasa.gov
1-650-604-0740

HECC User Services (Control Room)

Blaise Hartman
User Services Lead
blaise.hartman@nasa.gov

Here's What We Can Do For You

- **Tier-1 Support**
  - We are the first line of support for all the HECC groups.
  - We notify other HECC groups of user and system issues during off hours.
- **24x7 Concierge**
  - If we are unable to resolve your issues, we will contact the correct group for you.
  - We're available 24x7 to take your calls at (800) 331-8737.
- **Support@nas.nasa.gov**
  - Sending an email to this address will automatically create a ticket for you.
  - We take care of these tickets directly or redirect them to the appropriate group to address your questions and technical issues.

You Can Also Get In-Depth Support

- **Assistance Setting Up an HECC Passthrough**
  – A passthrough allows you to log into HECC systems without going through an SFE first.

- **Managing Archive Data**
  – We can archive the data of users that have already left your project.
  – We can move data into other group member's directories.

- **Data Restoration**
  – We can retrieve data that was accidentally deleted on a user's home filesystem.

- **Advanced File Sharing Assistance**
  – We can help make your data available to others.

Special Requests

- **Reservations and Special Queues**
  – We can setup special queues and reserve nodes for users and groups.
  – For this service please contact Support with your GID, the number of nodes, duration of required time, and a justification for the reservation or special queue.

- **Quota Modification**
  – We can increase your quotas on the home and nobackup filesystems as long as it is within reason.

- **In-Depth Analysis of SBU Usage**
  – We can help you identify why your allocation may be getting used up too quickly.

Some Success Stories

- **5,000 Simulations Reservation**
  - The Space Launch System team needed to create a database of parameters for booster separation. Each simulation was small, but they needed to run 5,000 simulations. We were able to help them by providing a reservation, allowing them to complete the database on-schedule.
- **Long Reservations**
  - A few users have made reservations for a month or more for critical projects or to meet critical agency deadlines.

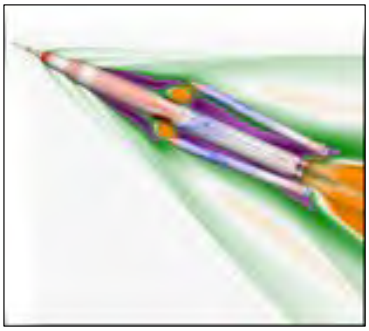

SLS Block 1 vehicle flight and wind tunnel geometries, showing both the flight flow field (left) and the wind tunnel flow field (right), simulated using NASA's FUN3D code. This vehicle will be used for Exploration Mission 1, the first SLS flight. The vehicle surface is colored by pressure contours, where blue is low and red is high. The green-white-orange colors represent low to high velocities. *Henry Lee, Stuart Rogers, NASA/Ames*

We listened: your feedback created changes

- **Scalability:**
  - A new queue was created, called “sky_wide,” for jobs needing up to half of the Skylake nodes. At least once a week, those jobs will be guaranteed to have enough nodes available to run.
  - Jobs between 1/4 and 1/2 of the Skylake nodes can submit to the “sky_wide” queue to run on Mondays. Jobs must be submitted to the queue by the previous Thursday at 5 PM.
    - » #PBS -q sky_wide
    - » #PBS -l select=NNN:model=sky_ele:ncpus=nn:aoe=sles12
- **Priority:**
  - While a reservation may not be warranted, if you have a critical deadline, you can request a priority increase with justification.

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time
Blaise Hartman
blaise.hartman@nasa.gov
support@nas.nasa.gov
1-650-604-2539

HECC User Services (Accounts and Allocations)

Emily Kuhse
Accounts Specialist
emily.kuhse@nasa.gov

Here's What We Can Do For You

- **Account Setup and Maintenance**
  - From getting new users on-boarded to meeting requirements to maintain accounts, we're here to help.
- **Allocation Requests and Management**
  - Assistance requesting a new allocation, renewing a project, or requesting additional resources.
- **RSA Token Support for Non-NASA Users**
  - Initial token setup, as well as replacement tokens and soft token support.
- **Representation**
  - We do our best to make sure that our remote users are represented when it comes to new or changing NASA policies.

Did You Know?

- **If you run out of time, you can request more!**
  - You can always request additional resources at any time, should you need them.
- **There is a rolling call for computing requests.**
  - While the majority of projects are allocated at the same time in the fall, you can submit requests for new computing projects at any time.
- **You can request reimbursement for SBUs lost due to system issues.**
  - We understand things happen and want to make it right whenever possible.
- **You can use myNAS to track your project's usage and view information on running, waiting, held, and recently completed jobs in near-real-time!**
  - Visit https://portal.nas.nasa.gov or download the myNAS app

But What If…

- **I didn't receive my full allocation request? I don't have enough time to complete my runs!**
  - You can always request more time. While we don't allocate time, we may be able to help you build a justification (for instance, misunderstandings about a project's SBU requirements).
- **I need to add new team members! What if they aren't US citizens?**
  - We can walk you and your new users through the process. Non-US citizens can absolutely get accounts. In fact, the process for non-US citizens has become much more efficient (down from a year to 2-3 weeks).

Some Success Stories

- **Maintaining Access**
  - HECC maintains a NASA affiliation agreement for remote users. If one of your group members is leaving and still needs access to your project, we can work with you to keep their NASA Identity active for supercomputing access.
- **Mission Critical GID**
  - One of our largest users needed a new GID set up for a mission critical test. The GID was created, allocated, and running in a reservation in under 4 hours!

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time
Emily Kuhse
emily.kuhse@nasa.gov
1-650-604-1687

HECC Application Performance & Productivity (APP)

Robert Hood
Manager, APP & Vis Groups
robert.hood@nasa.gov

Here's What We Can Do for You

- **Help with application-related issues, for example:**
  - Porting your code to Pleiades/Electra from other systems.
  - Basic performance analysis and optimization.
  - Setup for parallel execution of numerous serial jobs.
- **In-house developed tools for applications and job scheduling:**
  - Provided in the `/u/scicon/tools/bin/…` directory.
  - Info about your running jobs: `qtop.pl`, `qps`, `qsh.pl`
  - Node availability: `qs`, `node_stats`
  - Job startup: `mbind.x` (process binding)
    `several_tries` (resilience)
  - Memory monitoring: `gm.x`, `vnuma`

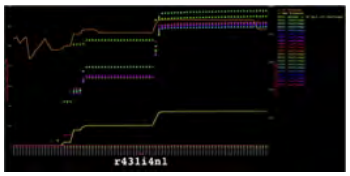

Screenshot from vnuma memory tool.

You Can Also Get In-Depth support

- **Detailed performance analysis and optimization.**
  - Can often improve run times by 10-60%.
- **OpenMP or MPI parallelization.**
  - As long as it's not too complicated; more extensive analysis and programming work may require external funding, based on the required level of effort.
- **Application development:**
  - Again, as long as it's not too complicated; external funding may be required.
- **We can teach your group how to do performance analysis.**
  - For several groups at Langley, we analyzed their codes and did some optimizations, then presented the methods and explained why they worked.

Did You Know…

- **You can access an extensive online Knowledge Base and Webinar series**
  - Often, the answer you're looking for is in the KB!
  - Tell us about topics you'd like in the KB or a Webinar.
- **You can use myNAS to get information about jobs even when you're not logged into HECC systems (see images on next page).**
  - Information provided by the myNAS web portal includes failure types for your jobs.
  - The myNAS mobile app can send notifications when a job starts or finishes.
- **You can get help building software packages if you need them.**
  - Software needed by multiple groups can be be made available on the HECC systems to all users; otherwise, you would get your own copy in your own directory.

myNAS Web Portal and Mobile App

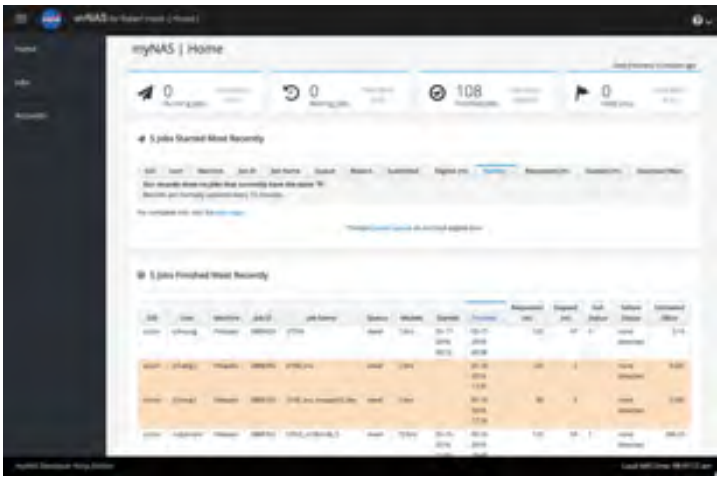

myNAS web portal:
https://portal.nas.nasa.gov

myNAS mobile app:

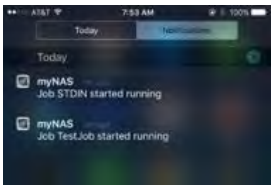

Did You Know …

- **We can help expedite your jobs.**
  - This is especially true if you've been affected by system issues.
  - Let us know if you have a hard deadline and you are running out of time (with manager approval).
- **We benchmark current computers and potential future systems.**
  - We set "SBU" rates for the node types.
  - If you have a code using a million SBUs a year, we'd like to talk to you about whether we need to develop a benchmark representing that code.
- **We monitor energy usage by jobs.**
  - We've found a way to identify some (simple-to-fix) cases of load imbalance.

Some Success Stories

- **Improved GEOSgcm job efficiency by 25%.**
  - APP performed in-depth analysis of memory requirements for the climate app.
  - Found that rather than leaving eight cores idle on each node, only the first node needed to do so.
  - Reduced the number of nodes required for jobs by 25% without an increase in run time.
- **Improved kernel of finite element solver by 85%**
  - FEMERA is a high performance finite element solver, implemented in C++ and OpenMP in a domain decomposition-based parallelization.
  - Speedup achieved through memory restructuring, use of vector intrinsics in critical loops, and flattening complex data structures.

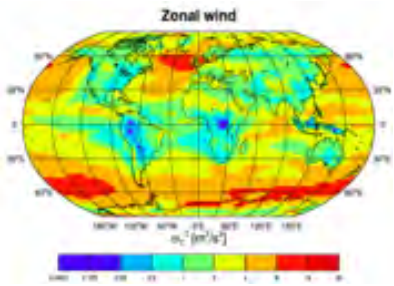

Climate models are the main tools used to simulate future climate conditions. However, the predictions made by these models are still subject to large uncertainties, and many computational runs with varying initial conditions are required to gauge the sensitivity of climate predictions to various model inputs.

More User Successes

- **Improved kernel of HyperSolve CFD solver.**
  - HyperSolve is part of Langley's T-infinity framework (CFD Vision 2030).
  - APP achieved improvement through compiler directives, memory layout changes, and code simplification.
  - The changes sped up the kernel by 2x, and the full application by 5–15%.
- **Enabled STAR-CCM+ to run on Skylake nodes.**
  - STAR-CCM+ was failing on Skylake nodes for one user.
  - APP determined the issue was that there are two types of Skylake nodes w.r.t. the IB network names, confusing the STAR-CCM+ MPI library.
  - APP was able to switch to HPE MPT by having the group use a private module performing app startup.

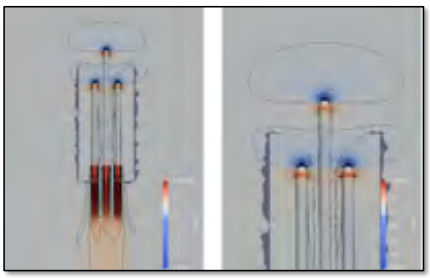

HyperSolve steady overset simulation of notional heavy-lift vehicle. The Mach contours on 3 domains are displayed on the left-hand side. Close up Mach contours are shown on the right-hand side. *Matthew O'Connell, Cameron Druyor, and Kyle Thompson, NASA/Langley*

We Listened: Since Our Last Visit

- **We helped a group at Glenn with a reproducibility issue.**
  - Initially, APP responded to a request to improve performance.
  - There was also a correctness issue (race condition?) several hours into a run.
  - We found the issue was due to use of MPI_Waitany; replaced with MPI_Waitall
- **Met with a GRC group (not current HECC users).**
  - They have multiphysics application that they want to improve by orders of magnitude.
  - They gave a team from APP an overview of their work and their long-term goals.
  - The APP team then:
    - Connected them to the Vis group who advised on file formats and efficient I/O patterns.
    - Outlined how optimization would work.
    - Gave pointers on approaches for meeting the long-term performance goals.

We're Listening Now …

- **Do you have suggestions for training webinars?**
  - GPU, optimization tools & techniques, …
- **Suggestions for myNAS?**
  - Are you using the mobile app? The web portal?
  - What would you like to see added?
    - » We will be adding Shift transfer status soon.
    - » Longer term: A future version will add Remedy ticket status.
    - » Would you like to see output from your job? Performance data?

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time:
Robert Hood
robert.hood@nasa.gov
1-650-604-0740

HECC Visualization & Data Analysis

Robert Hood
Manager, APP & Vis
robert.hood@nasa.gov

Here's What We Can Do For You

- **Scientific Visualization**
  - Create images and animations for research and presentations.
  - Apply scalar and vector field visualization techniques.
  - Handle all modern computational grid systems and point data.
  - Specialists in handling large and complex datasets.
- **Data Analysis**
  - Create derived quantities from computational primitives, including differentials.
  - Deploy feature extraction techniques, e.g., vortices.
  - Use topological methods for scalar and vector fields.
  - Provide time series analysis.

You Can Also Get In-Depth Support

- **Concurrent Visualization**
  - Instrument your code for direct access to data without disk I/O.
  - Create high-cadence visualizations (e.g., every timestep) with little or no overhead.
- **Scalable I/O**
  - Implement parallel I/O strategies for largescale codes.
  - Typically combined with concurrent visualization.
- **Big Data Management and Workflows**
  - Exploit software and hardware features for large data handling.
  - Develop file, filesystem, and archive strategies for scalable performance.
- **Custom Interactive Data Exploration Environments**
  - To use the hyperwall at NASA Ames, you must come to us.

Ocean scientists using the hyperwall with concurrent visualization at the NASA Advanced Supercomputing Division.

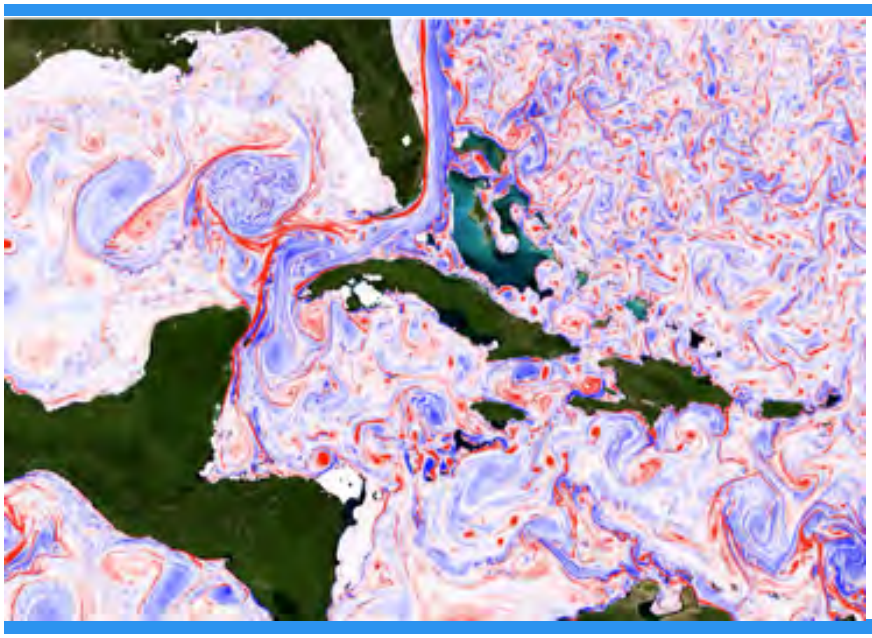

MITgcm Ocean Model:
22 billion gridpoints

Tools Developed for ECCO Data Analysis

The “multimovie” player:

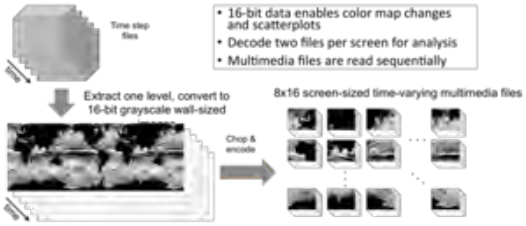

Combining 2 screens for comparisons:

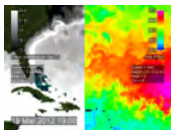

Showing 2 scalars in split screen

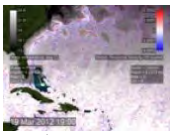

Overlaying extremes of one scalar on top of another

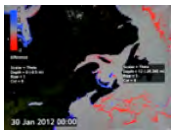

Subtracting two values, approximating a gradient

Scatterplots:

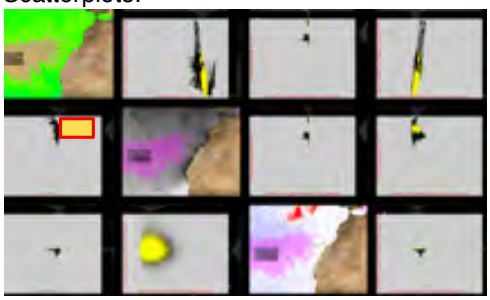

A 3 by 4 array of scatterplots with some map view displays. All of the displays show data for the same tile (geographical location), the eastern Atlantic ocean off North Africa. The red selection box is on the screen in the second row, first column. The selected points are shown in yellow on the scatterplot screens, and in light magenta on the map view screens.

These tools have also been used for analyzing Mars climate model data.

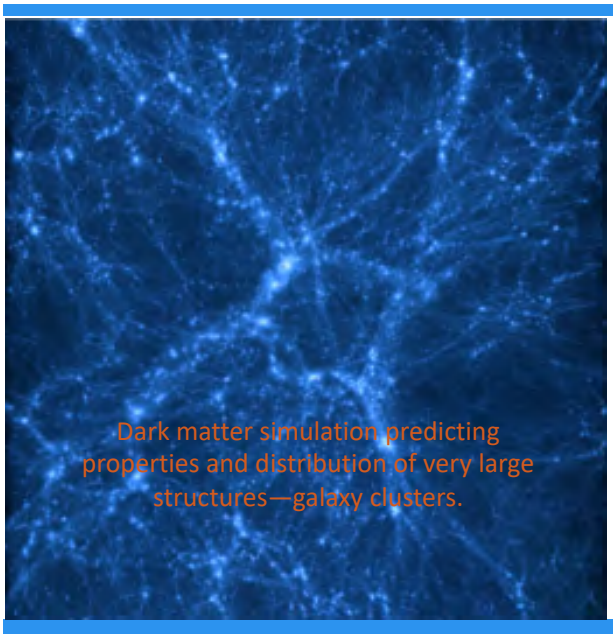

Bolshoi Cosmological Simulation:
10 billion points

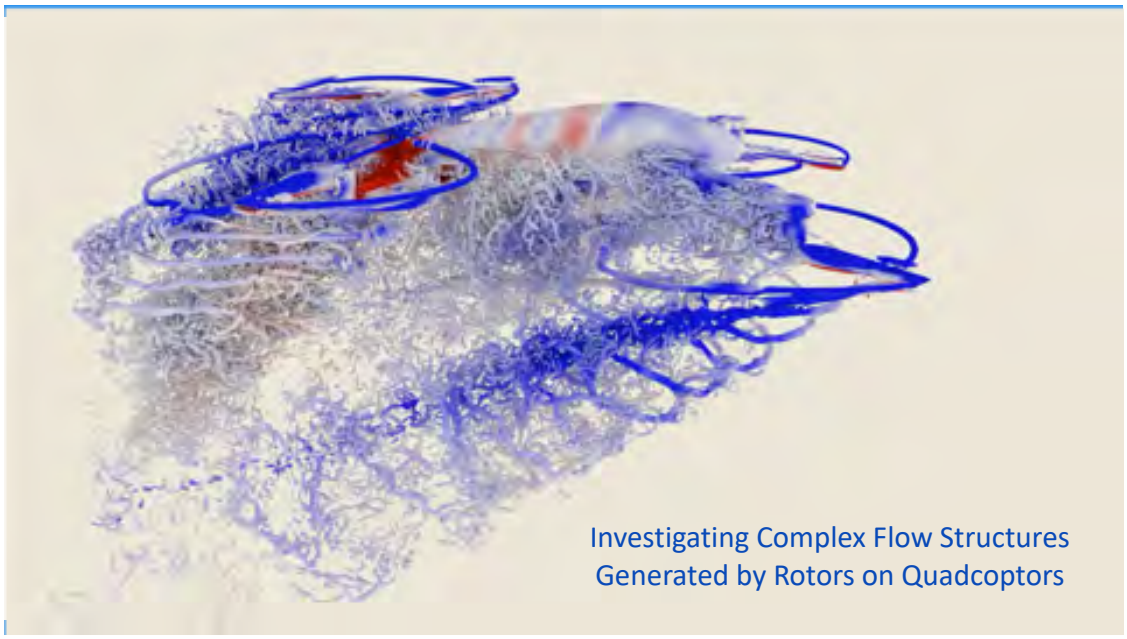

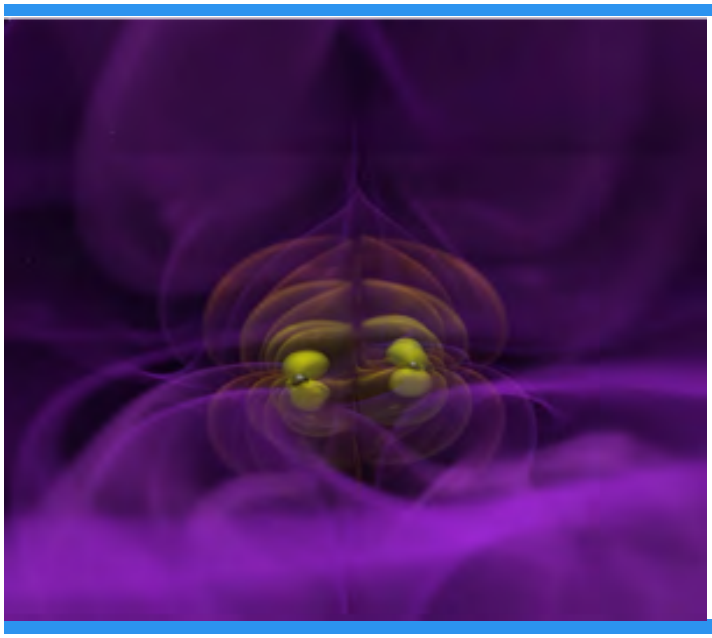

Gravitational Radiation
From Spiraling Black Holes:
LIGO First Detection

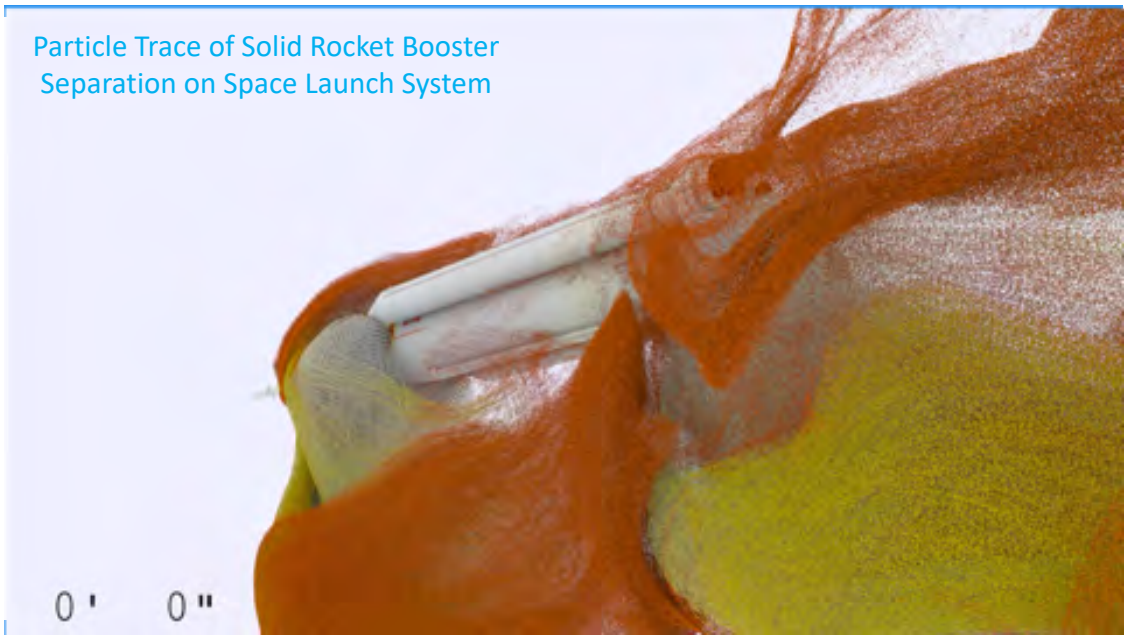

Unsteady Pressure-Sensitive Paint Project

- For several years a wind tunnel group at Ames Research Center has been requesting help with analysis of their Pressure-Sensitive Paint (PSP) experiments.
  - Goal is to do analysis in real time, rather than days or weeks later.
  - Security always posed a hurdle.
- Finally, this year the security concerns were addressed with network diodes.
- One key component of the project is data analysis.
  - Existing serial code processed one dataset in 14 hours.
  - The Vis team sped that up to 50 seconds—1000x—with:
    - Parallelization with MPI, OpenMP, and custom multi-threading
    - Using asynchronous I/O
    - Algorithmic improvements
- HECC hopes to begin delivering results in September 2019.

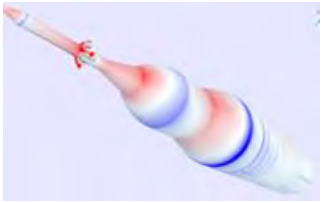

Rendering of the uPSP results, showing local pressure differences on the surface of an SLS model.

We're Listening!

Ideas
Suggestions
Questions

Contact us any time
Chris Henze
chris.henze@nasa.gov
1-650-604-3959

Here's What We Can Do For You

- **Assistance with using machine learning (ML) technology.**
  - Provide guidance with machine learning tools and techniques.
  - Provide support for running ML tools on HECC resources.
  - Help with TensorFlow, Jupyter Notebooks.
  - Work with NASA teams to help move to AI/ML technologies.
- **Assistance with data publication and discovery projects.**
  - Provide guidance in solving big data problems, using HECC resources.
  - Framework for sharing datasets stored at NAS with the public.

You Can Also Get In-Depth Support

- **Work with your pilot projects to develop a framework to address ML needs that can be applicable for multiple projects.**
  - Provide more focused support to infuse AI/ML and deep learning technology into projects.
  - Additional funding may be required depending on the required level of effort.
- **Big Data—Data Publication and Discovery.**
  - Get dedicated filesystem for data storage and/or sharing.
  - Web and data portal for sharing data stored on HECC resources with colleagues outside HECC and/or the public.

Some Success Stories: Data Portals

- Created a public repository for sharing large amounts of non-sensitive/non-proprietary data with colleagues.
- Share data in place using re-exporters providing required public access.
- Automated system to request sharing of datasets located at NAS to user community.
- Estimating the Circulation & Climate of the Ocean (ECCO) data portal – supports subsetting services.
- Heliophysics portal – query and download.
- Quantum Artificial Intelligence Lab (QuAIL) data portal.

More Successes: Metrics for Data Portal Usage (Peak Downloads over 100 K)

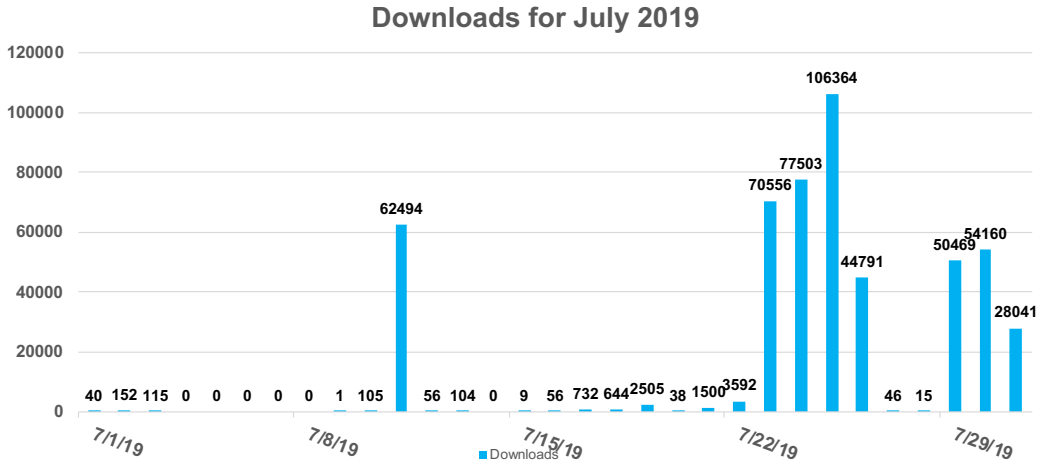

Yet More Successes: Machine Learning

- Created predictive models (Classification, Regression, LSTM) for ammonia concentration levels for Gas Chemical Sensors team at Ames.
- Photovoltaic Cell Characterization (GRC) - 2D Convolutional Neural Network model trained on ~2800 samples of differing chemistry was able to calculate the short circuit current and open circuit voltage with a 94% accuracy.
- 7x improvement in machine learning performance using GPUs for asteroid threat prediction calculation, to predict meteor characteristics based on measured light curve data.
- Looking to add one more pilot project.

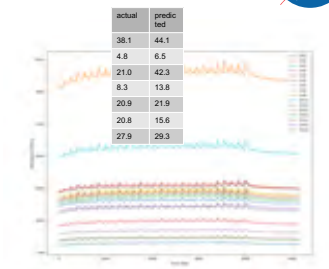

actual	predicted
38.1	44.1
4.8	6.5
21.0	42.3
8.3	13.8
20.9	21.9
20.8	15.6
27.9	29.3

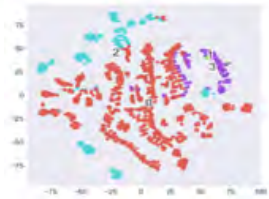

We Listened: Your Feedback Created Changes

- **GPU cluster expansion/upgrade.**
  - NVIDIA K40 – 64 nodes are available.
  - NVIDIA 4 X V100 – 14 nodes and 8 X V100 – 2 nodes are available.
- **Mechanism for data management/sharing.**
  - Data Portals.
  - Public repository for sharing large amounts of (non-sensitive/non-proprietary) data.
- **Assistance for moving into advanced analytics.**
  - TensorFlow is available in modules; includes both CPU/GPU versions.
  - Jupyter Notebooks with various TensorFlow environments.
  - Expert data scientists to assist with getting your ML/AI projects started.

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time
Shubha Ranjan
shubha.ranjan@nasa.gov
1-650-604-1918

HECC Publications and Media

Emily Kuhse
User Services
emily.kuhse@nasa.gov

Here's What We Can Do for You

- **HECC Knowledge Base: Detailed technical documentation.**
  - New user orientation, how-to articles, troubleshooting, optimizing, and more.
- **HECC News: The latest system information and announcements.**
  - Online notifications of system upgrades, scheduled downtimes, and other events.
- **Websites: Feature stories, resource info, system status, webinars.**
  - Details about HECC resources & services, webinar archive, success stories, and more.
- **Help with promoting your work to NASA and public audiences.**
  - In-depth feature stories about your projects, published on our websites.
  - Social media campaigns (Twitter, Facebook).
  - Technical conferences and public events.

Did You Know …

- **We want to tell your stories!**
  - When you have results you want to share with the HPC & science community, let us know!
  - We'll work with you to produce and promote:
    - » In-depth articles and short image features highlighting your work and visualization results.
    - » Demos at the annual SC Conference and other events.
- **We run Twitter and Facebook accounts for NASA supercomputing!**
  - Do you tweet about your work? Tag **@NASA_NAS** and we'll expand it to a wider audience!
  - Be sure to follow **@NASA_Supercomp** for all agency supercomputing news and events.
  - You can also find us on Facebook—just search for NASA supercomputing.
- **Send us your feedback on the Knowledge Base!**
  - Find a mistake? Something that needs updating? Have an idea for a new topic?

    Send a note to **michelle.c.moyer@nasa.gov** or use the **Ask a Question** feature in the KB.

Some Success Stories

- **We share your research with the world!**
  - After researchers presented their simulations of the Orion Launch Abort system in the NASA booth at SC17, we worked with them to develop a feature story for publication on our website.
  - We promoted the story and visualizations to a worldwide audience via our social media channels:
    - » **Twitter:** 17,003 media views; also shared by NASA Ames and ARMD accounts
    - » **YouTube** (NASA Ames): 15,979 views
    - » **Facebook** (NASA): 9,965 views
    - » Picked up by media including HPC Wire, Universe Today, Enterprise AI, European Space Agency

The SC Conference Series

- **SC19, SC20, and beyond!**
  - We organize, produce, and operate the NASA booth for the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC) each year. SC19 is our 31st year!
  - When you participate in the NASA booth, we will:
    - » Work with you to produce science-packed, public-friendly abstracts and posters, published at our NASA@SC websites.
    - » Provide opportunities to present your work to an international audience interested in HPC, science, and engineering.
    - » Work with you and agency public affairs experts to promote your science and research to a wider public audience via feature stories and social media.

We're Still Listening!

Ideas
Suggestions
Questions

Contact us any time

Jill Dunbar
jill.a.dunbar@nasa.gov
1-650-604-3534

Michelle Moyer
michelle.c.moyer@nasa.gov
1-650-604-2912

HECC Network Services

Celeste Banaag
Sr. Network Engineer
celeste.banaag@nasa.gov

Here's What We Can Do For You

- **End-to-End Networking Services**
  - The HECC user community is spread across the country and, in some cases, in other countries. Our users' ability to access resources is critical to the success of their computational projects. We believe that any problems you have in accessing HECC assets, or in moving data to or from our data center, is our problem to solve.
- **Multiple Methods for Data Transfer**
  - Secure Unattended Proxy (SUP), SSH Passthrough
  - Multiple file transfer protocols available for transferring large files:
    - » Shiftc
    - » bbftp/bbscp

You Can Also Get In-Depth Support

- **TCP Performance Tuning for WAN Transfers**
  - We work closely with you to optimize multiple aspects of end-to-end flows, select the most efficient transfer methods and protocols, and fine-tune your systems.
- **Network Troubleshooting**
  - Collaborate with remote network teams and system administrators, Communication Services Office (CSO), and wide-area network (WAN) service providers to identify and remove bottlenecks along the network path, and improve your flow throughput.
- **Performance Analysis**
  - Using Network Diagnostic Tool (NDT) to help identify last mile network issues, NPAD, and PerfSonar.
- **User Education**
  - Provide custom training, e.g., webinars, troubleshooting, diagnostic tools.

Some Success Stories

- **14x Transfer Rate Improvement to Langley Research Center**
  - In-house monitoring and analysis tool automatically notified network engineers of a user at NASA Langley getting a transfer rate of 12 megabits per second (Mbps) while transferring gigabytes worth of data.
  - Network staff immediately contacted the customer and, after reviewing system settings, determined an issue with the user's local desktop connection.
  - Langley network staff were then able to resolve a port negotiation problem and the user then achieved performance rates of more than 165 Mbps—yielding a 14x improvement in throughput rates.

More Success Stories

- **10x Transfer Speed Improvement to Johnson Space Center (JSC)**
  - HECC network engineers detected suboptimal transfer speeds from the JSC AeroLab to the NAS facility across a 2x1-gigabit (Gb) Ethernet link.
  - Before the upgrade, throughput for single- and multi-stream transfers averaged 125 Mbps.
  - The network team worked with JSC engineers to tune TCP window sizes and network stack for AeroLab hosts and upgrade links, resulting in throughput speeds up to 1.5 Gb/s for the bbSCP transfer method and 1.2 Gbps for the rsync method—about a 10x performance improvement.
  - By actively working with JSC engineers, the HECC team identified areas where we can improve the network and help the users understand the most effective way to use different file transfer tools.

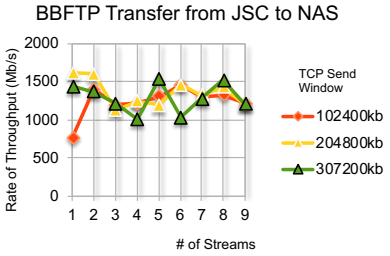

This chart shows data transfer rates (throughput) and the number of parallel streams for a single transfer between Johnson Space Center and NASA Ames using the bbFTP utility. HECC engineers achieved end-to-end performance tuning by adjusting the TCP window size (represented by each line on the graph.) TCP window size is simply the amount of data in bytes that a sender is willing to send at any point in time.

We're Still Listening!

Ideas
Suggestions
Questions

Contact me any time
Celeste Banaag
celeste.banaag@nasa.gov
1-650-604-2039

https://ntrs.nasa.gov/search.jsp?R=20190030673 2020-07-31T16:56:03+00:00Z

Industry Inputs into Future NASA Project Planning

September 5, 2019
Al Capps, ATD-2 Manager & Chief Engineer
Kevin Witzberger, ATD Deputy Project Manager
William N. Chan, ATM-X Project Manager
Dr. Bryan Barmore, ATM-X Deputy Project Manager for Technology

Motivation & Objective

Motivation

- ATD will complete in FY2020
  - ATD-1 completed in FY2018
  - ATD-3 completes in FY2019 (less than 30 days from now)
  - ATD-2 will complete in FY2020 (~ 12 months from now!)
  - ATM-X is planning for work for FY21 – FY25

- NASA seeks to maintain ATD-2 industry collaborations while transitioning ATD expertise to other NASA projects, including the ATM-X Project

Objective

- This breakout session is an opportunity for NASA to hear from industry to guide NASA AOSP planning

Planning Parameters

- Effort to begin executing in FY21 under the ATM-X Project
  - Requires NASA approval to execute, scheduled in summer 2020

- NASA envisions executing a series of joint NASA/Industry partnership evaluations from FY21 to FY25
  - Targets applications that industry (airline operators, airport operations, new entrants and vendors) expects to provide benefits
  - Complements current (TBFM, TFMS) and future (TFDM) FAA automation platforms as well as yet-to-be developed future systems
  - Service oriented technology to be validated through operational use to support commercialization potential by others in the aviation industry – expected deliverables include reference prototype, requirements, and associated documents

- Leverage the NASA/FAA/Industry partnership developed under ATD
  - Jointly develop future system requirements

Scope of Today's Discussion

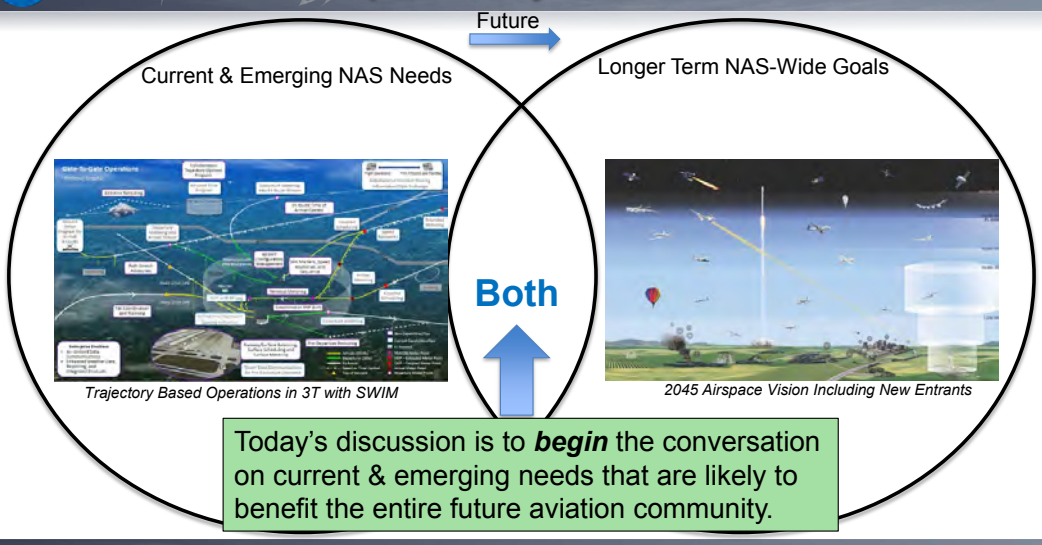

Follow up Webinar- Continuing the Dialog

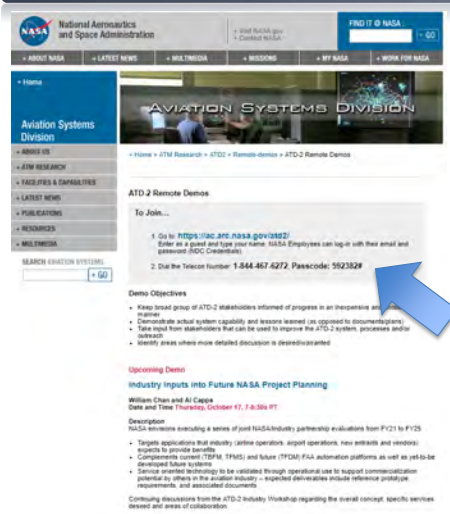

- A considerable amount of content will be presented today that benefits careful consideration
- In addition to today, follow up opportunities to provide input are welcome and available
- Feel free to join the currently scheduled follow-up Webinar on Oct 17th, 10-11:30 Eastern
  - https://www.aviationsystemsdivision.arc.nasa.gov/research/atd2/remote-demos/index.shtml

Applying End State Design Maxims to the Current Day System

- The community represented at this venue is very familiar with current and emerging needs
  - The information on this slide provides some additional necessary future system background
- The Unmanned Aircraft System (UAS) Traffic Management (UTM) system design paradigms are now commonly used worldwide in the new entrant field
  - The following slides walk through initial nearer-term examples of what applying paradigms from UTM might look like when projecting this onto the current day Air Traffic Management (ATM) system
  - The process of applying UTM paradigms to the current NAS has been called “UTM Inspired ATM” (see criteria below)
  - While these slides focus on nearer-term notional examples only, the larger objectives extend beyond NextGen timelines
- Criteria and design maxims applicable to the current day NAS that are being used in formulation
  - Maximize collaborative planning (the right user- taking the right action - at the right time)
  - Strong focus on early input from users to mitigate problems before they become a disruptive system event
  - Address known needs in the national airspace system (‘move the needle’)
  - Foster a streamlined development process (‘ecosystem for rapid innovation’)
- Drive toward a common set of aviation-wide services using greater data digital exchange in a Service Oriented Architecture (SOA) cloud environment

Operational Example- Maximizing Collaboration & Mitigating a Disruptive Event

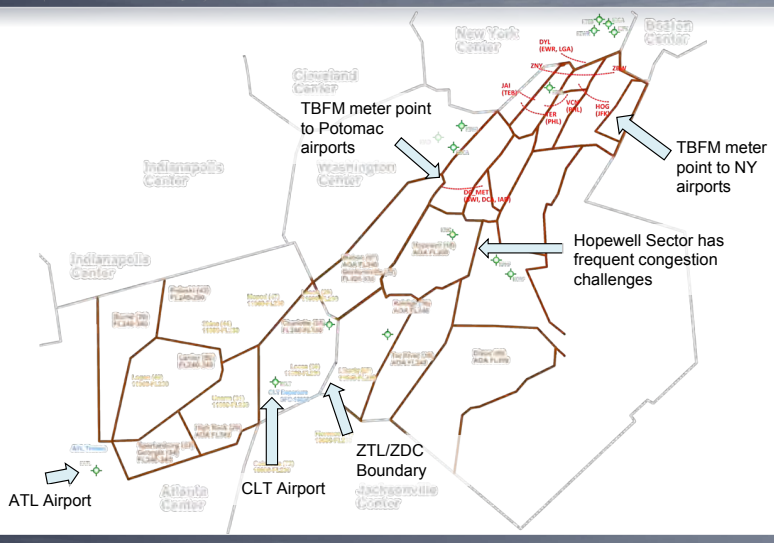

Operational Example- Maximizing Collaboration & Mitigating a Disruptive Event

The current day response to this sector demand/capacity imbalance (i.e.- congestion) is for the ATC to add a Miles in Trail (MIT) to flights that fly through this airspace

Implementing this ATC restriction may create substantial complexity. It may utilize a combination of TFMS, TBFM and/or procedural pass back delay to other Centers which impacts the airport surface.

Operational Example- Maximizing Collaboration & Mitigating a Disruptive Event

In the current day system, is the right user, taking the right action, at the right time?

In the future, might we be able to mitigate this disruptive congestion problem before it becomes an issue in the first place? If yes, what new services and procedures would allow for this?

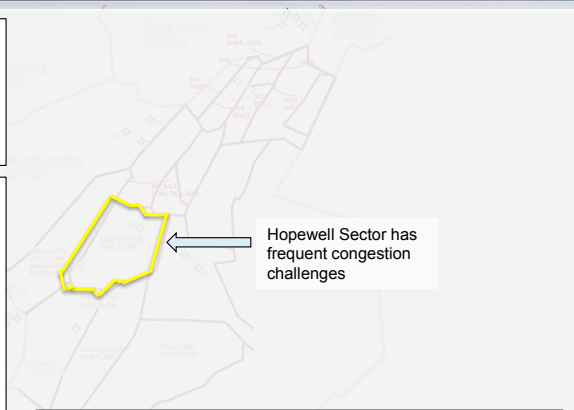

Note: This type of service is likely beneficial to both current day and new entrant airspace users.

How? Let's Start with Current Day SWIM Data Consumption

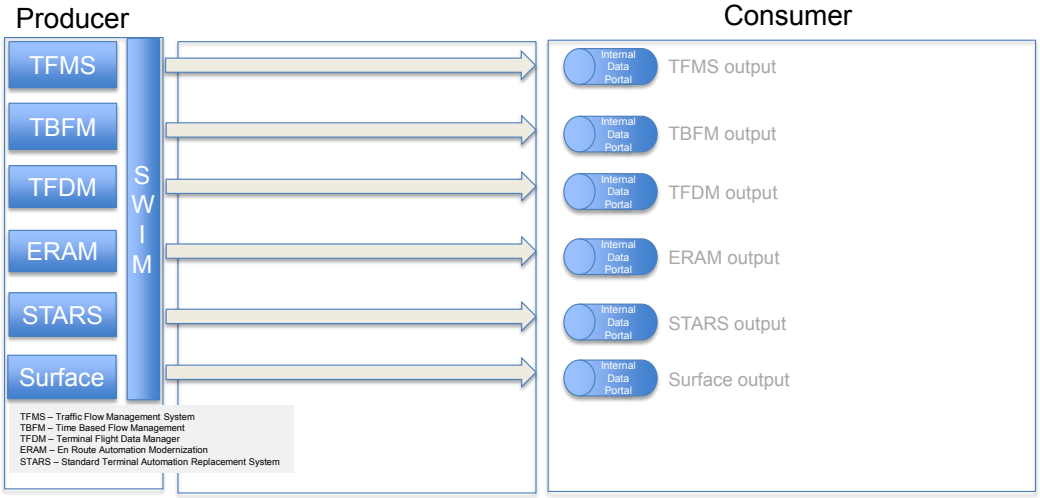

NAS Services in the Cloud- NASA Pathfinder

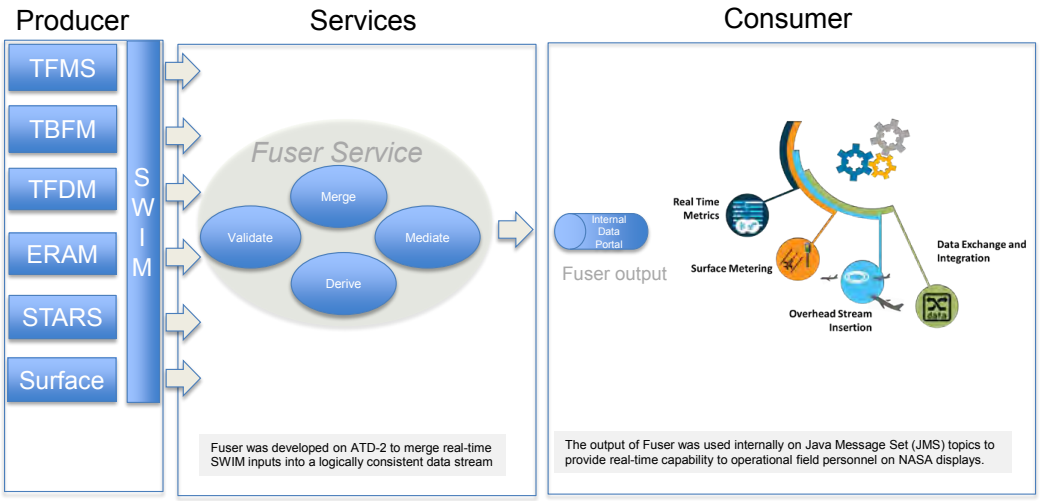

Building on the Pathfinder - NAS Services in the Cloud Examples

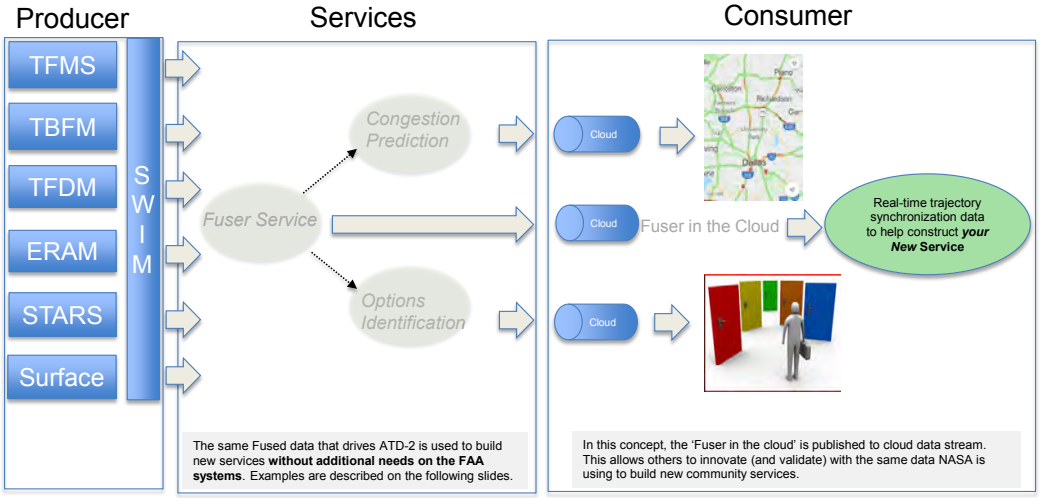

Example - Congestion Prediction Services

Problem

- Traffic congestion forecasting has become commonplace in our modern lifestyle. However, congestion forecasting (and monitoring) is not widely available in the U.S. aviation system.
- To predict congestion accurately, good estimates are required of both the capacity and the demand for the specific NAS resource (e.g. gate, runway, fix, etc.). Until recently, actual capacity or local demand was not accessible outside of traffic manager's visibility.

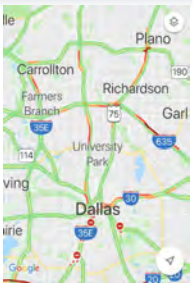

Workshop Traffic Prediction

How Data-Driven Services Will Help

- In recent years, new data sources and analytical methodologies have emerged that enable this capability. In this example, the NAS congestion prediction service would be a suite of foundational services that will provide real-time indications of congestion for NAS planning purposes and measure their accuracy.

Research Challenge

- Sector congestion. Can predictive analytics yield greater accuracy than purely deterministic algorithms?
- Surface runway congestion. Can new TFDM data help localized (e.g. flow dependent) predictions?
- Can TBFM metering delay pass back from the Center boundaries and airports be accurately predicted?
- Can new sensor data, not currently available to aviation, be used to improve capacity predictions?
- Note: The best algorithms for NAS resources ***may*** vary by domain and location. Can services be used to consolidate estimates of flight congestion by disparate domains into one useful representation?

Example - Flight Operator Options Identification Services

Problem

- A flight may be subject to multiple traffic restrictions and delay conditions which span ATC domains, Operators domains and mitigation strategies (e.g. substitute, re-route, fuel-efficient hold).
- From the perspective of flight operators, substantial experience exists with strategic restrictions (i.e. – substituting TFMS EDCTs). However, other flight alternatives are either new (surface substitution), invisible to operators (re-route alternative) or too complicated to work through quickly in operations.

FLT234 has 5 options (est. savings)

- TOS departure fix (12)
- TOS arrival fix (4)
- TOS low sector (8)
- Surface substitution (4)
- Strategic substitution (30)

How Data-Driven Services Will Help

- In this example, the flight operator options identification service would be in the unique position of being cognizant of all the delay producing conditions that exist in the NAS. This 'birds eye' view of the NAS has only recently become available with SWIM and other emerging initiatives.

Research Challenge

- What is the complete list of options available for a specific flight and the time window for decisions?
- Which options are mutually exclusive? (i.e.- can't take options A&B, options B&C can be together)
- TBFM. Can we route around high TBFM delay to improve predictability without unexpected impacts?
- TFDM. Which flights are exempt (unavailable) for surface substitution? Which have multiple delays?
- Would consistent, automated options recommendations lead to greater predictably/reliability?

Many Other Examples!!!

- Many other possibilities exist within this solution framework
  - Other examples are provided below (dozens exist). Solid progress is a key to success.
  - Ideas brought forward will be vetted with the aviation community in a collaborative fashion
  - We do not have time to discuss all the potentials today

- Input is desired on both the methodology and potential specific data-driven services

Leveraging Data Science

Platform for Predictive Analytics

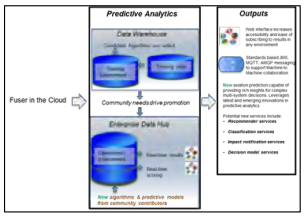

Innovative Pilot Tools

Pilot Communication Services

New sensors to improve planning data

Improving Capacity Estimation for Snow Affected Airports

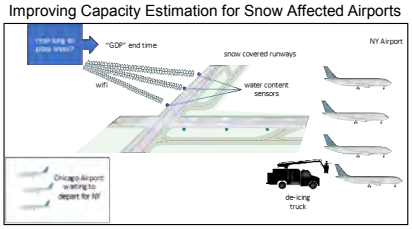

Formulation Input and Feedback During the Workshop

- Overall concept
- Specific services you would like to see
- Willingness to collaborate and good venues to do so
- Future discussion topics

Next Steps

- Collect formulative input from you during or after this conference
  - Please speak to any of the presenters here today after this breakout
  - If desired, we can schedule a follow up with your team
  - NASA will collect this formulative input, consolidate it, and provide a status of the formulation in the Webinar scheduled below

- Additional input can be sent to Al.Capps@nasa.gov

- NASA is hosting a follow up Webinar to continue the dialog
  - The information collected above will be discussed with the community
  - Oct 17th, 10-11:30 Eastern
  - https://www.aviationsystemsdivision.arc.nasa.gov/research/atd2/remote-demos/index.shtml

Thank You

Thank you for your input!

We are listening!

National Aeronautics and
Space Administration
NASA
Digital Transformation: AI and ML Strategy
Ed McLarney
11 June 2020
NASA Digital Transformation Team
3/30/20
HEC Needs Assessment Plan B
1

APMC Briefing Outline

- Introduction
  - Executive Introduction: Steve Jurczyk
  - Motivation for DT: Douglas Terrier, Jeffrey Seaton
  - DT Enterprise Framework: J.-F. Barthelemy
- Roadmaps and Implementation Plans
  - Collaboration Strategy: Yuri Gawdiak, Vanessa Wyche
  - Model-Based Anything Strategy: Tony DiVenti, Jill Marlowe
  - Data Strategy: Jason Duley
  - **Artificial Intelligence and Machine Learning Strategy: Ed McLarney**
  - Process Transformation Strategy: Beth DiGiovanni
  - Culture and Workforce Strategy: Nick Skytland
- Center Transformation
  - ACT2 LaRC: Jill Marlowe
  - Enterprise 2.0 JPL: René Fradet
- Discussion and Next Steps: Steve Jurczyk

Artificial Intelligence and Machine Learning Strategy

Ed McLarney, Nikunj Oza, Herb Schilling

June 2020

AI/ML enables machine capabilities to augment human intelligence: autonomous navigation and coordination, bot automation, analysis of vast information sources, intelligent assistants, technology horizon scanning, anomaly alerts, voice-driven interfaces, and more.

AI/ML – Augmenting Human Intelligence

VISION

*Amplify AI/ML mission and mission support benefit by investing in AI/ML leadership, education, research and development (R&D), practical application, and technologies. Serve as a **catalyst** for existing and emerging AI/ML work across the Agency.*

CHALLENGES & GAPS

- Slow and constrained AI/ML cloud access
- Efforts are ad hoc and multiplying
- Inaccessible data
- Insufficient internal AI/ML expertise
- Lack of AI/ML trust and understanding
- Unable to adapt to rate of change

GOALS

- AI/ML enables new NASA missions – **100x robots** for Moon and Mars; **2x sensors**; **50x data quantity** in science; **20-100x airspace capacity**
- Platforms and partners maximize AI/ML benefit – Efficiency of adapt/adopt vs. build **5x more efficient**; crowdsourcing multiplies brains to solve problems
- AI/ML helps optimize workforce effectiveness – Workforce able to do **25%+ more work** per person with AI/ML
- Vibrant NASA AI/ML community & expertise – **Redundant work decreases 3x** and sharing **increases 3x**
- NASA's data is AI/ML-ready – AI/ML-ready data can **reduce prep time by 75%**

AI/ML Success Example: GSFC Heliophysics Using Machine Learning

- Solar Dynamics Observatory satellite takes a high-resolution image of the sun every 1.3 seconds, over 200 Million images across 10 years
- Data is too massive for scientists to look at, let alone analyze
- Partnered with NVIDIA to develop customized deep learning algorithms to recognize solar features, track, and highlight trends
- Enables multiple complex Heliophysics questions to be explored
- This scale of analysis would be impossible without machine learning

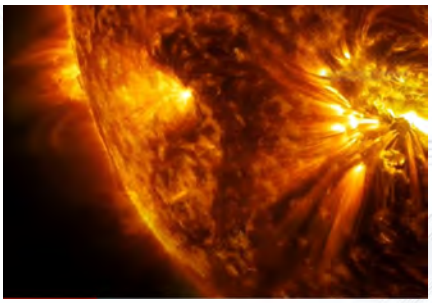

Technical Challenges (why ML was needed):
- Over 18 Petabytes of high-resolution solar images archived
- Expert human analysis only scratched the surface of the data leaving nearly all images unanalyzed

Approach (how ML was employed):
- Created original software to adapt NVIDIA hardware for image recognition using machine learning techniques
- Train the custom algorithms to identify solar features
- Verify and validate with subject matter experts and proven techniques

Principal Investigator: Michael Kirk, PhD, GSFC

Quantifiable Benefits:
- Enables exploring multiple complex Heliophysics questions, where a single problem outstripped human capacity
- Working toward predicting solar storms, in turn leading to early warning for space and terrestrial operations

Lessons Learned, Best Practices:
- Commercial techniques are available but must be curated to NASA use cases
- Incorporating physics into AI/ML development is necessary from step one for science success
- NASA requires data science, AI/ML expertise combined with domain expertise for optimum use of emerging capabilities

AI/ML Success Example: Machine Learning for Aeronautics

- Industry mostly uses rules over 2-3 variables to find safety issues
- Can only find known safety issues. Want to find unknown problems
- Developing data-driven methods for
  - Anomaly detection
  - Precursor identification: Finding precursors to particular safety issues
- Using data from aircraft operations (Flight Operations Quality Assurance—-FOQA), trajectory data on takeoffs and landings at airports and metroplexes

Technical Challenges (why ML was needed):
- Substantial time series data (150+ variables, many thousands of flights daily)
- Very few operationally significant anomalies---low false alarm rate difficult to achieve, but necessary

Approach (how ML was employed):
- Developed new algorithms for anomaly detection and precursor identification over heterogeneous data
- Developed active learning algorithm to learn from domain expert feedback to reduce false alarm rate
- Working with MITRE to deploy algorithms for use with 50+ airlines as part of Aviation Safety Information Analysis and Sharing (ASIAS) partnership

Principal Investigator: Nikunj C. Oza, PhD, ARC

Quantifiable Benefits:
- Have identified several anomalies deemed safety relevant and not identified by current rule-based methods. Used to modify rule-based method at one airline.
- Working with four airlines and MITRE to deploy algorithms for operational use

Lessons Learned, Best Practices:
- Allow substantial time for simple data analysis and cleaning before applying algorithms
- Existing commercial and open-source tools are almost never sufficient---machine learning algorithm R&D is required for application
- Machine Learning expertise and domain expertise are needed for both data preprocessing and results interpretation

DRAFT

AI/ML DT Capability Roadmap

Track	Work to Date	Near Term (2 years)	Long Term (5 years)	Future State
Tools	• 100+ AI/ML projects Agency-wide: Finding planets, EVA suit analytics, storm warning • HQ OCIO Development Environment • Agency cloud foundations (compute, storage, security)	• COI prioritizes/recommends enterprise-wide tools • Add AI/ML to existing cloud services • Establish library of trained ML models • Catalog of AI/ML tools • Expand teaming with internal users (ex. high-end computing, autonomy)	• AI/ML solutions permeate mission and mission support work • Establish integrated hybrid AI/ML platform (cloud, on-prem, multi-vendor, etc.) • AI/ML toolsets span mission lifecycles • NASA's data is AI/ML-ready	• Workforce optimized; bots conduct mundane tasks • AI-supported autonomous systems are reliable partners for human exploration • AI-powered decisions expand options for planetary missions • Virtual research partner scans for, analyzes, and integrates data • AI enhances air traffic and aircraft operations of the future
Partnering	• Google Cloud Platform AI/ML early win project • Experimentation with Microsoft and Amazon • Workshops with industry • Smart partner approach	• Public-facing NASA AI/ML COI • Add creative partners to COI (industry, academia)	• Well-established NASA, industry, OGA, academia COI • Robust tech and partner scanning for external AI/ML solutions	• Pervasive partnerships to seamlessly adapt external work • NASA recognized as leader in AI/ML in mission areas
Workforce	• Local training, speakers • Monthly Agency AI/ML COI telecon • Began cultivating Agency-wide AI/ML community by planning DT	• Scale training Agency-wide • Establish internal cadre of data scientists • Leverage collaborative learning with industry and academia (i.e. Frontier Development Lab) • Continue growing NASA AI/ML community • Policy for ethical AI/ML use	• Expand cadre of data scientists to external partners and citizen scientists • NASA AI/ML training, incorporated with overall transformation training • Easy, ethical AI/ML use	• Strong cadre of data scientists • Employees confidently rely on AI/ML tools for support

DRAFT

AI/ML Near-Term Roadmap

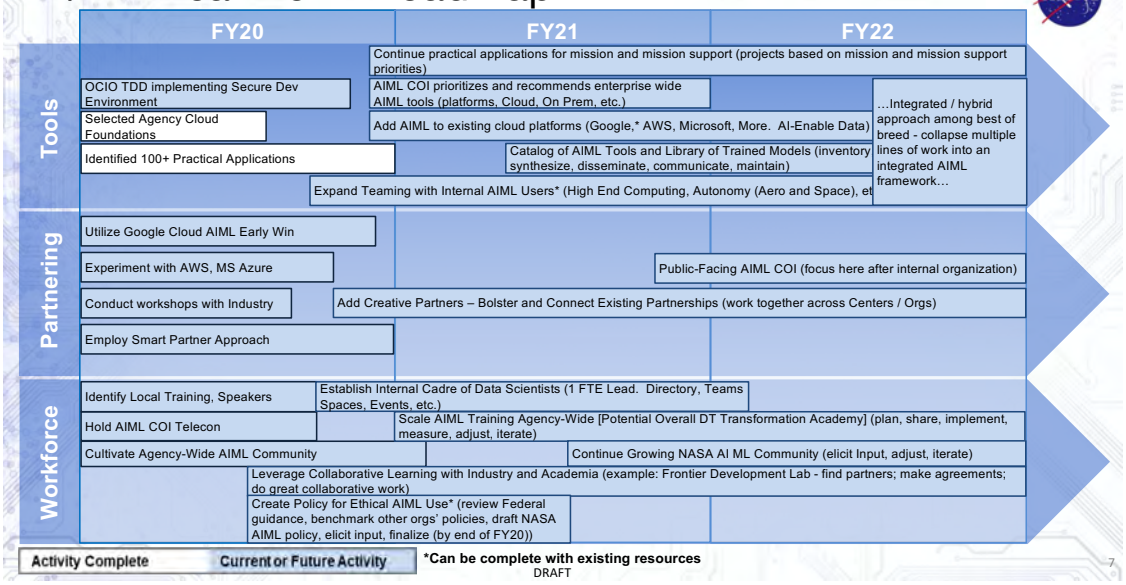

AI/ML Implementation Advances ALL Areas – Examples Highlighted

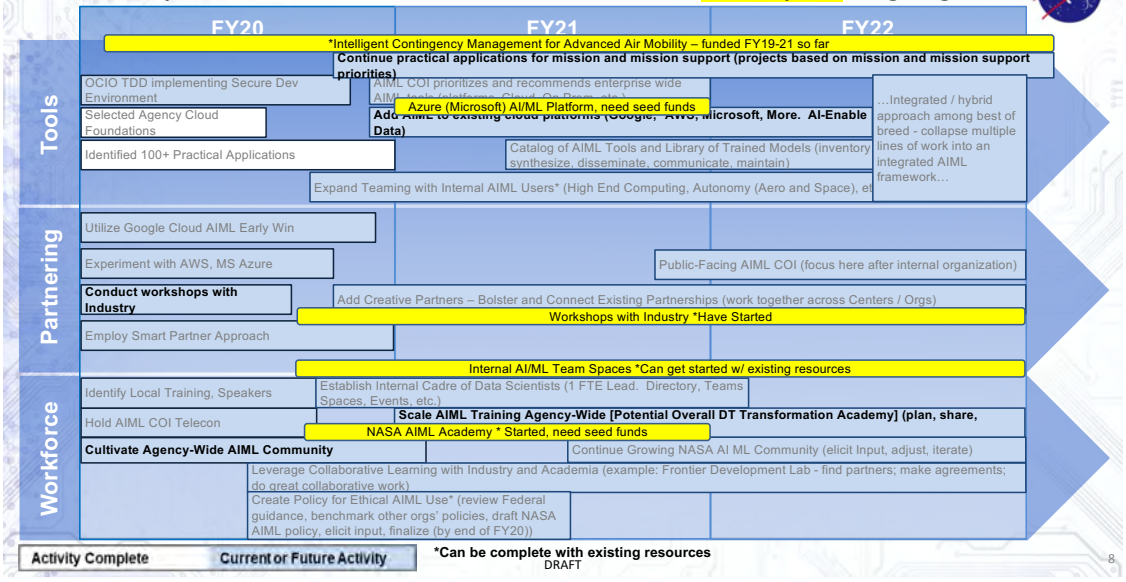

Initial AIML Practical Applications

Proposed via Center / Mission Reps – actively pursuing various funding sources

"Mission" is the relevant mission and does NOT mean the Mission has agreed to fully fund the project

These leverage the AIML foundation (next page)

Estimates in backups

Short Name	Mission	Funding
SLS Propulsion Green Run AIML Analysis (Do No Harm)	HEO	**Part Funded**
Measuring Human Performance Contributions to Aviation Safety	Aero	Proposed
FDL Challenge: Astronaut Human Performance + Nike Techniques	HEO	Proposed
Training Data Set for Earth Science Disasters	Science	Proposed
Hurricane Wind Speed Forecasting Using AI	Science	Proposed
Intelligent Design and Engineering of Aero Systems	Aero	Proposed
NLP Text Analytics for Requirements Optimization and Deconfliction	HEO	Proposed
OCHCO Human Capital IT Strategic Roadmap Transformation	MSD	Proposed
OCHCO AI Platform for Workforce Analytics	MSD	Proposed
NLP Text Analytics for Categorizing SES Job Applicants	MSD	Proposed
Intelligent Contingency Management for Advanced Air Mobility	Aero	**Funded**
SCADA for Hot Water Distribution	MSD	Proposed
Inverse Design of Materials - Infrastructure	All	Proposed
Additive Manufacturing Augmented by AI	All	**Part Funded**

DRAFT

Initial AIML Foundations

- Top list proposed as a portfolio of small core investments.
- Funding sources will depend on governance model

- Bottom list - larger projects for individual advocacy
- Again, funding source will depend on governance model

Activity	Funding
Google Cloud AI/ML	**Part Funded**
Leverage Secure Development Environment	**Part Funded**
Internal AI/ML Team Spaces	Proposed
Inventory of AI/ML Work Across NASA	Proposed
HEC Teaming	Proposed
AWS AI/ML	Proposed
GitLab / GitHub for AI/ML Code Sharing	Proposed
Establish an Overall AI/ML Lead	**Part Funded**
Continue Monthly AI/ML COI Talks	**Part Funded**
AI/ML Communication Program	Proposed
NASA AIML Academy	Proposed
Leverage Crowdsourcing for AI/ML	Proposed
Workshops with Industry	**Part Funded**
DT CONOPS Integration Workshop	Proposed
Externally-Facing AI/ML COE	Proposed
Dedicated High Performance Storage for AIML on NCCS (HEC)	Proposed
Transform On-Prem AI/ML Toolset (Dashlink - 500 existing users)	Proposed
AI-Enable NASA's Data – Teamwork with Data Thrust	Proposed
Azure Cloud AI/ML	Proposed
Combined AI/ML Cloud - AWS, Azure, GCP	Proposed 22
Trained Model Library	Proposed
Amplify Rearch Accelerators like FDL	Proposed
Expand to an Overall Virtual DT Academy	Proposed

DRAFT

Intelligent Contingency Management for Adv. Air Mobility

Project Underway: Use AI to safely fly a mission under all vehicle-allowable weather conditions, in a high-density airspace complex urban environment, including off-nominal situations without direct human intervention

Roadmap Alignment: Practical AIML Mission Application
Contributing Community Members:
- *LaRC Aeronautics (Irene Gregory) – LaRC D-316*
- *LaRC Data Science Team (Newton Campbell)*
- *Overall 8 SMEs involved ~4 person equivalent (FTE / WYE mix)*

KEY MILESTONES
- FY19: Initial framework and plan; Algorithm Demonstrations
- FY20: UAM Vehicle Simulation Developed; Contingency Management (CM) Tool Integration with UAM Simulation; Evaluation under well-defined hazard
- FY21: In-Simulation contingency management demonstration for UAM-relevant vehicle under hazard
- FY22: TBD

BENEFITS & RISKS
Benefits
- Enables robust autonomous UAS piloting, to include ability to react to situations beyond trained heuristics
- Contributes to safe UAS / UAM concepts

Risks
- Technical risk is moderate – this is new technical ground; mitigate via adoption of techniques from parallel fields
- Ethical risk is moderate – must train and refine models to operate UAS/UAM systems IAW safe, ethical guidelines. Mitigate via extensive modeling, real-world V&V, and digital forensics if accidents do occur

STAKEHOLDERS & FUNDING
- Aeronautics investing ~ 4 person years for 3-4 years
- Explicit AIML investment in PhD AI contract expert full time duration of project (~$250K / year)
- AIML intends to help support the team with selected cross-cutting enablers, such as expertise and access to AIML platforms described elsewhere in this brief
- Approach: Aero expert (Irene G) **augmented by AIML core (PhD AI expert)**

Google Cloud AIML Access

Obtain secure access to AIML services from Google Cloud Platform; provide initial training

Roadmap Alignment: Tools (Google is one of the major cloud AIML experts)

User Community: NASA-wide – data scientists, researchers, scientists, engineers, business analysts

KEY MILESTONES
- Sept 2019 – Feb 2020 – Early testing and planning
- Feb 2020 – IV&V assessment of GCP AIML services
- Mar-Jun 2020 – Security work
- July-Aug 2020 – Initial GCP AIML availability, training, hackathon

BENEFITS & RISKS
Benefits
- Secure, authorized method to get cloud AIML
- Project pays for itself if 10ea $50K AIML projects use the platform rather than building their own

Risks
- Risk to implementation: Availability of core GCP for NASA (storage, compute, etc.) – CSSO handling it!
- Risk of inaction in this area: Users invest in non-secure, redundant, expensive AIML tools in stovepipes

STAKEHOLDERS & FUNDING
- LaRC contributed $75K; Project partially funded
- LaRC OCIO investing .5 FTE all FY20
- Risk – Need $175K to scale availability
- Project and DT AIML leads to negotiate / recommend operationalization plan / owner for long-term O&M
- Project Lead: Charles Liles, in teamwork with MSFC, JSC, Agency TDD, and GSFC IV&V

Connections

Roadmap Contributors

- HQ: Tools & standards, Data de-duplication, ML Tags
- ARC: EVA analysis; Asteroid Threat Detection, Solar Flare Warning
- GRC: Jet Noise, Sat Signal Propagation, Propulsion Sys Design
- GSFC: Computer Aided Discovery, Helio; Proposed evaluation text analytics; Hyperspectral real time recalibration
- JSC: Google Cloud, ISS Object Detection, Astronaut Coach
- JPL: UAV Search/Rescue imagery, Galaxy formation, Mars 2020 planning; launch several ML pilots across JPL
- LaRC: Printed parts cert, Space Radiation, UAS coordination
- NSSC : First NASA Digital Employees (Washington Bot)
- HPC Community: Embracing / intersecting with AI / ML at many Centers
- MDs: SMD/AIST portfolio of $14M; ARMD/CAS

Key Implementation Stakeholders

- OCIO Transformation & Data Division (TDD) – on board, contributing co-lead
- ARMD – Potential interest in co-leading pending DT ESC governance decisions
- ARC – Co-lead personnel identified
- Funding depends on DT governance decision
  - Distributed = need co-lead contributions
  - Central = core from DT funding
- Other Key Stakeholders
  - AIML is relevant to all Missions, Centers, and Mission Support Orgs
  - OCIO Programs, OCIO MAP
  - DT Data, MBx, Collaboration, Culture/Workforce
  - Key partners in industry and academia

Backups

AIML Successes FY20

- Building an AIML community across NASA via DT teams
- Beginning to grow AIML cloud platform access: AWS, Google, Microsoft
- Started virtual AIML training agency-wide (TDD, GSFC, LaRC, IBM, more)
- Ongoing Practical applications at many Centers / Missions (a *small* sample):
  - LaRC - Aero Revolutionary Air Mobility / UAM intelligent contingency management
  - JSC – Autonomous Suit Fault Checking; Portable Life Support Simulator – ML analysis; IoT
  - JPL – Mars image recognition; Climate model analysis; Object discovery in sky surveys
  - MSFC – Coronavirus AI impact analysis on air traffic & shipping
  - GSFC – Pilot projects in Helio-Analytics and more
  - ARC and GSFC – investments in shared HEC AIML hardware, tools, and assistance

Technical Challenges (why DT was needed):
- Grass-roots adoption of AIML can be accelerated
- Insufficient AIML expertise within NASA

Approach (how DT was employed):
- Built teamwork and partnerships based on DT connections
- Got started by sharing IAW existing education resources
- Leveraging newly-made connections to share techniques, advice, scars, tools

Given all this, why invest in AIML as a DT Thrust?
- ***Amplify, accelerate, align, and scale NASA AIML impact***

Quantifiable Benefits:
- Baseline and forward metrics are part of implementation plan FY21…

Other Benefits:
- Less rework: AIML practitioners connect and share among projects
- Skills: Several shared virtual learning events, even during COVID
- Effectiveness: Early adopters excited about AIML platform testing

Lessons learned, Best Practices:
- Communicate, communicate, communicate!
- Start with available resources and grow
- Show value and advocate for shared, effective, efficient platforms
- Plenty of AIML to do; collaborate instead of competing
- Commercial AIML must be trained / adapted to NASA needs

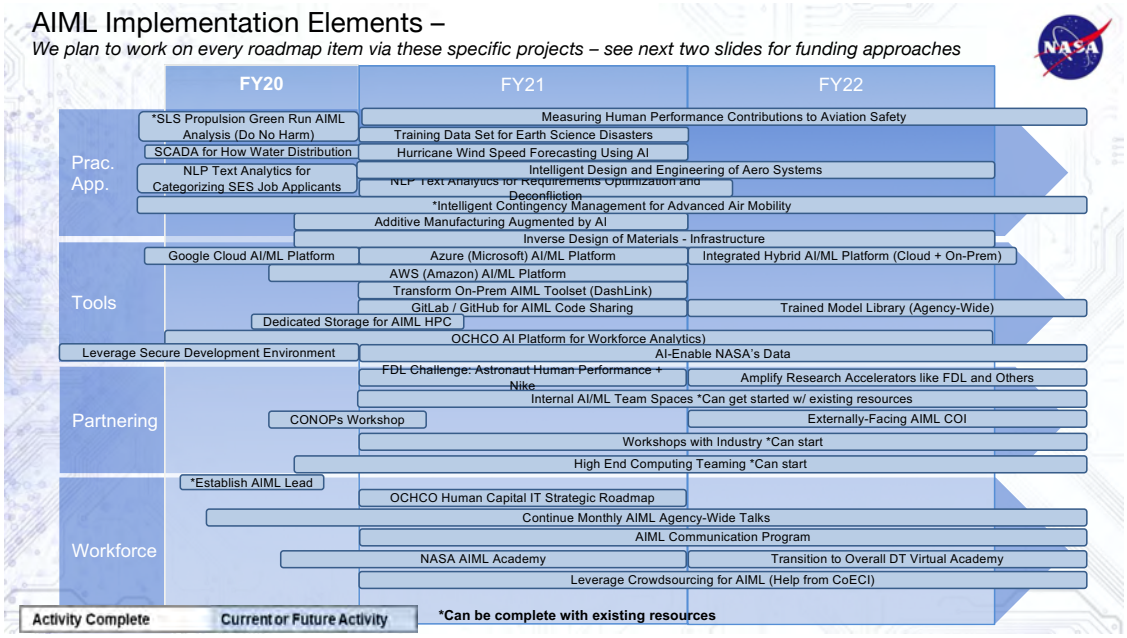
AIML Implementation Elements –
We plan to work on every roadmap item via these specific projects – see next two slides for funding approaches
FY20
FY21
FY22
Prac. App.
*SLS Propulsion Green Run AIML Analysis (Do No Harm)
SCADA for How Water Distribution
NLP Text Analytics for Categorizing SES Job Applicants
Measuring Human Performance Contributions to Aviation Safety
Training Data Set for Earth Science Disasters
Hurricane Wind Speed Forecasting Using AI
Intelligent Design and Engineering of Aero Systems
NLP Text Analytics for Requirements Optimization and Deconfliction
*Intelligent Contingency Management for Advanced Air Mobility
Additive Manufacturing Augmented by AI
Inverse Design of Materials - Infrastructure
Tools
Google Cloud AI/ML Platform
Azure (Microsoft) AI/ML Platform
Integrated Hybrid AI/ML Platform (Cloud + On-Prem)
AWS (Amazon) AI/ML Platform
Transform On-Prem AIML Toolset (DashLink)
GitLab / GitHub for AIML Code Sharing
Trained Model Library (Agency-Wide)
Dedicated Storage for AIML HPC
OCHCO AI Platform for Workforce Analytics)
Leverage Secure Development Environment
AI-Enable NASA's Data
Partnering
FDL Challenge: Astronaut Human Performance + Nike
Amplify Research Accelerators like FDL and Others
Internal AI/ML Team Spaces *Can get started w/ existing resources
CONOPs Workshop
Externally-Facing AIML COI
Workshops with Industry *Can start
High End Computing Teaming *Can start
Workforce
*Establish AIML Lead
OCHCO Human Capital IT Strategic Roadmap
Continue Monthly AIML Agency-Wide Talks
AIML Communication Program
NASA AIML Academy
Transition to Overall DT Virtual Academy
Leverage Crowdsourcing for AIML (Help from CoECI)
Activity Complete
Current or Future Activity
*Can be complete with existing resources

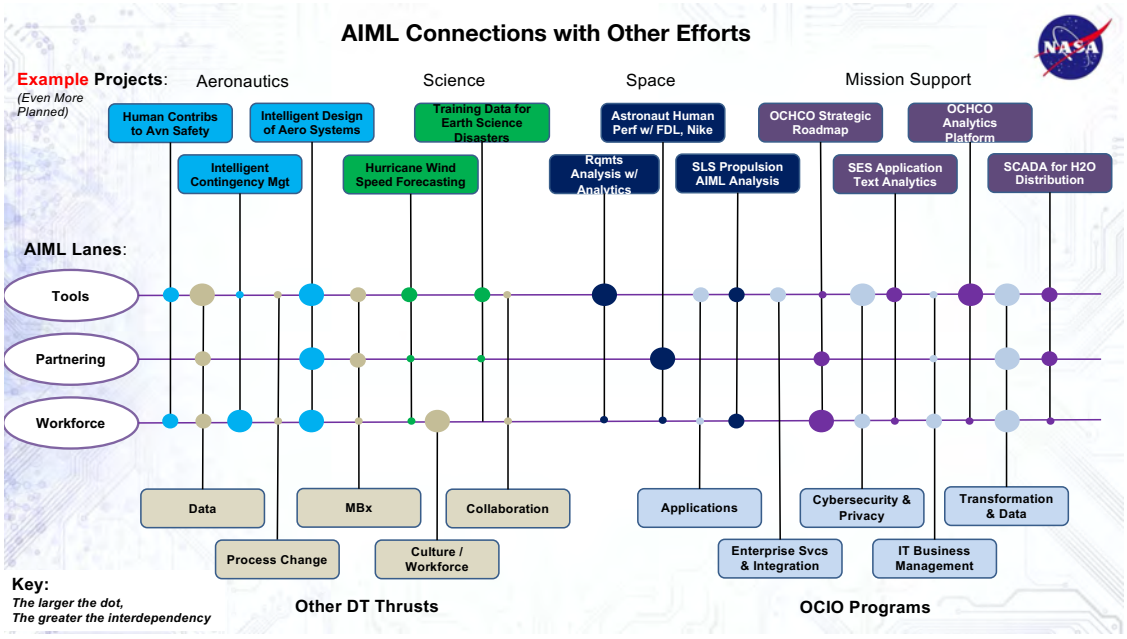
AIML Connections with Other Efforts
Example Projects:
(Even More Planned)
Aeronautics
Science
Space
Mission Support
Human Contribs to Avn Safety
Intelligent Design of Aero Systems
Training Data for Earth Science Disasters
Astronaut Human Perf w/ FDL, Nike
OCHCO Strategic Roadmap
OCHCO Analytics Platform
Intelligent Contingency Mgt
Hurricane Wind Speed Forecasting
Rqmts Analysis w/ Analytics
SLS Propulsion AIML Analysis
SES Application Text Analytics
SCADA for H2O Distribution
AIML Lanes:
Tools
Partnering
Workforce
Data
MBx
Collaboration
Applications
Cybersecurity & Privacy
Transformation & Data
Process Change
Culture / Workforce
Enterprise Svcs & Integration
IT Business Management
Key:
The larger the dot,
The greater the interdependency
Other DT Thrusts
OCIO Programs

NASA AI/ML Strategy – Elevator Pitch

Context:

- NASA is forming a Digital Transformation (DT) Strategy and Roadmap, led by the Office of Chief Technologist and Office of Chief Information Officer. This strategy includes AI/ML as one of six key strategic thrusts.
- NASA has a rich history of applying artificial intelligence (AI) to our hardest problems, such as autonomous behaviors in Mars rovers, deep analysis of space suit data, or image analysis to understand material strength. With the advent of powerful, plentiful, and affordable AI in business and industry, NASA is crafting a strategy to use AI as an accelerant for all NASA missions and business functions.

Strategy: As part of NASA's overall Digital Transformation, NASA's AI strategy includes:

- **Apply:** Solve relevant mission and mission support problems via AI / ML.
- **Teamwork:** Lead and synchronize NASA AI/ML via an open Agency AI / ML community.
- **Reskill:** Expand AI training, education, hiring, and retention across the workforce.
- **Tools:** Assess, recommend, and establish AI / ML platforms for NASA-wide adoption.
- **Data:** AI-enabled! Establish secure, authoritative access to the right data.
- **Outreach:** Make selected data and problems available for public / partner AI / ML work.
- **Adapt:** Leverage industry AI / ML work and adapt it to NASA use rather than reinventing.
- **Scale:** Plan to promote selected AI / ML capabilities from pilot to production operations.

The AI/ML team is from **across** the Agency with over 35 active members; additional contributors are always welcome.

Contact:
Ed McLarney (edward.l.mclarney@nasa.gov), **Nikunj Oza** (nikunj.c.oza@nasa.gov) or **Herb Schilling** (hschilling@nasa.gov)

Planned Outcomes from Near Term Roadmap

- Benefits
  - Amplify NASA AIML capabilities to enable planned missions and maximize existing workforce
    - Aeronautics, Space Exploration, Science missions embrace AIML as core enabler; DT AIML as jump-start with Missions
    - AIML skills advancement in leaders, data scientists, researchers, scientists, engineers, business analysts
    - Growing set of AIML core toolsets available, affordable, standardized, interoperable
    - NASA's Data is on the road to being AI-ready
  - What's the yield.
    - ~3 Interoperable Standardized AI platforms vs ~500 by-project custom AIML stacks. At least 5x more cost efficient to make common investment rather than have each project stand up all its tools
    - 50% of workforce learns AIML relevant to their job; on path to >25% greater overall effectiveness
    - Automate the most redundant 10% of business processes
    - Reduce data prep time from 80% to 50%... On the way to 10-20%
- After 3 years, NASA is well on the way to making AIML impactful to all NASA's missions and mission support functions. The workforce is knowledgeable of AIML and NASA has become excellent at adopting its new methods

DRAFT

AI/ML Team Members

Ed McLarney – LaRC
Nikunj Oza – ARC
Herb Schilling - GRC
Dwayne Armstrong - SSC
Kofi Burney
Mark Carroll - GSFC
Marshall Cottrell - AppDAT
Dan Crichton - JPL Programmatic data science lead
Gary Crum – GSFC - Science data processing
Marcus Friske - AppDAT
Martin Garcia - JSC
Yuri Gawdiak – ARMD A/Dir for airspace ops and safety
Phyllis Hestnes - GSFC SW engineering division
Dave Kelldorf – CTO IT, JSC
Mike Little - HQ Earth Science, at GSFC. AIST PM for Earth Science
Rodney Martin – ARC deputy for data science group
Manil Maskey, MSFC AI/ML SME
Chris Mattman - JPL
Nargess Memarsadeghi – GSFC - Science data systems
David Meza - Agency WF
Lien Moore, MSFC AI/ML SME
Leigh Anne McMahon, MSFC AI/ML SME
Shubha Ranjan – ARC – Data Science Lead, HECC
Laura Rogers. LaRC / ESTO analyst / expert
Tom Soderstrom - JPL
Anthony Shoemaker - AppDAT
John Sprague – Acting Agency CTO-IT
William Thigpen - ARC - Manager, High-End Computing Capability
Brian Thomas – Agency data scientist and strategist
Heather Thomas - JSC
Barbara Thompson – GSFC Data Science
Lui Wang, JSC AI/ML SME
Prentice Washington - HQ
Shan Zeng - LaRC Science
Keith Krut - HQ OCHCO

AIML Team Recommendation to APMC

- ARMD and ARC, with OCIO and AIML Strategy Team support, to develop an initial Agency implementation plan for the **Artificial Intelligence and Machine Learning Strategic Thrust area**, aligned with agency level roadmaps; this should build on accomplishments and planned content.
  - Interested members from the AIML Strategy Team coordinating / advocating for volunteers to take implementation plan ownership:
    - ARMD due to AIML necessity in future air vehicles and air traffic control
    - ARC due to enduring Center work and expertise in AIML
    - OCIO TDD as OCIO's focal point for digital innovation
  - Due: April 2020

AIML Implementation Co-Leads

- Roles, Personnel, Notes
  - ARMD: SWS, ATMX, AAM – Impl should help organize AIML for these; cross-pollination among them and other parts of NASA; engagement with external partners; tech transfer into and out of NASA. FY21 – ask for $250K as part of crowd-funding. See Yuri's chart. Tiffany: would like to see OCIO investment too. Maybe for cyber too... Splunk, Collab, also mission-embedded.
  - ARC: Trey Smith, Jeremy Frank, Justin Hash. Nikunj: $250K might be hard, but some fraction realistic; could also contribute Agency abstraction of AIML services. (In-kind resources)
  - OCIO TDD: Ed. FY21 $250K investment. Must get detailed plans in place to help make advocacy

Green Run Post-Test Data Analysis AIML Analysis Quad Chart

Proposal: Proof-of-Concept to parallel post-test data analysis of SLS Green-Run Test in Summer 2020

Roadmap Alignment: AIML Practical Application;

Contributing Community:

- *MSFC Data Scientists, Software Test Engineers and Engine Test Engineers*
- *External Expert Partner*
- *Agency-Wide AIML Volunteer Army*

KEY MILESTONES

- March/April 2020 – Refine concept; OK with core analysis team; build team
- April– July 2020 – Define AIML approach, assemble AIML toolsets, train tools on related data (prior Engine tests, etc.)
- July – Green-Run performed at Stennis Space Center
- Late Summer/Early Fall – perform analysis work, in parallel to defined SLS activity on a non-interference basis.
- Schedule may change with virus

BENEFITS & RISKS

Benefits

- Provides additional insight into trends or anomalies; complements traditional analysis…NOT on critical path
- Identify areas for greater focus in traditional analysis

Risks

- Risk to implementation: Novel volunteer approach
- Inaction: Potential to miss anomaly in massive data

STAKEHOLDERS & FUNDING

- Pursuing external expert Partner, with history of AIML tool development at GSFC
- MSFC Data Science and Software Test Engineers
- AIML Volunteer Army ~$125K worth of brainpower
- Potentially Seeking seed funds to enable teamwork and tool access for the overall team
- Lead: David Reynolds/MSFC; PI Brandon Steele/MSFC

Measuring Human Performance Contributions to Aviation Safety

Proposal: Conduct multi-data analysis of human actions that keep flights within safe operating parameters

Roadmap Alignment: AIML Practical Application, MBx

Contributing Community:

- *TTT, SWS*
- *LaRC / D318, Agency Data Science, DT AIML*
- *ARC / Code TH, Code TI*
- *Partnerships: American Airlines, The LOSA Collaborative*

BENEFITS & RISKS

Benefits

- Traditional flight safety analysis focuses on what went wrong in rare incidents / accidents
- This analysis seeks to learn from what went right in a huge set of safe flights, with the intent of eventually automating selected currently-human tasks

Risks

- Risk to implementation: Synthesis across multiple forms of data is complex; AIML can help mitigate
- Inaction: Failure to enable key safety elements in emerging UAS/UAM environment

KEY MILESTONES

- March-September 2020. Plan a multi-year project
- FY21 – V1.0 - Refine problem, develop analysis approach, collect initial data, clean data, integrate models / sims, initial analysis, human interpretation / model tuning
- FY22 - V2.0 - Refine all aspects: data collection, analysis, models/sims, data cleaning pipeline, analysis, human tuning
- FY23 – V3.0 - Final refinements, V&V, Scale / Inject into key UAS / UAM work

STAKEHOLDERS & FUNDING

- FY20 funding for planning via participating organizations' core budgets
- FY21-23 funding primarily from TTT, SWS. Potential seed investment from Agency DT. Anticipated yearly ROM: On-order of $XYZ; details TBD
- Lead: Jon Holbrook (LaRC/D-318)

FDL Challenge for Autonomy and Astronaut Exploration Medical Capability Quad Chart

Capitalize on inaugural Nike Innovation workshop with FDL held in January 2020

Roadmap Alignment: AIML Practical Application;

Contributing Community Members:

- *Yvonne Cagle (NASA Astronaut)*
- *Leveraging Previous FDL challenges on the topic of*
  - *Generation of Simulated Biosensor Data*
  - *Harnessing AI to support medical care in space*

BENEFITS & RISKS

Benefits

- Gain further insight in exploring cognitive technologies to improve human health for the future of space exploration.
- Integrate applied apparel into human adaptation both in space and on Earth
- Leveraging analogue training of elite athletes as proxy authoritative data source,
- Refine spin-off benefits for galvanizing the masses to train like the humans who aspire and endeavor for space

Risks

- Risk to implementation: Relies on establishing formal partnering arrangement with Nike
- Potential data access issues due to sensitivity associated with Medical PII

KEY MILESTONES

- Key sub activities and timeline
- May 2020 – Refine concept; OK with core analysis team; build team
- June – Define AIML approach; request data feeds; assemble AIML toolsets
- July – Recruit participants and stage outreach campaign
- August - Demo

STAKEHOLDERS & FUNDING

- Seeking ~$250K seed funds to enable teamwork and tool access for the overall team;
- Possible sources of alternate funding (listed in order of probability of success)
- *Contacts being solicited by Yvonne Cagle (see notes section for details)*
- *TRISH Catalyst proposal (Generally less than $100k, but can be up to $350k)*
- *Nike: ~$???K (pending negotation using FDL context)*
- *NIH: ??*
- Idea from: Yvonne Cagle, Rodney Martin
- Potential beneficiaries: HEOMD (Human Health and Life Support), STMD (Various TAs)
- Addresses both HEMOD/STMD technology gaps for a broad variety of physiological space flight hazards
- Compact wearables address HEOMD-prescribed SWaP constraints

Supervisory Control and Data Acquisition (SCADA) for Hot Water Distribution

Intended to supplement [NASA@work Crowdsourcing Contender submision]

Will use real-world time-series and apply open source code bases to develop a Supervisory Control and Data Acquisition (SCADA) system

Roadmap Alignment: AIML Practical Application;

Contributing Community Members:

- *NASA@work*
- *NASA's Smart City Initiative*

KEY MILESTONES

- Key sub activities and timeline
- June 2020 – Refine concept; OK with core analysis team; augment team with WYE developers and engineers
- July – Define approach; finalize data feeds; assemble toolsets
- August – Execution and implemention
- September – Demo to stakeholders

BENEFITS & RISKS

Benefits

- Addresses the question of what procurements or developments could be made that would enable NASA field centers to evolve into Smart Campuses.
- Ultimately aimed at offering each Center an opportunity to participate in the definition of an agency wide Smart Campus concept and technology challenge

Risks

- Identified need and interest in enabling technologies of type that this proposal seeks to investigate (e.g., from ARC's and MSFC's facilities groups).
- ARC's previous experience with piloting new technologies were never funded, so gaining buy-in for participation as a stakeholder will be a heavy lift.
- Unless funds are sufficent to move towards adoption and support from an on-going program, it's hard for them to justify the time required to support R&D initiatives.

STAKEHOLDERS & FUNDING

- Jennifer Groman from OSI identified Condition-Based Manitenance (CBM) as a target/candidate investment area (potentially up to ~$250k)
- Aligns with activities being lead out of LaRC and JSC on: IoT AI/ML Anomaly Detection, Other Mission Support Opportunities, e.g. IoT Architecture and NASA Enterprise Remote Data Collection and Management System(s)
- Aligns with the federal smart city program's vision and priorities (advocacy from Yuri Gawdiak, Agency rep for GCTC - Global City Teams Challenge)
- Dr. Sokwoo Rhee (GCTC lead, Associate Director of Cyber-Physical Systems Program at NIST, and expert in entrepreneurial endeavors of the type that we're embarking upon) expressed interest in offering expertise in the shape in of in-kind support
- Idea from: Rodney Martin, Dan O'Neil, Brent Garber, Keith Swearingen
- Potential beneficiaries: OSI (Office of Strategic Infrastructure), Agency facilities groups and initiatives (e.g. Ames Campus of the Future Initiative, etc.)

Google Cloud AIML Access Quad Chart

Obtain secure access to AIML services from Google Cloud Platform; provide initial training

Roadmap Alignment: Tools (Google is one of the major cloud AIML experts)

User Community: NASA-wide – data scientists, researchers, scientists, engineers, business analysts

KEY MILESTONES

- Sept 2019 – Feb 2020 – Early testing and planning
- Feb 2020 – IV&V assessment of GCP AIML services
- Mar-Jun 2020 – Security work
- July-Aug 2020 – Initial GCP AIML availability, training, hackathon

BENEFITS & RISKS

Benefits

- Secure, authorized method to get cloud AIML
- Project pays for itself if 10ea $50K AIML projects use the platform rather than building their own

Risks

- Risk to implementation: Availability of core GCP for NASA (storage, compute, etc.) – CSSO handling it!
- Risk of inaction in this area: Users invest in non-secure, redundant, expensive AIML tools in stovepipes

STAKEHOLDERS & FUNDING

- MSD funded DT & LaRC $250K for this project
- LaRC OCIO investing .5 FTE all FY20
- Risk – MSD funding not yet received (Jill Marlowe re-connecting with funding sponsors for this and other earmarked early win funds)
- Project and DT AIML leads to negotiate / recommend operationalization plan / owner for long-term O&M
- Project Lead: Charles Liles, in teamwork with MSFC, JSC, Agency TDD, and IV&V

Large-scale training dataset for Earth science disaster events

Biggest bottleneck in adoption of AI/ML (supervised learning) is availability of large-scale training datasets.

Proposal: Create well curated training datasets focused on disaster events (Floods, Cyclones, etc.). These datasets will be used for data science competitions to develop models to detect the disaster events.

Roadmap Alignment: Training Data

User Community: NASA, FEMA, insurance industry, data scientists, researchers, scientists, engineers, business analysts

KEY MILESTONES

- Aug 2020: identification of disaster events for training data
- Dec 2020 – creation of training data for the first disaster event
- April 2021 – creation of training data for the second disaster event
- July 2021 – Data science competition

BENEFITS & RISKS

Benefits

- High quality benchmark datasets generated by experts
- Allows development and evaluation of subsequent models
- Implications to disaster response
- Conduct at least 2 data science challenges to build/evaluate models using the benckmark datasets

Risks

- Risk to implementation: availability of domain experts for the tedious process.

STAKEHOLDERS & FUNDING

- Funding required ~$60K
- Seed funding from IEEE $20K, seeking additional $20K from ESD, will be seeking $20K from DT
- Project will be conducted as part of the student summer project
- Domain experts from the NASA/MSFC and University of Alabama in Huntsville
- Project Lead: Manil Maskey - MSFC

Hurricane wind speed forecasting using AI

MSFC team has developed and evaluated an AI-based hurricane wind speed estimation system (http://hurricane.dsig.net/) (https://www.youtube.com/watch?v=OJUvyAo_QqY&t=3s). The fully automated system monitors the storm and start estimating the wind speed using just satellite images. This proposal is to implement forecasting capability to the existing system to augment current forecasting capabilities.

Roadmap Alignment: Tools, Applications

User Community: NASA, FEMA, NWS, insurance industry, data scientists, researchers, scientists, engineers, business analysts

KEY MILESTONES

- Aug 2020: Data preparation, initial model
- Dec 2020 – Evaluation of models
- March 2021 – Transition model to production
- July 2021 – Ready to evaluate 2021 hurricane season

BENEFITS & RISKS

Benefits

- Automated and objective means of forecasting
- Forecasting at higher temporal frequency
- Implications to disaster response
- Potential to make forecasts of wind speed every 15 minutes compared to every 6 hours

Risks

- Risk to implementation: dependent of large-scale training data, occurrence pf future storms to evaluate against

STAKEHOLDERS & FUNDING

- Total funding required $200K
- Will be seeking $100K from solicitation and $50K from ESD. Will request $50K from DT
- Project will leverage existing code and datasets
- Project will be conducted in collaboration with hurricane science team at MSFC
- Domain experts from the NASA/MSFC and University of Alabama in Huntsville
- Data and code will be made available
- Planning to collaborate with JPL (D. Crichton)
- Project Lead: Manil Maskey - MSFC

IDEAS - Intelligent Design and Engineering of Aerospace Systems

Glenn Research Center has developed PeTaL and is working on developing IDEAS to automate and simplify propulsion system design using AI. The tool will link natural systems design and material discovery to propulsion system design workflow through ML/AI. The tool will also capture knowledge and connect users to resources and potential partnering opportunities.

Roadmap Alignment: Tools, Applications, workforce, teamwork, partnering

User Community: NASA, DOE, DOD, Propulsion industry, data scientists, researchers, scientists, engineers, business analysts

KEY MILESTONES

- Oct 2019: IDEAS workshop with industry+academia to ID NASA role, gaps
- Sept 2020 – turbine design workflow complete
- Sept 2021 – multifunctional AI-generated TMS demo
- Sept 2022 – integrated TMS system demo

BENEFITS & RISKS

Benefits

- Design workflow automation
- 10X reduction in time to discover, design and build
- Enable data driven multi-objective, multi-parameter decision making at system level

Risks

- Lack of infrastructure (cloud, database, security) to scale up and interface with industry and academia

STAKEHOLDERS & FUNDING

- Seeking $285K to hire a PIF; multiple potential sponsors; possible DT co-investment
- Project will leverage existing code and datasets
- Project will be conducted in collaboration with AATT, TAC-P, SBIR, STTR and NASA fellowships
- Natural system design and material discovery tool links from GRC.
- In-house and industry expertise (UTRC, IBM, GE)
- Tools are open source
- Project Idea: Vikram Shyam/NASA GRC

Natural Language Processing (NLP) for Requirements Quad Chart

Proposal: Develop NLP system to parse documents, requirements, specifications for conflicts, duplications, incomplete requirements, and classifications

Roadmap Alignment: Tools, Applications

Contributing Community:

- *MSFC Systems Engineering branches*
- *Datakata LLC*
- *Langley Research Center Data Scientist*
- *NASA Proposal Writing and Evaluation Experience (NPWEE) (Team 11)* Academia
- *MSFC Software Engineering, Strategic office, Chief Engineers office*

KEY MILESTONES

- March 2020 – Presented to MSFC Office of Chief Technologist
- April 2020 – Submission of NPWEE proposal
- ~~May 2020 – Possible TDD awarding of funding~~ *Cancelled*
- May 2020 – Submission of Collaboration Agreement Notice (CAN)
- October 2020 – Target CAN start date

BENEFITS & RISKS

Benefits

- Consistency in definitions/faster reviews
- Streamlining and Augmenting insight functions for projects
- Feedback mechanism to improve for all projects
- Scalable for all projects/agency
- Leverage already created documents
- multiple applications

Risks

- Things that are not Natural Language, such as Tables, Equations
- Ingesting of non-standard documents

STAKEHOLDERS & FUNDING

- Human Lander Systems, Mars Ascent Vehicle, Space Launch System – donating documents, and insight support
- Datakata LLC
- LaRC Data Scientist – IBM WEX technology expertise
- MSFC office of Chief Technologist – 10K (Target)
- NPWE -- 10K and Team members time (Target)
- ~~NASA TDD Labs – 25K (Target)~~ *Cancelled*
- NASA MSFC CAN – 120K (plus company matching) (Target)
- Donated FTE from projects/branches (~.1-.2 per)
- Potential collaboration with JSC, JPL, LaRC
- All NASA program/projects to be able to use
- Looking for 20-35K more funding for resources
- Return on Investment: (see Benefits), Cost avoidance bycatching problems earlier, Faster change tracking for Systems engineers, Standardization, Reduce labor for reviews
- Lead/PI: Scott Tashakkor/MSFC

OCHCO Human Capital IT (HCIT) Strategic Roadmap Overview

OCHCO is developing a high-level, long-term roadmap and business case to modernize Human Capital Information Technology (HCIT) that enables a vision for integrated HC management and an enhanced experience.

To achieve desired business outcomes, OCHCO needs a HC platform powered by emerging technology, flexible to adjust with the speed of innovation, and integrated across HR and enterprise applications and workflows.

Roadmap Alignment: Tools & Architecture, Workforce & Education

User Community: NASA-wide

CASE FOR CHANGE

- HCIT modernization is critical to ensure that NASA can strategically manage human capital
  - NASA's data silos and limited data governance inhibit an integrated talent experience and prevent leaders from having the data needed to make talent decisions and support an enterprise view of mission needs
  - Managers and employees seek a more personalized and simple employee development experience
- HC application stovepipes and data silos create risk of inefficient and ineffective spending on localized technology solutions, adversely impacting NASA's budget and ability to strategically manage a diverse workforce

KEY MILESTONES

- Oct 2019: Project kickoff
- Dec 2019: Visioning workshop with Agency, Center, Mission Support and Mission Directorate leaders
- Jan 2020: Opportunities workshop with functional owners and SMEs to define new ways of doing work
- Mar 2020: Complete an HCIT As Is & Gaps Report
- Apr 2020: Build HCIT Strategic Roadmap & Business Case
- FYs 21-26: Implement HCIT Strategic Roadmap

PLANNED OUTCOMES

- Make the case for investment in a modern HC platform and with an open and flexible architecture to create agility, embed security, and increase ability to use data across the employee lifecycle
- Create a plan to reduce risk of having inaccurate and inaccessible data when making HC decisions by creating a trusted HR data landscape for leadership and managers, as well as reducing time spent on reporting and recording HR data for customers
- **Identify a multi-year approach to incorporate advanced technologies (such as AI/ML) into HC operations and workflows**

OCHCO Capabilities AI Platform for Workforce Analytics

OCHCO People Analytics team is developing a methodology and architecture to map NASA workforce to necessary Knowledge, Skills and Abilities(KSA). The system will implement a combination of machine learning and deep learning models to identify Agency strengths and gaps in KSA.

Current use cases:

- *Skills Analysis for OCIO Map transformation*
- *Predict emerging skills & forecast skill gaps*
- *AIML assisted assignment recommendations*
- *Assist in meeting Cyber Workforce identification requirement*
- *Talent Market place AIML assisted search*

Roadmap Alignment: Practical Application, Tools & Architecture

User Community: NASA-wide

KEY MILESTONES

- Dec 2019 – Jan 2020 – Early testing and planning
- Jan 2020 – April 2020 – Cyber workforce identification
- Jan 2020 – Dec 2020 – OCIO Map Transformation
- Feb 2020 – Apr 2020 – Graph Database Development
- May 2020 – Jun 2020 – Interface Development
- July-Dec 2020 – Pilot testing
  - Skill prediction and forecast
  - AI assisted assignment recommendation
  - Workforce predictions

BENEFITS & RISKS

Benefits

- Greater understanding and alignment of workforce
- Improve workforce mobility
- Provide recommendations for training and career opportunities

Risks

- Risk to implementation: technology stack unavailable for operationalizing the platform. Access to data is cumbersome and difficult

STAKEHOLDERS & FUNDING

- Seeking $300K for graph database, visualization application, and analytic cloud resources
- OCHCO provide FTE
- HQ ITCD Server support for analysis and visualization
- OCIO support for using cloud analytics
- Risk – Limited funds for cloud services, HQ ITCD servers not implemented
- Project and DT AIML leads to negotiate / recommend operationalization plan / owner for long-term O&M
- Project Lead David Meza

NLP/ML Tool for Categorizing SES Job Applicants

The purpose of this project is to automate the human review workflow in the SES hiring process. Today applicants write lengthy essays describing how they meet the Executive Core Qualifications (ECQs) and Technical Qualifications (TQs) for the position. Staffers and Executives manually review the essays and categorize applicants as Not Qualified, Qualified, or Highly Qualified. The prototype, developed from archival data analyzed with NLP and ML, will be an algorithmic human-assisted decision-making tool that can review the essays and accurately assign applicants to a qualification category..

Roadmap Alignment: AIML Practical Application

User Community: NASA-wide, panel reviewers, HR specialist, data scientist, business analyst

KEY MILESTONES

- Mar 2020 – Data collection and exploration
- Apr 2020 – AI Model development
- May 2020 – Model validation and review
- June 2020 – Deployment

BENEFITS & RISKS

Benefits

- HR Modernization
- $200K savings per year
- Decrease time to hire
- Possible use in other use cases

Risks

- Risk to implementation: technology stack unavailable for operationalizing the platform.

STAKEHOLDERS & FUNDING

- Seeking $50K for analytic, storage and database cloud resources
- OCHCO
- HQ ITCD Server support for analysis and visualization
- OCIO support for using cloud analytics
- Risk – Limited funds for cloud services, HQ ITCD servers not implemented
- Project and DT AIML leads to negotiate / recommend operationalization plan / owner for long-term O&M
- Project Lead: James Illingworth, in teamwork with MSFC, JSC, Agency TDD, and IV&V

NCCS AIML Resources Quad Chart

Proposal: Purchase dedicated, shared fast storage (e.g. NVMe) for NCCS AIML environment

Roadmap Alignment: Team with HEC AIML community and AIML Practical Application

Contributing Community:

- NASA Center for Climate Simulation (NCCS)
- Solar Data Observatory (SDO)

KEY MILESTONES

- Configure 22 node AIML cluster to augment existing 3 nodes – currently on going within the NCCS
- Research storage solutions, including NVMe
- Draft requirements, request quotes, procure the solution
- Integrate the solution into the NCCS AIML environment
- Establish process for identifying data for this storage
- Transfer identified datasets

BENEFITS & RISKS

Benefits

- Reduce the effort by scientists to transfer and manage data for AIML research
- Create a high-speed repository of training data to accelerate research
- High performance storage will shorten training times to allow for more model iterations and development

Risks

- Risk to implementation: Since the storage will be shared, multiple users may put a high demand on the service and reduce its effectiveness
- Inaction: Scientists spend time wrangling data instead of doing research

STAKEHOLDERS & FUNDING – ask for

- NCCS – Scientific Computing recently purchased $1 million 22 GPU cluster, augmented with funding from Heliophysics
- Science Mission Directorates at Goddard, e.g. Heliophysics, Astrophysics, Hydrology, Climate Science
- Seeking ~$150K for dedicated, shared fast access solid state storage to lower the barrier for data access and accelerate training, e.g. NVMe

DaSHlink modernization toward AI/ML Consultancy

DRAFT Proposal: Modernize DaSHlink platform and underlying code base to use web-based tools to facilitate collaboration and re-use of machine learning workflows

Roadmap Alignment: Tools and Architecture;

Contributing Community Members:

- *Current DaSHlink members (we will choose most active subset among the ~500 users)*
- *ARC Data Sciences Group and KBR/Wyle system administrators)*
- *Innovation Portal Team*

KEY MILESTONES

- May 2020 – Determine best subset of NASA platforms with which to connect, identify other code bases to bring in for integration into DaSHlink (e.g., Jupyter notebooks)
- June – Connect DaSHlink with appropriate NASA platforms
- July – Integrate Jupyter notebooks and invocation of containers that run on user's local machine or cloud instance
- August – Build prototype recommender engine that connects users' uploaded content
- September – Integrate all code and demonstrate

BENEFITS & RISKS

Benefits

- Enable domains to easily set up modern web-based platforms for machine learning/data mining and customize it for their work and community.
- Facilitate sharing machine learning workflows within communities by enabling capture of workflows involving dataset preparation, execution of machine learning algorithms, and results preparation
- Connect with existing, complementary websites such as data.nasa.gov, analytics.nasa.gov, and Innovation Portal.

Risks

- Connection with other platforms requires an uncertain amount of other teams' time and effort

STAKEHOLDERS & FUNDING

- Seeking ~$250K seed funds to prepare demonstration platform and set stage for expansion. Several potential co-investors; possible DT seed funding;
- Made TDD Labs submission for $25k – this year's call was cancelled
- ARMD augmentation fund request for ~$100k - We will submit an ARMD augmentation request next FY, won't be possible for the current FY.
- Will explore other possibilities such as CIF , IRaD & may consider NASA@work challenge, etc.
- Idea from: Rodney Martin, Nikunj Oza
- Potential beneficiaries: ARMD (DaSHlink), SMD (NEX). DaSHlink and NEX have same underlying code base developed by Data Sciences Group.

Intelligent Contingency Management for UAS

Project Underway: Use AI to safely fly a mission under all vehicle-allowable weather conditions, in a high-density airspace complex urban environment, including off-nominal situations without direct human intervention

Roadmap Alignment: Practical AIML Mission Application

Contributing Community Members:

-LaRC Aeronautics (Irene Gregory) – need official org

-LaRC Data Science Team (Newton Campbell)

-Overall 8 SMEs involved ~4 person equivalent (FTE / WYE mix)

KEY MILESTONES

• FY19: Initial framework and plan; Algorithm Demonstrations
• FY20: UAM Vehicle Simulation Developed; Contingency Management (CM) Tool Integration with UAM Simulation; Evaluation under well-defined hazard
• FY21: In-Simulation contingency management demonstration for UAM-relevant vehicle under hazard
• FY22: TBD

BENEFITS & RISKS

Benefits

- Enables robust autonomous UAS piloting, to include ability to react to situations beyond trained heuristics
- Contributes to safe UAS / UAM concepts

Risks

- Technical risk is moderate – this is new technical ground; mitigate via adoption of techniques from parallel fields
- Ethical risk is moderate – must train and refine models to operate UAS/UAM systems IAW safe, ethical guidelines. Mitigate via extensive modeling, real-world V&V, and digital forensics if accidents do occur

STAKEHOLDERS & FUNDING

- Aeronautics investing ~ 4 person years for 3-4 years
- Explicit AIML investment in PhD AI contract expert full time duration of project (~$250K / year)
- AIML intends to help support the team with selected cross-cutting enablers, such as access to AIML platforms described elsewhere in this brief

Inverse Design of Materials Infrastructure

AI/ML enhanced infrastructure *to enable the rapid discovery, optimization, qualification, and deployment of* ***fit-for-purpose materials****. A metal detector for a needle in the haystack problem.*

Roadmap Alignment: Practical Application, Tools

User Community: NASA, aerospace, automotive, and tech industries, scientists, engineers

KEY MILESTONES

- May 2020: Complete ML software for automatic analysis of microstructure on several pilot materials (funded)
- Oct 2020: ML enhanced multiscale modeling pilot project completed (funded)
- Dec 2020: complete AMSII data ML pilot project (ingest test data, ML analysis, data storage in GrantaMI) (unfunded)
- Jan 2021: Begin integrating tools (unfunded)

BENEFITS & RISKS

Benefits

- Individual parts have concrete benefits (powerful predictive models, data storage and reusability, automatic analysis, design tools, visualization)
- Integrated ecosystem will accelerate materials discovery and deployment
- ROI: 4x reduction in development/deployment time of new materials (20 years to 5 years)

Risks

- Risk to implementation: standards for data storage and tool integration, security and cloud infrastructure

STAKEHOLDERS & FUNDING

Total Resources Required: $250K.

Current funding: Project (TTT) providing 1 FTE, GRC providing LERCIP, 2020 Fellowship ($30-60K), possible NSF providing $60K (NSF INTERNship)
DT funds needed: $60K

Project expenditures

0.2 FTE / 0.3 WYE to lead tool integration (internal and external tools) and develop ecosystem; Fellowships and faculty collaboration to develop tools; LERCIP Internships; Software Licenses

Project idea: Joshua Stuckner, Steve Arnold

Additive Manufacturing Data AI

Proposal: Develop a standardized data architecture compatible with machine learning and AI techniques to maintain additive manufacturing data populated with a data set for AI development.

Roadmap Alignment: AIML Practical Application; Tools & Architecture

Contributing Community:

- *MSFC Materials & Processes Laboratory*
- *Auburn University*
- *Langley Research Center*

KEY MILESTONES

- March/April 2020– Refine concept and scope; build team
- Q4 FY20 – Proposal Award & ATP
- ATP+ 0m - Kick-off with collaborators and stake holders
- ATP+ 5m – Database Schema finalized
- ATP+ 10m – Collation of Auburn Data Complete
- ATP+10m – Completion of Database Framework Complete
- ATP+12m- Data Ingestion Complete

BENEFITS & RISKS

Benefits

- Enables implementation of AI/ML on data of emerging technology of high interest.
- Provides a database framework to capture AM data from a variety of sources.
- Captures historical data from Auburn University Center of Excellence to increase the value of heritage data.

Risks

- Testing costs may prohibit the development of a large, complete data set from the outset.

STAKEHOLDERS & FUNDING

- Pursuing CAN agreement with Auburn University for core architecture development.
- MSFC Materials and Processes Lab
- Langley Research Center
- Other potential funding sources are being evaluated.
- Lead: Auburn University/Nima Shamsaei; NASA POCs Christopher Roberts/Teresa Miller

Soil Property Object Classifier (SPOC)

In Progress: Machine learning Moon / Mars terrain analysis – process hundreds of thousands of images vs. ~200 / day currently – inform enhanced navigation..

Roadmap Alignment: AIML Practical Application

Contributing Community:

- *JPL*
- *Hiro Ono*

BENEFITS & RISKS

Benefits

- Process 100-1000x more images / day for robust rover operation
- Faster operations by processing more on-board rover via labelled images

Risks

- NA.

STAKEHOLDERS & FUNDING

- JPL
- Funded
- Hiro Ono, Chris Mattman, Dan Crichton
- JPL Software of the Year

Complex Data Explorer (CODEX)
"See What's In The Data"

Objective:

- *Develop automated tool using machine learning techniques to quickly scan mission science and engineering data to reveal structures and problems*
- *Can flag anomalies and/or select best subset of data for later, more detailed analysis*

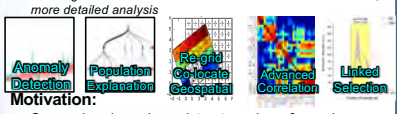

Motivation:

- Ops, engineering, science datasets are large & complex
- Time-critical need to see data contents & quality
- Data problems not immediately noticed since buried in data
- Custom data monitors limited; much analysis still manual

Major Milestones:

• Develop interactive machine learning workflows to support the following areas for exploration:

- Clustering ("How many groups?)"
- Regression ("Predict this from that")
- Correlation ("What relates to what?")
- Anomaly Detection ("What occurs when?")
- Template Scanning ("Find more like this")
- Data Quality Exploration ("What looks unexpected?")

• Provide a fast and intuitive user interface, which allows the user to rapidly explore their data.

• Provide export capabilities for down selected data, plots and code framework to reproduce the exploration outside CODEX.

Benefits:

- Risk Reduction: Rapidly recognize problems in data
  - Missing values, unexpected correlations
  - Changes in data relationships, bi-modality
- Identify multiple populations: require statistics analysis
- Isolate anomalies: for later study or omission
- Search by example: "events like this" in entire dataset
- Faster Research Results: Weeks of initial survey removed
- Rigorous Stats: Understand incoming data distributions
- Broad Applicability: All high-dimensional data consumers

Team:

Lukas Mandrake (Concept Formulation & Vision)
Jack Lightholder (Principal Investigator & Server Implementation)
Rob Tapella (User Interface Design & User Experience)
Josh Rodriguez (Front End Development)

Funding Profile: $600K (FY18-20)

Alignment with Strategic Plan:

- Reduce mission operations cost and risk
- Fast data problem and anomaly identification
- Leverage technological advances in computing
- Generalizes to Support Future Missions
- High performance queries (s/c data) enabling support for faster operational decisions
- System applicable to operations, science data analysis, engineering data analysis and multiple mission phases

Mission Relevance:

All planetary/Earth missions with large amounts of data downlink (Europa, M2020, SWOT, OCO3, etc.)

Infusion of Astronomical Source Vetting and Variable Star Classification Pipelines at MIT Lincoln Laboratory

Task: Distinguish Between…

Real Astronomical Sources | Bogus Artifact

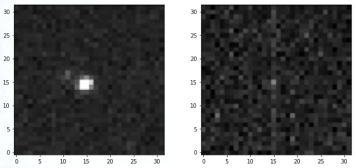

Filter artifacts from downstream processing pipelines

Task Objectives/Outcomes:

Objectives

- **Infuse** ML Pipeline Module into Space Surveillance Telescope Data Pipeline at MITLL that filters data artifacts of astronomical imaging

Outcomes

- Transfer JPL / Caltech best practices from astronomical survey experience to MITLL
- Improved catalog quality by filtering out bogus artifacts and non-detections using a machine-learned vetting system
- Classification of "light curves" emanating from a single astrophysical source in order to create a catalog of classified variable stars

Background

- Pilot a task that demonstrates collaborative potential between JPL/Caltech and MITLL
- Leverage complementary expertise
  - JPL/Caltech: machine learning, time domain astronomy
  - MITLL: astronomical data processing, SST survey data, high performance computing
- Pilot DSP initiated in FY18 ($50K) with follow-up in FY19 ($15K)

Identified task: JPL and MITLL collaborate on development of a data science pipeline for the Space Surveillance Telescope (SST)

Team Members & External Collaborators

- Umaa Rebbapragada (JPL)
- Ashish Mahabal (Caltech)
- Deborah Woods (MIT LL PI)
- Jacob Varburg (MIT LL)
- Brian Cook (MIT LL)

Importance to NASA/Mission Applications

- Fosters collaboration between JPL and MITLL
- Solidifies JPL's leadership in the building and deployment of machine learning systems for astronomical data science pipelines, adding SST to the following list: PTF, ZTF, V-FASTR at the Very Long Baseline Array (VLA), and the Dark Energy Survey (DES).

JPL Data Science Pilot Program

Information Technology for Cancer Research Program, National Cancer Institute

Meaningful Representative Graphics

Background

Develop an international knowledgebase linking data from across several programs to support research in cancer biomarkers and gene mutations.

Provide a roadmap forward for NCI to build a Cancer Biomarker Aggregator for the community.

Demonstrate emerging data science capabilities for linking diverse data models and distributed data sources.

Team Members & External Collaborators

George Washington University (GWU), University of Delaware, Swiss Bioinformatics Institute

Task Objectives/Expected Accomplishment:

Integration of diverse ontology models for cancer biomarkers, mutation and gene expression data

Development of a data science ecosystem for access and computation of biomarker-related data across highly diverse datasets including international partners.

Integration of Natural Language Processing (NLP)

Development of a discovery analytics framework governed by prevalence of data in the literature, statistical significance, evidence codes, and other pre-defined criteria.

Key Milestones	TRL: X to Y
Development of an ontology model that extends the EDRN	4 to 6
Development of APIs for remote access to data and computational services	4 to 6
Integration of NLP capabilities for integration publications	4 to 6

FY20 Investment: $150K ($450K - 3 years)

Importance to NASA and Medical Partner:

- Automated text extraction of data from scientific publications [NASA TA 11.4.3.1]
- Automated integration of multiple ontological models and associated databases [NASA TA 11.4.3.3 Semantic Bridge Framework].
- Search analytics technologies for integrating, visualizing and presenting scientific data results across distributed data sources [NASA 11.4.2.4 Intelligent Search and Mining]

(NCI/NIH) Apply data science capabilities in support of the NIH Data Science Strategy (June 2018) to support paradigm changes in biomedical research. Support integration of complex data sources to enhance data discovery.

Improving Hurricane Intensity Forecast with Machine Learning of NASA Satellite Data

***Meaningful Representative Graphics:* The ML-HIFS - A Machine Learning Hurricane Intensity Forecast System**

Task Objectives:

Leverage IBM Watson Studio to streamline data collection, feature extraction, ML modeling, performance evaluation and visualization

Establish the end-of-end data analytics framework (ML-HIFS) for NASA data applications into TC forecast improvements

Outcomes:

Created ~20-years training datasets for ML modeling of TC intensity change

Demonstrated the use of ML can drastically increase the predictive skills for TC Rapid Intensification (RI) over the state-of-the-art

Showcased the significant impact of NASA satellite data on improving TC forecasts (Su et al., 2020, in review)

Background

Tropical cyclone intensity forecast, especially for rapid intensity change, has been a notorious challenge for years.

Many NASA satellite measurements of TC environmental conditions and internal structures are not utilized operationally.

Intelligent use of these data with machine learning (ML) techniques has potential to make break-throughs in TC forecast.

Team Members & External Collaborators

Hui Su, Longtao Wu, Svetla Hristova-Veleva, Peyman Tavallali, Brian Knosp (JPL/Caltech), Raksha Pai (IBM), Alex Liu (RMDS Lab)

Importance to NASA/Mission Applications

Exemplifies rapid application of NASA data with advanced technology for societal benefits

Strengthens collaborations between NASA and NOAA operational forecasting centers

Establishes partnership between NASA/JPL and commercial companies such as IBM

Demonstrates the need and benefits of connecting science domain experts with data science specialists

Carbon Event Detection and Analytics System (CEDAS)

Meaningful Representative Graphic: Cross-scale, CH4 point-source detection

Yadav, Thorpe, Duren *et al*

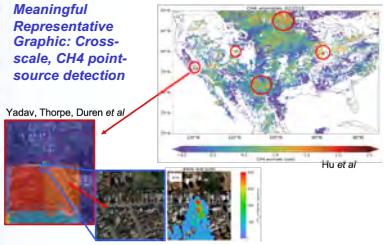

Hu *et al*

Task Objectives/Outcomes:

Pilot project to develop automated multi-scale CH4/CO2 event/anomaly detection and classification that:

- Leverages and extends key data science technologies: Smart Inferencing, On-demand Services, Data Integration, Machine Learning, Scalable Computation
- Carbon Specific Use-Case
  - Automates multi-scale CH4/CO2 event/anomaly detection and classification
  - Ingests L1B/L2 data & ancillary from multiple earth observing platforms
  - Runs algorithms for plume detection and attributes to source
  - Makes data available for analysis and download

Background

- Methane growth rate: causes are poorly understood and currently incompatible with greenhouse gas mitigation goals
- We have tiered Observing & Analysis Strategy
- Lack an integrated data system and automated CH4 point source detection & classification using remote sensing data, which results in long latency
- Lack of on-demand services including QC
- Methods have been demonstrated over limited pilot regions across California

Team Members & External Collaborators

Natasha Stavros, Brian Bue, Michael Garay, Andrew Thorpe, Winston Olson-Duvall, Rob Tapella, Kevin Gill, Vineet Yadav, Dan Cusworth, Riley Duren

Importance to NASA/Mission Applications

- *Tracking and characterizing the mechanisms of environmental change* such as methane is an objective in NASA's Strategic Plan
- **Contributes to the US Carbon Cycle Science Plan** Goal-1 and Goal-6
- **Responds to NASA's Carbon Cycle and Ecosystems** focus by reducing uncertainty in Goal-1 and Goal-3
- Reduce risk, cost, and time for delivering products from current and future Earth Science missions as **highlighted by the 2017 Earth Decadal Survey**: Earth System Explorer, Greenhouse Gas thrust
- **Responds to NASA Applied Sciences Program** as it is endorsed by public and private sector stakeholders indicating interest and strong potential for infusing the technologies

Information Technology for Cancer Research Program, National Cancer Institute

Meaningful Representative Graphics

Background

Develop an international knowledgebase linking data from across several programs to support research in cancer biomarkers and gene mutations.

Provide a roadmap forward for NCI to build a Cancer Biomarker Aggregator for the community.

Demonstrate emerging data science capabilities for linking diverse data models and distributed data sources.

Team Members & External Collaborators

George Washington University (GWU), University of Delaware, Swiss Bioinformatics Institute

Task Objectives/Expected Accomplishment:

Integration of diverse ontology models for cancer biomarkers, mutation and gene expression data

Development of a data science ecosystem for access and computation of biomarker-related data across highly diverse datasets including international partners.

Integration of Natural Language Processing (NLP)

Development of a discovery analytics framework governed by prevalence of data in the literature, statistical significance, evidence codes, and other pre-defined criteria.

Key Milestones	TRL: X to Y
Development of an ontology model that extends the EDRN	4 to 6
Development of APIs for remote access to data and computational services	4 to 6
Integration of NLP capabilities for integration publications	4 to 6

FY20 Investment: $150K ($450K - 3 years)

Importance to NASA and Medical Partner:

- Automated text extraction of data from scientific publications [NASA TA 11.4.3.1]
- Automated integration of multiple ontological models and associated databases [NASA TA 11.4.3.3 Semantic Bridge Framework].
- Search analytics technologies for integrating, visualizing and presenting scientific data results across distributed data sources [NASA 11.4.2.4 Intelligent Search and Mining]

(NCI/NIH) Apply data science capabilities in support of the NIH Data Science Strategy (June 2018) to support paradigm changes in biomedical research. Support integration of complex data sources to enhance data discovery.

The NASA AI & ML Academy: AI/ML Skill Development Resources

GSFC has a testbed program in FY2020 called the Goddard Machine Learning Academy. It consists of seven assets available to GSFC staff that provide training and "upskilling" assets.

The Machine Learning Academy consists of online and in-person training courses to learn:

- Machine learning and deep learning
- Data science programming essentials
- Cloud and GPU-based computing
- Collaborative computing and data science architectures
- Fundamentals of AI

OCHCO's Learning and Development Strategy Team is investigating expanding the program to all NASA centers.

KEY MILESTONES

- October 2019: Testbed program at Goddard commences, seven learning assets including O'Reilly Learning, LinkedIn Learning
- December 2019: Learning fair at Goddard has exhibits by Google, AWS, NVIDIA, O'Reilly and more. Attracts numerous participants from every GSFC directorate.
- March 2020: OCHCO Learning and Development Strategy Team begins investigation into other centers. Learning Academy expands to include participants from MSFC and other centers
- August 2020: OCHCO and Learning Academy leadership will evaluate the FY2020 testbed and determine best course of action. Possible transition to Digital Tranformation Academy.

BENEFITS & RISKS

Benefits

- Staff have convenient, cost-effective access to resources necessary for skill development and proficiency maintenance of essential data science skills
- Learning jointly with NASA colleagues establishes a common knowledge base
- Supports networking and community of practice that can identify new collaborations and strategic partnerships

Risks

- There is no single solution that answers to all learner needs. The diverse learning assets are from different vendors and providers, so there is some redundancy of services and difficulty in coordinating across platforms.

STAKEHOLDERS & FUNDING

- GSFC orgs leading development: OCHCO (100), GSFC Library (272), Software Engineering Division (580) Sciences and Exploration (600) and Information Technology (700), AI Center of Excellence.
- Funding is for internal FTE labor support and external resources (such as learning seats in a virtual learning environment)
- Additional potential stakeholders are all NASA orgs that have proficiency requirements in data science, machine learning, and advanced methods

NASA AI Center of Excellence (AICOE)

Proposal: *Develop, enhance, and maintain AI Center of Excellence website with up-to-date content for NASA's AI/ML community to* ***facilitate partnerships (both within the agency and with academia and industry), innovation, scientific discoveries, AI/ML training, workshops, conferences, and receiving R&D resources and funds.*** *An existing baseline model under development is GSFC's AICOE site: https://ai.gsfc.nasa.gov and Team's page.*

Roadmap Alignment: 1) Teamwork and Community, 2) Skills and Education.

External Expert Partner: *academia (e.g. EPSCoR program grantees) and industry (e.g. NVIDIA, Google, Amazon)*

Key Milestones

April 2020: Identify NASA AI/ML Training Resources
April 2020: Identify existing NASA AI/ML Computational Resources
May 2020: Identify NASA AI/ML projects and activities.
June 2020: participate in administration and management of AI/ML component of NASA EPSCoR program (academic partnerships)
September 2020: Design & preliminary development of the agency AICOE site and Teams page with above content.
October 2020: initial release of the agency's AI/ML site in Drupal platform.
Ongoing: Test, update, maintain, and enhance AICoE site
Ongoing: Organize and Support AI/ML training on center and for the agency in partnership with GSFC ML Academy
November 2020: Help with organization of the next NASA AI workshop

Benefits

- Provide a central catalogue of all agency's AI/ML projects, computational resources, training opportunities, conferences, and funding sources.
- Prevent re-inventing the wheel via providing inventory of existing AI/ML software solutions, labeled data, models, ..
- Facilitate partnerships both within and outside the agency

Risks

- Large size of the organization
- Too many internal and external stakeholders
- Lack of similar efforts across many centers to facilitate agency level implementation.

Seeking a total of $160K as follows:

Civil Servant Project Manager (0.5 FTE ~ $100K)

- Continuously gather data/input/content from stakeholders across the agency, academic and industrial partners and curate content for the AICOE site.
- Lead development of the AICOE website.
- Form partnerships for R&D AI/ML projects (e.g. manage components of EPSCoR program), facilitate AI/ML R&D proposals (team building, proposal writing, …).
- Attend agency AI/ML and DT meeting, workshops,…

Web developers and graphic artist ($50K)

Travel ($10K) to AI/ML workshops/conferences, or other centers.

•**PI:** *Nargess Memarsadeghi*, NASA/GSFC, Code 586

Using Machine Learning to Support Earth Science Mission Designs:

Deep exploration and creation of active remote sensing science library via integration of LIDAR techniques

PROBLEM TO SOLVE:

- Emerging Science missions are complex and involve many tradeoffs among quality of science, cost, etc.; traditional analysis approaches work, but become expensive and / or time consuming.
  - Can ML efficiently help design new instruments for different levels of missions?
  - Can ML help to bridge science, engineering and management?
  - How to keep a leading position in LIDAR techniques in long-term? (best way to adopt new techniques, improve old ones and invest them in appropriate and efficient way)
  - Can ML prepare for even bigger mission integration? (space ground mission, passive active remote sensing mission, etc)

BENEFITS & RISKS

Benefits

- Like AlphaGo can defeat top human players, a well-designed ML searching tool can take advantage of computing power to find a best solution in system design with a faster speed and could analyze more design options. It will enhance human's ability to design LIDAR instruments.
- ML design tools can be integrated (grow as our digital brain) long-term impact is high. ML can help build and accumulate an accessible LIDAR library on science parts that can further support for mission design and operation. On the road ahead, this trained library can help for integration with engineers, management, earth-space science, etc. I can also help with post-mission analysis such as satellite imagery retrievals)

Risks

- High technical risk, but potential impact is also high. Novel approach, computing power challenging

IDEA

We propose to start to integrate LIDAR techniques and build LIDAR science library by integrating machine learning and LIDAR simulations. Full integration is a long-term goal and we start from the development of a simple ML framework for a single LIDAR technique integration. We will grow the system over time to integrate other LIDAR and other ML techniques. The results can be used for mission design in the future.

STEPS

- Running Lidar simulations. The simulations build the LIDAR science value library, a multiple dimensional Look-up-Table between LIDAR designs and science values.
- Due to the large volume of the library built through simulation, it is not easy for quick access and impossible for integration. Machine learning surrogate models train the library in a second step to make simulation results easy and quick to access. ML method build a one-to-one link between LIDAR designs and science value without any complicated middle processing and this trained library is now possible for integration.
- As the library for integration is built, another ML (RL model) method can be used to search into the library for design, optimization and decision making by combining engineering model and cost model. ML takes an agent position as scientist, dig into the LIDAR science library to look for the best option for combination of LIDAR design and explore their science values.

Funding

- Cost estimate for initial experiment: $K; Mission funding: $K; Other funding: $K
- Team: LaRC Science; POC Shan Zeng

High End Computing Capability (HECC) Project

HECC Hardware Assets

5 Compute Clusters

- Pleiades 158 Racks / 11,292 nodes / 7.42 PF
- Electra 24 Racks / 3,422 nodes / 8.32 PF
- Aitken 4 E-Cells/ 1,152 nodes / 3.69 PF
- Merope 56 ½ Racks / 1,792 nodes / 252 TF
- Endeavour 3 Racks / 2 nodes / 32 TF

1 Visualization Cluster 245 million pixel display /128 node /703 TF

10 Lustre File Systems 46.0 PB

6 NFS File Systems 1.5 PB

Archive System 1,000 PB

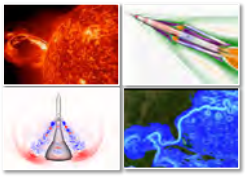

D-Wave 2000Q quantum system

- Whistler processor with 2031 qubits

NAS Facility Extension (under development)

- An acre site with 30 MW power to house HPC systems in modules

HECC Services

HECC provides a suite of complimentary services to the user community to enhance the scientific and engineering results obtained from the hardware assets.

- Systems – Customized solutions including compute and storage solutions to meet specific project or mission requirements.
- Application Performance and Productivity – Software solutions provided to research/engineering teams to better exploit installed systems.
- Visualization and Data Analysis – Custom visualization during traditional post-processing or concurrent during simulation to understand complex interactions of data.
- Network – End-to-end network performance enhancements for user communities throughout the world.
- Data Analytics – Exploitation of data sets through neural nets and emerging new techniques.
- Machine Learning – Custom environments to enable learning through advanced data techniques.
- Custom Data Gateways – Custom data portals to support diverse programs and projects

GPU/AI/ML Addition to HECC

- HECC Data Science team has setup Machine Learning environments and tools to help NASA users to utilize their data to perform machine learning tasks
- 3 HPE Racks with 16 Apollo GPGPU
  - 2 Racks with 6 Apollo Systems each with the following configuration
    - 2 Intel Xeon 6154 processors (20-core, 3.0 GHz, 384 GB memory)
    - 2 EDR IB 2-port cards
    - 4 NVIDIA V100 32 GB GPGPU
    - 2 1.6TB SAS SSD
  - 1 Rack with 4 Apollo Systems with the following configuration
    - 2 Apollo Systems with the following configuration
      - 2 Intel Xeon 6154 processors (20-core, 3.0 GHz,
      - 384 GB memory)
      - 2 EDR IB 2-port cards
      - 4 NVIDIA V100 32 GB GPGPU
      - 2 1.6TB SAS SSD
    - 2 Apollo Systems with the following configuration
      - 2 Intel Xeon 6154 processors (20-core, 3.0 GHz,
      - 384 GB memory)
      - 3 EDR IB 2-port cards
      - 8 NVIDIA V100 32 GB GPGPU
      - 2 1.6TB SAS SSD

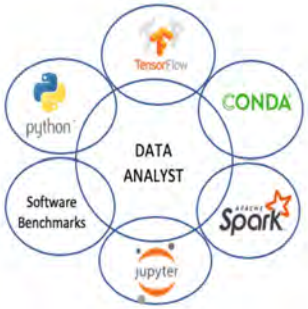

Future Outlook

- Next steps beyond initial implementation plan
  - Large community "CONOPS" session to move from initial planning within current resources to recommend the investment levels and methods to truly maximize AIML's benefit to NASA
- How to scales/extend beyond these activities
  - Initial AIML activities for practical application, workforce AIML training, AIML platform access, AI-Ready Data, and more all form the early foundations for even greater AIML benefit to NASA missions in the long term
- Long term vision of this implementation
  - AIML permeates everything NASA does, in secure, ethical manner
  - AIML as the core enabler to making NASA's hardest missions feasible

AI/ML Integration/Alignment Input

- Key lessons, challenges, best practices
  - Concern that our current resource-constrained approach to AIML is insufficient, especially as other federal orgs are making major AIML investments (e.g., $50M DoE AI center of excellence, major DoD investments, etc.)
- Key touch points with other ST roadmaps (or OCIO roadmaps)
  - We have met with Data and need to continue the dialog
  - Great to see AIML on Cloud and Info Mgt roadmaps; 'have been coordinating with these OCIO teams; need to do more
- Integration/alignment with your Implementation Plan team
  - AIML began implementation planning before IIPs were named. We are coordinating with ARC, Aeronautics, and OCIO TDD leaders as possible co-sponsors of AIML implementation formalization.
  - In the meantime, we will continue planning with the existing team
- Other coordination/integration challenges
  - We recommend early consideration of a full-time AIML lead catalyst FTE for the Agency; potentially within OCIO TDD as a sister position to Chief Data Scientist; have been discussing with TDD lead (Ron T)

AI/ML DRAFT Roadmap / Implementation Sketch **(Medium DRAFT - still open to lots of input and refinement. DT editors will make it look prettier)**

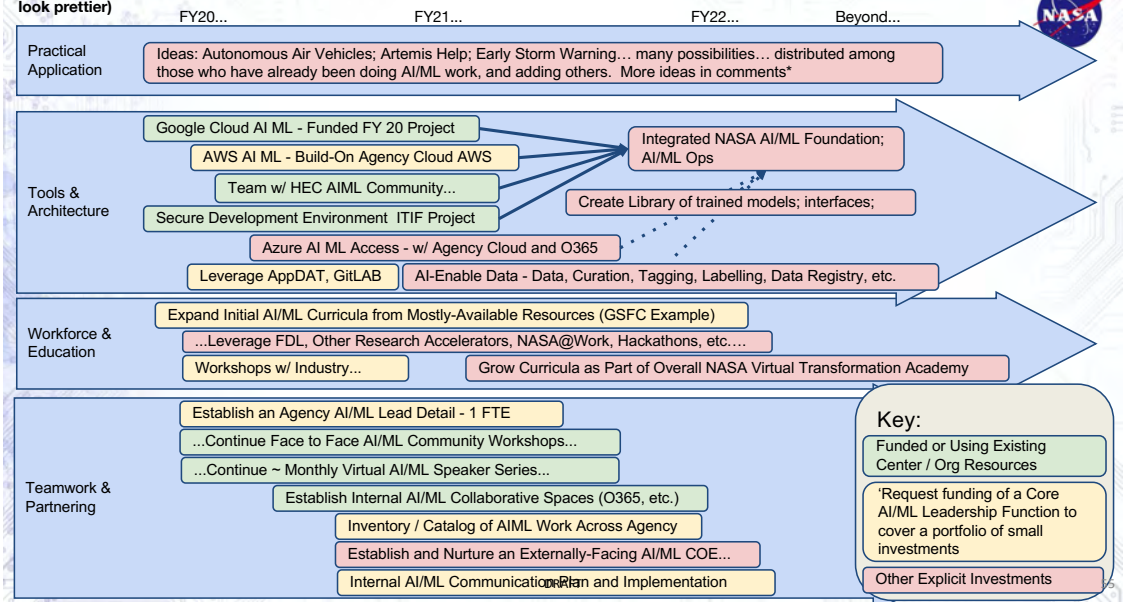

AI/ML Roadmap Items for Initial Implementation (1 of 2)

Name	Description	Rationale	Lead / Partners
AI/ML Mission Applications	High-priority AI/ML help to Missions. Coordinated as-needed with Missions. Grow supporting AI/ML teams across NASA	The point of AI/ML is to act as an accelerant.	Agency or Centers AI/ML Teams, as needed by Missions. $$ by Project
AI/ML Msn Support Apps	High-priority AI/ML help to Mission Support. Coordinated as-needed with Support Orgs. Grow supporting AI/ML teams across NASA	The point of AI/ML is to act as an accelerant.	Agency or Centers AI/ML Teams, as needed by Support Orgs. $$ by Proj
Core Agency AI/ML Support Team and Investment Catalyst	$750K / year and 1 FTE to act as AI/ML Strategic Lead and catalyst for NASA. Combine catalyst funding with existing Mission and Center AI/ML spending to steward a portfolio of core AI/ML platforms, policies, training, education, communication, community of interest, etc.	- Much AI/ML work is already happening across NASA; need a focal point to coordinate, advocate, synchronize. - Able to maximize existing spending via small catalyst (~$750K)	New AI/ML Strategic Lead within OCIO TDD (Proposed) by 1 October 2020 *$750K covers most "yellow" arrow items
Google Cloud AI/ML Early Win	Achieve ATO for Google Cloud Platform AI/ML tools, building on Agency Cloud, JSC, MSFC core GCP work already underway. Document and share availability. Train, encourage use.	- Cloud vendors are delivering greatest bang for buck in AI/ML platforms… better than building hundreds of unique, non-interoperable AI/ML stacks	LaRC, MSFC, AAO, JSC, HQ OCIO. $250K already funded from MSD
Secure Dev Environment ITIF – Leverage!	OCIO-funded IT Investment Fund project to create a secure, modern SW development environment; leverage for AI/ML uses	- Excellent forward-leaning project; use it to jump-start modern AI/ML development	Agency Data Scientist - already funded for ~$1.5M total in FY18 and 19

AI/ML Roadmap Items for Initial Implementation (2 of 2)

Name	Description	Rationale	Lead / Partners
AI-Enable NASA's Data	Team with the Data Team (and others) to create, share, and encourage standard methods for data curation, tagging, labeling, registry, security, etc. Make NASA's data easy for AI/ML services to access and make sense of.	Data prep is the most time-consuming current step in the AI/ML pipeline	DT Data Team, TDD, AI/ML; ~$2M / year for 3 years
MS Azure AI/ML Services	Team with Agency Cloud office to get Azure cloud AI/ML services authorized for use at NASA. Document and share available services. Train, encourage use.	MS Azure is one of the big three cloud providers, potential O365 integration	AI/ML Team, Agency Cloud Team, O365 Team; $250K
Integrated AI/ML Foundation	Unify and standardize AI/ML cloud platforms among AWS, Microsoft (Azure), Google, and any others we onboard. Emphasize portability among platforms and longevity of solutions for longer-term NASA missions.	Let turbulent market stabilize and then consolidate among the winners	AI/ML Team, TDD, Data; $500K initial + $300K yearly ops
Library of Trained AI/ML Models	Create and curate a library of trained ML models to be shared across NASA, rather than each data scientist having to train their own.	Share trained models rather than costly, redundant training	AI/ML Team; $200K + $100K yearly ops
Amplify Research Accelerators	Increase the use of the FDL and other research accelerators. Includes NASA workforce AND partners	Leverage external expertise and partnerships	AI/ML Team, Missions; scalable $$TBD
Virtual Transformation Academy	Make AI/ML curriculum part of an overall NASA Virtual Transformation Academy; see if this idea makes sense with the culture / workforce team	Contribute to overall DT skills	Culture / WF team, OCHCO, AI/ML Team; $500K / year for AI/ML
Public AI/ML COI	Establish and nurture an externally-facing NASA AI/ML COE; see example of major announcements by other federal orgs.	Provide external partners a focal point for engagement	OSAC Offices, AI/ML Team; $TBD by FY22

NASA CFD Vision 2030 Study

A Path to Revolutionary Computational Aerosciences

Mujeeb R. Malik
Aeronautical Sciences Project

Presentation at NASA HQ
Aeronautics Research Mission Directorate
September 30, 2014

Outline

- **Executive Summary**
- **CFD Vision 2030 Study**
  - **Vision**
  - **Findings**
  - **Recommendations**
  - **Grand Challenge Problems**
  - **Notional Technology Development Plan**
- **The Way Forward**

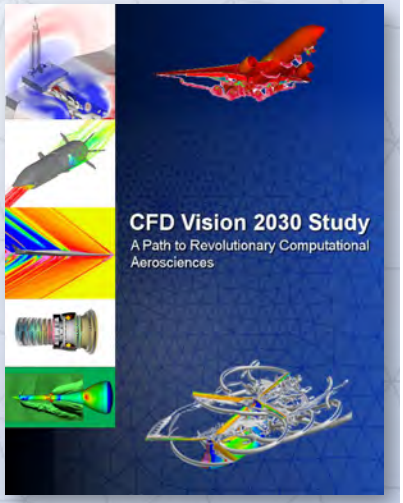

NASA CR 2014-218178

Report available at: http://www.aeronautics.nasa.gov/fap/tech_highlight_aero6.html

Executive Summary

CFD within NASA Mission Directorates

- **Aeronautics Research Mission Directorate (ARMD):**
  - CFD is a cross-cutting technology that supports four of the ARMD strategic thrusts, and the associated "outcomes"
  - Plays an important role in subsonic and supersonic civil aircraft and rotorcraft technology development
  - **Basic computational tool development**
    - **Overflow, CFL3D, ARC3D, Wind-US, Vulcan …**
    - **FUN3D, USM3D, CART3D…**
- **Human Exploration and Operations (HEOMD):**
  - Development of Space Launch System, Orion
- **Science (SMD):**
  - Planetary entry systems (MSL/Curiosity)
  - Climate, weather, environment

Executive Summary CONTINUED

Impact of CFD-1

"Since the first generation of jet airliners, there has been about a 40% improvement in aerodynamic efficiency and a 40% improvement in engine efficiency … about half of that has come from CFD."

(Robb Gregg, BCA Chief Aerodynamicist)

During Panel Discussion Session "Transformative Aerospace System Analysis, Design and Certification: A Vision for CFD in 2030" at AIAA Aviation 2014
http://new.livestream.com/AIAAvideo/Aviation2014/videos/54061902

Executive Summary CONTINUED

Impact of CFD-2

"Using CFD in the design … has resulted in 75% reduction in the number of tests to develop a combustor, and 60% reduction in emissions."

(Steve Morford, P&W Chief Engineer for Systems and Aerodynamics)

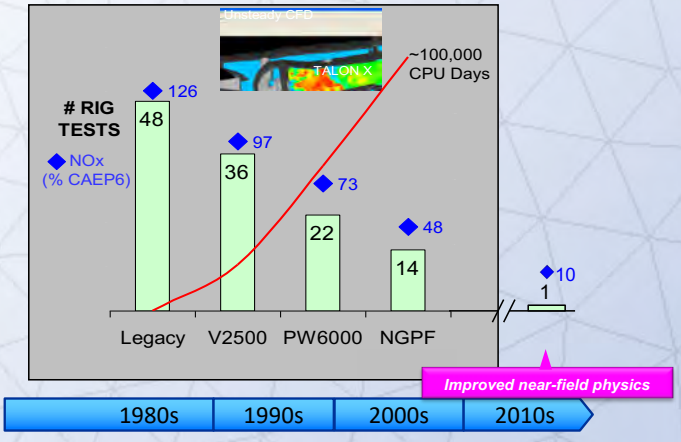

Executive Summary CONTINUED

CFD can not reliably predict separated flows

- **CFD has drastically reduced testing for cruise design**
  - Attached flow, well predicted by current turbulence models
- **Testing still required for off-design (e.g., high-lift) conditions, even for conventional configurations**
  - Flow separation is the key issue
- **Increased testing will be required for innovative configurations**
  - Prediction of flow separation and transition are key physics issues

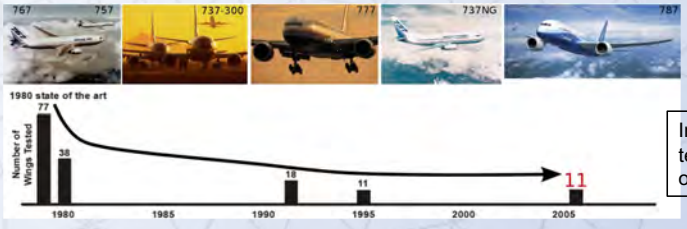

Inability to further reduce number of tests due to deficiency in modeling of turbulent flow physics

Executive Summary CONTINUED

Breakthroughs and Innovation by High Performance Computing

- **New computational technologies can fundamentally change the design space**
  - Provide better understanding of the design space, leading to innovation
- **Improved simulation capabilities enable**
  - Superior/more capable designs
  - Reduced development cycle time/cost/risk
  - Industrial competitiveness
- **Current CFD capability is not what it should be**
  - Potential to significantly improve prediction accuracy, speed and reliability/robustness
  - Need investment to advance the state-of-the-art
  - These investments will eventually payback by reducing physical testing infrastructure

Decadal Survey of Civil Aeronautics (NRC): ***"…an important benefit of advances in physics-based analysis tools is the new technology and systems frontiers they open"***

Overview

- **CFD Vision 2030, an "N+3" Class Study:**
  - "…provide a **knowledge-based forecast** of the future computational capabilities required for **turbulent, transitional, and reacting flow simulations**…"
  - "…and to lay the foundation for the **development of a future framework/environment** where physics-based, accurate predictions of complex turbulent flows, including flow separation, can be accomplished routinely and efficiently in cooperation with **other physics-based simulations to enable multi-disciplinary analysis and design**."

Overview CONTINUED

Vision 2030 NRA Team Members

NASA Technical Monitors – **Mujeeb Malik**

Jeffrey Slotnick
Principal Investigator
Boeing Research & Technology
jeffrey.p.slotnick@boeing.com

Abdi Khodadoust
Project Manager
Boeing Research & Technology
abdollah.khodadoust@boeing.com

Juan Alonso
Stanford University

David Darmofal
Massachusetts Inst. Of Technology

William Gropp
National Center for Supercomputing Applications

Elizabeth Lurie
Pratt & Whitney – United Technologies

Dimitri Mavriplis
University of Wyoming

Extended Vision 2030 Team:

- **Joerg Gablonsky, Mori Mani, Robert Narducci, Philippe Spalart, and Venkat Venkatakrishnan** – *The Boeing Company*
- **Robert Bush** – *Pratt & Whitney*

Overview CONTINUED

- **Approach of the study effort:**
  - Define and develop **CFD requirements**
  - Identify the most critical **gaps and impediments**
  - Develop and execute a **community survey** and **technical workshop** to solicit input and gain community consensus for future CFD development needs and potential
    - Over 150 survey participants (Industry 40%, Academia 24%, Government 36%) – SMEs, Managers
    - 70 workshop participants
  - Synthesize input and develop a detailed **research roadmap** to
    - capture anticipated technology **trends** and future technological **challenges**,
    - and provide **focus** to the broader CFD community for future research activities

Vision of CFD in 2030

- **Emphasis on physics-based, predictive modeling**
  Transition, turbulence, separation, chemically-reacting flows, radiation, heat transfer, and constitutive models, among others.
- **Management of errors and uncertainties**
  From physical modeling, mesh, natural variability, lack of knowledge in the parameters of a particular fluid flow problem, etc.
- **A much higher degree of automation in all steps of the analysis process** Geometry creation, meshing, large databases of simulation results, extraction and understanding of the vast amounts of information generated with minimal user intervention.
- **Ability to effectively utilize massively parallel HPC architectures that will be available in the 2030 time frame**
  Capacity- and capability-computing tasks in both industrial and research environments.
- **Seamless integration with multi-disciplinary analyses**
  High fidelity CFD tools, interfaces, coupling approaches, etc.

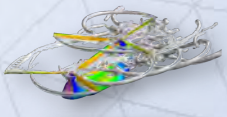

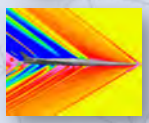

Predictive and automated CFD tools for timely analysis/design of novel configurations.

Findings

1. **Investment in technology development for simulation-based analysis and design has declined significantly in the last decade and must be reinvigorated if substantial advances in simulation capability are to be achieved.**
   - NASA's Revolutionary Computational Aerosciences (RCA) sub-project is a step in the right direction **and should be maintained and expanded**
2. **High Performance Computing (HPC) hardware is progressing rapidly**
   - Many CFD codes and processes do not scale well on petaflops systems
   - CFD codes achieve only 3-5% of peak theoretical machine performance
   - NASA poorly prepared for exaflops (10^{18} flops) revolution
3. **The accuracy of CFD in the aerospace design process is severely limited by the inability to reliably predict turbulent flows with significant regions of separation**

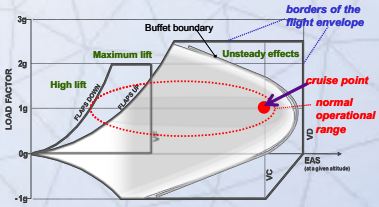

CFD accurate only near cruise point

Findings CONTINUED

4. **Mesh generation and adaptivity continue to be significant bottlenecks in the CFD workflow, and very little government investment has been targeted in these areas.**
   - Goal: **Make grid generation invisible to the CFD analysis process** → Robust and optimal **mesh adaptation methods** need to become the norm
5. **Algorithmic improvements will be required to enable future advances in simulation capability.**
   - Robust solution convergence for complex geometries/flows is lacking
   - Improved scalability on current and emerging HPC hardware needed
   - Develop "optimal" solvers, improve discretizations (e.g., high-order)
6. **Managing the vast amounts of large-scale simulations data will become increasingly complex due to changing HPC hardware.**
7. **In order to enable multidisciplinary simulations, for both analysis and design optimization purposes, several advances are required: CFD solver robustness/automation, standards for coupling, computing and propagating sensitivities and uncertainties.**

Recommendations

1. **NASA should develop, fund and sustain a technology development program for simulation-based analysis and design.**
   - Success will require **collaboration** with experts in computer science, mathematics, and other aerospace disciplines
2. **NASA should develop and maintain an integrated simulation and software development infrastructure to enable rapid CFD technology maturation.**
   - **Maintain a world-class in-house simulation capability**
     - Critical for understanding principal technical issues, driving development of new techniques, and demonstrating capabilities
3. **HPC systems should be made available and utilized for large-scale CFD development and testing.**
   - **Acquire HPC system access for both throughput (capacity) to support programs and development (capability)**
     - *improved software development, implementation, and testing is needed*
   - **Leverage national HPC resources**

Comparison of HPC at NASA and DoE

- **Typical NASA simulation: 1000's of cores**
- **State-of-the art: 1.6 million cores**
- **NASA capability trailing by 3 orders of magnitude**

DOE and NASA HPC Systems Among the Top 100 Worldwide

Machine / Agency	Speed
2. Titan	11.7
3. Sequoia	11.5
5. Mira	5.7
9. Vulcan	2.9
15. Cascade	1.7
18. Edison	1.1
21. Pleiades	1.0
32. Cielo	0.7
34. Hopper	0.7
47. Zin	0.5
54. Cetus	0.5
64. Gaea C2	0.4
87. Red Sky	0.3
94. HPCEE	0.3

Department of Energy
13 systems, 4.2 million cores capable of 57 petaflops

NASA
1 system, 96 thousand cores capable of 1.5 petaflops

- *Machines with billions of cores are just around the corner*
- *Speed is relative to that of Pleiades*

Pleiades Divided Among Hundreds of Users, allowing "capacity" computing only

Hierarchical Approach at DOE

Recommendations CONTINUED

4. **NASA should lead efforts to develop and execute integrated experimental testing and computational validation campaigns.**
   - High quality experimental test data for both **fundamental, building-block** *and* **complex, realistic configurations**, coupled with careful computational assessment and validation, is needed to advance CFD towards the Vision 2030 goals
     - Experiments to provide data for development of advanced turbulence models/prediction capability
   - **NASA is uniquely positioned** to provide key efforts in this area due to the availability of world-class experimental test facilities and experience, as well as key expertise in benchmarking CFD capabilities

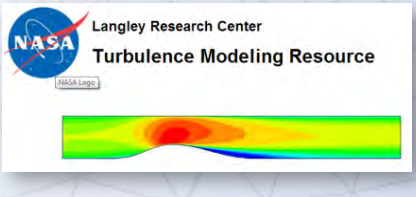

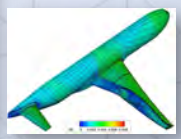

Recommendations CONTINUED

5. **NASA should develop, foster, and leverage improved collaborations with key research partners across disciplines within the broader scientific and engineering communities.**
   - Emphasize funding in **computer science and applied mathematics**
   - Embrace and establish **sponsored research institutes** → provides centralized development of cross-cutting disciplines.

CTR

6. **NASA should attract world-class engineers and scientists.**
   - **Success in achieving the Vision 2030 CFD capabilities is highly dependent on obtaining, training, and nurturing a highly educated and effective workforce**
     - Expand **fellowship programs** in key computational areas
     - Encourage and fund long-term **visiting research programs**

Grand Challenge Problems

- Represent critical **step changes** in engineering design capability
- May **not be routinely achievable** by 2030
- Represent key elements of **major NASA missions**

1. Large Eddy Simulation (LES) of a powered aircraft configuration across the full flight envelope
2. Off-design turbofan engine transient simulation
3. Multi-Disciplinary Analysis and Optimization (MDAO) of a highly-flexible advanced aircraft configuration
4. Probabilistic analysis of a powered space access configuration

Notional Technology Development Roadmap

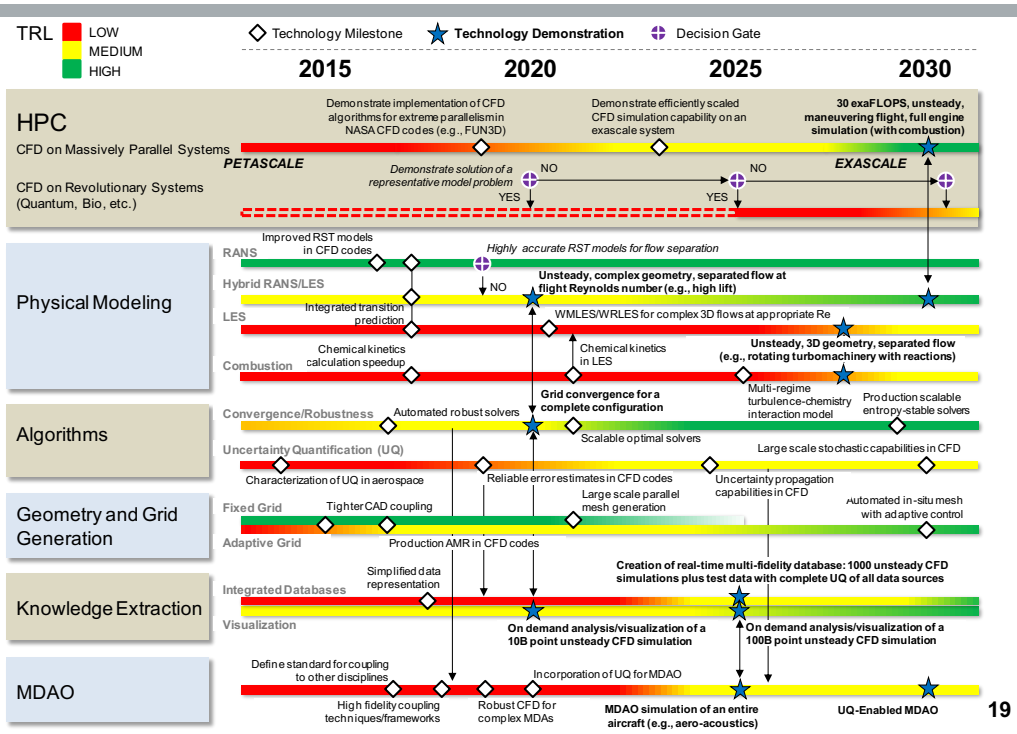

Summary of the CFD Vision 2030 Study

- Although CFD has become an integral part of modern aerospace vehicle design and analysis, many technical difficulties in accuracy and modeling ***still*** persist.
- In the future, to effectively utilize CFD for ***all phases*** of the design process and in ***all corners*** of the design space, significant resource investment, coupled with a collaborative, multi-disciplinary research environment, is critically needed.
- A comprehensive study to define a vision of CFD in the year 2030 has been completed → Study recommendations will help guide research investment in CFD development over the next two decades to enable ***transformational computational analysis***.
- Once available, the transformational CFD capability will significantly reduce non-recurring product design and development costs, and enable ***certification by analysis*** as well as ***innovation***.

Proposed NASA Way Forward

- **Study Recommendations that would Need Investments**
  - Access to HPC for "capability" computing
  - Validation Experiments
  - Collaborations with academia/research institutes
    - Note: Current Annual RCA NRA budget is $2.3M. FY15 onwards, NRA budget for the entire TTT project will be $2.5M/year
  - Grid-generation/solution adaptive grids (human out of the loop)
  - High-fidelity multidisciplinary analysis and design optimization
  - Software development infrastructure/frameworks
- **Program in Simulation-Based Engineering**
  - Comprehensive implementation of all recommendations would require a new program ($50-$60 Million). Note DoE ASC/ASCR budget is about $1B
    - Acquire new hardware for capability computing
  - Possible only if across mission directorates because of investment level
  - Will address disciplines relevant to the three NASA missions
- **An "ARMD only" Solution**
  - Augment RCA budget
  - Accomplish study recommendations by leveraging outside resources

Proposed NASA Way Forward CONTINUED

- **Notional Proposal**
  - Access to HPC for "capability" computing: $1M
  - Validation Experiments: $4M
  - Collaborations with academia/research institutes: $5M
    - Maintain current level of RCA NRA investment
    - Fund research institutes for collaborative/multidisciplinary research
    - Engage subject matter experts in computer science, applied mathematics
  - Grid-generation/solution adaptive grids: $1M
- **Collaborations will aim to solve GC Problems**
  - Suggested by the CFD Vision 2030 Study
  - Or, those in consultation with ARMD focused projects

The Way Forward

- **Study Recommendations that would Need Investments**
  - Access to HPC for "capability" computing
  - Validation Experiments
  - Collaborations with academia/research institutes
    - Maintain current level of RCA NRA investment
    - Fund research institutes for collaborative/multidisciplinary research
    - Engage subject matter experts in computer science, applied mathematics
    - Expand graduate student fellowship program
  - Solution adaptive grids (human out of the loop)
  - High-fidelity multidisciplinary analysis and design optimization
  - Software development infrastructure/frameworks
- **Collaborations will aim to solve Grand Challenge Problems**
  - Suggested by the CFD Vision 2030 Study
  - Or, those in consultation with ARMD focused projects

The Way Forward CONTINUED

- **Access to "Capability" Computing**
  - Leveraging DoE resources?
  - High level "hand-shake" to provide access for NASA scientists?
  - In 2013, INCITE Program awarded 2.8 billion hours on ALCF and 1.8 billion hours on OLCF (Academia 46%, Govt. 37%, International 12%)
- **An Unsuccessful NASA INCITE Proposal to DoE**
  - Joint AS/FW proposal for aeroacoustic simulations of open rotor, isolated and installed on HWB
  - Declined because FUN3D does not have extensive experience with Graphical Processing Units (GPUs) and Titan is GPU heavy (20K GPUs)

- **Some Successful non-NASA INCITE Proposals**
  - DES of high-lift wing with active flow control; U. Colorado
  - DNS of Ignition in Engine-like Geometries; ETH(Zurich)
  - LES of Combustor Liner Flows; GE

- **Collaborate with Academia to gain access to DoE HPC**
  - Solve NASA challenge problems
  - Gain experience with "Capability" computing
  - Leveraging not free, but will cost much less than acquiring new hardware

The Way Forward CONTINUED

What would be Gained by Research in HPC?

- **An Order of Magnitude Speed up, even on Existing Hardware**
  - Current CFD codes perform at 3-5% of theoretical peak machine performance
  - It should be possible to achieve 30-50%
- **Gordon Bell Prize 2013**
  - Won by a team from ETH (Zurich), TU-Munich, IBM Research and LLNL
  - 'Simulation of Cloud Cavitation Collapse' on DOE's Sequoia (IBM Blue-Gene) using 13 trillion cells including 15000 bubbles
  - Team increased performance by 20x for their application
- **Teaming with Computer Scientists/Algorithm Specialists is Essential**

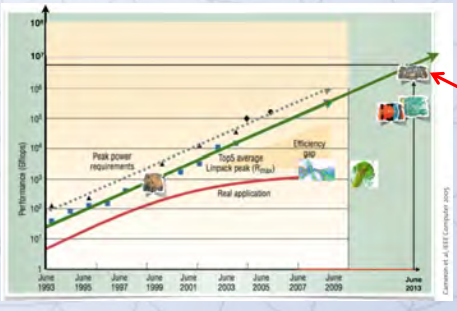

Using 1.6 M cores, ETH Team Achieved: 70% of Linpack; 55% of theoretical peak

The Way Forward CONTINUED

- **Validation Experiments**
  - Validation data critical for advancing CFD prediction capability to complex flows (separated flows, shock/BL interaction, etc.)
  - Foundational experiments with detailed measurements of turbulence quantities, skin-friction, etc. for overcoming turbulence model deficiencies
    - Unique NASA role; too foundational for industry

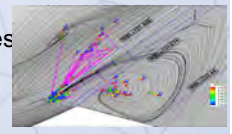

- **Juncture Flow Experiment (JFE)**
  - Informed by Drag Prediction Workshop series: prediction of side-of-body separation an issue for CFD codes
  - JFE to allow 4 conditions (no separation, incipient, small and massive separation)
  - Outer mold line currently under exploratory design
  - Innovative instrumentation will be used
  - Experiment currently not funded

JFE Model

- **Several Other Experimental Ideas in the 'Hopper'**
  - Building block BL Transition and Turbulent Flow Experiments
  - Aeroelasticity
  - Aircraft Boat-Tail Flow
  - Shock/Boundary-Layer Interaction

The Way Forward CONTINUED

Why Investment in CFD is Needed?

"***Rapid Physics-Based Virtual Testing Capabilities*** **are needed to reduce the cost to design and develop new aviation systems. As sophisticated as modern design, development, test and evaluation practices are, there is still a high reliance on data-driven, empirical models, which eventually require significant amounts of physical testing to validate the performance and safety of design decisions. In order to speed development cycle-time, there must be more 'first principles' physics-based models that can directly reveal the impact of design decisions. Such 'virtual' testing capability would decrease development cycle times and costs, enabling greater innovation and responsiveness to dynamic transportation needs."**

(An excerpt from 2012 version of Bob Pearce's White Paper)

CFD Vision 2030 Study strongly recommends to develop this capability.

The Way Forward CONTINUED

Conclusions

- **Physics-based computational tools have tremendous potential**
  - Multidisciplinary tools with accuracy, speed and reliability to design future aerospace vehicles with little or no ground testing
  - Advance the cause of ARMD "Outcomes"
    - e.g., "Improved Vehicle Efficiency and Environmental Performance in 2035"
  - Cross-cutting tools that impact 3 NASA mission directorates
- **Current NASA investment not sufficient**
  - Potential adverse impact on NASA missions (in terms of innovation, cost, schedule)
  - Potential gains for US industry will be lost
- **NASA needs to invest in Simulation-Based Engineering**
    - Provide access to capability computing
    - Perform foundational experiments needed to expand the range of applicability of computational tools
    - Collaborative research to solve NASA Grand Challenge Problems

Simulation-Based Engineering is increasingly becoming important and will be the way of the future.

Back up Charts

CFD 2030 Study Scope and Requirements

- Focus on air and space vehicles within the realm of the NASA mission
- Identify simulation requirements for specific flow phenomena associated with fixed and rotary wing vehicles, engine propulsion, and space systems, from static conditions to hypersonic speeds. Consider CFD requirements in three categories:
  - Solution Accuracy
  - Effective Use of HPC
  - Technology Robustness
- Identify current shortcomings in CFD technology and key impediments in achieving the 2030 vision

CFD Requirements – Solution Accuracy

- Predict **performance/efficiency/noise/loads/emissions** characteristics with certifiable (e.g. FAA) accuracy over the complete flight envelope.
- Routinely calculate both **steady-state** and **time-dependent** flows including simulations with **dynamically deforming geometries** and **relative body motion** with possible changes in topology (e.g., real-time high-lift system deployment, aeroelastic wing response, rotor/airframe interaction, store separation, etc.).
- Be applicable to **all Mach number** ranges from freestream static conditions to subsonic to hypersonic flows, from low to high Reynolds numbers.
- Routinely simulate flows with all types of **flow separation** (e.g., transonic buffet, etc.) and other **complex flow physics** (e.g., chemically reacting flows, etc.)
- Routinely model laminar to turbulent **flow transition** of all modalities (T-S waves, cross-flow, Görtler, and 2nd-mode instabilities; natural and bypass) including effects of surface roughness.
- Perform simulations on **accelerating vehicle flowfields** (i.e. time-dependent Mach number)
- Enable **quantification of various error sources** including discretization (both spatial and temporal), algebraic and modeling errors.
- Provide automated capability for **simulating to overall error tolerances**.
- Provide (as standard output) full quantification of numerical errors, sensitivity information, and computational uncertainty for specified quantities.

CFD Requirements – Effective Use of HPC

- Petascale predictions (and what actually happened …)
- Exascale computing
  - Architecture
  - Programming
- HPC forecast in 2030
  - Programming
- Alternative Computer Architectures
  - Quantum computing

CFD Requirements – Technology Robustness

- Employ an **integrated, fully automated CFD process** from pre- to post-processing, including CAD incorporation, grid generation, and solution adaptive techniques for entire vehicle and propulsion simulations with minimal user intervention.
- Provide an **intuitive parameter-free interface** enabling optimal use for a wide range of problems while minimizing the required user learning curve.
- Accept both **epistemic and aleatory probabilistic inputs** and return suitable outputs for the purposes of quantifying uncertainty.
- Operate across **multi-platform** computing environments.
- Be seamlessly integrated into **visualization and data mining techniques** that make efficient use of results from time-resolved, physically complex flowfield simulations.
- Provide flexible linkages with ground-based and flight test datasets to build **integrated aerodynamic databases** with prescribed confidence intervals throughout the database.
- Enable **coupling** with all other **disciplines in computational mechanics** (e.g., structures, thermal, electromagnetic, etc.) and with flight control system simulations for steady and time dependent trim and maneuver simulations.
- Enable robust simulation capability across all flow regimes, in particular due to nonlinear and transient effects, **without the need for users to perform application-specific tuning**.
- Provide **fault-tolerant simulation** execution, particularly for use with aerodynamic optimization workflows

CFD Requirements – MDA/O and Multi-Physics

- CFD in 2030 is expected to be a critical element of an **integrated, multi-disciplinary analysis and design system** that will be routinely used throughout all phases (conceptual, preliminary, and detailed) of the vehicle design process, and as an essential part of accurate and efficient computational **multi-physics capabilities** coupling fluid dynamics to many other physical processes.
- Focus is on the use of **higher-fidelity**(higher accuracy) CFD tools
- **Mix of loosely- and tightly-coupled CFD tools** with multi-disciplinary and multi-physics capabilities will be required to address the broad range of engineering applications.
- Enable **system-level optimization** of the integrated platform (aero, propulsion, structures, etc.)
- Enable optimization of platform systems (e.g. airframe, propulsion, etc.) **across partnering OEMs**

CFD Requirements – Shortcomings and Impediments

- Separated flows
- Boundary layer transition
- Complex flow physics
- Geometry and grid generation
- Error quantification
- CFD process automation and robustness
- CFD validation
- CFD interaction with MDO and multi-physics
- CFD databases
- Software development for HPC
- Education of the CFD workforce
- Coordinated development efforts between Academia, Government, and Industry

LES of a Powered Aircraft Configuration Across the Full Flight Envelope

- **Assess the ability to use CFD over the entire flight envelope, including dynamic maneuvers**
- **Assess the ability of CFD to accurately predict separated turbulent flows**
  - Monitor increasing LES region for hybrid RANS-LES simulations
  - Evaluate success of wall-modeled LES (WMLES)
  - Determine future feasibility of wall-resolved LES (WRLES)
- **Assess the ability to model or simulate transition effects**
- **Enable future reductions in wind tunnel testing**

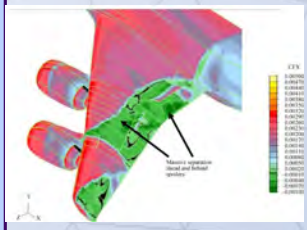

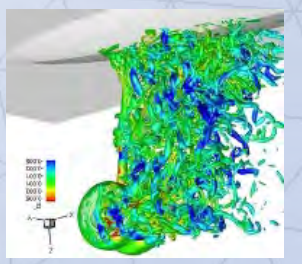

Off-Design Turbofan Engine Transient Simulation

- **Measure progress towards virtual engine testing and off-design characterization**
- **Assess the ability to accurately predict:**
  - Separated flows
  - Secondary flows
  - Conjugate heat transfer
  - Rotating components, periodic behavior
- **Potential to demonstrate industrial use of WRLES for lower Re regions**
- **Assess progress in combustion modeling and prediction abilities**

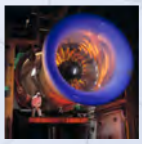
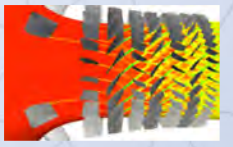
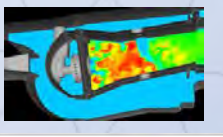

MDAO of a Highly-Flexible Advanced Aircraft Configuration

- **Ultimate utility of CFD for aerospace engineering is as component for MDAO**
- **Future vehicle configurations to be highly flexible**
- **Assess progress in analyzing the important multidisciplinary problems: Science of coupling**
  - Time dependent
  - Aero-structural
  - Aero-servo-elastic
  - Aerothermoelastic
- **Assess multidisciplinary optimization capabilities**
  - Availability of sensitivities
  - Performance of optimizations
  - Optimization under uncertainties

Probabilistic Analysis of a Space Access Vehicle

- **Opening up new frontiers in space vehicle design hinges on development of more capable high-fidelity simulations**
- **Assess specific relevant capabilities**
  - Separated turbulent flows
  - High-speed/hypersonic flows
  - Aero-plume interactions
  - Aerothermal predictions
- **Emphasis on reducing risk through uncertainty quantification techniques**
  - Unique configurations
  - Limited experience base
  - Difficult conditions for ground-based testing

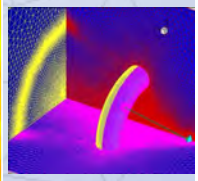

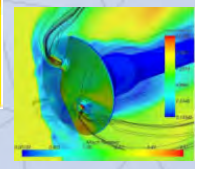

Supporting Safe, Robust Aerospace Structures Through HPC Enhanced Computational NDE

Elizabeth Gregory & Bill Schneck

Computational Nondestructive Evaluation

- Nondestructive Evaluation (NDE) of Aerospace Structures: Detect, quantify, and characterize damage inside aerospace materials/structures through any method that leaves the structure intact and unchanged.
- Computational Nondestructive Evaluation (CNDE): Simulation, modeling, and data analysis tools to aid in the NDE process.

Value to NASA:

- Modeling

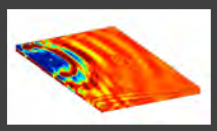

  - Provides NDE information during the design stage leading to inspectable and optimal designs.
  - Makes it feasible to study a large number of damage and flaw cases to establish confidence in inspectability.
  - Cost effective method to develop optimized inspection techniques.

- Data Analysis

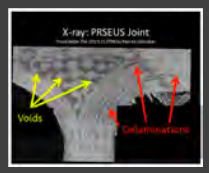

  - Reduces inspection time and cost.
  - Illustrates non-obvious trends and relationships.
  - Can remove subjectivity of human interpretation.
  - Connects with modeling to inform decision making.

Ultrasound Inspection Simulation CompCell EFIT

- Elastodynamic Finite Integration Technique optimized for Intel Xeon Phi Knights Landing processors
- Cannot handle anisotropic heterogeneous media (e.g. composites)

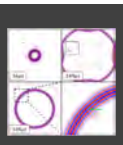

X-Ray CT Ray Tracing

- A simplified interface layer for scientific ray-tracing has been implemented as an available library for additional 'app' development.

RSG

- 3D staggered grid time-domain finite difference
- C++ code with MPI parallelization, CPU kernel validated with GPU kernel under development
- Natively handles 3D fully anisotropic heterogeneous media

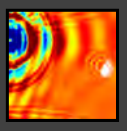

Data Fusion and ML/AI

- Automated Analysis of many GB of thermographic data from Automated fiber placement (AFP), quickly.
- Tensorflow with python real-time processing.

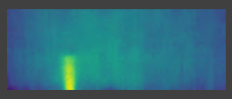

Ultrasound Simulation

Ultrasound Testing Simulation for NDE

Ultrasonic elastic waves excited using transducer, piezoelectric wafers, pulse laser, or other methods. → Elastic waves interact with defect/damage, propagating with damage information. → Detection of these waves using a piezoelectric transducer, fiber optic sensor, non-contact laser Doppler vibrometer, etc.

Simulation can:

- **Provide insight** into underlying interaction phenomena and energy paths.
- **Predict and quantify** energy-damage interaction.
- **Identify features** correlated with relevant parameters.
- **Optimize** inspection setup.
- **Provide confidence** in inspection results.

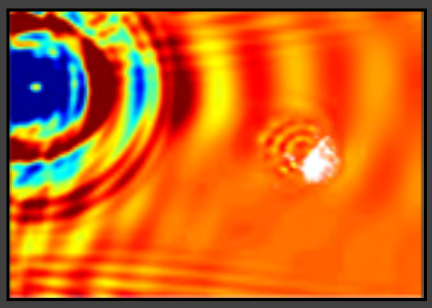

EFIT CompCell Ultrasound Simulation Code

- EFIT - Elastodynamic Finite Integration Technique: 2nd order Finite Difference scheme
- CompCell: 3D modular code version of EFIT aimed at complex geometries and written in C++
- Works well for isotropic and orthotropic materials
- Used to simulate a crack in a stringer of an Aluminum fuselage panel in collaboration with the FAA.
- Scaled well for medium to large problems.
- Not appropriate for anisotropic materials or small problems that result in too much communication.

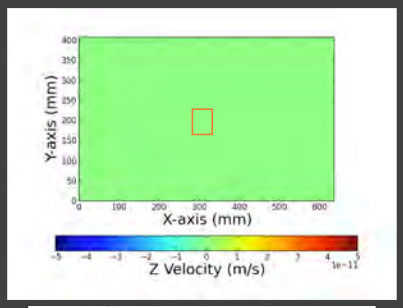

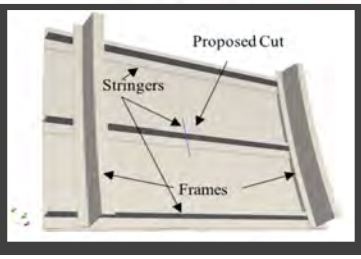

Rotated-Staggered Grid Ultrasound Simulation Code

- Development began under the Advanced Composites Project (ACP).
- Computational tool that models the physics of an ultrasound (UT) inspection
  - 3D staggered grid time-domain finite difference
  - C++ code with MPI parallelization, CPU kernel validated with GPU kernel under development
  - Natively handles 3D fully anisotropic heterogeneous media
- Solves stress-velocity elastodynamic equations, 2nd order accurate in time and space.

$$\rho \partial_t v_i = \partial_{x_j} \sigma_{ij} + f_i^v$$

$$\partial_t \sigma_{ij} = c_{ijkl} \partial_{x_l} v_k + f_{ij}^\sigma$$

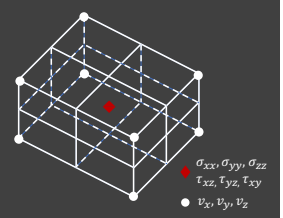

Repeating unit cell for ultrasound simulation algorithm – rotated staggered grid (RSG) finite difference

General Capabilities

- MPI parallelized, highly optimized solver
- Verified and validated for a range of general elastodynamic cases

Additive Manufacturing (AM) Applications

- Exploration of flaw parameter space and microstructural effects on UT inspection
- Guiding the emergence of novel UT techniques for AM part inspection
- Help define limits of inspectability as relevant to effect-of-defect analysis

Benchmark Case – Delaminated Composite Plate

- Experimental, FEM, and in-house code benchmark for composite laminate
  - Delamination between 2nd and 3rd ply
  - Experimental data full wavefield capture via laser Doppler vibrometry
- Quantifiable velocity and wavenumber changes above delamination observed via simulation and experiment

v_z surface wavefield, RSG simulation code

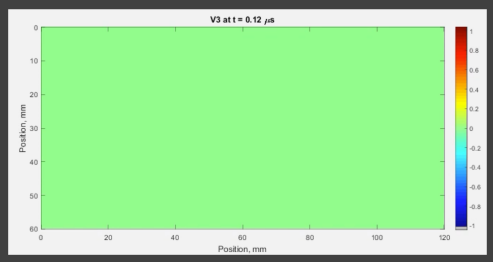

Experimental Wavefield data

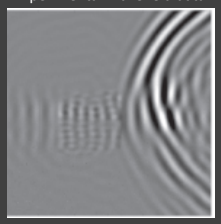

Run time comparison with COMSOL, Abaqus, ANSYS

	T_{run} (hours)	# DoF (10^6)	Memory (GB)
COMSOL	19.5	7.77	266
ABAQUS Implicit	40	1.7	28
ABAQUS Explicit	53	26	30
ANSYS	170	10	16
RSG	0.15	16.6	4.6

Between ~130-1100 times faster simulation time for converged model, using comparable compute resources

Leckey, Cara AC, et al. "Simulation of guided-wave ultrasound propagation in composite laminates: Benchmark comparisons of numerical codes and experiment." Ultrasonics 84 (2018): 187-200.

Data Fusion, Analysis, and Machine Learning for In Situ Thermographic Inspection

In Situ Thermographic Inspection System (ISTIS) of Automated Fiber Placement

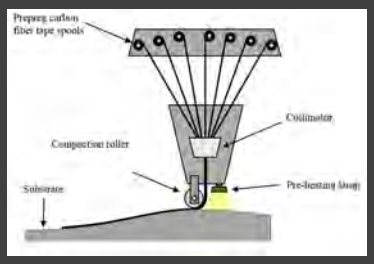

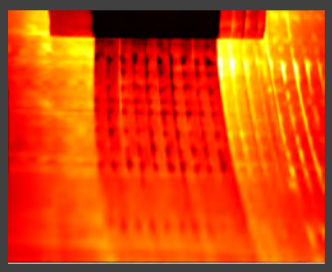

- Faster than traditional hand layup methods, less scrap waste
- Better repeatability
- Advances in robotic automation can enable rapid prototyping and more complex structures

- Added IR camera is minimally invasive.
- Uses AFP head for power and communication.
- Results in 100s of GB of data for a single part.

- Shows:
  - Tow Overlaps and Gaps.
  - FOD
  - Peel-up
  - Tow Twists

Data Registration and Fusion

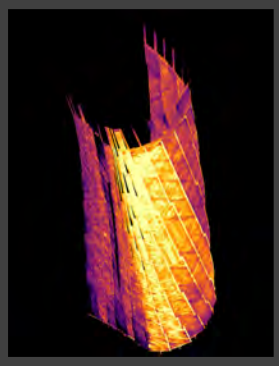

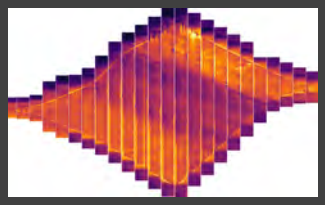

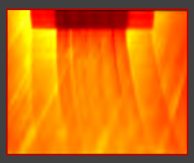

- Large scale data (~100GB) processing required for real-time analysis
- Optimized, rapid code is required
- Frames
- Courses
- 2D Plies
- 3D Plies
- Snapshots
- Time histories
- Registration in Space and Time

Machine Learning for Labeling and Localization

- Machine Learning Techniques
  - Convolution Neural Networks for Labeling
  - Clustering for Segmentation
  - Encoder and Decoder for Segmentation and Labeling
  - Edge Detection
  - Background Subtraction
- Optimized code required for real-time analysis.
- Data Labelling services
  - AWS SageMaker Ground Truth
  - Google Cloud Data Labelling

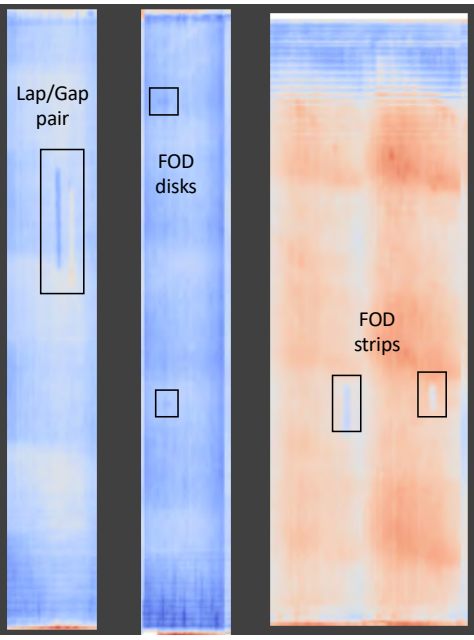

Impact of Fast HPC Code:

- CNDE Physics Simulations (ultrasound & X-ray mini-apps)
  - Fast code enables use of simulation tools to identify inspection solutions for NASA's unique inspection challenges
    - Rapid simulations (running in less than 60 seconds) is a requirement for many optimization methods (such as HPC based Monte Carlo)
  - Quick iterations enable robust exploration of the problem space, reducing risk
  - Parameter exploration enables solving the 'inverse problem' for improved damage/defect characterization
- Large Scale Data Analysis
  - Data-informed decisions are enabled by ability to rapidly analyze very large datasets to yield quick and robust identification of trends and features.
  - Provides insight at the time of data collection instead of hours or days later.
  - Machine learning is only possible with efficient HPC based code

The Path Forward

- X-Ray CT: Continuing model development under the Convergent Aeronautics Solutions (CAS) - Sensor-based Prognostics to Avoid Runaway Reactions & Catastrophic Ignition (SPARRCI) project to model Lithium metal battery X-Ray inspections.
- Ultrasound Simulation: Transformational Tools and Technologies (TTT) is using ultrasound simulation code (RSG and components of CompCell) to simulate inspection of additively manufactured (AM) parts to enable faster certification of AM parts.
- ISTIS: Advanced Air Transportation Technology (AATT) is using ISTIS for AFP process development and AFP part inspection.

Lessons Learned

Trade names and trademarks are used in this report for identification only. Their usage does not constitute an official endorsement, either expressed or implied, by the National Aeronautics and Space Administration

Collaboration and Portability and Continuous Integration and Testing

- Use version control to collaborate and share.
- Use workflows to prevent conflict and maintain functionality.

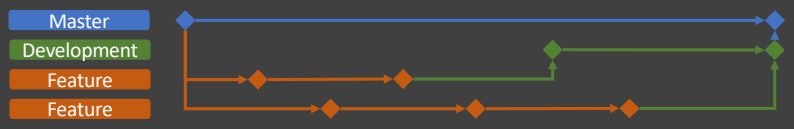

- Establish style guidelines for the team (inherit from another guide).
- Interface control documentation for components.
- Select a testing framework or combination of testing frameworks.
  - CMAKE® (of Kitware, Inc.)
  - PYTEST℠ (of H. Krekel, et al.)
  - GOOGLETEST℠ (of Google, LLC)
- If you can't test it, it is bad code.
- Test often.

Machine Learning and AI Tools

- Use Virtual Environments in python on K-cluster to maintain package compatibility. (matches Pleiades environment)
- Use cloud computing resources when applicable to leverage their ML/AI environments.
- Services:
- KERAS℠ (of F. Chollet, et al.)
- PyTorch
- TENSORFLOW® (of Google, LLC)
- GOOGLE® (of Google, LLC) Cloud
- AMAZON WEB SERVICES® (of Amazon Technologies, Inc.)

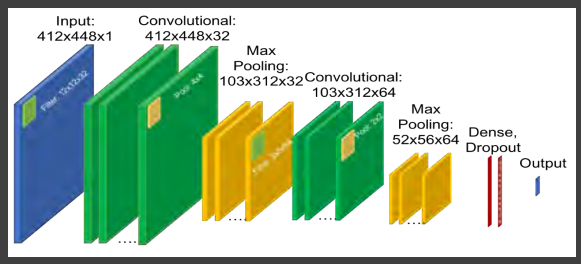

Selected Publications

- W.C. Schneck III, E.D. Gregory, C.A.C. Leckey, H. Reed, *Optimized Manycore Ultrasonic Simulation Code with Sequential Monte-Carlo Based Parameter Estimation*, QNDE 2018.
- E.D. Gregory, W.C. Schneck III, C.A.C. Leckey, P. Swindell, *Simulation Assisted Guided Wave Structural Health Monitoring of Aerospace Structures*, QNDE 2018.
- H. Reed, W.C. Schneck III, E.D. Gregory, C.A.C. Leckey, *Sequential Monte Carlo Based Parameter Estimation for Structural Health Monitoring*, QNDE 2018.
- E.D. Gregory, W.C. Schneck III, C.A.C. Leckey, *A Versatile Simulation Framework for Elastodynamic Modeling of Structural Health Monitoring*, MODSIM World 34, 2018.
- W.C. Schneck III, E.D. Gregory, C.A.C Leckey, *Optimization of Elastodynamic Finite Integration Technique on Intel Xeon Phi Knights Landing Processors*, Journal of Computational Physics, submitted March 2018.

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

ECCO-ACCESS: a Cloud-Native Storage and Data Analysis System for Ocean Climate Research

Dr. Ian Fenty (Co-I, JPL)

Ian.Fenty@jpl.nasa.gov

Project Team:
- Dr. Patrick Heimbach (PI, UT Austin)
- Thomas Huang (Co-I, JPL)
- Duncan Bark (*intern*, UCSC)
- Maya Debellis (JPL)
- Eamon Ford (JPL)
- Kevin Gill (JPL)
- Dr. Joe Jacob (JPL)
- Kevin Marlis (*intern*, CSLA)
- Dr. Ou Wang (JPL)
- Elizabeth Yam (JPL)

© 2020 California Institute of Technology. Government sponsorship acknowledged.
Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsements by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Outline

- ECCO Background
- Big Data Challenges
- ECCO-Cloud System
- Future Directions

Goal of ECCO: *to make the best possible estimates of the time-evolving ocean circulation in support of climate research*

- ECCO is a *multi-platform, multi-instrument, multi-variable* synthesis product that integrates ocean and ice observations with models
- ECCO Consortium is comprised of an international group of scientists across several institutions

NASA Program Support

Physical Oceanography (PO)

Cryosphere

Modelling, Analysis, and Prediction (MAP)

Advancing Collaborative Connections for Earth System Science (ACCESS)

ECCO is a *multi-platform, multi-instrument, multi-variable* synthesis product that integrates ocean and ice observations with models

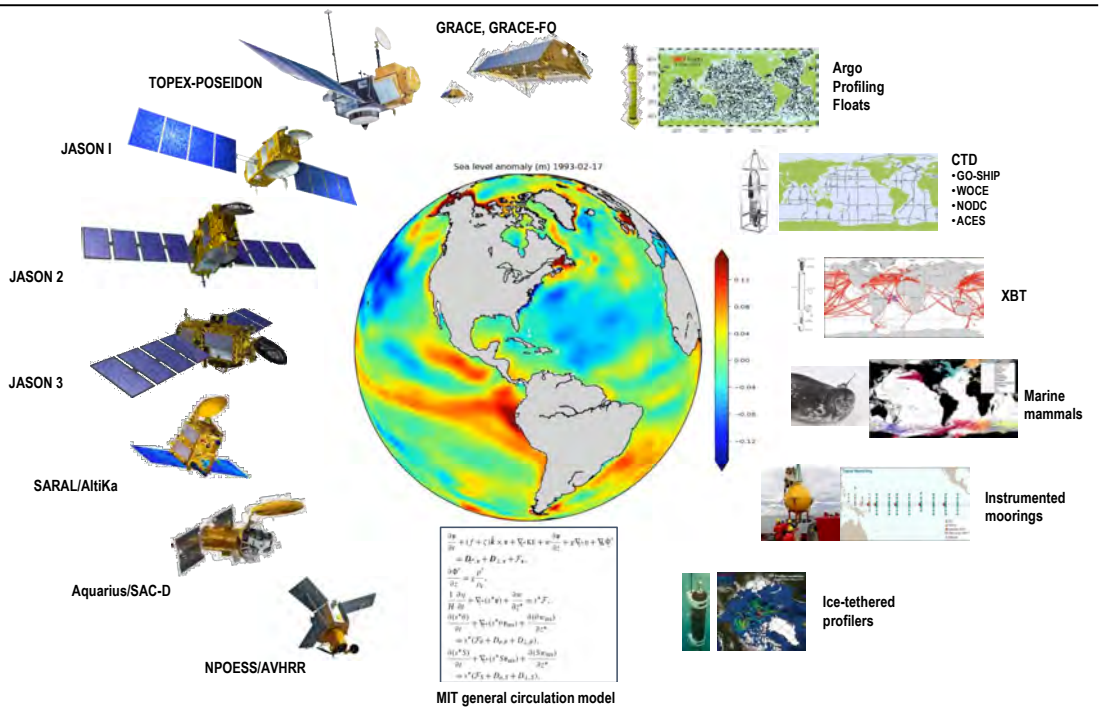

Goals and approach of Ocean State Estimation are distinct from those of Numerical Weather Prediction

Numerical Weather Prediction: a "filtering" problem

- Goal: optimal forecasting
- find initial conditions which produce best possible forecast
- ***dynamical consistency and property conservation NOT required for prediction***

Ocean state estimation: a "smoothing" problem

- Goal: understand the state of the ocean and its dynamical evolution from past to present
- find initial and boundary conditions such that a free-running ocean general circulation model (OGCM) evolves consistently with observations
- ***dynamic consistency and property conservation ESSENTIAL for climate research***

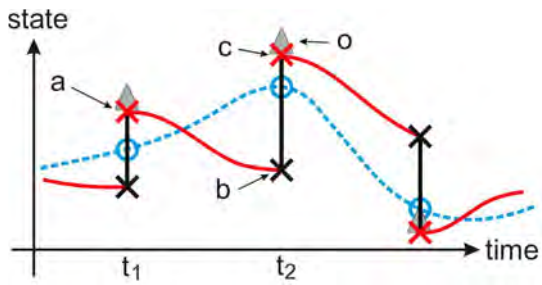

x : observation (data)
— : forecast (prediction)
x : forecast at observation time
--- : state estimate
o : state estimate at observation time

Wunsch and Heimbach (Physica D 2007)
Forget et al. (Geosci. Model Dev. 2015)
Stammer et al. (Annu. Rev. 2016)
Heimbach et al. (Front. Mar. Sci. 2019)

FENTY/JPL © 2020. All rights reserved.

ECCO Central Estimate Product: Version 4, Release 4

Period: 1/1992-12/2017

Resolution: 0.3˚~1˚, 50 levels

Ocean + sea-ice fields

- *T, S, u, v, w, η, ρ, Φ*
- Sea-ice and snow *h* and *c*
- Lateral and vertical fluxes of heat, salt, volume, and momentum

Atmosphere fields

- T, q, |u|, τ, long- and radiative fluxes
- Air–sea-ice–ocean fluxes of heat, moisture, energy, and momentum

Subgrid-scale mixing parameters

- 3D GM κ and Redi κ
- 3D vertical diffusivity

Monthly & select daily mean fields = 7 TB

ECCO fields are provided on two grids

0.5° lat-lon

"lat-lon-cap 90"

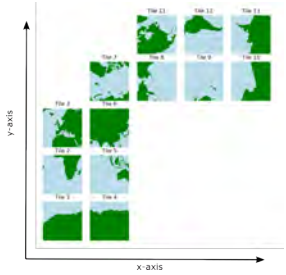

13 tiles of 90x90x50

FENTY/JPL © 2020. All rights reserved.

Outline

- ECCO Background
- Big Data Challenges
- ECCO-Cloud System
- Future Directions

Big Data Challenges: Data Discovery and Acquisition

- NASA distributes **100s of ocean and ice datasets,** many of which could be integrated into ECCO
- NASA DAACs continually add new datasets and update existing datasets, *with new data arriving at an accelerating rate.*
  - new missions
  - new data from ongoing missions
  - updated/reprocessed data
- **Searching for new ocean and ice data is time-consuming** → *latest datasets are rarely used.*
- Instead of "***Earthdata Search by CMR***" we rely more on "***Earthdata Search by serendipitous hallway conversations and conference posters***".
- Datasets contain >10^3-10^4 granules for ECCO period (1992-present)
  - long transfer times
  - large local disk storage requirements

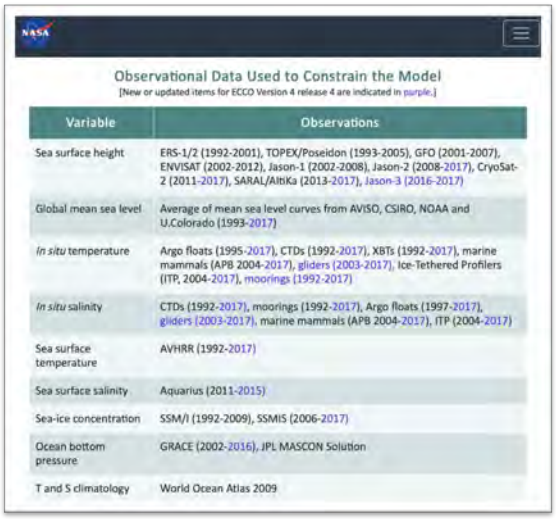

Observational Data Used to Constrain the Model

[New or updated items for ECCO Version 4 release 4 are indicated in purple.]

Variable	Observations
Sea surface height	ERS-1/2 (1992-2001), TOPEX/Poseidon (1993-2005), GFO (2001-2007), ENVISAT (2002-2012), Jason-1 (2002-2008), Jason-2 (2008-2017), CryoSat-2 (2011-2017), SARAL/AltiKa (2013-2017), Jason-3 (2016-2017)
Global mean sea level	Average of mean sea level curves from AVISO, CSIRO, NOAA and U.Colorado (1993-2017)
In situ temperature	Argo floats (1995-2017), CTDs (1992-2017), XBTs (1992-2017), marine mammals (APB 2004-2017), gliders (2003-2017), Ice-Tethered Profilers (ITP, 2004-2017), moorings (1992-2017)
In situ salinity	CTDs (1992-2017), moorings (1992-2017), Argo floats (1997-2017), gliders (2003-2017), marine mammals (APB 2004-2017), ITP (2004-2017)
Sea surface temperature	AVHRR (1992-2017)
Sea surface salinity	Aquarius (2011-2015)
Sea-ice concentration	SSM/I (1992-2009), SSMIS (2006-2017)
Ocean bottom pressure	GRACE (2002-2016), JPL MASCON Solution
T and S climatology	World Ocean Atlas 2009

ECCO would benefit from regular updates about new and updated datasets and from new approaches for acquiring and keeping track of downloaded granules.

Big Data Challenges: Data Transformation

- Data granules must be ***transformed to a common grid*** before ECCO assimilation
- Data products have ***dazzling variation:***
  - spatial grid resolution and array orientation
  - geographic grid projection
  - temporal frequency or averaging period (daily, 7-day, monthly, etc.)
  - uncertainty, quality flags, etc.
- Developing transformation codes for each unique dataset is time consuming, ***limiting the number of assimilated data products***.
- Datasets require ***10s of hours*** to process.
- Data transformation routines are ***ad-hoc***, residing on the local machines of Co-I's.
- Data provenance from source (EOSDIS) to final assimilation ***is not logged***

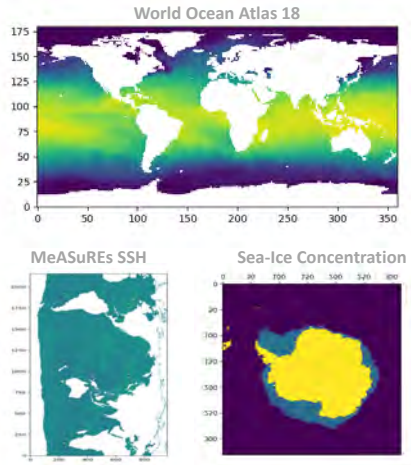

ECCO would benefit from new approaches for transforming dataset granules and from a better system of recording data provenance from "acquisition to assimilation".

Big Data Challenges: Data Distribution

- ECCO provides ***80+ geophysical fields***
- Historically ECCO fields are distributed as ***monthly averages*** because of unaffordable storage requirements **for us** *(as data provider)* and for users, although ***hourly*** fields are available!
- ECCO metadata is totally ad-hoc and non-conforming to modern conventions
  - **CF**: *Climate and Forecast*
  - **ACDD**: *Attribute Convention for Data Discovery*
- Preparing ECCO granules for distribution is labor intensive resulting in infrequent updates (***yearly or longer***), although ***quarterly*** updates are possible
- As ECCO spatial resolution and product time period increase, user requirements for data transfer and local storage ***grow exponentially***.
- Next generation ECCO product (expected q4 2020) will be **70+ TB**

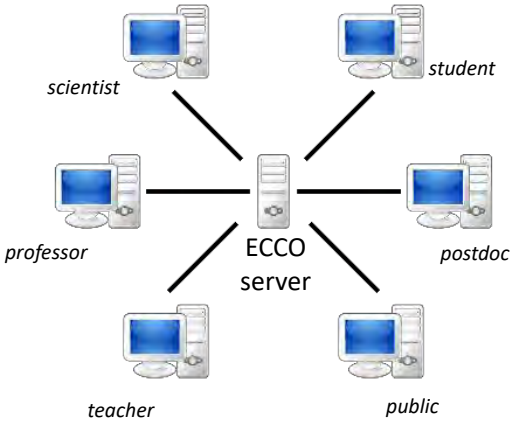

The ECCO user community would benefit from more frequent, more complete, and more standardized distribution of our data product

Big Data Challenges: Data Analysis

- ECCO products can provide unique insights into the mechanisms and dynamics of the 3D ocean and sea-ice state
- Typical analyses ***require O(100 GB)***
  - transfer to local disk
  - local I/O
  - local calculations
- Access to machines with large local storage and memory ***is not universal***
- Analyses with ECCO lend themselves to parallelization, ***but user machines have few cpus and are I/O time limited***.
- As ECCO product evolves, ***local analysis becomes increasingly difficult***.
- ***Users migrate to simpler, less accurate data products*** to avoid local computational and storage limitations.

The ECCO user community would benefit from new approaches to analyze data products without excessive local computational and storage requirements.

FENTY/JPL © 2020. All rights reserved.

Outline

- ECCO Background
- Big Data Challenges
- ECCO-Cloud System
- Future Directions

FENTY/JPL © 2020. All rights reserved.

ECCO-Cloud System

A *production-ready cloud-native storage* and *data analysis system* created to achieve the following goals:

(1) expand and accelerate the assimilation of NASA Earth System data into ECCO through improved *data discovery* and cloud-based *data preprocessing and transformation*

(2) radically streamline integration of ECCO products into NASA's Earth Observing System Data and Information System (EOSDIS), specifically the Physical Oceanography Distributed Active Archive Center (PO.DAAC)

(3) facilitate and expand the scientific utilization of NASA remote sensing data integrated into ECCO by the growing community of interdisciplinary researchers in the oceanographic, sea-ice, sea-level rise, and climate fields.

We conceived of the ECCO-Cloud system after learning about new search, data interoperability, and cloud-based analysis technologies developed within the ACCESS program.

FENTY/JPL © 2020. All rights reserved.

ECCO-Cloud System Architecture

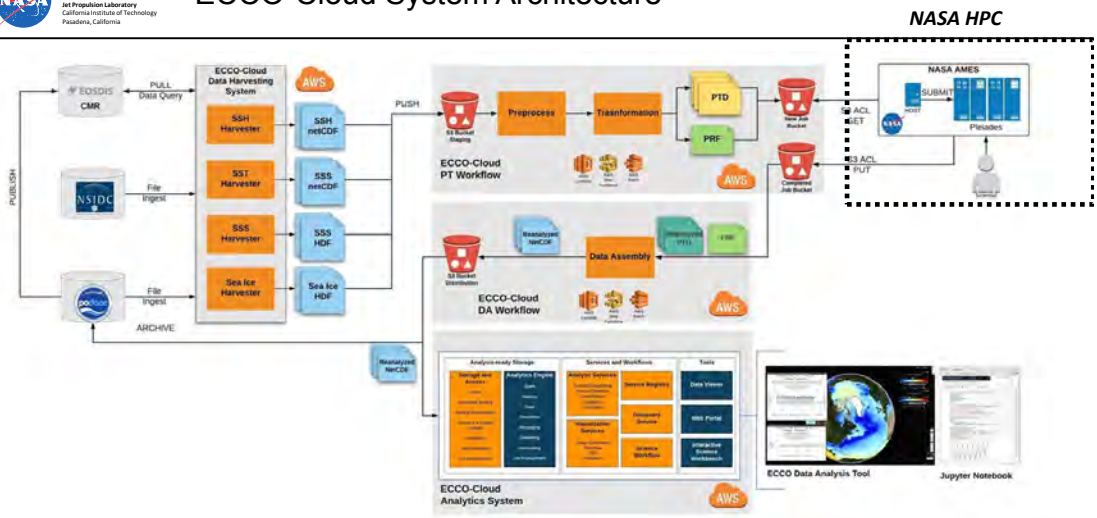

- **Data Harvesting System:** Queries Earthdata CMR for relevant data products, ingest them from the designated DAAC, and stage the ingested product to Object Store to trigger Product Transformation Workflow
- **Product Transformation Workflow**: preprocess the harvested data optimized for HPC with provenance and job tracking metadata
- **Data Assembly Workflow**: takes the reanalysis output from HPC and generates CF-compliant netCDF output for long-term archive
- **Analysis System**: Ingest the CF-compliant netCDF product for cloud-optimized data analysis services

FENTY/JPL © 2020. All rights reserved.

ECCO-Cloud System Architecture

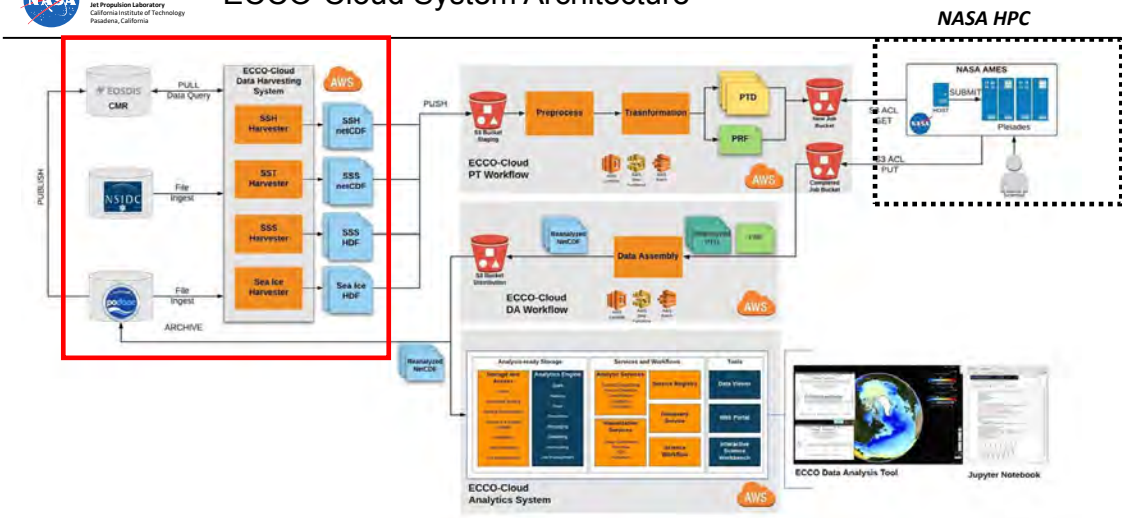

- **Data Harvesting System:** Queries Earthdata CMR for relevant data products, ingest them from the designated DAAC, and stage the ingested product to Object Store to trigger Product Transformation Workflow
- **Product Transformation Workflow**: preprocess the harvested data optimized for HPC with provenance and job tracking metadata
- **Data Assembly Workflow**: takes the reanalysis output from HPC and generates CF-compliant netCDF output for long-term archive
- **Analysis System**: Ingest the CF-compliant netCDF product for cloud-optimized data analysis services

FENTY/JPL © 2020. All rights reserved.

Data Discovery Tool

Big Data Challenge: Data Discovery

- Queries Earthdata CMR for relevant data sets given a set of ***GCMD*** or ***free keyword*** search terms
- Generates regular reports highlighting potentially relevant datasets for ECCO scientists, including direct links to product pages on DAACs
- Uses ECCO's cloud-based registry (SOLR) to compare current vs. past state of Earthdata CMR

```
CMR Data Discovery Tool Report: 2020-04-09
==========================================================

Query Dates: current v.s reference (delta days)
----------------------------------------------------------
2020-04-09  vs.  2020-04-08  (1 days)

New and updated products report filenames
----------------------------------------------------------
1. new_product_2020-04-09_vs_2020-04-08.json
2. updated_product_2020-04-09_vs_2020-04-08.json
3. unchanged_product_2020-04-09_vs_2020-04-08.txt
4. unchanged_product_2020-04-09_vs_2020-04-08.json

Query keywords and # hits (# reference hits)
----------------------------------------------------------
  1. sea_surface_temperature    200 (199)
  2. sea_ice    200 (200)


----------------------------------------------------------
NEW/UPDATED PRODUCTS SUMMARY (reason for difference)
----------------------------------------------------------
New Products
----------------------------------------------------------
  1. OSCAR_L4_OC_1deg     (new product)
    https://doi.org/10.5067/OSCAR-10D01
      keywords: sea surface temperature


Updated Products and Updated Field(s)
----------------------------------------------------------
  1. G02202     (EndingDateTime)
    https://doi.org/10.7265/N59P2ZTG
      keywords: sea ice
                EndingDateTime
                        2018-12-31T23:59:59.999Z (current)
                        2017-12-31T23:59:59.999Z (reference)


Unchanged Products
----------------------------------------------------------
for a list of unchanged products, see unchanged_product_2020-04-
for full details of the unchanged products, see unchanged_produc
```

query interval

search terms

links to new and updated products

FENTY/JPL © 2020. All rights reserved.

Data Harvester Tool

Big Data Challenge: Data Acquisition

- ***Harvester transfers individual granules for each dataset*** from DAAC to cloud or local storage using CMR + OPeNDAP, or other method if necessary
- Creates detailed entries into a SOLR database for **each dataset** and **granule**.
- Database entries are initial record of ***data provenance***
- Determination of new or updated granules made by ***comparing database entries with CMR queries***
- Only new or updated data is transferred from DAAC and staged for transformation.

database **DATASET** entry


```
"response":{"numFound":1,"start":0,"docs":[
  {
    "type_s":"dataset",
    "dataset_s":"sst_AVHRR",
    "last_checked_dt":"2020-07-07T20:01:18Z",
    "years_updated_s":"1992",
    "start_date_dt":"1992-01-01T00:00:00Z",
    "end_date_dt":"1993-01-01T00:00:00Z",
    "last_download_dt":"2020-07-07T20:01:18Z",
    "status_s":"harvested",
    "source_s":"https://podaac.jpl.nasa.gov/ws/search/granule/?formats
    "fields_s":"analysed_sst, analysis_error",
    "id":"00b6d83c-a3d8-473f-b9cc-d73712a37fe1",
    "llc90_factors_path_s":"/Users/kevinmarlis/Developer/JPL/src/grids/
    "llc90_factors_stored_dt":"2020-07-07T20:06:56Z",
```


dataset name

download date

fields

database **HARVESTED GRANULE** entry


```
"response":{"numFound":366,"start":0,"docs":[
  {
    "type_s":"harvested",
    "date_s":"1992-12-31T00:00:00Z",
    "dataset_s":"sst_AVHRR",
    "source_s":"https://podaac-opendap.jpl.nasa.gov/opendap/a
    "modified_time_dt":"1992-12-31T00:00:00Z",
    "checksum_s":"836ffa7f7d5b7bccda02d34d7e20c867",
    "pre_transformation_file_path_s":"/Users/kevinmarlis/Deve
    "harvest_success_b":true,
    "filename_s":"19921231120000-NCEI-L4_GHRSST-SSTblend-AVHR
    "file_size_l":1061945,
    "download_time_dt":"2020-07-07T20:01:18Z",
```


granule date

transfer status

ECCO-Cloud System Architecture

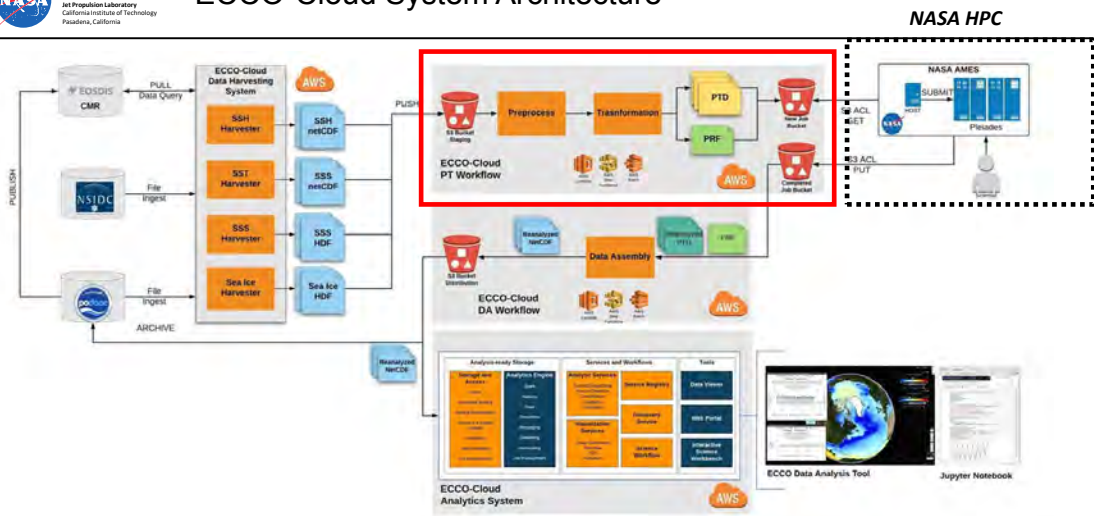

- **Data Harvesting System:** Queries Earthdata CMR for relevant data products, ingest them from the designated DAAC, and stage the ingested product to Object Store to trigger Product Transformation Workflow
- **Product Transformation Workflow**: preprocess the harvested data optimized for HPC with provenance and job tracking metadata
- **Data Assembly Workflow**: takes the reanalysis output from HPC and generates CF-compliant netCDF output for long-term archive
- **Analysis System**: Ingest the CF-compliant netCDF product for cloud-optimized data analysis services

Product Transformation Pipeline

Big Data Challenge: Data Transformation

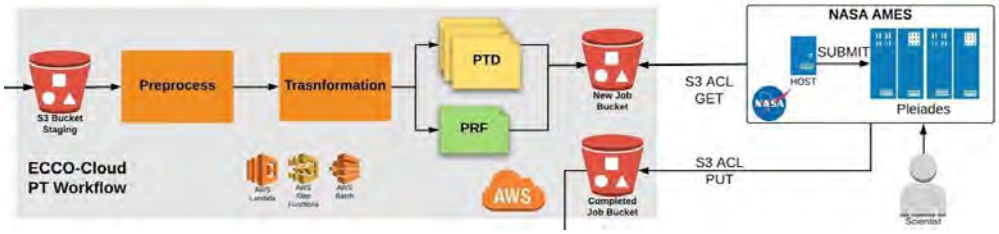

- Prepares granules for assimilation into ECCO
- Queries SOLR database to determine granules requiring attention
- Uses consistent and reproducible routines to transform granules to common ECCO grid
- Aggregates transformed granules, saves in **netCDF** (with metadata) and **flat binary** (for assimilation) formats, stages for transfer to NASA Pleiades supercomputer
- Records all transformations and aggregations to SOLR database, extending data provenance record
- ***Informs ECCO scientists that new data is ready!***

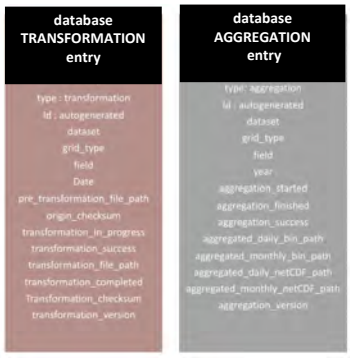

ECCO-Cloud System Architecture

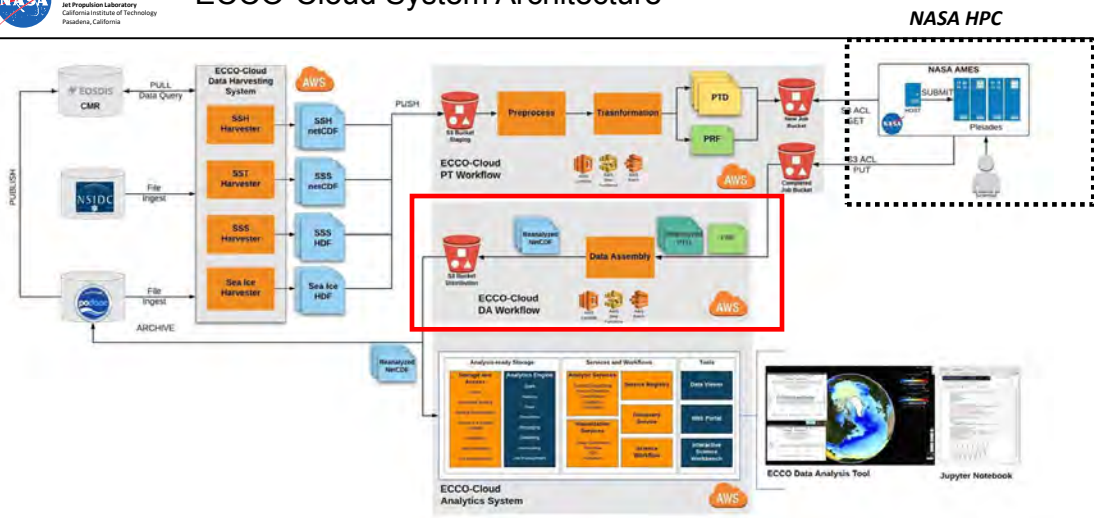

- **Data Harvesting System:** Queries Earthdata CMR for relevant data products, ingest them from the designated DAAC, and stage the ingested product to Object Store to trigger Product Transformation Workflow
- **Product Transformation Workflow**: preprocess the harvested data optimized for HPC with provenance and job tracking metadata
- **Data Assembly Workflow**: takes the reanalysis output from HPC and generates CF-compliant netCDF output for long-term archive
- **Analysis System**: Ingest the CF-compliant netCDF product for cloud-optimized data analysis services

Product Assembly Pipeline

Big Data Challenge: Data Distribution

curvilinear model grid

regular lat-lon grid

- Transforms ECCO model output from **flat binary** to **NetCDF**
- Remaps fields to a regular latitude-longitude grid from curvilinear model grid
- Adds metadata conforming to netCDF Climate and Forecast Metadata Conventions (**CF**) and Attribute Convention for Data Discovery **(ACDD)** including
  - standard name and units (for interoperability)
  - geographical and temporal extents (for search)
- Groups related parameters together to reduce total number of ECCO Datasets (from 80+ parameters to ~16 datasets)
- Creates new entries in SOLR database
- Stages files to "distribution bucket" for PO.DAAC

FENTY/JPL © 2020. All rights reserved.

Product Assembly Pipeline

Big Data Challenge: Data Distribution

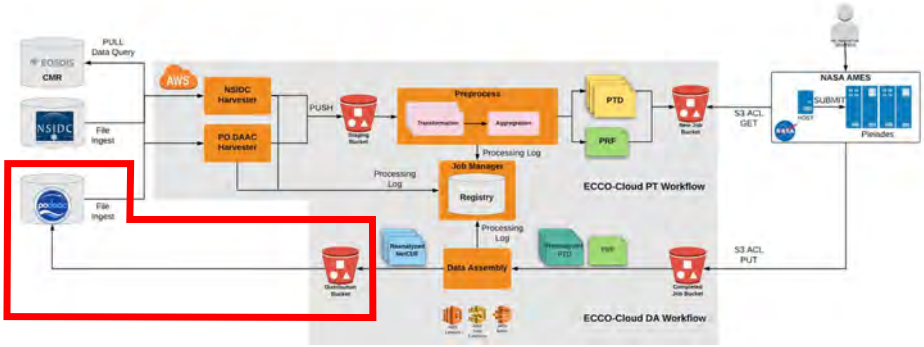

- Technical Interchange with PO.DAAC Management, System Engineer, and Data Engineer
- **Product Delivery Interface**: Read-only S3 bucket
- **Product Packaging**: CF-compliant netCDF
- **Information Model**: UMM-C and UMM-G
- **Time Resolution**: Daily and Monthly
- Grid: latitude-longitude and native model grid
- GCMD Keyword mapping

NetCDF filename syntax

vvvvvvvv_ppp_tttt_gggggg_YYYY-MM-DD-ECCO_V4r3_rrrrrr.nc

vvvvvvvv - ECCO parameter or group name
ppp – "day" or "mon" averaging period
tttt – "mean" or "snap" (mean or instantaneous snapshot)
YYYY - year, MM - month, DD - day
Vxry – ECCO version 'x' release 'r'
gggggg – grid type, "latlon" for latitude-longitude or "native" for model
rrrrrr- grid resolution (for regular "latlon" grid) or native (e.g., llc90)

Example: *Slarea_day_mean_latlon_1992-01-01_ECCO_V4r4_0p50deg.nc*
Example: *Slarea_mon_mean_native_1992-01_ECCO_V4r4_llc90.nc*

FENTY/JPL © 2020. All rights reserved.

ECCO-Cloud System Architecture

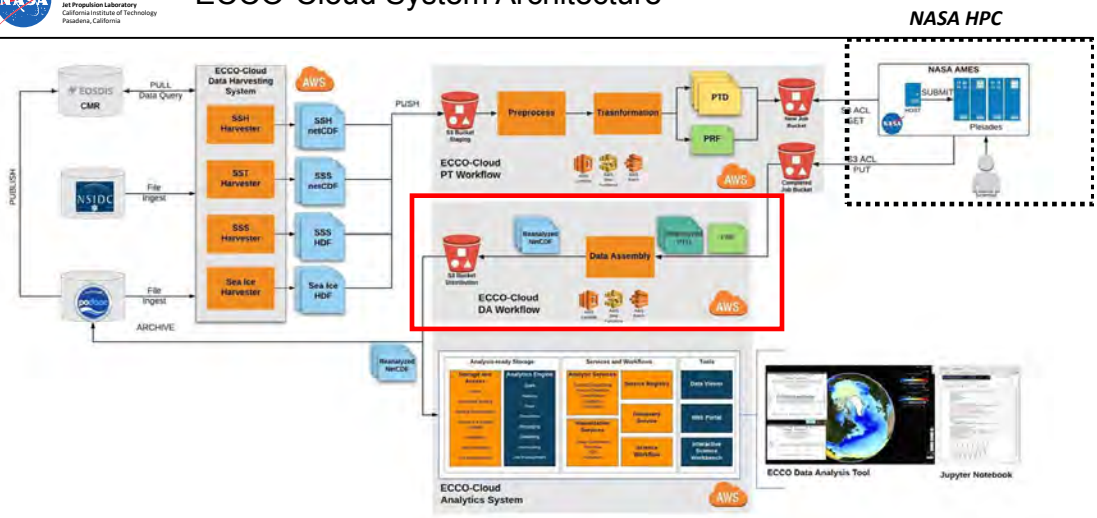

- **Data Harvesting System:** Queries Earthdata CMR for relevant data products, ingest them from the designated DAAC, and stage the ingested product to Object Store to trigger Product Transformation Workflow
- **Product Transformation Workflow**: preprocess the harvested data optimized for HPC with provenance and job tracking metadata
- **Data Assembly Workflow**: takes the reanalysis output from HPC and generates CF-compliant netCDF output for long-term archive
- **Analysis System**: Ingest the CF-compliant netCDF product for cloud-optimized data analysis services

ECCO-Cloud Analytics System

Big Data Challenge: Data Analysis

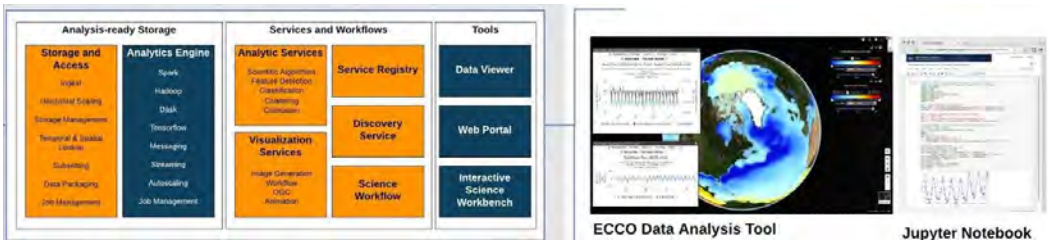

- ECCO fields are ingested to the "**ECCO-Cloud Analytics System**", an open-source solution based on **SDAP** (the Apache Science Data Analytics Platform)
- ECCO-Cloud Analytics System offers ***cloud-based data management and parallel computing***
- Researchers ***can focus on analysis***, ***not data wrangling***
- Analysis is done "close to the data": no need to transfer ECCO fields to local machines
- Access to ECCO fields via simple "webservice" calls using any programming language (e.g. Python, IDL, MATLAB, C/C++, JavaScript)
- ***EXAMPLES***
  - ***(1) ECCO Data Analysis Tool***
  - ***(2) Python and Jupyter Notebook***

ECCO-Cloud Analytics System: **Data Analysis Tool**

Big Data Challenge: Data Analysis

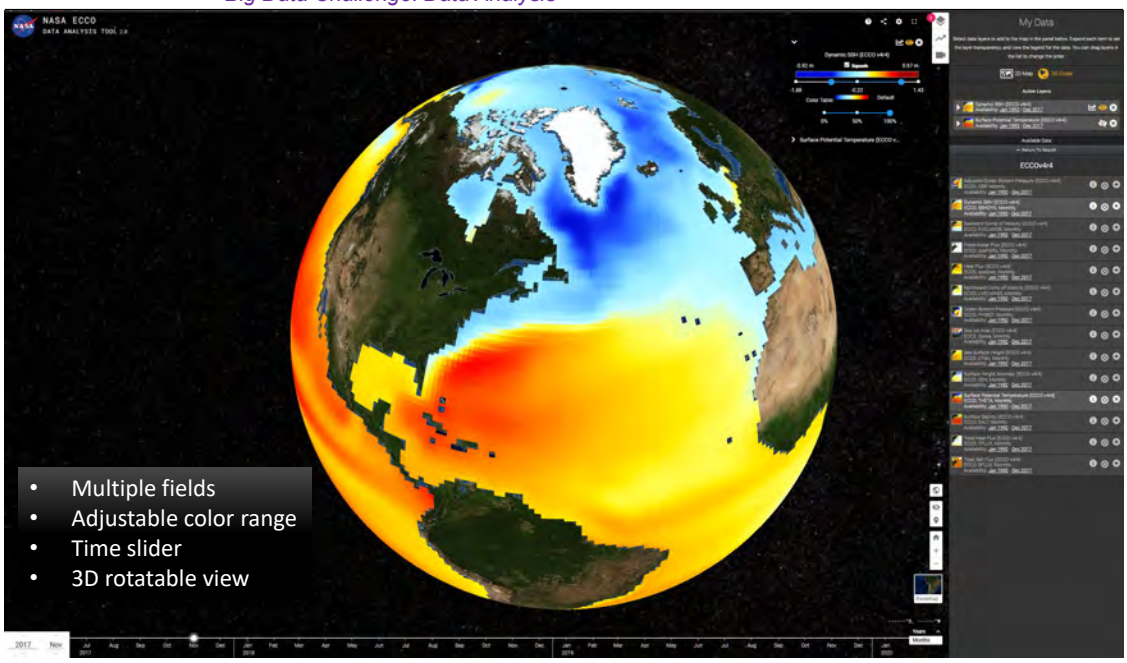

Extends technology deployed on the NASA Sea Level Portal
https://sealevel.nasa.gov/data-analysis-tool/

ECCO-Cloud Analytics System: **Data Analysis Tool**

Big Data Challenge: Data Analysis

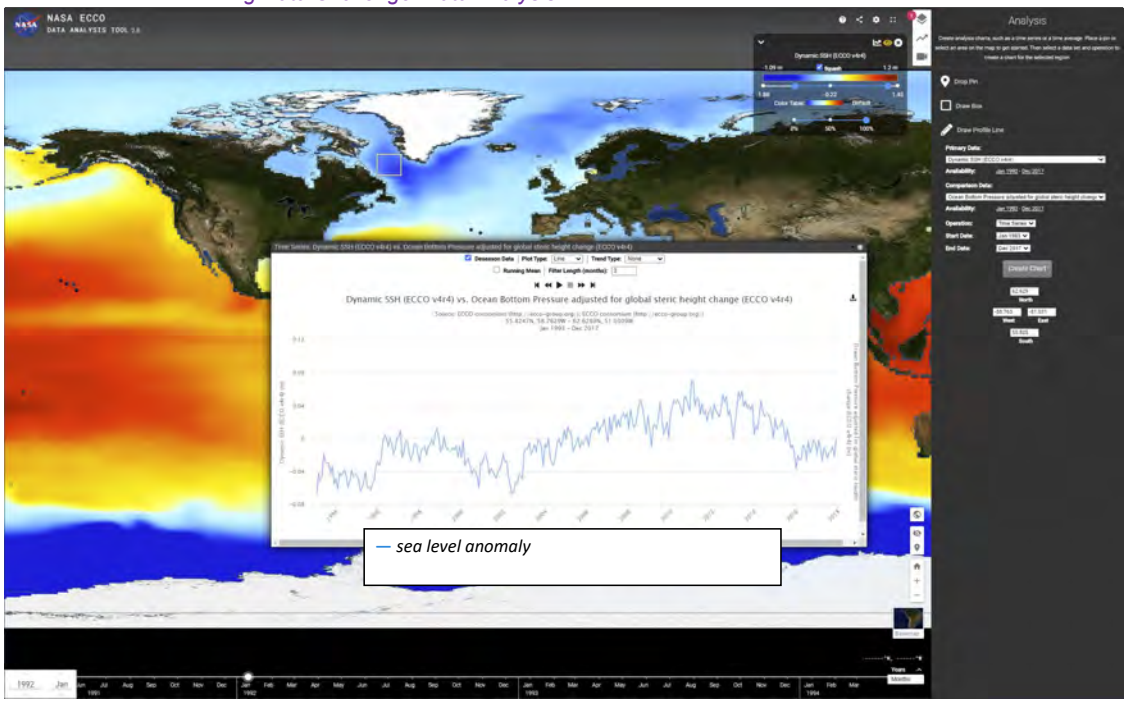

ECCO-Cloud Analytics System: **Data Analysis Tool**

Big Data Challenge: Data Analysis

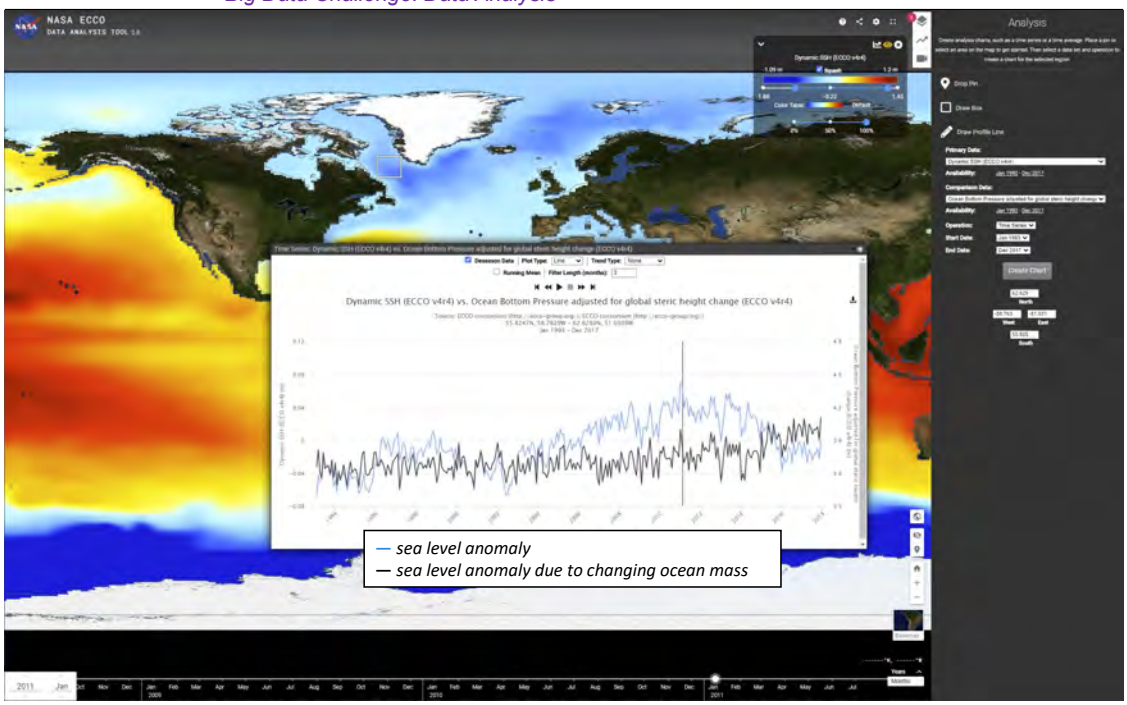

ECCO-Cloud Analytics System: **Data Analysis Tool**

Big Data Challenge: Data Analysis

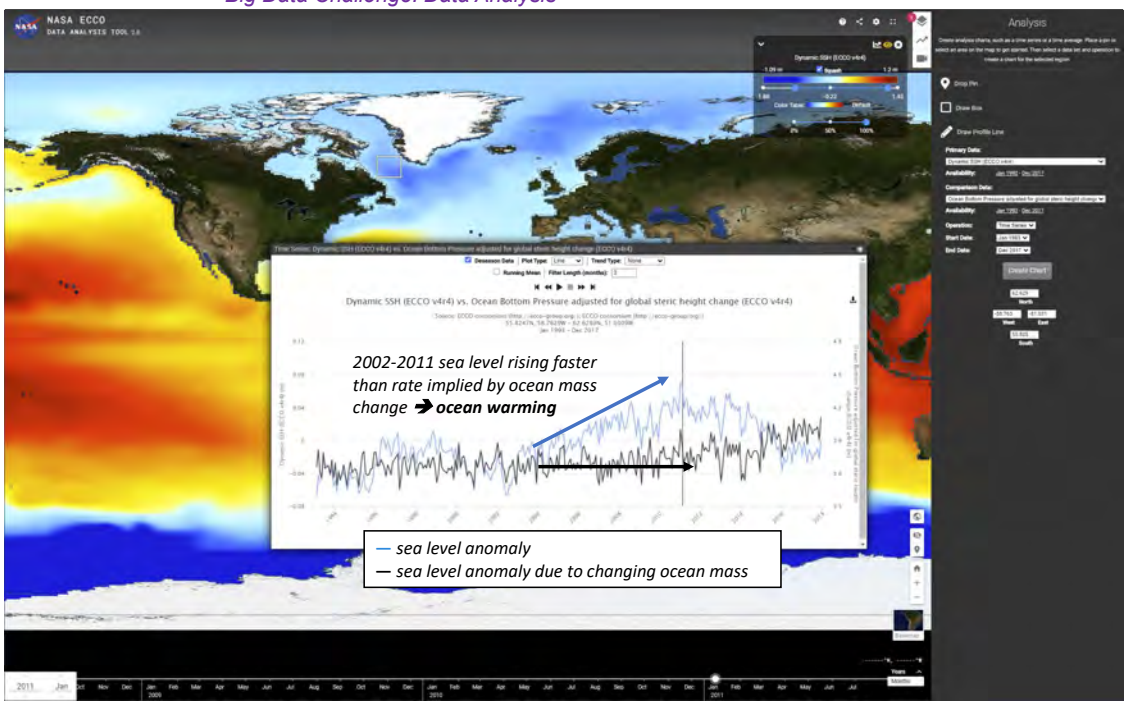

ECCO-Cloud Analytics System: **Data Analysis Tool**

Big Data Challenge: Data Analysis

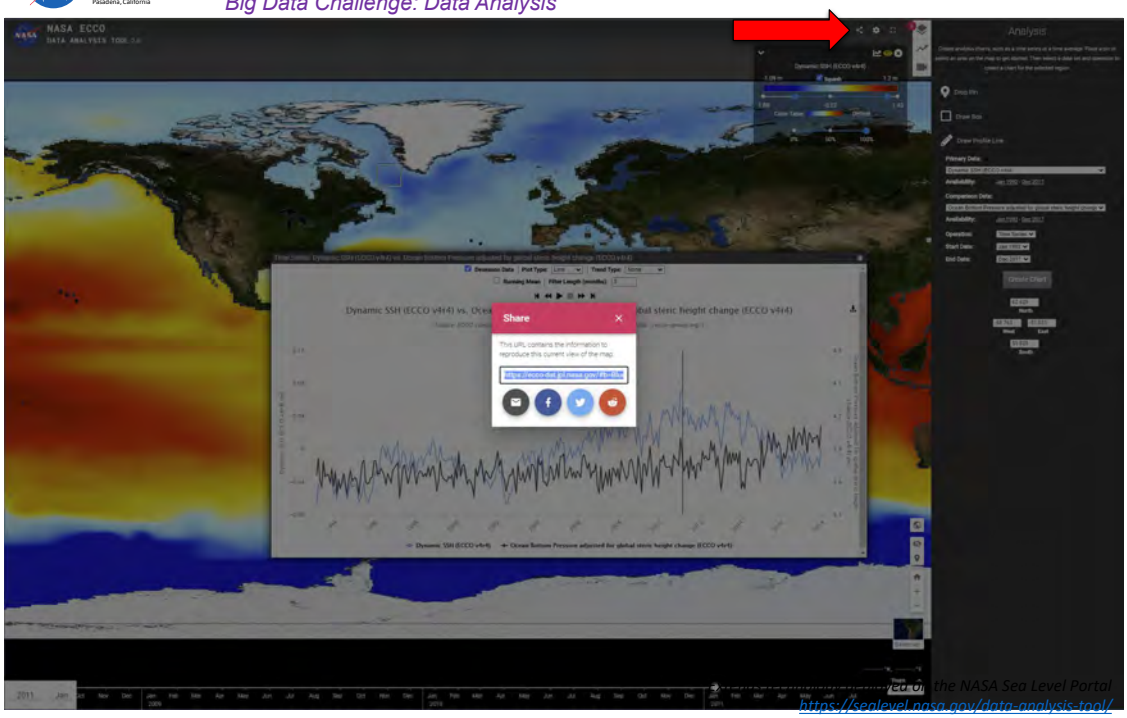

Accessible through the NASA Sea Level Portal
https://sealevel.nasa.gov/data-analysis-tool/

ECCO-Cloud Analytics System: **Python + Jupyter Notebook**

Big Data Challenge: Data Analysis

Analysis: *Cloud-based compute & store for analysis close to the data*

```
In [16]: # Compute area-averaged time series using the SDAP/NEXUS web/HTTP inte
         rface
         # Specify the SDAP name of the dataset
         dataset = "ECCO_v4_r4_SSHDYN"

         url = '{}/timeSeriesSpark?ds={}&minLon={}&minLat={}&maxLon={}&maxLat={
         }&startTime={}&endTime={}'.\
             format(ecco_cloud_base_url, dataset, *bbox.bounds,
                    start_time.strftime(dt_format), end_time.strftime(dt_format
         ))

         # Display some information about the job
         print(url); print()

         # Query SDAP to compute the time series
         print("Waiting for response from SDAP on AWS..." + dataset)
         start = time.perf_counter()
         ecco_ssh_ETP_ts_json = requests.get(url).json()
         print('Time series took {} seconds'.format(time.perf_counter
         ))

         https://ecco-cloud.jpl.nasa.gov/nexus/timeSeriesSpark?ds=EC
         SSHDYN&minLon=-125.0&minLat=-8.0&maxLon=-76.0&maxLat=8.0&st
         993-01-01T00:00:00Z&endTime=2016-12-31T00:00:00Z

         Waiting for response from SDAP on AWS...ECCO_v4_r4_SSHDYN
         Time series took 4.476107846000001 seconds
```

Plot SSH ECCO vs. MEaSUREs

```
In [24]: plt.figure(figsize=(15,6))
         plt.plot(measures_ssh_ETP_ts_result_a.time, measures_ssh_ETP_ts_result
         _a.values,color='Green',linewidth=2);plt.grid();
         plt.plot(ecco_ssh_ETP_ts_result_a.time, ecco_ssh_ETP_ts_result_a.value
         s,color='Maroon',linewidth=2);plt.grid();

         plt.ylabel(ecco_ssh_ETP_ts_result.units)
         plt.grid();
         plt.title('Sea Surface Height Anomalies\n' + box_str);
         plt.legend(('MEaSUREs','ECCOv4r4'));
```

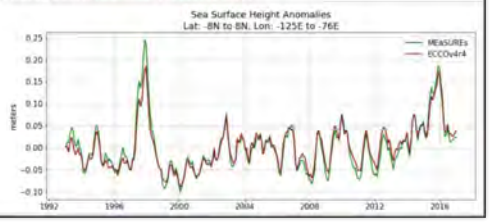

ECCO-Cloud Analytics System: **Python + Jupyter Notebook**

Big Data Challenge: Data Analysis

Analysis: *Cloud-based compute & store for analysis close to the data*

```
In [16]: # Compute area-averaged time series using the SDAP/NEXUS web/HTTP inte
         rface
         # Specify the SDAP name of the dataset
         dataset = "ECCO_v4_r4_SSHDYN"

         url = '{}/timeSeriesSpark?ds={}&minLon={}&minLat={}&maxLon={}&maxLat={
         }&startTime={}&endTime={}'.\
             format(ecco_cloud_base_url, dataset, *bbox.bounds,
                    start_time.strftime(dt_format), end_time.strftime(dt_format
         ))

         # Display some information about the job
         print(url); print()

         # Query SDAP to compute the time series
         print("Waiting for response from SDAP on AWS..." + dataset)
         start = time.perf_counter()
         ecco_ssh_ETP_ts_json = requests.get(url).json()
         print("Time series took {} seconds".format(time.perf_counter
         ))

         https://ecco-cloud.jpl.nasa.gov/nexus/timeSeriesSpark?ds=ECC
         SSHDYN&minLon=-125.0&minLat=-8.0&maxLon=-76.0&maxLat=8.0&sta
         993-01-01T00:00:00Z&endTime=2016-12-31T00:00:00Z

         Waiting for response from SDAP on AWS...ECCO_v4_r4_SSHDYN
         Time series took 4.476107846000001 seconds
```

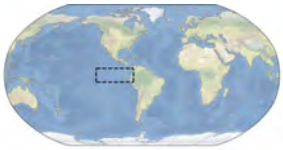

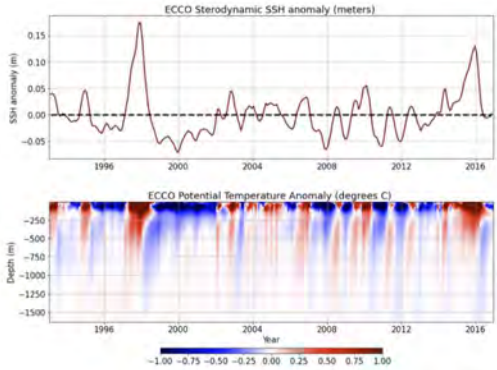

FENTY/JPL © 2020. All rights reserved.

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

ECCO-Cloud Analytics System: **Python + Jupyter Notebook**

Big Data Challenge: Data Analysis

Analysis: *Cloud-based compute & store for analysis close to the data*

```
In [16]: # Compute area-averaged time series using the SDAP/NEXUS web/HTTP inte
         rface
         # Specify the SDAP name of the dataset
         dataset = "ECCO_v4_r4_SSHDYN"

         url = '{}/timeSeriesSpark?ds={}&minLon={}&minLat={}&maxLon={}&maxLat={
         }&startTime={}&endTime={}'.\
             format(ecco_cloud_base_url, dataset, *bbox.bounds,
                    start_time.strftime(dt_format), end_time.strftime(dt_format
         ))

         # Display some information about the job
         print(url); print()

         # Query SDAP to compute the time series
         print("Waiting for response from SDAP on AWS..." + dataset)
         start = time.perf_counter()
         ecco_ssh_ETP_ts_json = requests.get(url).json()
         print("Time series took {} seconds".format(time.perf_counter
         ))

         https://ecco-cloud.jpl.nasa.gov/nexus/timeSeriesSpark?ds=ECC
         SSHDYN&minLon=-125.0&minLat=-8.0&maxLon=-76.0&maxLat=8.0&sta
         993-01-01T00:00:00Z&endTime=2016-12-31T00:00:00Z

         Waiting for response from SDAP on AWS...ECCO_v4_r4_SSHDYN
         Time series took 4.476107846000001 seconds
```

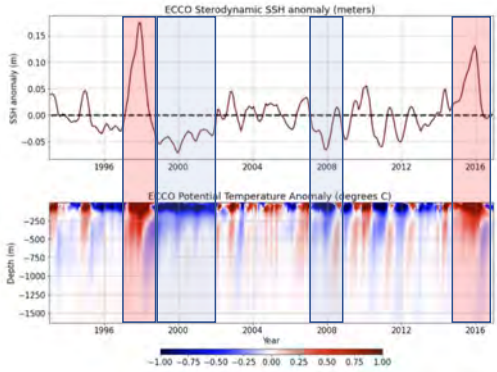

***Red**: periods of positive sea level → warm upper 1000m*
***Blue**: periods of negative sea level → cool upper 1000m*

FENTY/JPL © 2020. All rights reserved.

Open Source: Building Solutions for the Community

ECCO Serverless Workflow Subsystem
TBD License: MIT / Apache v2.0
Future Repo: https://github.com/ECCO-GROUP/

Apache Science Data Analytics Platform
License: Apache v2.0
Repo: https://sdap.apache.org

Common Mapping Client
License: Apache v2.0
Repo: https://github.com/nasa/common-mapping-client

OnEarth
License: Apache v2.0
Repo: https://github.com/nasa-gibs/onearth

MIT General Circulation Model: MITgcm
License: MIT
Repo: https://github.com/MITgcm/MITgcm

ECCOv4-py
License: MIT
Repo: https://github.com/ECCO-GROUP/ECCOv4-py

From: "Morrison, Brian A (US 1761)" <brian.a.morrison@jpl.nasa.gov>
Date: Thursday, October 3, 2019 at 2:00 PM
To: "Jacob, Joseph C (US 398M)" <joseph.c.jacob@jpl.nasa.gov>, "Huang, Thomas (US 398M)" <thomas.huang@jpl.nasa.gov>
Subject: APPROVED: Re: [EXTERNAL] [Open Source] Cloud-based Data Processing System for ECCO (NTR-51406) V1 between Caltech/JPL, Brian Morrison, THOMAS HUANG, and 2 more is Signed and Filed!

Hi Team,

**** APPROVED ****

The software known as the Cloud-based Data Processing System for ECCO (NTR-51406) V1 has been approved for open source release. You may upload the code to a known repository when you are ready to do so.

All the best,

Brian Morrison
Software Release Authority
Advanced IT Research and Open Source Program (1761)
Commercial Programs - Office of Technology Transfer (2450)
Caltech - JPL

FENTY/JPL © 2020. All rights reserved.

Outline

- ECCO Background
- Big Data Challenges
- ECCO-Cloud System
- Future Directions

FENTY/JPL © 2020. All rights reserved.

More Big Data Challenges to Tackle: ***REPRODUCIBILITY***

Computational Costs to reproduce ECCO solutions

ECCO **v4** llc90	96 CPUs	**12 hr**
ECCO **v5** llc270	787 CPUs	**36 hr**
ECCO **v6** llc1080	10821 CPUs	**28 days** (without tides)
ECCO **v6** llc1080	10821 CPUs	**75 days** (with tides)

How can we ensure reproducibility for researchers without access to large, dedicated supercomputer resources?

FENTY/JPL © 2020. All rights reserved.

More Big Data Challenges to Tackle: ***REPRODUCIBILITY***

- *ECCO is not only the first (to our knowledge) NASA modelling system to leverage cloud capabilities in terms of **production** and **analysis,** it is also the first NASA modelling system that can be **RUN in the cloud!***
  - Users can fully run the state estimate simulation on their own
  - Customize output (e.g., add variables, change output frequency, etc.)
  - Run experiments (e.g., new atmospheric forcing, mixing parameterizations, sub-models)
- Cost: ECCO v4 ~$40 on AWS virtual cluster with 6 nodes with 16 cores per node
- Recipes exist to reproduce ECCO runs on AWS
  - “Snapshots” of pre-configured machine images available shared with users.
- ***we are happy to share lesson learned with others***

Vinogradova, N.T., Shiffer, M., Forget, G., & Hill, C. (2017). Cloud-based solutions for distributed climate modeling.

FENTY/JPL © 2020. All rights reserved.

Next Steps

- Finalize the “Transformation” and “Product Distribution” workflows
- Add support for the ECCO-Cloud Analytics System to calculate on the model grid
- Integrate the “Data Analysis Tool” into the ECCO website
- ***Transition ECCO-Cloud from development to operation***
  - PO.DAAC could be the long-term home for ECCO-Cloud.
  - We need guidance from the ACCESS Program and NASA EOSDIS on how we should engage the data center.
  - PO.DAAC is interested in the ECCO-Cloud project, but their current main interest is archiving and distributing standard ECCO products
- Project officially ends in FY20, but I anticipate a few aspects of the system will require more resources to button up.
- Integrate ECCO-Cloud’s *Preprocessing and Transformation* and *Product Distribution* workflows with NASA Cumulus (a multi-year effort to develop a cloud-based framework for data ingest, archive, distribution, and management.)

FENTY/JPL © 2020. All rights reserved.

High-Level Issues

- NASA’s new restrictions on new public-facing website / webservices could affect our ability to make ECCO-Cloud analytic service available to the public
- NASA data centers should consider supporting distributing datasets on complex grids.
  - latitude-longitude or projected ‘x-y’ grids are common among data centers.
  - Some ECCO analyses require calculations be performed on the model grid (curvilinear multi-faceted grid akin to cubed-sphere or lat-lon-cap)
- Not all datasets on CMR contain links to granule files. We had to develop DAAC-specific interface to harvest relevant granule files

FENTY/JPL © 2020. All rights reserved.

ECCO Website: https://ecco-group.org

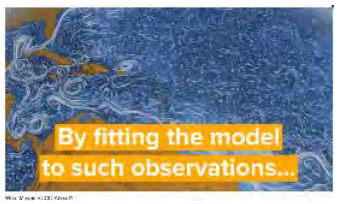

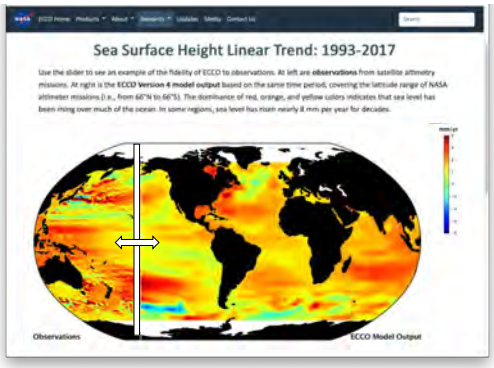

FENTY/JPL © 2020. All rights reserved.

Thank you – ECCO & Friends!

Estimating the Circulation and Climate of the Ocean
ECCO
National Aeronautics and Space Administration

Serverless Workflow Architecture

Variable	Data Product	Level
Sea Surface Temperature	GHRSST Level 4 AVHRR_OI Global Blended Sea Surface Temperature Analysis (GDS version 2)	4
Sea-ice Concentration	Near-Real-Time NOAA/NSIDC Climate Data Record of Pas. Mic. Sea-ice Conc, V1.	4
Ocean Bottom Pressure	JPL GRACE Mascon Ocean, Ice, and Hydrology Equivalent Water Height JPL RL05M.1 Version 2	3

Example: Sea Ice State Diagram

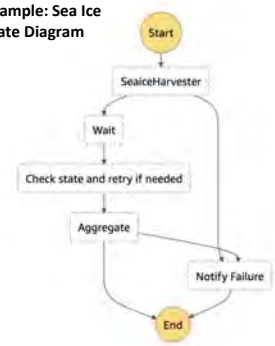

Tracking dataset and granule

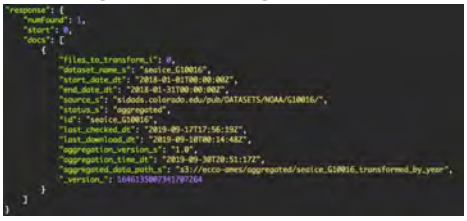

- **Pre-processing**
  - Harvesting from CMR and datacenters (e.g. NSIDC, PO.DAAC)
  - Serverless workflow using AWS Step Functions to integrate science algorithms for data transformation
  - Package and publish job to NASA Ames Pleiades
- **Post-processing**
  - Generic serverless architecture to take binary outputs from Pleiades and package into CF-compliant NetCDF products for long-term archive, visualization, and analysis

FENTY/JPL © 2020. All rights reserved.

NASA's Drive Towards Exascale Computing

Dr. Piyush Mehrotra

Chief, NASA Advanced Supercomputing (NAS) Division
NASA Ames Research Center in Silicon Valley

July 14, 2020

2018 Strategic Plan Overview

NASA's 64-page Strategic Plan is organized around 4 themes and their related Strategic Goals.

Vision

To discover and expand knowledge for the benefit of humanity.

Mission

Lead an innovative and sustainable program of exploration with commercial and international partners to enable human expansion across the Solar System and bring new knowledge and opportunities back to Earth. Support growth of the Nation's economy in space and aeronautics, increase understanding of the Universe and our place in it, work with industry to improve America's aerospace technologies, and advance American leadership.

Strategic Goals

- **DISCOVER** — EXPAND HUMAN KNOWLEDGE THROUGH NEW SCIENTIFIC DISCOVERIES.
- **EXPLORE** — EXTEND HUMAN PRESENCE DEEPER INTO SPACE AND TO THE MOON FOR SUSTAINABLE LONG-TERM EXPLORATION AND UTILIZATION.
- **DEVELOP** — ADDRESS NATIONAL CHALLENGES AND CATALYZE ECONOMIC GROWTH.
- **ENABLE** — OPTIMIZE CAPABILITIES AND OPERATIONS.

High-End Computing Capability (HECC) Project

NASA's premier Supercomputer Center
(hosted by the NAS Division at the Ames Research Center)

Charter: to meet the supercomputing needs of all NASA's Mission Directorates

Over 600 science & engineering projects with more than 1,600 users across the nation & the world

Pleiades: 7.85 PF
11,215 Xeon nodes
16 GPU nodes
245,000+ cores

Electra: 8.32 PF
3456 Xeon nodes
124,416 cores

Aitken: 3.69 PF
1152 Xeon nodes
46,080 cores

Global Storage Capacity :
Online: ~50 PB
Archival: 1 EB

NASA's Diverse HPC Requirements

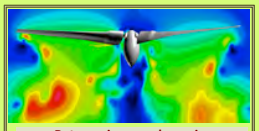

Rotary wing aerodynamics

Leadership / capability Research requires HEC resources that can handle high-fidelity long-running large-scale computations to advance theoretical understanding.

High throughput / capacity
Engineering requires HEC resources that can handle large ensembles of moderate-scale computations to efficiently explore design space.

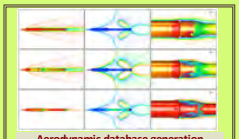

Aerodynamic database generation

High availability
Time-sensitive mission-critical applications require HPC resources readily available on demand.

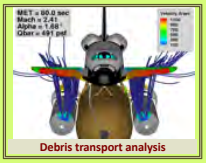

Debris transport analysis

HECC Resource Usage by NASA Mission Directorates

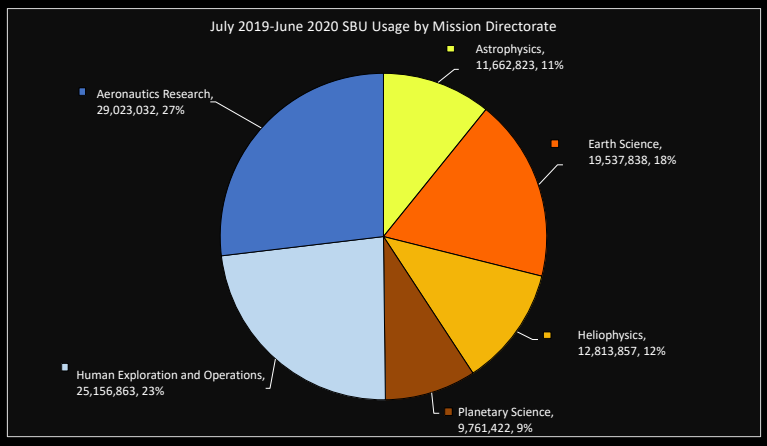

An SBU (Standard Billing Unit) represents the amount work that can be performed on a dual-socket Intel Xeon Broadwell node in 1 hour.

HECC Resource Usage by Job Size July 2018-June 2020

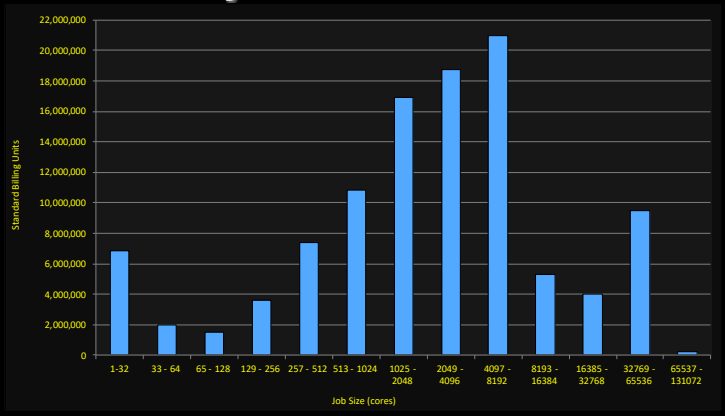

Total Usage: ~108 Million SBUs

An SBU (Standard Billing Unit) represents the amount work that can be performed on a dual-socket Intel Xeon Broadwell node in 1 hour.

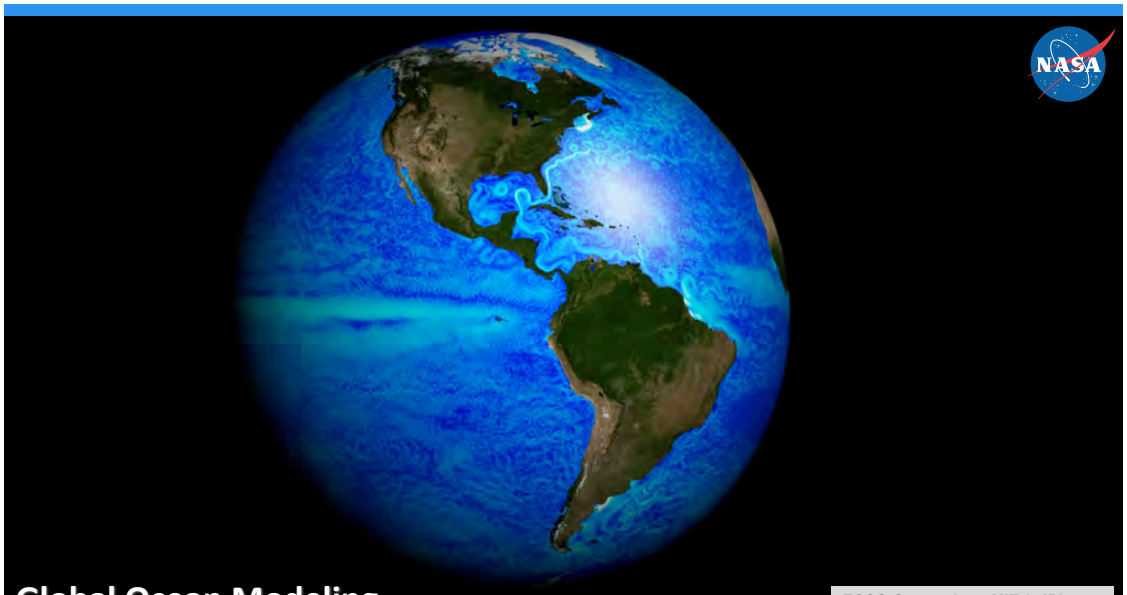

Global Ocean Modeling

ECCO Consortium: MIT & JPL

Coupled Global Atmosphere/Ocean Modeling GEOS & MITGcm

2020 Capability		by 2025		by 2030	
Resolution	*Cores/Data*	*Resolution*	*Cores/Data*	*Resolution*	*Cores/Data*
Atmosphere 6 km *137 Levels*	O(10,000) 500 TB per simulation	**Atmosphere** 3 km *181 Levels*	O(100,000) >5 PB per simulation	**Atmosphere** 1.5 km *>200 Levels*	O(1 million) >100 PB per simulation
Ocean *4 km* *90 Levels*		***Ocean*** *2 km* *100 Levels*		***Ocean*** *1 km* *150 Levels*	

Bill Putman: NASA Goddard Space Flight Center

Retropropulsion for Mars entry (FUN3D)

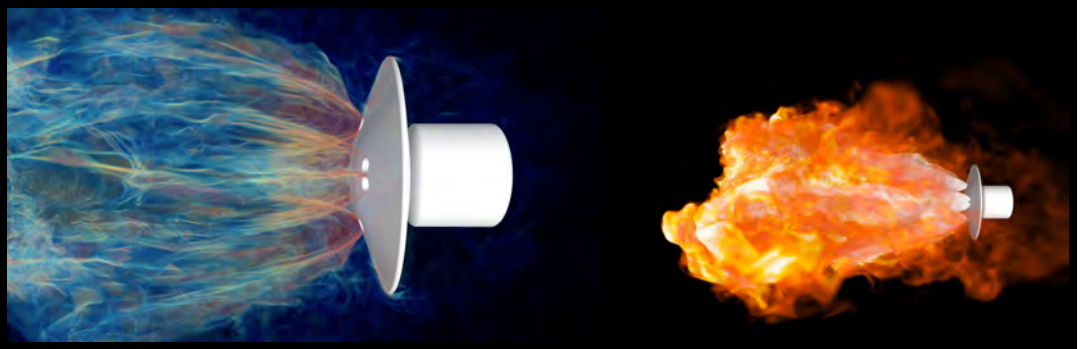

Eric Nielsen: NASA Langley Research Center

Retropropulsion for Mars Entry

2020 Capability		by 2025		by 2030	
Complexity	*GPGPUs**	*Complexity*	*GPGPUs**	*Complexity*	*GPGPUs**
Reacting gas chemistry addition	O(30,000) for 30 days	*Aero-database generation*	O(3M) for 30 days	Multiphysics at nominal and off-design conditions	O(300M) for 30 days

2019 Run required 3300 GPGPUs for 5 days – would have preferred a 30 day run to resolve the temporal statistics

* GPGPUs or equivalent processors

Eric Nielsen: NASA Langley Research Center

Launch Environment Analysis

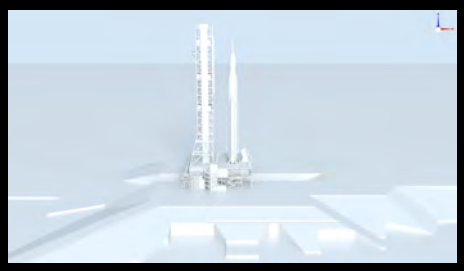

Visualization of geometry used in LAVA Cartesian simulation

Temperature cutting plane passing through an SRB centerline. Plume is clipped. Green people shown for scale.

Cetin Kiris: NASA Ames Research Center

Launch Environment Analysis

2020 Capability		by 2025		by 2030	
Resolution	*Cores/Data*	*Resolution*	*Cores/Data*	*Resolution*	*Cores/Data*
Adaptive Meshing 1 Billion Cells *2cm Δx*	O(10,000) 100 TB per simulation	**Adaptive Meshing** 10 Billion Cells *1cm Δx*	O(100,000) 1 PB per simulation	**Adaptive Meshing** 100 Billion Cells *0.5cm Δx*	O(1 million) 10 PB per simulation
Physics *Multispecies* *Unsteady* *Acoustics*		***Additional Physics*** *Multiphase* *WMLES*		***Additional Physics*** *Chemistry* *6-DOF liftoff* *Full Flight Analysis*	

Cetin Kiris: NASA Ames Research Center

CFD Vision 2030 Technology Development Roadmap

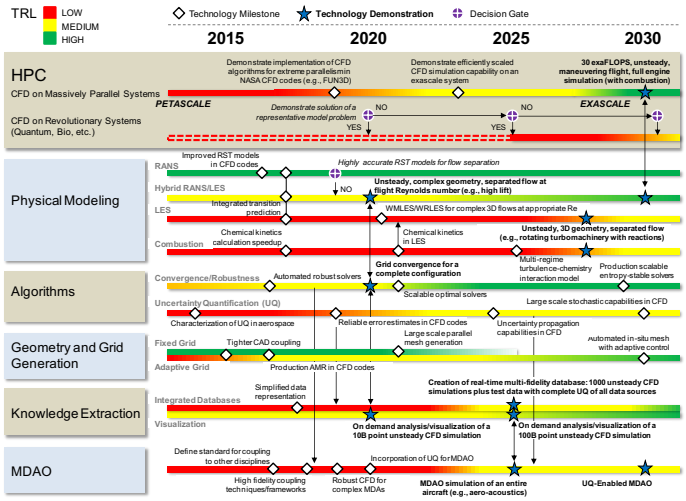

Slotnick, J., et al. "CFD vision 2030 study: a path to revolutionary computational aerosciences." NASA/CR–2014-218178 (2014).

Timeline of NAS and NASA CFD Towards Exascale Computing

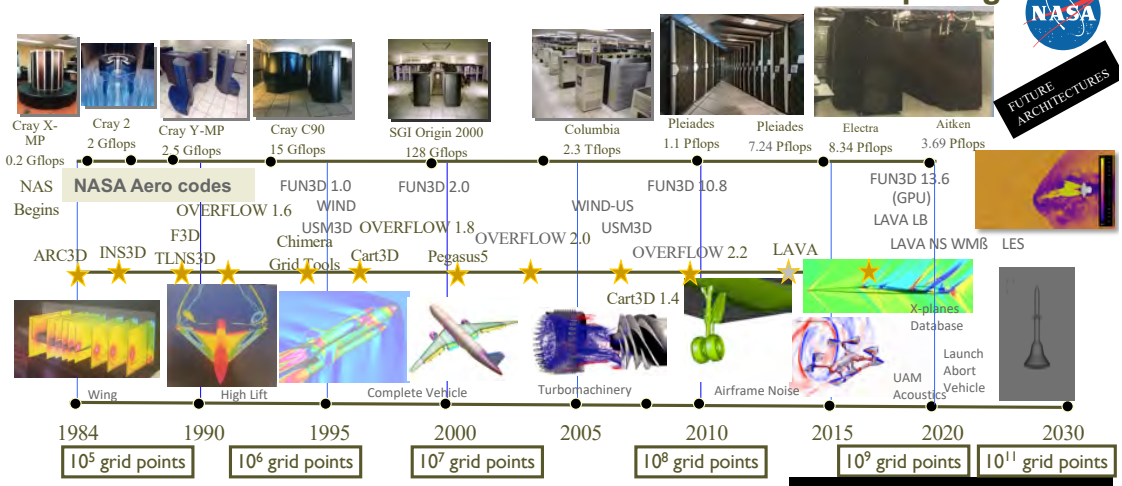

Cetin Kiris: NASA Ames Research Center

Wind tunnel runs of SLS model using pressure sensitive paint to estimate pressure/loads

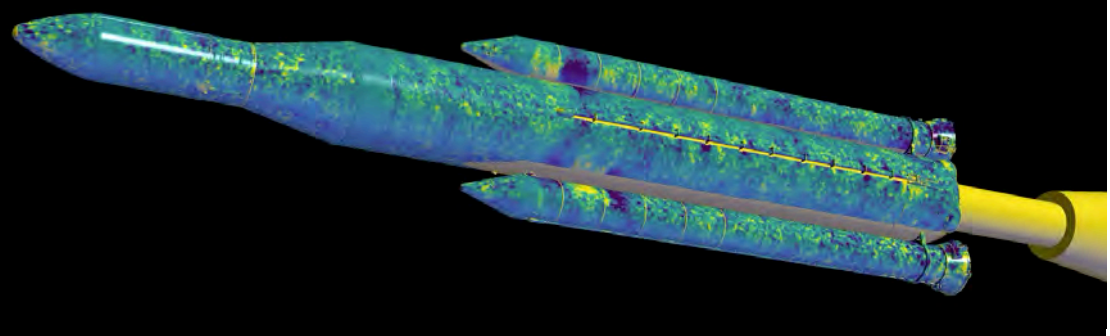

Near real-time analysis of test data using Supercomputer resources with potential for computer-guided data acquisition

Challenges to Exascale (1 of 2)

- NASA HPC budget essentially flat, resulting in a modest growth of capacity over the next several years
- Project needs to cater to throughput computing also making it difficult to support capability runs
- Limited funding for multi-disciplinary commitment to exploit current and emerging architectures
  - A few scaling efforts but mostly the focus is on science results
  - Limited focus on algorithmic development necessary for exploiting new architectures
- Resulting in:
  - Only a few petascale level codes
  - Very few codes successfully ported to accelerators e.g., GPGPUs

Challenges to Exascale (2 of 2)

Issues with utilizing DOE systems

- **Proposal-based system for allocation very competitive**
  - High-risk with a high level of effort required
  - Unpredictability implies reliance on DOE resources cannot be in the critical path of a program
- **Cannot propose ITAR work on available systems**
- **NASA programs generally not aligned with DOE objectives (except potentially green aviation) for ALCC allocation**
- **Ground rules for DOE system access require jobs at least as wide as 20% of machine**
  - Guidelines do not preclude ensemble runs, but proposers have had difficulty advocating ensembles which can increase scientific value
  - No hierarchy of mid-range systems to systematically step through
    - Chasm will be increasingly immense when Aurora/Frontier/El Capitan arrive
  - Exceedingly difficult for new/young/small development teams to catch up with leadership class

Questions?

piyush.mehrotra@nasa.gov

http://www.nas.nasa.gov/

QuakeSim: Multi-Source Synergistic Data Intensive Computing for Earth Science

PI: Andrea Donnellan, JPL

Objective

- Update QuakeSim services to integrate and rapidly fuse data from multiple sources to support comprehensive efforts in data mining, analysis, simulation, and forecasting for earthquakes
- Extend QuakeSim infrastructure to include tiered publishing mechanisms and data provenance, trust and history tracking
- Develop and deploy a cloud computing architecture to access and analyze large and heterogeneous data products and integrate them with earthquake models and simulations in collaboration with the NASA Earth Science Division - Earth Surface and Interior focus area

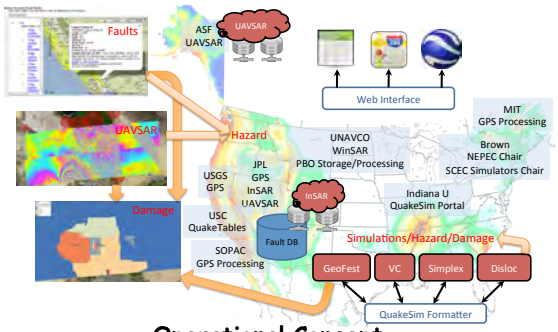

Operational Concept

Approach

- Integrate multi-source data from NASA, USGS, NSF and others through bridging services
- Support fault model optimization by integrating multiple data types in a Cloud Computing framework
- Integrate model contribution, provenance, version tracking, commenting, and rating of fault models produced by the optimization framework

Co-Is/Partners: J. Parker, R. Granat, M. Glasscoe, JPL; J. Rundle, UC Davis; L. Grant, UC Irvine; D. McLeod, USC; G. Fox, M. Pierce, Indiana Univ.; T. Tullis, Brown Univ.

Key Milestones

- Update fault models with change notification service — 12/12
- Implement data provenance, trust and history tracking — 05/13
- Complete statistical algorithms and tools for simulated and real data — 05/13
- Implement Cloud architecture leveraging NASA CMAC Cloud-Enabled Scientific Collaborative Research Environment (CESCRE) radar processing capability — 09/13
- Deploy publishing mechanism for trusted sources — 01/14
- Integrate high data rate and multi-source GPS time series analysis — 06/14
- Demonstrate pattern analysis of fused and cross-correlated data — 11/14
- Complete earthquake cloud computing system — 06/15

TRL_{in} = 2 $TRL_{current}$ = 4

ESTO Earth Science Technology Office

Appendix D9 HEC Monthly Status Report Project Highlights

ID	Mission Directorate	Division/ Program	HEC Project Title
D9.1	HEOMD	Commercial Collaboration	Simulating Dream Chaser ® Spaceplane Aerodynamics
D9.2	SMD	ESD	Simulations Probe the Sun's Effect on Climate
D9.3	SMD	ESD	Upgrade Markedly Improves Skill of NASA Seasonal Prediction System
D9.4	ARMD		Predicting High-Altitude Relight in Aircraft Engines
D9.5	SMD	ESD	NCCS Hosted Models Look at Fires on Two Continents
D9.6	ARMD		New CFD Methods for Predicting Quadcopter Drone Noise
D9.7	SMD	ESD	Realistic Simulations of the Coupled Atmosphere-Ocean-Ice System
D9.8	HEOMD		Building Booster Separation Aerodynamic Databases for Artemis II
D9.9	SMD	ESD	GMAO Visualizes Smoke Transport from Australian Fires
D9.10	SMD	ESD	Applications Performance & Productivity Team Helps Improve Ice Sheet Model Performance
D9.11	ARMD		Minimizing Sonic Boom Through Simulation-Based Design: The X-59 Airplane
D9.12	SMD	ESD	NCCS Accelerates Simulation of Atmospheric Chemistry Using Machine Learning
D9.13	ARMD		Simulating a Full-Scale, Large Airliner Landing
D9.14	SMD	ESD	NCCS Accelerates Fine-Scale Forest Modeling Across the North American Boreal Zone
D9.15 (2 pg)	SMD	Helio	Researchers Model Superflare from Sun-Like Star at NCCS
D9.16	STMD		Using Retrorockets for Human Exploration of Mars
D9.17	SMD	ESD	OCO-2 and GEOS Team Up to Produce a New View of Carbon Dioxide
D9.18	SMD	ESD	Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability
D9.19	SMD	ESD	Enabling NASA Earth Exchange (NEX) Petabyte Data Production Systems
D9.20	SMD	ESD	GMAO and NCCS Begin ORACLES Support
D9.21	SMD	ESD	Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability
D9.22	SMD	ESD	Using Katrina and Sandy Data to Improve Hurricane Prediction Tools
D9.23	SMD	ESD	Cloud Resolving Models – Cloud Library
D9.24	SMD	ESD	Projecting Sea Level Rise by Modeling the Evolution of Ice Sheets

D9.25	SMD	ESD	Using GEOS-5 Nature Run Results to Characterize Impact of Sampling Frequency on DSCOVR-Derived Energy Budget of the Earth
D9.26	SMD	ESD	KORUS-AQ Field Campaign Support
D9.27	SMD	ESD	Estimating Trees and Shrubs in the Sub-Sahara With Satellite Imagery
D9.28	SMD	ESD	GISS Simulations Show Farms Are a Major Source of Air Pollution
D9.29	SMD	ESD	High Resolution Particulate Matter Studies Are Fueled by 1-km MODIS AOD, Generated from MAIAC Algorithm on Discover
D9.30	SMD	ESD	Virtual Earth Observing via GMAO's GEOS-5 Global 1.5-km Simulation
D9.31	SMD	ESD	30-Year Global Fire Weather Analysis Based on MERRA and Ground Station Observations Sets Stage for Contributions from GPM, TRMM, and GPCP
D9.32	SMD	ESD	Estimating Woody Biomass on South Side of the Sahara at the 40-50 cm Scale Using Amazon Web Services (AWS): Progress for Intel Challenge "Head in the Clouds" Project
D9.33	SMD	ESD	March 2016 Field Campaign Support (ENRR, AfriSAR)
D9.34	SMD	ESD	Discover Supports Twice-Daily 7-Day GEOS-5 Experimental 6-km Global Forecasts
D9.35	SMD	ESD	Pleiades-Enabled Reprocessing by OCO-2 Provides a Consistent Ten-Month Data Record
D9.36	SMD	ESD	GEOS-5 Chemistry Climate Model Simulation of Stratospheric Ozone Intrusion Featured at SIGGRAPH 2015 Conference
D9.37	SMD	ESD	Modeling the Carbon Balance in Arctic Ecosystems
D9.38	SMD	ESD	Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring
D9.39	SMD	ESD	Pleiades' computational power has enabled improved understanding of methane (CH_4) fluxes
D9.40	SMD	ESD	Discover Expansion Powers Progress in Dynamical Downscaling Studies with the NASA Unified-WRF Model
D9.41	SMD	ESD	NASA Releases Detailed Global Climate Change Projections
D9.42	SMD	ESD	NASA Releases Detailed Global Climate Change Projections
D9.43	SMD	ESD	Climate Model Data Services (CDS) Support NASA's Response to Nepal Earthquake
D9.44	SMD	ESD	Real-Time Assimilation of ISS-RapidScat Observations
D9.45	SMD	ESD	Pleiades-Run Simulations Provide A Virtual Telescope on the World's Oceans and Sea Ice

D9.46	SMD	ESD	CISTO Climate Data Services Staff Cited in Federal Lab Consortium's 2015 Interagency Partnership Award for UV-CDAT
D9.47	SMD	ESD	HECC is Used to Clear 5 Years of UAVSAR Data Processing Backlog and Reduce Latency
D9.48	SMD	ESD	1.5km Global Simulation with GEOS-5 on SCU10 at NCCS
D9.49	SMD	ESD	GEOS-5 Aerosol and Chemistry Forecasts Support CalWater 2/ACAPEX Field Campaigns
D9.50	SMD	ESD	Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring
D9.51	SMD	Astrophysics	FOGGIE: Simulating the Cosmic Fog Around Galaxies

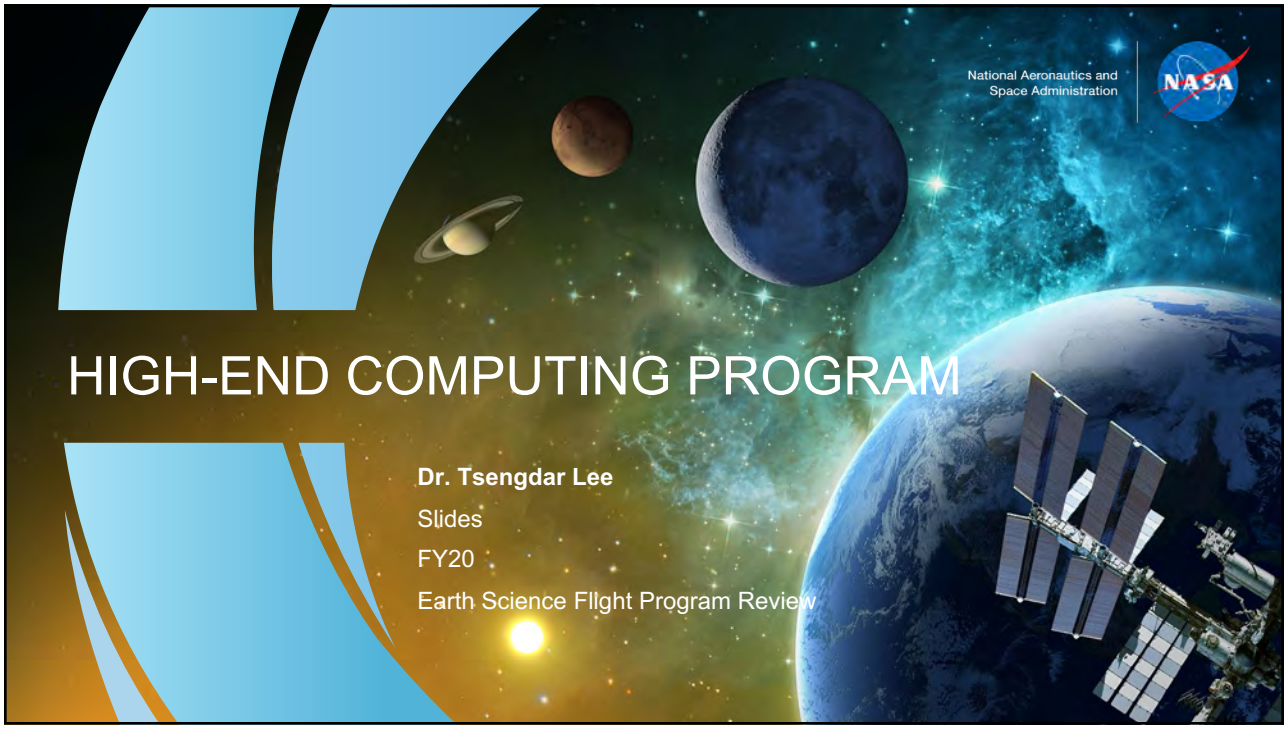

Simulating Dream Chaser® Spaceplane Aerodynamics

- Aerodynamics engineers at Sierra Nevada Corporation (SNC) ran CFD simulations on the Pleiades and Endeavour supercomputers to develop aerodynamic and aerothermal preflight databases for SNC's reusable Dream Chaser spacecraft.
  - These databases are used by SNC's Guidance, Navigation and Control for preflight predictions and control analysis and development of the Thermal Protection Systems (TPS) used to protect the vehicle during reentry.
- Using wind tunnel data as anchoring data, the SNC team developed and expanded the Dream Chaser preflight databases using thousands of aerodynamic and aerothermal predictions, and NASA-developed CFD packages, including FUN3D, CART3D, and DPLR (Data Parallel Line Relaxation).
- Access to HECC's computing resources resulted in approximately 20,000 unique high-fidelity CFD simulations to calculate the aerodynamic forces and moments acting on the vehicle, control surface interactions, and aerodynamic increments by firing Reaction Control System thrusters throughout the Entry, Descent, and Landing trajectory.

** HECC provided supercomputing resources and services in support of this work.*

IMPACT: Vital to the success of the Dream Chaser spacecraft is the accurate characterization of aerodynamic forces on the airframe and control surfaces during atmospheric flight. Improvements to spacecraft will benefit NASA's payload capability to and from the ISS.

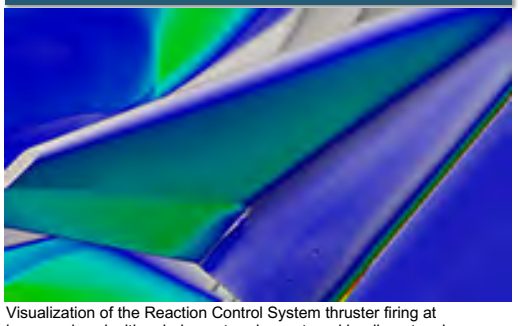

Visualization of the Reaction Control System thruster firing at hypersonic velocities during entry, descent, and landing at various angles of attack of the Dream Chaser spacecraft. As the angle of attack increases, the thruster plume impinges on the wing, changing both the pressure distribution and the vehicle's aerodynamics.
Matt Opgenorth, Sierra Nevada Corporation

Simulations Probe the Sun's Effect on Climate

- To investigate complex, remaining questions about how radiation from the Sun affects Earth's climate, Duke University and NASA GISS scientists ran century-long simulations under a variety of solar conditions at the NASA Center for Climate Simulation (NCCS)
- After a 100-year control run to bring the ocean to equilibrium, the team ran two sets of simulations on the NCCS Discover supercomputer.
- One entailed three simulations covering 110 years (10 solar cycles) and the other ran four simulations covering 160 years
- The computations revealed that long-term changes in the Sun's output cause clear effects on Earth's climate that can be seen from the surface to the upper stratosphere. In contrast, the 11-year solar cycle has clear impacts in the stratosphere but relatively weak effects on surface climate that are similar in magnitude to natural variability
- The study also showed that the tropics and high latitudes respond to solar variability at different wavelengths, especially during winter

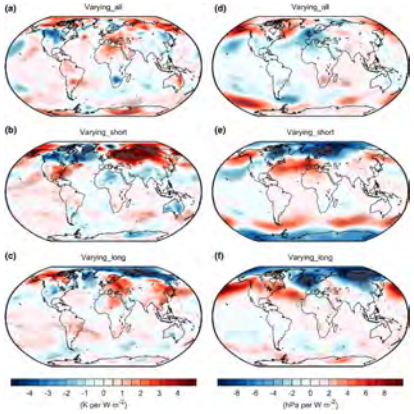

IMPACT: Changes in the Sun's output can play an important role in climate change over centuries or longer timescales but have relatively small effects on decadal and shorter timescales. This study shows how output changes at visible, near-infrared, and ultraviolet wavelengths affect our planet differently depending on latitude and season.

The panels show December–February surface temperature (left; a–c) and sea-level pressure (right; d–f) responses in three distinct fixed solar condition simulations. Dots indicate statistical significance at 90% confidence. Figure by Shindell et al.

Upgrade Markedly Improves Skill of NASA Seasonal Prediction System

- NASA's upgraded seasonal prediction system shows substantial improvement in performance, infrastructure, and forecast skill over its predecessor, per a broad set of experiments carried out at NCCS
- GEOS-S2S-2 includes a coupled Atmosphere-Ocean General Circulation Model, an Ocean Data Assimilation System, and a weakly coupled Atmosphere-Ocean Data Assimilation System.
- Compared to GEOS-S2S-1, the upgrade uses newer versions of the models, doubles atmosphere model resolution to 50 km, adds the Goddard Chemistry Aerosol Radiation and Transport model, and implements a revamped ocean data assimilation that employs an advanced Kalman Filter model ensemble and assimilates along-track sea-level data for the first time.

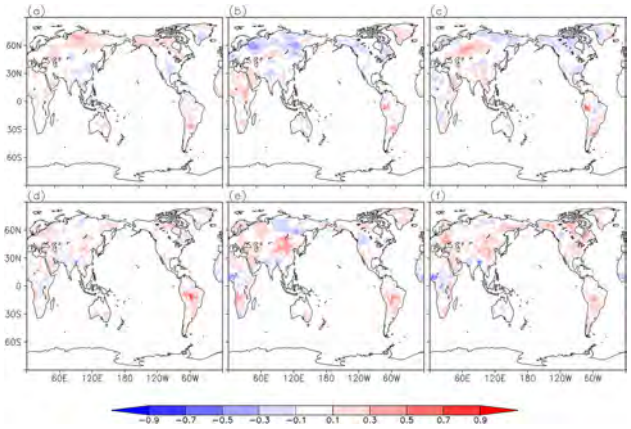

IMPACT: Improved GEOS-S2S-2 seasonal predictions raise the forecasting skill of the North American Multi-Model Ensemble and other national and international multi-model ensembles in which NASA participates. These multi-model ensembles play vital roles in tailored regional prediction and decision support by policy-makers.

GEOS-S2S-2 notably improves forecasts of near-surface temperatures, as indicated by red shading in the maps. The top row shows seasonal mean temperatures for GEOS-S2S-2 minus its predecessor GEOS-S2S-1 for (a) June-July-August at 1-month lead, (b) 4-month lead, and c) 7-month lead. The bottom row shows the same values for December-January-February.

Predicting High-Altitude Relight in Aircraft Engines

- Researchers at the University of Michigan ran simulations on Pleiades using their new HPC framework, which provides an end-to-end simulation platform to predict high-altitude relight of aircraft engines
  - Low ambient air pressures and temperatures at high altitude can lead to engine flame-out and hamper relight attempts. In high-altitude relight, an aircraft engine must ignite within a certain period after its initial flameout.
  - Although aircraft fuels are tested to evaluate their ignition characteristics at different operating conditions, it is difficult to experimentally replicate all the physical parameters that affect ignition
  - Detailed computational models are needed to provide insight into the complex relight process
- The researchers ran simulations to generate enough samples to represent statistical effects of parameters such as turbulent flow field and initial kernel energy on ignition outcome. Outputs were used to reconstruct ignition probability, using uncertainty quantification.
- Results showed that the modeling framework can efficiently generate abundant high-fidelity data of turbulent forced-ignition processes, which can be applied as input for techniques used to study patterns of altitude relight problems, benefitting future engine design.

IMPACT: The ability to re-ignite or relight an aircraft engine at high altitude—critical for safety—is an FAA certification requirement for new aircraft designs. Simulations run on HECC resources help researchers understand the critical physical processes that control ignition.

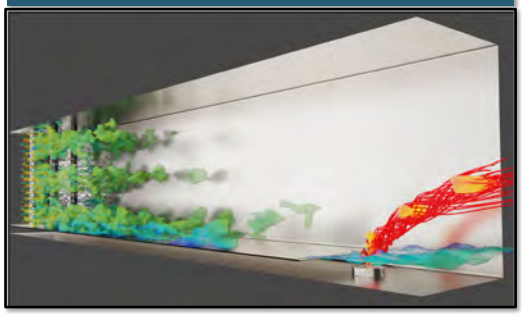

Image of a model aircraft engine combustor that represents high-altitude relight. Shown are upstream fuel mixing (isosurface of mixture fraction colored by velocity magnitude) and downstream non-premixed spark ignition (red lines tracing a spark-induced ignition kernel). *Yihao Tang, Venkat Raman, University of Michigan*

NCCS Hosted Models Look at Fires on Two Continents

- As massive fires wreaked havoc in California and Australia over recent months, computer models combined with observations were probing the dynamics and global effects of these fires
- Model forecasts supported FIREX-AQ flights collecting comprehensive, detailed measurements of trace gas, optical properties, and emissions from wildfires and agricultural fires
- FIREX-AQ daily flight planners and modeling teams had continual access to near real-time results from 12 models on the NCCS Data Portal and NCCS/NASA Global Modeling and Assimilation Office (GMAO) FLUID visualization portal
- "For FIREX-AQ, NASA supercomputing resources have played a significant role in providing computational supports and storage resources, which guaranteed timely forecasts and facilitated visualizations of results to tremendously support the science team during flight planning." — Xinxin Ye, UCLA
- Continuing GEOS-FP forecasts revealed the extent and impact of the Australian bushfires during the Dec 2019–Jan 2020 peak intensity period
- In GEOS-FP, deep vertical transport into the upper troposphere and lowermost stratosphere accelerated movement of smoke across the Southern Ocean. Smoke plumes then circumnavigated the globe back around to Australia and became particularly pronounced across the southern Pacific Ocean

IMPACT: Observations from FIREX-AQ supported by model forecasts will significantly advance our understanding of fire — including emissions, plume rise, chemical transformation, and air quality impacts.

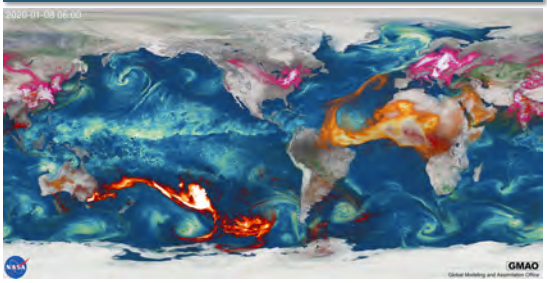

NASA's GEOS-FP data assimilation system incorporating NASA MODIS observations captures the global distribution of aerosols from the prominent bushfires in Australia. Aerosol species are dust (orange), sea-salt (blue), nitrates (pink),and carbonaceous (red); brighter regions correspond to higher aerosol amounts. Visualization by William Putman, GSFC.

New CFD Methods for Predicting Quadcopter Drone Noise

- NASA Ames engineers ran computational fluid dynamics (CFD) simulations on Electra to predict the aeroacoustic noise generated by a Straight Up Imaging quadcopter drone in forward flight using the Lattice Boltzmann Method (LBM)
  - To produce the simulations, the team used a Lattice Boltzmann solver within NASA's Launch Ascent and Vehicle Aerodynamics (LAVA) solver framework
  - The blades were represented on the automatically generated Cartesian grid through an in-house algorithm to complete the streaming step for lattice links broken by the geometry
- Coupled with far-field acoustic propagation in the LAVA framework, LBM successfully predicted tonal noise levels and captured broadband noise trends accurately for the first time
  - In detailed comparisons with test data, the predicted mean thrust was within 5% of the value measured experimentally
  - No manual mesh generation was required; the model used adaptive mesh refinement based on running statistics of pressure fluctuations and turbulent kinetic energy
  - 30-rotor rotations were simulated with 568 million cells in under 125 hours
- The solver's computational efficiency—and the absence of labor-intensive manual mesh generation—are key to making routine aeroacoustic analysis of urban air taxis and drones from first principles possible.

IMPACT: Up to ten times faster than traditional approaches to analyzing quadcopter aeroacoustics, new CFD methods—enabled by the Electra supercomputer—help design engineers address noise, which is a major roadblock to community acceptance of autonomous air taxis and drones.

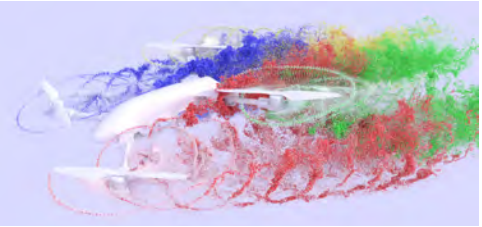

Passive particles tracing the complex flow structures generated by four sets of blades spinning at up to 4,937 RPM, in this simulation of a Straight Up Imaging quadcopter in forward flight. *Francois Cadieux, Timothy Sandstrom, NASA/Ames*

* HECC provided supercomputing resources and services in support of this work.

Realistic Simulations of the Coupled Atmosphere-Ocean-Ice System

- Earth science researchers are successfully running groundbreaking coupled ocean-atmosphere simulations, using two flagship data-assimilating NASA models, in order to accurately represent and understand coupled air-sea exchange processes
  - The Goddard Earth Observing System (GEOS), which is used for global cloud-resolving simulations with horizontal grid spacing as small as 1.5 kilometers (km)
  - The Estimating the Circulation and Climate of the Ocean (ECCO) model, which is used for global-ocean, submesoscale and internal-wave-admitting simulations with horizontal grid spacing as small as 1–2 km
- Initial analysis of the coupled-model output revealed three-to-six-day oscillations of sea surface temperatures and surface wind anomalies—a phenomenon appearing in observational records and reanalyses, but previously unseen in ocean-only simulations
- These simulations are also used to guide the utilization of existing satellite observations and development of new instruments
- The simulations are available to the scientific community on NAS and NCCS data portals

* HECC provided supercomputing resources and services in support of this work.

IMPACT: NASA supercomputing resources are revolutionizing Earth science studies by enabling increasingly realistic simulations of the coupled ice-ocean-atmosphere system that are shared with and used by the scientific community.

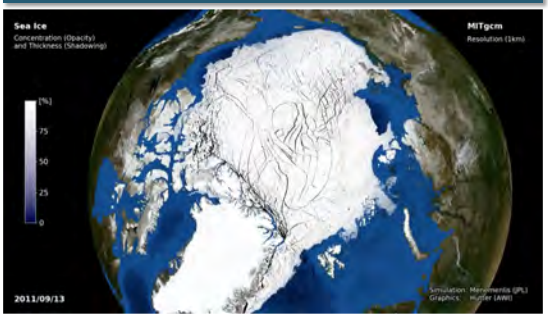

Arctic sea ice concentration & thickness from an internal-wave-admitting Estimating the Circulation and Climate of the Ocean (ECCO) simulation. The small-scale deformation of the ice causes the sea ice cover to be divided into several floes separated by strips of open ocean. Dimitris Menemenlis, NASA/JPL; Nils Hutter, Alfred-Wegener Institute

Building Booster Separation Aerodynamic Databases for Artemis II

- Aerospace engineers at NASA Ames ran computational fluid dynamics (CFD) simulations on the Pleiades and Electra supercomputers to formulate the booster separation databases for NASA's Artemis II flight configuration
- Created using NASA's FUN3D flow solver, the calculations include the aerodynamic effects of 22 different plumes during the booster separation event and 13 independent variables. Additionally, separate sets of computations were run to simulate possible effects caused by a core-stage engine failure
- These databases will be used by the Guidance, Navigation, and Control group at NASA Marshall to create dynamic simulations of the booster separation event for Artemis II to ensure the boosters can separate successfully without reconnecting to the core under all possible flight conditions
- The CFD simulations that construct these databases allow engineers to model flight conditions and complex interactions that are difficult to test in an experimental setting

IMPACT: The results of these simulations, run on HECC resources, will be critical in reducing the risk to the crew of the Artemis II missions during booster separation.

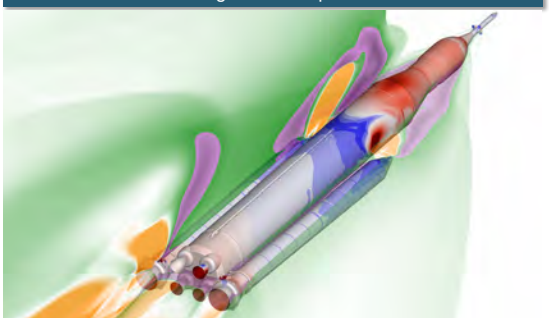

Isometric view of the Artemis II vehicle, simulating the effect of a failure in a core stage engine with the boosters four feet downstream from their original, attached position. The vehicle surface is colored by pressure contours, where blue is low and red is high. The green and orange colors represent low and high Mach numbers, respectively. *Stuart Rogers, Jamie Meeroff, NASA/Ames*

* HECC provided supercomputing resources and services in support of this work.

GMAO Visualizes Smoke Transport from Australian Fires

- The local impacts of the Australian bushfires have been devastating to property and life in the country, while producing extreme air quality impacts throughout the region
- As smoke from the massive fires has interacted with global weather, the transport of smoke plumes around the globe has accelerated into the upper troposphere and even the lowermost stratosphere, leading to long-range transport around the globe
- The smoke from these bushfires traveled across the Southern Ocean, completing a global circumnavigation back around to Australia, and is particularly pronounced across the southern Pacific Ocean out to South America
- This visualization, created from NASA's GEOS-FP data assimilation system uses HEC resources at Goddard's NASA Center for Climate Simulation (NCCS)
- Global observations of aerosol optical depth (AOD) from the NASA MODIS instruments on the Terra/Aqua satellites were assimilated to constrain aerosols in the GEOS-FP system
- Active fires, also detected by the MODIS instruments, are used in a fire-emissions module that is incorporated in GEOS-FP to constrain the carbon aerosol emissions

IMPACT: NCCS resources enable the Global Modeling and Assimilation Office (GMAO) to generate products like this one

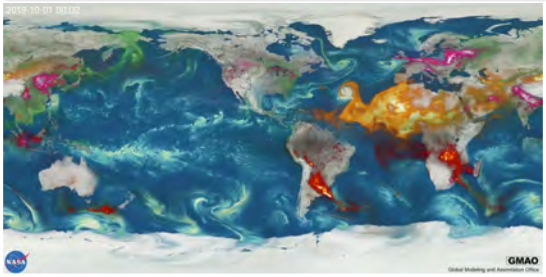

NASA's GEOS-FP data assimilation system captures the global distribution of aerosols. Different aerosol species are highlighted by color, including dust (orange), sea-salt (blue), nitrates (pink) and carbonaceous (red), with brighter regions corresponding to higher aerosol amounts. NASA's MODIS observations constrain regions with biomass burning as well as the aerosol optical depths in GEOS, capturing the prominent bushfires in Australia and transport of emitted aerosols well downstream over the South Pacific Ocean.

* NCCS provided supercomputing resources and services in support of this work.

Applications Performance & Productivity Team Helps Improve Ice Sheet Model Performance

- The Ice Sheet System Model (ISSM) is a new-generation thermo-mechanical ice sheet model to improve projections of sea level rise over the coming centuries
- The ISSM team relied heavily on HECC resources and expertise to perform different simulations for both Greenland and Antarctica, as part of the Ice Sheet Model Intercomparison Project within the Coupled Model Intercomparison Project – phase 6 (CMIP6)
- The Applications Performance & Productivity (APP) experts confirmed the impact of those modifications with a new round of profiling
- The collaboration went through many rounds and resulted in an **11.9x speedup**
  - The primary benefit came from reduced use of dynamic memory allocation and deallocation
  - Improved compiler optimizations and load balancing also contributed significantly

IMPACT: Software performance improvements achieved through collaboration between HECC applications experts and code developers allows more simulation scenarios to be included in the ISSM ensemble statistics. This improves uncertainty bounds for sea level rise.

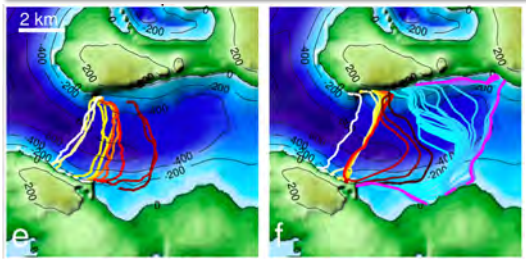

Observed (left) and modeled (right) ice front position for Upernavik Isstrøm C under current conditions. Warm colors are for 2007 to 2017 and cold colors are the model projections for 2017 to 2100. Mathieu Morlighem, University of California, Irvine

Minimizing Sonic Boom Through Simulation-Based Design: The X-59 Airplane*

- Aerospace engineers at NASA Ames ran high-fidelity computational fluid dynamics (CFD) simulations on Pleiades, Electra, and Endeavour to help shape the design of NASA's upcoming X-59 Quiet SuperSonic Technology X-plane
- NASA's production-level Cart3D simulation package for CFD was used to determine the pressure field near the aircraft and evaluate the ground noise carpet of each major design evolution of the X-59
- Cart3D was coupled with an atmospheric propagation solver to estimate the noise level on the ground; and with uncertainty quantification tools to provide uncertainty estimates in the pressure signatures due to variations in the aircraft's operating conditions and configuration
- The simulations contributed to many design improvements such as reducing the noise generated by the nose of the aircraft, instrumentation probes, and secondary-air-systems inlets. Cart3D is also used to support supersonic wind tunnel tests

**HECC provided supercomputing resources and services in support of this work.*

IMPACT: These simulations, run on HECC resources, support one of NASA's six Strategic Thrusts for the agency's Aeronautics Research Mission Directorate: "Innovation in Commercial Supersonic Aircraft."

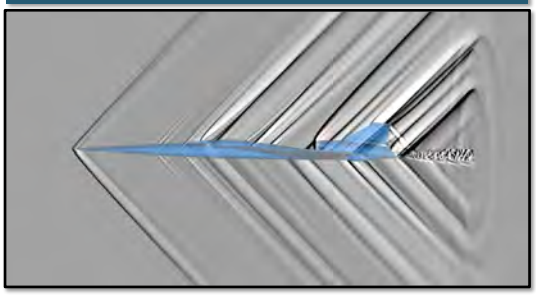

Video from a Cart3D simulation showing the complex shock system of NASA's X-59. Dark and bright regions represent shockwaves and expansions, respectively. Weaker shocks propagate from the lower surface of the aircraft, quieting sonic booms to sonic thumps on the ground. Marian Nemec, Michael Aftosmis, NASA/Ames

NCCS Accelerates Simulation of Atmospheric Chemistry Using Machine Learning

- To understand the origins, quality, and extent of harmful air, NASA scientists are combining satellite observations and computer simulations
- Researchers employed the NCCS Discover supercomputer and Advanced Data Analytics Platform (ADAPT) to simulate atmospheric chemistry and train a machine learning (ML) model
- This new method simulates atmospheric chemistry using ML with the potential to be **at least two times faster** than the conventional method
- Researchers trained the ML model using data from the NASA GEOS model with full chemistry
  - Once trained, the ML model can be inserted into existing air quality models to **provide global air quality predictions at half the computational cost**

"Thanks to advances in computing, we can now simulate global air pollution at unprecedented scale. Machine learning offers exciting new opportunities to simulate air pollution even faster and in ever greater detail." - Christoph Keller, USRA/GMAO

IMPACT: ML is a promising tool to simulate atmospheric chemistry. The ML-enhanced GEOS model has several applications, including creating ensemble air quality predictions and better combining NASA observations and atmospheric models.

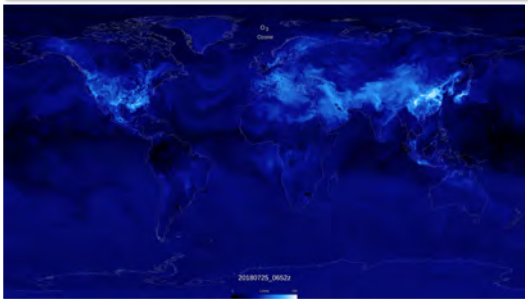

This visualization of a GEOS atmospheric simulation shows 96 chemical species. Visualization by Greg Shirah (lead) and Trent Schindler, NASA's Scientific Visualization Studio.

Simulating a Full-Scale, Large Airliner Landing

- NASA Langley scientists ran simulations on Pleiades to compute the far-field noise signature of a full-scale, large civil transport airliner in landing configuration - airframe noise contributes significantly to the total noise generated by aircraft during approach and landing
  - Simulations were run on low-res meshes to establish best practices for attaining the numerical, spatial, and temporal resolution required to capture the key noise-producing flow features
  - In partnership with Boeing, scientists used Dassault System's Lattice Boltzmann code PowerFLOW to investigate the complex, unsteady flow field around the airliner to better understand the prominent noise-generation mechanisms and locate major noise sources
  - Validated predicted results against measured acoustic data obtained from flight tests of the same aircraft
- Results from preliminary comparisons of the noise spectrum on the ground with flight test data indicate there will be good agreement. This bodes well for the suitability of their current simulation approach to solving this airframe noise grand challenge for the aerospace community.
- Scientists will next perform the simulations on a finer spatial grid to capture the high-frequency segment of the far-field noise spectrum

IMPACT: Simulation- based prediction of airframe noise is essential for the design of effective noise-reduction strategies that can substantially improve the quality of life in communities near major airports

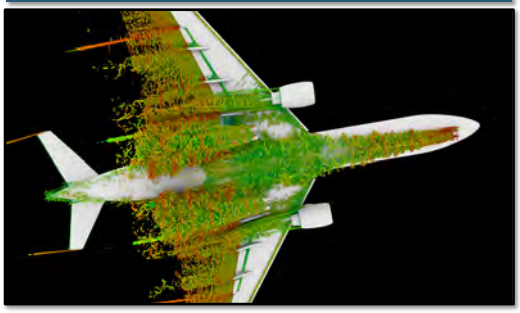

Visualization of the simulated flow field around a large airliner. Of note: complex, vortical, unsteady flow features generated by the wing high-lift devices and aircraft landing gear. Benedikt Koenig, Dassault Systemes; Patrick Moran, NASA/Ames

**HECC provided supercomputing resources and services in support of this work.*

NCCS Accelerates Fine-Scale Forest Modeling Across the North American Boreal Zone

- Scientists updated an individual-tree-based forest model - the University of Virginia Forest Model Enhanced (UVAFME) - for the boreal region
- Research used NCCS Advanced Data Analytics Platform (ADAPT) and ADAPT's Arctic Boreal Vulnerability Experiment (ABoVE) Science Cloud
- The updated model was applied across the Tanana River Basin
- UVAFME simulated forest dynamics under two climate change scenarios: Representative Concentration Pathway (RCP) 4.5, a more conservative prediction for climate change; and RCP 8.5, a more extreme prediction.
- For both scenarios, UVAFME predicted overall biomass decline, with some areas of increasing biomass in historically very cold, moist areas
- UVAFME also predicted increasing deciduous forest cover across the region

IMPACT: Increase understanding of the potential futures of the Alaskan boreal forests. Demonstrates the importance of fine-scale models like UVAFME for simulating future forest states.

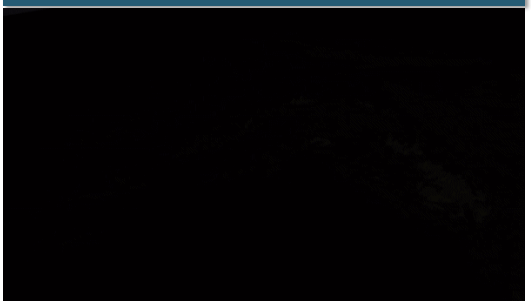

Put in presentation mode to view animation. *This animation of UVAFME model output initially shows total biomass across all interior AK. Dark greens represent high biomass, and yellows represent areas with low biomass. The movie then "zooms in" to a single site and shows how the model simulates individual tree growth, the response to wildfire, and re-establishment following fire.*

Researchers Model Superflare from Sun-Like Star at NCCS: 1 of 2

- Employing NASA Center for Climate Simulation (NCCS) resources, university and NASA researchers modeled a gigantic superflare and coronal mass ejection from Kappa Ceti—a Sun-like star 29 light-years away from Earth.
- The simulated explosion released ~700,000 times the energy used by all humans on Earth over an entire year, an event as powerful as our Sun's famous 1859 Carrington storm. A solar storm of this magnitude today would widely damage communications and electrical power infrastructure, with $40 billion in daily economic losses in the U.S. alone. Astoundingly, observations show superflares 10 times more powerful erupting from Kappa Ceti.
- The computations used the ARMS 3D magnetohydrodynamics code developed by NASA Goddard Space Flight Center's Rick DeVore, Spiro Antiochos, and collaborators.
- The published ARMS simulation ran on 256 cores of the NCCS Discover supercomputer for ~200 hours. It produced 300 gigabytes of data later moved to NCCS local storage for analysis and visualization, yielding another 300 gigabytes of data.

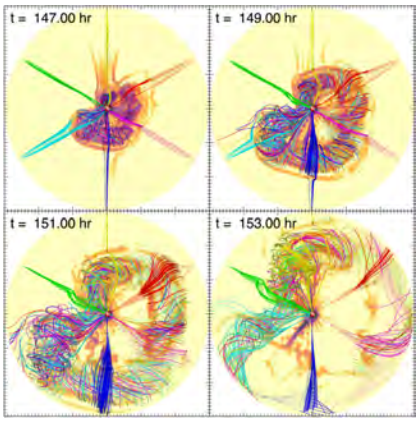

Evolving representative magnetic field lines during eruption of a superflare and coronal mass ejection from the star Kappa Ceti. The magnetically closed stellar atmosphere is sufficiently energized that the sheared and twisted fields become unstable and erupt, creating a massive 3D magnetic flux rope ejection encircling the entire star. Viewpoint is above the north pole of the star (orange sphere at image center).

IMPACT: Kappa Ceti is the same size and mass as our Sun but only 700 million years old, so studying it shows what solar activity was like early in our solar system's history. Moreover, understanding the origin and character of stellar storms and space weather affecting exoplanets guides our search for extraterrestrial life.

Researchers Model Superflare from Sun-Like Star at NCCS: 2 of 2

- These results show that global-scale shear can energize the closed magnetic field lines of the stellar corona (atmosphere) sufficiently to power an eruptive superflare from Kappa Ceti releasing as much energy as the massive Carrington event.
- The global-scale eruption represents the most extreme stellar space weather event possible within the observed constraints on the surface magnetic flux and sets a baseline for comparison with observations of superflares on Kappa Ceti and similar stars.
- "From a theorist and modeler's perspective, the availability of high-performance computing infrastructure like NCCS Discover is absolutely necessary for advancing our first-principles understanding of complex, multi-scale physical systems such as solar and stellar atmospheres. These resources dramatically increase our capacity for space weather forecasting through the ability to model the dynamic, time-dependent evolution of our own Sun-to-Earth system and to facilitate the application of these tools to more exotic astrophysical environments."
  — Ben Lynch, Space Sciences Laboratory, University of California, Berkeley

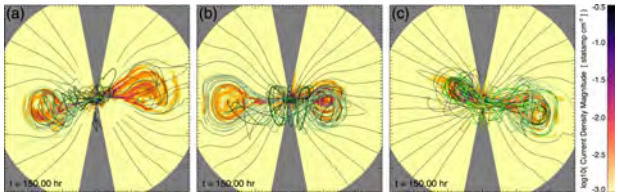

Magnetic field lines and current density magnitude (color shading) in the 3D magnetic flux rope ejection encircling Kappa Ceti. Viewpoints are from three different directions in the equatorial plane of the star.

More Information:

- Lynch, B.J., V.S. Airapetian, C.R. DeVore, M.D. Kazachenko, T. Lüftinger, O. Kochukhov, L. Rosén, and W.P. Abbett, 2019: Modeling a Carrington-scale Stellar Superflare and Coronal Mass Ejection from Kappa Ceti. *Astrophysical Journal,* **880**, 97, doi:10.3847/1538-4357/ab287e.
- Kerry Hensley, "Launching a Stellar Superflare," *AAS Nova*, 8/16/19.

Using Retrorockets for Human Exploration of Mars*

- The Descent System Studies team at NASA Ames is running computational fluid dynamics (CFD) simulations on Pleiades to study retropropulsion—a method for Entry, Descent, and Landing (EDL) that may be used in human missions to Mars.
- Aerospace engineers are considering using retropropulsion for these missions because EDL methods used in previous Mars missions are not scalable for delivering the massive payloads necessary for humans to survive on Mars.
- The simulation results help the researchers study the complex interactions between rocket plumes and incoming flow, providing important insights for analyzing the feasibility of design concepts for future Mars missions.
- CFD provides a valuable tool for analyzing different concept vehicles quickly and at relatively low costs. Engineers can change design parameters parameters to study the controllability and stability of a spacecraft using retropropulsion for powered descent and landing.

IMPACT: The knowledge and insights gained from these simulations are being applied to refine the design and architecture of future human exploration missions to Mars.

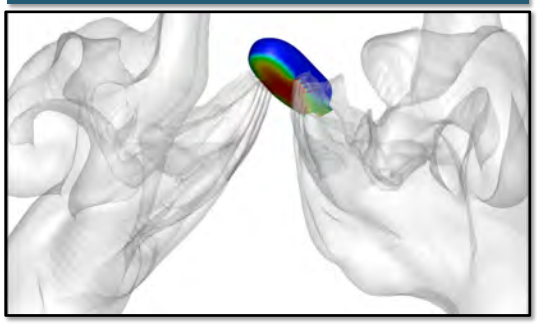

Visualization of the flow around a spacecraft with a mid lift-to-drag-ratio, with eight rocket engines firing to decelerate the vehicle. Colors indicate pressure (red is high, blue is low). The simulation illustrates how the expanding rocket plumes interact with the supersonic freestream, changing the surface pressure. *Chun Tang, NASA/Ames*

** HECC provided supercomputing resources and services in support of this work.*

OCO-2 and GEOS Team Up to Produce a New View of Carbon Dioxide

Combining observations from NASA's Orbiting Carbon Observatory 2 (OCO-2) satellite with information from the GEOS modeling system, GMAO scientists have created a new data-driven view of atmospheric carbon dioxide.

OCO-2 is NASA's first mission dedicated to observing atmospheric carbon dioxide from space, providing approximately 100,000 observations of column-averaged mixing ratio per day.

Assimilating OCO-2 data enables the GMAO to extend GEOS capabilities to studying greenhouse gases and air quality.

A 3D visualization by the SVS reveals in startling detail the complex patterns in which carbon dioxide in the atmosphere increases, decreases, and moves around the globe from September 2014 to September 2015.

The simulation ran on 336 cores of the NCCS Discover supercomputer for roughly 22 days.

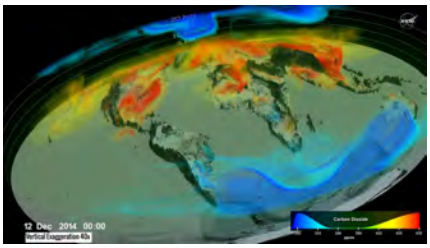

In this new global view of carbon dioxide provided by OCO-2 satellite data and the GEOS model, the height of Earth's atmosphere and topography have been vertically exaggerated and appear approximately 40 times higher than normal to show the complexity of the atmospheric flow.

Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability

Susan Strahan (Code 614, Atmospheric Chemistry and Dynamics Laboratory) and colleagues used the NCCS Discover system to investigate behavior of total column ozone in Arctic winter and spring during years with and without Sudden Stratospheric Warming (SSW) events.

The study used the Global Modeling Initiative (GMI) chemistry transport model, employing the Modern-Era Retrospective Analysis for Research and Applications (MERRA) meteorological fields, also generated on Discover, and compared results with 2005 – 2015 observations from Aura's Microwave Limb Sounder (MLS) ozone and N2O data to show GMI credibly represents transport processes and the net heterogeneous chemical loss of O3.

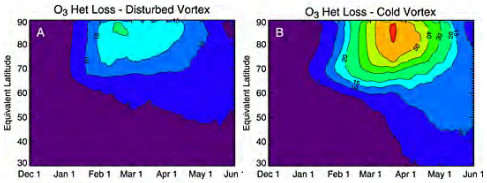

Average column ozone depletion (Dobson Units) for (top) 6 years with Sudden Stratospheric Warmings (SSWs) and (bottom) 5 years with no major or minor SSW before mid-February. Maximum ozone depletion in years without a midwinter warming is roughly 3 times greater than years with a warming.

- The investigation showed that winters with SSW events (where the Arctic stratosphere is warmer, the polar vortex is more disturbed, and more atmospheric meridional transport occurs) have less than half the ozone depletion of years without one.
- The analysis showed that 2/3 of the observed large variability of Arctic total column ozone each spring resulted from interannual dynamical variability – i.e., dynamic transport (resupply or reduction) by atmospheric waves or eddies – while only 1/3 was due to ozone loss by ozone depleting substances.

Impact: *Studies such as this one demonstrate the value of the Discover cluster in providing the environment for simulations and analyses that enable researchers to investigate important phenomena such as the attribution of behavior of Arctic ozone variability .*

Enabling NASA Earth Exchange (NEX) Petabyte Data Production Systems *

- To support a growing number of projects that are developing new, large-scale science data products, the NEX team developed a hybrid HPC-cloud data production system during the last year.
- The NEX system represents both the production rules and the execution of the rules as a directed graph, from which user can analyze and reproduce scientific results.
- The system makes it possible for small science teams to execute projects that would not have been possible previously. Examples include:
  - Production pipeline for Landsat and NASA Global Imagery and Browse Service – 37 steps, 7 petabytes of data, and 670,000 processor-hours.
  - NEX Downscaled Climate Projections (NEX-DCP30) and NEX Global Daily Downscaled Projections (NEX-GDDP) pipeline – 40 terabytes (TB) of data and 320,000 processor-hours.
  - Estimating biomass by counting trees across the United States by using machine learning and computer vision – 40 TB of data/100,000 processor-hours.
- Over the last year, NEX projects used more than 9 million processor-hours on Pleiades and processed more than 3 petabytes of data.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: The Pleiades supercomputer, combined with HECC's massive data storage and high-speed networks, enables the NASA Earth Exchange (NEX) team to engage large scientific communities and provide large-scale modeling and data analysis capabilities not previously available to most scientists.

NEX production pipeline example: Preliminary national forest disturbance map from the North American Forest Dynamics (NAFD) team, developed using 25 years of Landsat data (1985–2010) and over 200,000 compute hours. Product is a wall-to-wall forest disturbance map for the entire U.S. at 30-meter spatial resolution.

GMAO and NCCS Begin ORACLES Support

ObseRvations of Aerosols above CLouds and their intEractionS (ORACLES) is a 5-year investigation designed to study key processes that determine the climate impacts of African biomass burning aerosols. It will provide airborne observations over the complete vertical column of the key parameters that drive aerosol-cloud interactions in the Southeast Atlantic, an area with some of the largest inter-model differences in aerosol forcing assessments on the planet.

The first of three Intensive Observation Periods began August 30, 2016. Subsequent month-long deployments are scheduled for 2017 and 2018.

ORACLES instruments sit along the top of a NASA P-3B aircraft. Photo by Samuel LeBlanc.

- The GMAO and NCCS are providing GEOS-5 model forecasting support on the Discover supercomputer to aid in aircraft mission planning for ORACLES. Data access is provided through web-based maps, FTP, and OPenDAP.
- The GMAO's Arlindo da Silva is an ORACLES Science Team member and serves as a liaison.

Chemical and Dynamical Impacts of Stratospheric Sudden Warmings on Arctic Ozone Variability

Susan Strahan (Code 614, Atmospheric Chemistry and Dynamics Laboratory) and colleagues used the NCCS Discover system to investigate behavior of total column ozone in Arctic winter and spring during years with and without Sudden Stratospheric Warming (SSW) events.

The study used the Global Modeling Initiative (GMI) chemistry transport model, employing the Modern-Era Retrospective Analysis for Research and Applications (MERRA) meteorological fields, also generated on Discover, and compared results with 2005 – 2015 observations from Aura's Microwave Limb Sounder (MLS) ozone and N2O data to show GMI credibly represents transport processes and the net heterogeneous chemical loss of O3.

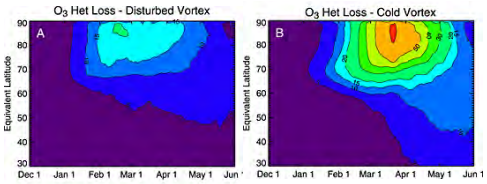

Average column ozone depletion (Dobson Units) for (top) 6 years with Sudden Stratospheric Warmings (SSWs) and (bottom) 5 years with no major or minor SSW before mid-February. Maximum ozone depletion in years without a midwinter warming is roughly 3 times greater than years with a warming.

- The investigation showed that winters with SSW events (where the Arctic stratosphere is warmer, the polar vortex is more disturbed, and more atmospheric meridional transport occurs) have less than half the ozone depletion of years without one.
- The analysis showed that 2/3 of the observed large variability of Arctic total column ozone each spring resulted from interannual dynamical variability – i.e., dynamic transport (resupply or reduction) by atmospheric waves or eddies – while only 1/3 was due to ozone loss by ozone depleting substances.

Impact: *Studies such as this one demonstrate the value of the Discover cluster in providing the environment for simulations and analyses that enable researchers to investigate important phenomena such as the attribution of behavior of Arctic ozone variability .*

Using Katrina and Sandy Data to Improve Hurricane Prediction Tools *

- Using data from Hurricanes Katrina (2005) and Sandy (2012), researchers at San Diego State University (SDSU) ran large-scale simulations on Pleiades to study the role of multiscale processes associated with environmental flows and improve prediction capabilities for high-impact tropical cyclones (TCs).
- The SDSU team deployed NASA's Coupled Advanced global Modeling & Visualization (CAMVIS) framework and Multiscale Analysis Package (MAP) to obtain the following:
  - Accurate 5-day track and intensity predictions and high-fidelity 3D visualizations for Katrina, showing multiscale interactions with an approaching upper-level trough that may have led to the storm's intensification prior to landfall.
  - Realistic 7-day track and 6-day genesis simulations of Sandy; multiscale analysis and high-fidelity visualizations revealed the multiscale processes of Sandy's formation and movement.
  - Multiscale analysis with MAP's parallel ensemble empirical mode decomposition capability revealed the role of tropical waves, westerly wind belt, and a Madden-Julian Oscillation in Sandy's initial formation.
- The SDSU team is now running multi-TC genesis simulations and analyzing results to understand how downscaling and upscaling processes may help predict TCs in numerical models.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: HECC supercomputing technologies and integrated services make it feasible for scientists to improve large-scale analysis tools for better understanding and prediction of hurricanes, which in turn may save lives and reduce damage costs.

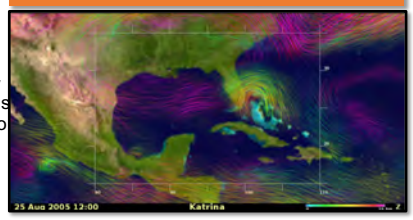

This four-dimensional (X-Y-Z-time) visualization from a multiscale simulation for Hurricane Katrina (2005) shows the horizontal phasing of an approaching jet stream (upper left corner) and Katrina's southwesterly outflow (to the southeast of the jet) prior to landfall. The phasing further strengthens the upper-level anticyclonic flow over the hurricane (dense, streamlines in expanding pink at right), enhancing Katrina's development and creating strong, deep convections. *Bron Nelson, NASA/Ames*

Cloud Resolving Models – Cloud Library

The NCCS is supporting Dr. Wei-Kuo Tao's AIST project

Empowering Data Management, Diagnosis, and Visualization of Cloud-Resolving Models (CRM) by Cloud Library upon Spark and Hadoop

Science objectives

Diagnose key processes for cloud-precipitation

Provide tools for the intercomparison of cloud models

NU-WRF with the Goddard Cumulus Ensemble (GCE) model

High-resolution simulations are being run on Discover specifically looking at deep convective cells in the Midwest

The results will be accessible through a Super Cloud Library (SCL) which contains a database management, distribution functions, visualization, subsetting, and evaluation built on top of a Hadoop framework

Project funded by ESTO AIST

Impact: Accelerating the development and adoption of cloud resolving models in weather and climate research.

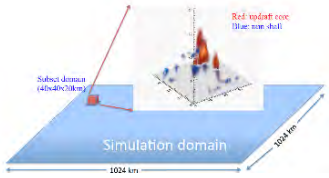

Representative subsetting example of overlaying multiple variables to study their relationships. Shujia Zhou

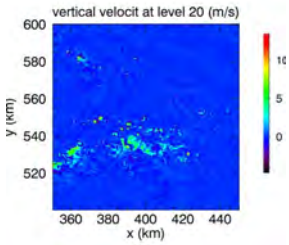

250 meter resolution regional model simulation using NU-WRF showing a subdomain plot of the vertical velocity. Xiaowen Li (GSFC)

Projecting Sea Level Rise by Modeling the Evolution of Ice Sheets *

- A team at JPL and the University of California at Irvine (UCI) use Pleiades extensively to accurately model the evolution of ice sheets and understand how freshwater fluxes to the ocean will impact sea level rise in the coming decades.
- Using the open-source Ice Sheet System Model (ISSM) software, developed at JPL/UCI, the team's results have had a widespread impact on the cryosphere scientific community. Among their contributions:
  - ISSM's automatic differentiation feature;
  - Optimization of bedrock topography data, and creation of physically consistent bedrock topography for the entire Greenland ice sheet; and
  - Improvements in modeled physical processes.
- While ISSM can be deployed in a wide range of HEC environments, the JPL team's reliance on Pleiades is of particular importance—they regularly complete large simulation runs involving thousands of computational nodes on the system.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: Understanding and projecting the evolution of polar ice sheets is a priority for the global science community. By leveraging HECC resources, scientists can complete jobs in hours, compared to weeks on conventional computing resources.

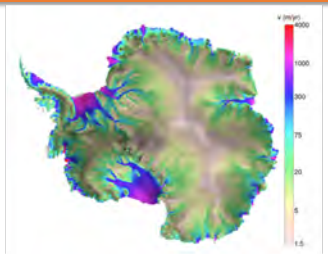

This image from the Ice Sheet System Model represents the computed velocities for the entire continent of Antarctica. Brown and green represent relatively slow ice surface velocities (fractions of meters per year to a few meters per year). Blue and purple represent much faster velocities, ranging up to 3,000 meters per year.

Using GEOS-5 Nature Run Results to Characterize Impact of Sampling Frequency on DSCOVR-Derived Energy Budget of the Earth

Daniel Holdaway (GMAO) and Yuekui Yang (Climate and Radiation Laboratory) used the NCCS Discover cluster to analyze time series of outgoing radiation from the 7-km GEOS-5 Nature Run (G5NR) data, employing it as a surrogate for 2 years of observations from DSCOVR's EPIC and NISTAR instruments.

Their study, published as part 1 of 2, characterized effects of temporal sampling frequency on the resulting Earth radiation budget derived from DSCOVR observations, measuring skill using correlations and normalized errors.

A particular focus was to analyze whether a given sampling frequency of outgoing shortwave and longwave radiation of Earth's sunlit side would provide an accurate measure over both shorter and longer time scales.

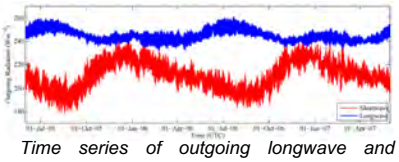

Time series of outgoing longwave and shortwave radiation produced by the GEOS-5 Nature Run (June 1, 2005 to June 1, 2007).

Left: a July 5, 2015 image from DSCOVR's Earth Polychromatic Imaging Camera (EPIC). Right: artist's rendering of the DSCOVR spacecraft.

Evaluations revealed that, for global radiation budget assessments, sampling frequency coarser than every 4 hours resulted in significant error, and that more frequent observation rates were particularly important for the Arctic region around the beginning and end of the polar night.

- Part 2 of the study will examine impacts of temporal sampling on DSCOVR observations of cloud cover.

Studies such as this one demonstrate the value of the Discover cluster in powering the G5NR simulation, and in providing the environment for critical analyses that can assist researchers in optimizing both accurate observations and efficient instrument sampling frequency.

KORUS-AQ Field Campaign Support

KORUS-AQ (April–June 2016) is an international cooperative air quality field study focused on the South Korean peninsula and surrounding waters to better understand the factors controlling air quality across urban, rural, and coastal interfaces.

NASA's Global Modeling and Assimilation Office (GMAO) is providing aerosol and carbon monoxide forecasting support and participating in flight planning discussions.

The NCCS Discover supercomputer is hosting the GEOS-5 analyses and forecasts and the Data Portal is delivering forecast imagery and data at: http://gmao.gsfc.nasa.gov/field_campaigns/KORUS-AQ

GMAO scientist Arlindo da Silva gave a forecast briefing to the campaign team.

NASA Earth Expeditions video (starting at 0:14): https://www.youtube.com/watch?v=KxefWn_b-rE

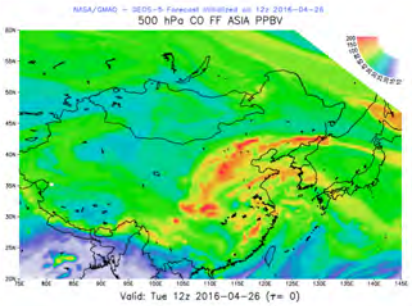

This image shows a GEOS-5 forecast of carbon monoxide (CO) over Asia computed for the KORUS-AQ field campaign. Interactive forecast maps and data are available through the NCCS Data Portal.

Impact: Enabled by the NCCS computational and data distribution resources, the GMAO provides critical aerosol forecasts for KORUS-AQ flight planning.

Estimating Trees and Shrubs in the Sub-Sahara With Satellite Imagery

NASA Goddard and University of Minnesota scientists are estimating the carbon stored in vegetation across a coast-to-coast swath of Africa from the Sahara south to the savanna zones.

This vast region includes an estimated one billion trees and shrubs that are visible in commercial satellite sub-meter resolution imagery acquired through the National Geospatial Agency.

The ultimate result will be a biomass estimate for the entire study region. The team is currently assembling mosaics using panchromatic (i.e., black and white) imagery and will soon be using multispectral data, which they believe will give a higher degree of accuracy.

The team of scientists and computer scientists have innovatively combined the NCCS Advanced Data Analytics Platform (ADAPT) with the Amazon Web Services cloud to divide tens of thousands of satellite images into smaller tiles for analysis using machine learning algorithms.

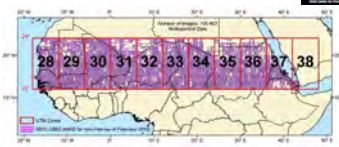

The project encompasses Universal Transverse Mercator (UTM) zones 28 to 38.

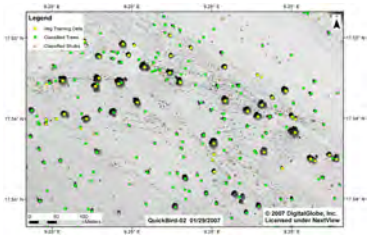

Vegetation classifier algorithms identify trees and bushes in UTM Zone 32. Image by Jamon Van Den Hoek, Gerasimos Michalitsianos, and John L. David, GSFC.

Impact: The NCCS has teamed with NASA Goddard and University of Minnesota scientists to use high-resolution commercial satellite imagery to create an unprecedented measure of biomass over sub-Saharan Africa.

GISS Simulations Show Farms Are a Major Source of Air Pollution

Recently published Goddard Institute for Space Studies (GISS) research says that emissions from farms outweigh all other human sources of fine-particulate air pollution in much of the United States, Europe, Russia, and China.

The research is based upon simulations using GISS Earth System ModelE2 (2º by 2.5º resolution, 40 layers) coupled to the aerosol microphysical scheme MATRIX (Multiconfiguration Aerosol TRacker of mIXing state).

ModelE2 simulations on the NCCS Discover supercomputer included a base run plus experiments where anthropogenic emissions are set to zero and where only agricultural emissions are set to zero.

For coming decades most model projections show that if industrial emissions decline, fine-particle pollution will go down even if fertilizer use doubles as expected.

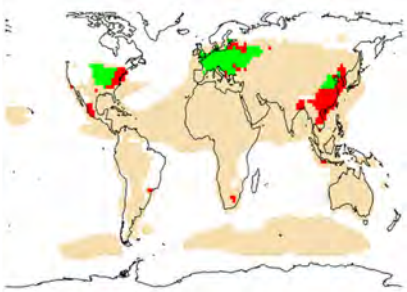

A world map shows main sources of fine-particulate pollution: natural (brown), industrial (red), and agricultural (green). Image from Bauer et al., Geophysical Research Letters, 2016.

High Resolution Particulate Matter Studies Are Fueled by 1-km MODIS AOD, Generated from MAIAC Algorithm on Discover

- 1-km MODIS Aerosol Optical Depth (AOD) data, generated on the Discover cluster using the innovative Multi-Angle Implementation of Atmospheric Correction (MAIAC) algorithm, are used in several new studies as a key input for high-resolution estimates of Particulate Matter (PM).

The studies show daily PM estimates based on the MAIAC AOD compare well to 10+ years of observations from monitoring sites.

The studies also show high resolution PM estimates from MAIAC AOD to be particularly useful in areas such as Mexico City and Israel, where bright land surface and complex terrain prove challenging for satellite monitoring; and in rural areas where monitoring sites are sparse.

Using the MAIAC algorithm, Terra and Aqua's long MODIS records can be used for high-resolution PM estimates to reconstruct both long- and short-term exposure to PM for epidemiological studies and other purposes.

While current MODIS L2 AOD products are produced at 10-km resolution, MODIS data processing using the MAIAC algorithm is slated to become operational soon.

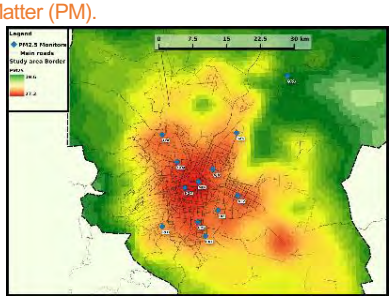

Estimated using MAIAC 1-km MODIS Aerosol Optical Depth for years 2004–2014, the overall mean of daily 1-km fine particulate matter (PM2.5) predictions across the Mexico City region (from Just et al., doi: 10.1021/acs.est.5b00859). MAIAC Point of Contact: alexei.i.lyapustin@nasa.gov)

Powered by the Discover cluster, MAIAC proof-of-concept processing on large MODIS selections has produced high resolution results that are being used by the environmental science community as a springboard for new capabilities.

Virtual Earth Observing via GMAO's GEOS-5 Global 1.5-km Simulation

- Using the NASA Hyperwall at the Fall 2015 AGU meeting, GMAO's Bill Putman presented a four-day GEOS-5 global model simulation run at an unprecedented 1.5-km resolution.
- Employing more than 200 million grid cells, the four-day simulation used more than 28,000 cores on Discover's 1.2-PetaFLOP SCU10 for throughput of one simulation day per wall-clock day.
- Such high resolution is expected to be used in operational weather forecasts about the year 2030, but would require a staggering 10 million cores with current technology to the meet operational throughput requirement for 10-day simulations every 6 hours.
- This non-hydrostatic, high resolution simulation incorporated a new microphysics scheme, and spawned very realistic weather features such as large severe thunderstorm complexes with over-shooting cloud tops and intricate stratocumulus clouds blanketing the oceans.

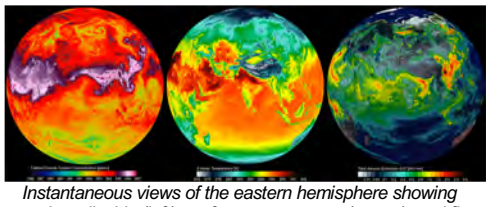

Instantaneous views of the eastern hemisphere showing carbon dioxide (left), surface temperature (center), and fine particles (right) from a 1.5 kilometer (~1 square-mile-grid cell) GEOS-5 global atmospheric model simulation for June 16, 2012. William Putman, NASA/Goddard

The first ultra-high-resolution global simulation to include interactive aerosols and carbon, its visualizations show fine aerosol particles streaming from industrial plants and fires around the globe, and waves of carbon dioxide drifting from their sources and engulfing the world.

GMAO scientists and software engineers trimmed GEOS-5's memory footprint to fit more than 5 terabytes of global data into the 128 gigabytes of available memory per node.

The powerful combination of Discover's capacity and environment with GMAO's scientific and computational expertise enables pioneering advancements with GEOS-5 to create a more efficient and robust modeling capability to meet the demands of NASA's current and future Earth science missions.

30-Year Global Fire Weather Analysis Based on MERRA and Ground Station Observations Sets Stage for Contributions from GPM, TRMM, and GPCP

Robert Field (GISS/Columbia University) has **developed a 30-year Global Fire WEather Database** (GFWED) that employs global gridded data from the Modern Era Retrospective Analysis for Research and Applications (MERRA) and gridded estimates derived from rain gauge measurements.

Using the Canadian Forest Fire Weather Index (FWI), GFWED integrates the effects of the different weather factors which influence the likelihood of a vegetation fire starting and spreading.

Use of reanalysis data **enables a historical view of fire weather that is continuous in space and time**, an advantage compared to use of measurements from rain gauges and ground stations alone.

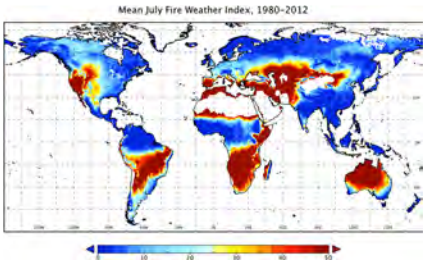

Mean July Fire Weather Index from 1980-2012, based on NCEP 0.5-degree resolution daily precipitation estimate over land.

- This research forms the basis for a new effort supported by the Precipitation Measurements Mission, to apply NASA satellite precipitation products from GPM, TRMM and GPCP into GFWED.
- GFWED data, computed on the Discover Cluster, is distributed via the NCCS Dataportal.

Enabled by NCCS computational and data distribution resources, GFWED provides a 30 year record of global fire weather indices, and a platform for incorporating NASA satellite precipitation measurements.

Estimating Woody Biomass on South Side of the Sahara at the 40-50 cm Scale Using Amazon Web Services (AWS): Progress for Intel Challenge "Head in the Clouds" Project

Project Science Goal - Using National Geospatial Agency (NGA) imagery data, estimate tree and bush biomass over the entire arid and semi-arid zone on the south side of the Sahara to:
Estimate carbon stored in trees and bushes; and

- Establish carbon baseline for later research on expected CO_2 uptake.

Initial Test Runs in Amazon Web Services (AWS) Cloud

Executed on 1/3 of Universal Transverse Mercator (UTM) Zone 32 with data (8 TBs preprocessed at NCCS) from 1 of 4 total satellites.
Used 200 virtual machines (VMs) for 6-7 hours via low-cost AWS spot instances (processing cost est. for test: ~$80).

- Results suggest improvements in surface soil moisture, root zone soil moisture and total runoff.

Work scales linearly, estimate to finish the entire study domain: (11 UTM Zones) * (4 Satellites) * (600 VMs/UTM Zone) = 26,400 VMs for 6-7 hours

Proved total dataset of 80 TBs (10x initial test run) could be completed in same amount of time as 8 TBs.

– Estimated total compute cost, AWS spot instance: ~$11,000

***Received 2015 HPCwire Readers' Choice Award for Best HPC Collaboration Between Government and Industry**

Zone of study for the Estimating Woody Biomass "Head in the Clouds" project.

Researcher Martin Brandt at a verification test site.

Following extensive testing and fine-tuning, NCCS and partners are demonstrating the potential power, for suitable data analysis projects, of a high performance computing workflow that includes "cloud bursting" to commercial clouds.

March 2016 Field Campaign Support

The NCCS Discover cluster provided computational resources for the GMAO's operational research GEOS-5 analyses and forecasts, with data and forecast images delivered via the NCCS Dataportal, including support for the following Field Campaigns:

Campaign	Description	Time Frame
NOAA's ENRR El Niño Rapid Response http://www.esrl.noaa.gov/psd/enso/rapid_response/	Ground, seaborne, and airborne observations measured how the very strong El Niño influences the atmosphere. Included NASA's HIWRAP, HAMSR, LIP instrument packages on board NASA's Global Hawk.	01/2016 - 03/2016
AfriSAR-NASA 2016 Africa Synthetic Aperture Radar http://www.nasa.gov/feature/goddard/2016/nasa-partner-space-agencies-measure-forests-in-gabon	NASA collaboration with European Space Agency for very high-resolution 3D measurement of ecosystem structure over Gabon, in support of global carbon budget studies, using NASA's Land, Vegetation, and Ice Sensor (LVIS, a precursor for GEDI (Global Ecosystem Dynamics Investigation) lidar) on board NASA B-200, and UAVSAR (NASA-ISRO Synthetic Aperture Radar precursor) on board NASA C-20A.	02/19/2016 - 03/08/2016

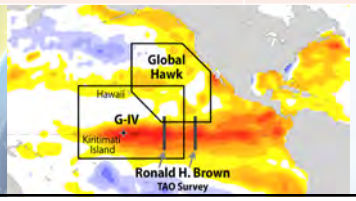

Left: El Niño Rapid Response mission map.

Discover Supports Twice-Daily 7-Day GEOS-5 Experimental 6-km Global Forecasts

Beginning in mid-March, the GMAO's Bill Putman has been using the Discover cluster to run twice-daily experimental global 7-day forecasts using the 6-km GEOS-5 system.

This is believed to be the highest-resolution global forecast run with this near-operational frequency.

Putman engineered his code to use a configurable number of processors, with the effect that adequate resources remain available for other Discover jobs while the 6-km forecasts run.

Forecast jobs using 18,000 cores (643 28-core Haswell nodes) typically complete in 5 ½ to 6 hours.

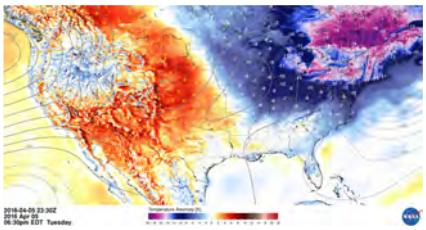

Temperature anomaly forecast for April 5, 2016 from an experimental 6-km GEOS-5 global forecast initialized on April 1, showing unusually cold temperatures in the northeastern US.

- Selected images, currently generated by automated plotting jobs on Discover, are served via the NCCS Dataportal ftp download service.
- To support the large volume of high-resolution forecast image data, the NCCS is using its special-purpose read-only export service to provide specifically designated data on Discover disk to the Dataportal, avoiding the additional network and disk space burden.

Using NCCS Discover cluster and Dataportal resources, GMAO scientists are demonstrating the feasibility of near-operational generation of NASA research forecasts at trailblazing resolutions.

Pleiades-Enabled Reprocessing by OCO-2 Provides a Consistent Ten-Month Data Record

The OCO-2 team at JPL recently completed a major reprocessing campaign covering ten months of the mission, enabled by the extensive use of NASA's Pleiades supercomputer. **About 60% of the campaign was pushed through Pleiades, including 18.6 million instrument soundings (using 1.4 million CPU-hours).**

The effort utilized the highly-reliable transfer tools and up to 500 Haswell nodes concurrently. Because the Full Physics retrieval process operates on the soundings independently to extract X_{CO2}, it is well suited to the Pleiades architecture.

The added computing resources **allowed the team to complete the campaign approx. 3 months sooner than initially planned and get a consistent, well-calibrated and comprehensive data set** out to the public quickly.

Science team members and the larger community are now mining the data record at both global (see maps at right) and regional scales.

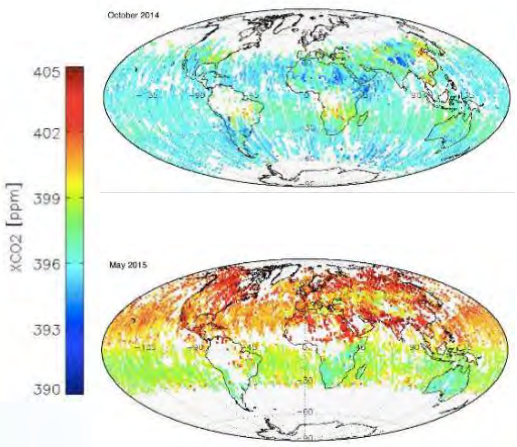

Global XCO2 maps for October 2014 and May 2015

GEOS-5 Chemistry Climate Model Simulation of Stratospheric Ozone Intrusion Featured at SIGGRAPH 2015 Conference

Stephen Pawson, Andrea Molod and colleagues are using the Discover cluster for investigations focusing on "Coupled Predictions of Ozone and Climate."

Team member Eric Nielsen's simulations using the Goddard Earth Observing System Chemistry Climate Model (GEOS-5 CCM) were featured in visualizations by Trent Schindler (USRA, for the Science Visualization Studio) at the SIGGRAPH 2015 conference, August 9-13, in Los Angeles.

The "still" and movie visualizations demonstrated that stratospheric ozone intrusions are captured accurately by simulations run at 28-km resolution (roughly ¼ degree latitude).

The team is currently running higher-resolution simulations with updated model components for historical simulations beginning with 1950, and eventually extending to projections up to 2100.

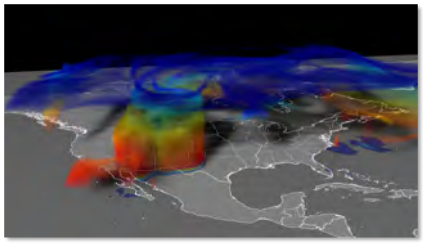

This visualization shows the 28-km GEOS-5 CCM simulation of a stratospheric ozone intrusion event that occurred on April 6, 2012. (Simulation point of contact: Jon.E.Nielsen@nasa.gov; image credit: Trent Schindler, USRA, Science Visualization Studio, Code 606.4)

The 28-km simulations used for the visualizations employed 1,700 to 2,500 cores on the NCCS's Discover Linux cluster.

Discover cluster's capacity enables researchers' use of NASA observations with coupled chemistry climate simulations to investigate interactions of natural phenomena and anthropogenic emissions on the physical climate for the historical period, and evaluate scenarios and uncertainties for longer term projections.

Modeling the Carbon Balance in Arctic Ecosystems *

- Meteorologists supporting NASA-sponsored research, including the Arctic Boreal Vulnerability Experiment (ABoVE), ran simulations to improve understanding of the current balance of carbon dioxide (CO_2) and methane (CH_4) in the Arctic. Using Pleiades, the researchers:
  - Generated high-resolution reanalyses of air flow patterns over the Arctic during time periods when CO_2 and CH_4 observations were being made.
  - Computed backward trajectories for hundreds of thousands of air particles to identify the pathways of gases released from the Earth's surface.
  - Computed the potential contribution from these upwind surface-to-air gas fluxes to the concentration measurements at "receptor" locations.
- Pleiades enables modeling of time periods spanning multiple years at a resolution that is an order of magnitude higher than prior studies.
- Researchers will continue to use Pleiades to refine their Weather Research and Forecasting/Stochastic Time-Inverted Lagrangian Transport (WRF-STILT) model, and incorporate current, future and historical observations to provide new insights into the vulnerability and resilience of ecosystems and society to the changing Arctic environment.

HECC provides supercomputing resources and services in support of this work

Mission Impact: HECC resources enable researchers to achieve fine-scale detail in model fields of the Arctic ecosystem at a resolution that is an order of magnitude greater than that of prior research studies.

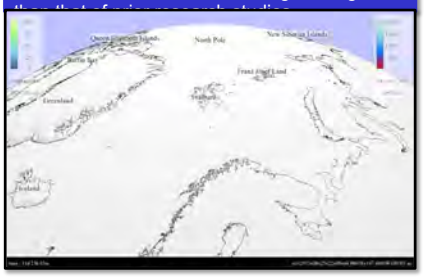

Visualization of the trajectories of 500 air particles ultimately converging at a receptor (yellow marker) over a 10-day period. Colored line segments denote the positions of particles at high altitudes (white) and low altitudes (red) over 4 hours of that period. Particles passing into the planetary boundary layer contribute to the footprint field (shaded) that accumulates over the 10 days. *Timothy Sandstrom, NASA/Ames; John Henderson, Atmospheric and Environmental Research*

Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring

Teppei Yasunari (Hokkaido University) and colleagues used the Discover cluster for simulations examining the impact on Northern Hemisphere snowpacks of three types of light-absorbing aerosols – dust, black carbon (BC) and organic carbon (OC) – including their effects on snow amount and heating on the ground in spring.

Utilizing the GOddard SnoW Impurity Module (GOSWIM), the team added snow darkening effects to the GMAO's GEOS-5 climate model, and incorporated aerosol deposition from the Goddard Chemistry, Aerosol, Radiation, and Transport (GOCART) model from Goddard's Atmospheric Chemistry and Dynamics Laboratory (Code 614).

Running two sets of simulations – with, and without, the dust, BC, and OC constituents in snow – scientists analyzed results from 2002 to 2011, to quantify the impact of the constituents on boreal spring climate.

Simulations including effects of light absorbing aerosols absorbed more of the sun's energy, with the darkened snow surface temperatures up to 10 degrees Fahrenheit warmer compared to simulations without snow darkening effects.

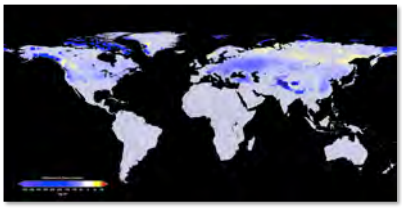

Reduced albedo of snow from dust, black carbon and organic carbon depositions ("snow darkening effect") alters springtime snow water equivalent (SWE, in kg/m^2) through increased springtime melt compared to simulations with no snow darkening effect. (Image credit: Cindy Starr/GST, Science Visualization Studio)

For much of the Northern Hemisphere snowpack, the higher surface temperatures led to reductions in the springtime snow water equivalent compared to simulations without snow darkening effects.

The power of the Discover cluster not only enabled ten-member ensembles for both snow darkening and non-snow darkening cases, but also provided capacity for additional constrained "replay" simulations (using MERRA for certain atmospheric states) to allow for contextual comparisons against the free-running simulations.

Pleiades' computational power has enabled improved understanding of methane (CH_4) fluxes

Step 1: Collect atmospheric CH_4 observations – from aircraft and tall tower platforms.

Step 2: Run weather model and "influence footprint" simulations on NASA Pleiades, using massively parallel architecture. These footprints estimate the fraction of emitted gas from any location that flowed to the observation site:

[Modeled concentrations] = [footprints] x [estimated emissions]

Step 3: Optimize emissions to match modeled and observed concentrations.

Recent Results

Wetland fluxes:

Chang et al. (2014):
Estimated a CH_4 budget for Alaska. Found no signs of rising Arctic CH_4.

Miller et al. (2014, 2015):
All biogeochemical models overestimate CH_4 from Canadian wetlands.

Anthropogenic emissions:

Miller et al. (2013):
National emissions from agriculture and the oil/gas industries exceed inventory estimates.

McKain et al. (2015):
Estimated leaks from Boston's natural gas infrastructure. Leaks ~2.5x inventory estimates.

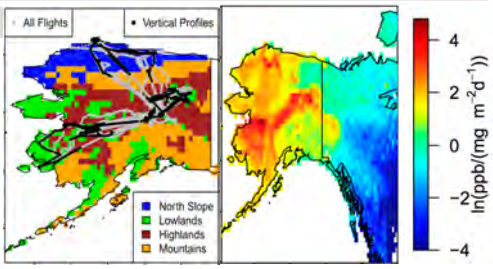

Figure 1: Observations (left) and cumulative footprints (right) from the NASA-lead CARVE aircraft campaign in Alaska. Footprints have units of concentration per flux.

Discover Expansion Powers Progress in Dynamical Downscaling Studies with the NASA Unified-WRF Model

A team led by PI Christa Peters-Lidard is investigating dynamic downscaling, employing the NASA Unified-Weather Research Forecasting (NU-WRF) regional Earth System model.

A key objective is to establish a NASA framework for assessing the credibility of downscaled climate projects.

Evaluation is focusing on three phenomena: northeastern US winter storms, mid-continent summertime mesoscale convective systems, and west coast wintertime atmospheric rivers.

Experiments to date are investigating the 11-year period 1999 to 2010, using the GMAO's MERRA-2 six-hour analysis as initial and boundary conditions, on a 3-level nested grid at 24, 12, and 4 km, with 41 vertical levels.

Intermediate progress was presented at the June

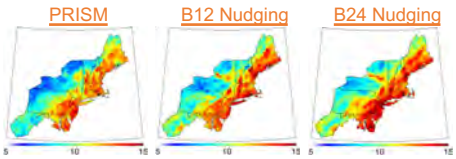

1999 to 2010 northeast US median winter storm frequency from, left, Oregon State University's PRISM Climate Group's climate observations dataset; center, NU-WRF Downscaling from 12-km domain B simulations employing spectral nudging; and right, 24-km domain B simulations with nudging. Intensity is defined by maximum daily precipitation during a storm (mm/day).

The capacity of Discover's recent 28-core Haswell additions has enabled significant progress, especially for the 12- and 4- km domains.

The NCCS also established a special 8-day runtime "quality of service" for the NU-WRF downscaling work.

The expanded capacity of the Discover cluster makes it feasible to explore science questions requiring unprecedented spatial resolutions and long time integrations, such as performance comparisons of parameterized convection vs. resolved convection for impactful weather phenomena.

NASA Releases Detailed Global Climate Change Projections

http://www.nasa.gov/press-release/nasa-releases-detailed-global-climate-change-projections

NASA has released data showing how temperature and rainfall patterns worldwide may change through the year 2100 because of growing concentrations of greenhouse gases in Earth's atmosphere. This NASA dataset integrates actual measurements from around the world with data derived from the General Circulation Model (GCM) runs conducted under the Coupled Model Intercomparison Project Phase 5 (CMIP5) ~~from climate simulations created by the international Fifth Coupled Model Intercomparison Project~~. These climate simulations used the best physical models of the climate system available to provide forecasts of what the global climate might look like under two different greenhouse gas emissions scenarios: a "business as usual" scenario based on current trends and an "extreme case" with a significant increase in emissions. The NASA climate projections provide a detailed view of future temperature and precipitation patterns around the world at approximately 15.5 mile (25 kilometer) resolution, covering the time period from 1950 to 2100. The 11-terabyte dataset provides daily estimates of maximum and minimum temperatures and precipitation over the entire globe. The dataset is the latest product available from the NASA Earth Exchange (NEX), and was released in conjunction with the launch of the Climate Services for Resilient Development Partnership, designed to strengthen developing counties' national resilience against the impacts of climate change. The high-resolution data, which can be viewed on a daily timescale at the scale of individual cities and towns, will help scientists and planners conduct climate risk assessments to better understand local and global effects of hazards, such as severe drought, floods, heat waves and losses in agriculture productivity.

NEX is a collaboration and analytical platform that combines state-of-the-art supercomputing, Earth system modeling, workflow management and NASA remote-sensing data. Through NEX, users can explore and analyze large Earth science data sets, run and share modeling algorithms and

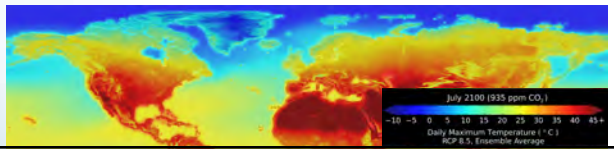

NASA Releases Detailed Global Climate Change Projections

On June 9, in response to a request from WH/OSTP, NASA released the NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP); available to the public via OpenNEX.

Goddard's Climate Model Data Services supported the data publication request, including: downloading, capturing metadata, reprocessing, and organizing 2TB of data for publication/distribution; standing up TDS, FTP, and CREATE-V data distribution access points; and creating an information/landing page as an access point for data and the official White House press release.

The dataset provides the 5th Coupled Modeling Intercomparison Project **downscaled climate model data (daily max/min temperature and precipitation) for the globe at 25km resolution from 1950–2100** under 2 different greenhouse gas emissions scenarios, RCP 4.5 and RCP 8.5.

It was **released in conjunction with the launch of the Climate Services for Resilient Development Partnership**. The high-resolution data, which can be viewed on a daily timescale at the scale of individual cities and towns, aids climate risk assessments enabling improved understanding of local and global effects of hazards, such as severe drought, floods, heat waves and losses in agriculture productivity.

Above: Chief Scientist Dr. Ellen Stofan presented a talk during launch of the Climate Services for Resilient Development Partnership describing NASA's NEX-GDDP release, "Visualization of Climate Services for Development."

Below: NEX-GDDP forecasts how global temperature and precipitation might change up to 2100 (shown here for the RCP 8.5 scenario).

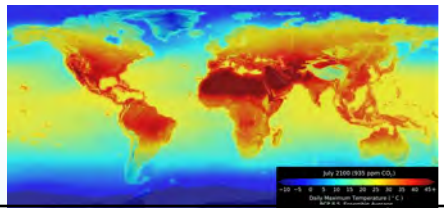

Climate Model Data Services (CDS) Support NASA's Response to Nepal Earthquake

CDS Support for NASA's Response to Nepal Earthquake With National Geospatial-Intelligence Agency (NGA) Very-High Resolution Commercial Satellite Imagery

CDS Support Included:

- Rapid response coordination between NASA HQ Applied Sciences (David Green), and the Polar Geospatial Center (PGC) at the University of Minnesota
- Established path forward under existing CDS, PGC, NGA infrastructure connections to retrieve, format, and present pre/post earthquake to NCCS ADAPT system, with no additional cost
- Coordination with Ohio State University for rapid creation of Digital Elevation Maps
- Retrieved and provided access to 1.4TB and 3,173 of Mono/Stereo Nepal NGA data on ADAPT system in approximately 5 days

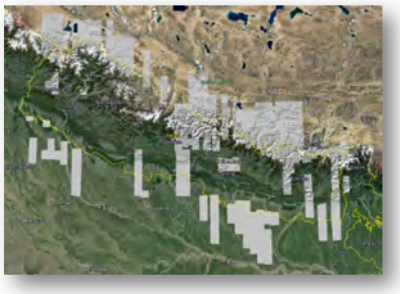

Pre-earthquake Digital Elevation Map of Nepal using NGA/Digital Globe Very-High Resolution Commercial Satellite Imagery

NASA Center for Climate Simulations

Real-Time Assimilation of ISS-RapidScat Observations

Increase of Surface Wind Observations by Assimilating RapidScat

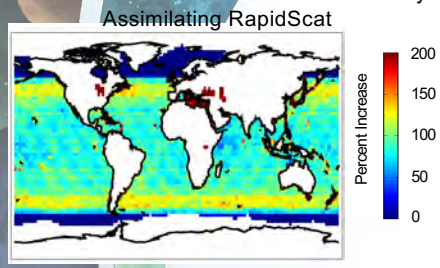

Flying on the International Space Station, RapidScat is providing surface wind observations between ±56° Latitude

These observations are available with low latency from JPL RapidScat Team and are *assimilated in real-time* at GMAO

The ISS orbit provides data coverage complementary to EUMETSAT ASCAT, greatly improving global sampling

Observation impacts, which quantify the reduction of short-term forecast error in observation space, show that ***RapidScat*** is outperforming ***ASCAT*** in our system

The metric also shows that, per observation, the RapidScat observations are nearly on-par with atmospheric motion vectors derived from geostationary satellite imagery

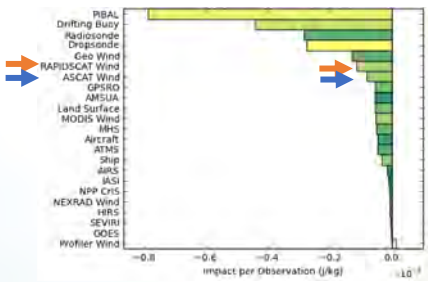

Pleiades-Run Simulations Provide A Virtual Telescope on the World's Oceans and Sea Ice *

- Researchers with the Estimating the Circulation and Climate of the Ocean (ECCO) consortium are running simulations on Pleiades to produce global "maps" of Earth's ocean and sea-ice system at an unprecedented resolution (~1 km horizontal grid).
- Simulations are produced with the Massachusetts Institute of Technology general circulation model (MITgcm) on up to 70,000 Pleiades cores, and are compared to observational data from NASA satellites, ocean sensors, and ship-borne and mooring data.
- Researchers can use this modeling tool to:

  Investigate fundamental questions such as how the circulation, chemistry, and biology of the ocean collectively interact with atmospheric carbon.

  Determine how a pollutant plume or debris field might spread from a particular ocean location, or where and when heat is absorbed by or released from the ocean.
- As HECC storage capabilities continue to increase, this work can potentially develop into a transformative strategy for understanding and predicting the impact of global ocean circulation on climate.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: HECC resources enable high-resolution global simulations that help researchers monitor ocean, sea-ice, and atmospheric systems to learn how they interact and evolve, and to better understand their impact on climate.

The full display resolution (25,600 x 9,600 pixels) of the NAS facility's hyperwall is required to navigate through multiple layers of the simulation to examine scalar and vector fields at various ocean depths. *David Ellsworth and Chris Henze, NASA/Ames*

High-End Computing Capability Project

CISTO Climate Data Services Staff Cited in Federal Lab Consortium's 2015 Interagency Partnership Award for UV-CDAT

CISTO Climate Data Services (CDS) staff collaborating in the Ultrascale Visualization Climate Data Analysis Tools (UV-CDAT) effort were included in the UV-CDAT Project's 2015 Federal Laboratory Consortium for Technology Transfer's (FLC) Interagency Partnership Award.

The Interagency Partnership Award is one of the FLC's highest honors, and it recognizes the efforts of teams who have collaboratively accomplished outstanding work in the process of transferring a technology.

The UV-CDAT Project is a consortium of government agencies, educational institutions, and companies dedicated to delivering robust computational infrastructure for interactive analysis and visualization of extreme-scale climate data.

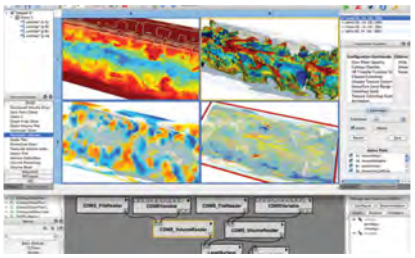

A sample UV-CDAT screen showing 3D data visualization, workflow, and widgets.

- UV-CDAT integrates more than 70 disparate scientific software packages and libraries for large-scale data analysis and visualization, capturing independent workflows and provenance for enhancing reproducibility and repeatability.

***IMPACT:** CDS staff's key contributions to UV-CDAT, including tightly integrated advanced interactive visualization tools for climate scientists, three-dimensional (3D) data visualization, and UV-CDAT tutorials, aid climate researchers in solving their most complex data analysis and visualization problems.*

NASA Center for Climate Simulation

HECC is Used to Clear 5 Years of UAVSAR Data Processing Backlog and Reduce Latency

Impact: Enabled by NASA's Pleiades supercomputer, UAVSAR was able to clear a 5-year processing backlog to deliver surface deformation products to scientists for research.

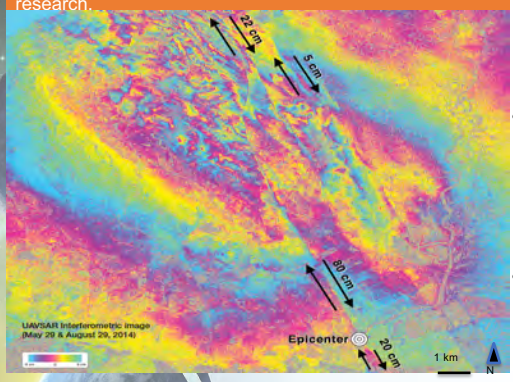

- Scientists are using repeat-pass interferometric (InSAR) data from Uninhabited Airborne Vehicle Synthetic Aperture Radar (UAVSAR) to study centimeter-scale surface deformation of earthquake faults, volcanoes, landslides, and glaciers.
- The computationally intensive InSAR processing code was ported to Pleiades Supercomputer at Ames to take advantage of the large number of processing nodes, each with more than 32 GB RAM, and ample data storage.
- With Pleiades and processor automation, HECC was able to clear a 5-year InSAR processing backlog in 6 months and reduce processing latency to 2 weeks.

UAVSAR Interferometric image (May 29 & Aug 29, '14) of the M6.0 South Napa Earthquake, CA. Colors in the image represent the amount of ground motion between the two flights from the radar's point of view. Linear discontinuities in the colors indicate locations where a surface rupture is highly likely.

Andrea Donnellan, NASA/JPL

Global Modeling and Assimilation Office

1.5km Global Simulation with GEOS-5 on SCU10 at NCCS

- *The highest resolution simulation performed with any US global model*
  - Executed on 30,000 cores of the NCCS 'Discover' SCU10 cluster over three days
  - SCU10 is the first PetaFlop system at the NCCS
  - This non-hydrostatic, high resolution simulation incorporated a new 2-moment microphysics scheme, interactive aerosols, and produced very realistic weather features
  - This engineering demonstration led to rapid development of the infrastructure of GEOS-5 to support high-resolution global downscaling applications for climate and weather
  - Pushing toward these explicit cloud resolving resolutions provides valuable insight into the strengths and weaknesses of physics parameterizations as global models evolve into the 'grey-zone' of parameterizations where sub-grid scale processes are partially resolved

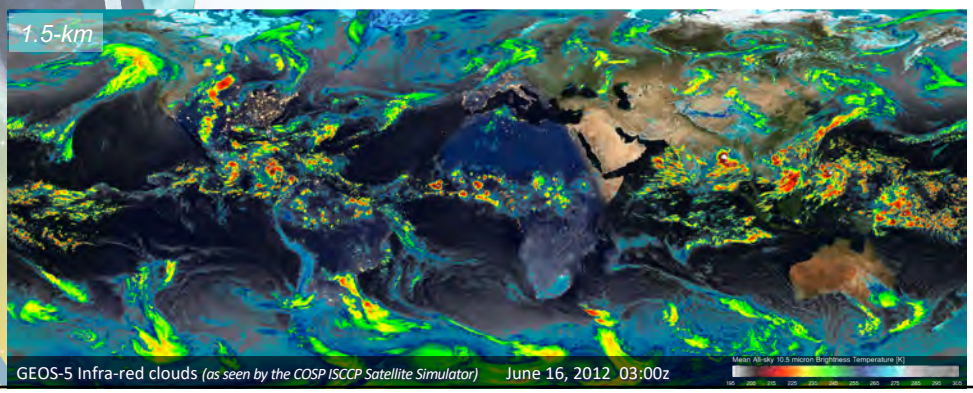

GEOS-5 Infra-red clouds *(as seen by the COSP ISCCP Satellite Simulator)* June 16, 2012 03:00z

50

GEOS-5 Aerosol and Chemistry Forecasts Support CalWater 2/ACAPEX Field Campaigns

The multi-agency CalWater 2 - ACAPEX Field Campaign, conducted in January and February, focused on two phenomena (atmospheric rivers and aerosols from local sources and transported from remote continents) that play key roles in the water supply and the incidence of extreme precipitation events along the US West Coast.

The GMAO's GEOS-5 Chemical and Aerosol analyses and forecasts are providing global, high resolution long-range transport forecasts.

GEOS-5 uses the Goddard Chemistry Aerosol Radiation and Transport (GOCART) module of AeroChem, which simulates major tropospheric aerosol components, to make near-real-time estimates and forecasts of aerosols, CO and CO_2 tracers.

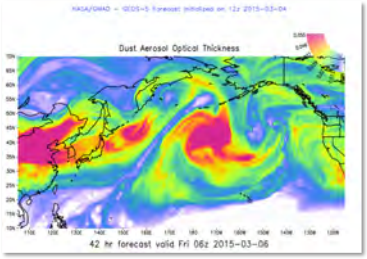

GEOS-5 Dust Aerosol Optical Thickness for a 42-hour forecast over the CalWater domain.

The GEOS-5 forecasts run on the NCCS Discover cluster, and data and images are shared with the research community via the NCCS Dataportal.

IMPACT: Discover's computational power and proximity to the NCCS Dataportal enable the GMAO to provide timely analyses, forecasts, and image to flight planners and researchers in the field.

NASA Center for Climate Simulation

Simulations with Light Absorbing Aerosols Yield Higher Surface Temperatures, Reduced Snow Water Equivalent for Boreal Spring

Teppei Yasunari (Hokkaido University) and colleagues used the Discover cluster for simulations examining the impact on Northern Hemisphere snowpacks of three types of light-absorbing aerosols – dust, black carbon (BC) and organic carbon (OC) – including their effects on snow amount and heating on the ground in spring.

Utilizing the GOddard SnoW Impurity Module (GOSWIM), the team added snow darkening effects to the GMAO's GEOS-5 climate model, and incorporated aerosol deposition from the Goddard Chemistry, Aerosol, Radiation, and Transport (GOCART) model from Goddard's Atmospheric Chemistry and Dynamics Laboratory (Code 614).

Running two sets of simulations – with, and without, the dust, BC, and OC constituents in snow – scientists analyzed results from 2002 to 2011, to quantify the impact of the constituents on boreal spring climate.

Simulations including effects of light absorbing aerosols absorbed more of the sun's energy, with the darkened snow surface temperatures up to 10 degrees Fahrenheit warmer compared to simulations without snow darkening effects.

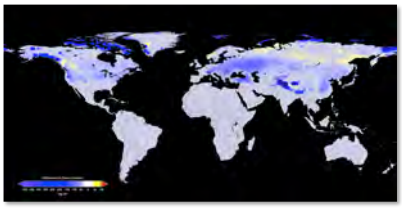

Reduced albedo of snow from dust, black carbon and organic carbon depositions ("snow darkening effect") alters springtime snow water equivalent (SWE, in kg/m^2) through increased springtime melt compared to simulations with no snow darkening effect. (Image credit: Cindy Starr/GST, Science Visualization Studio)

For much of the Northern Hemisphere snowpack, the higher surface temperatures led to reductions in the springtime snow water equivalent compared to simulations without snow darkening effects.

The power of the Discover cluster not only enabled ten-member ensembles for both snow darkening and non-snow darkening cases, but also provided capacity for additional constrained "replay" simulations (using MERRA for certain atmospheric states) to allow for contextual comparisons against the free-running simulations.

FOGGIE: Simulating the Cosmic Fog Around Galaxies*

- Researchers from the Space Telescope Science Institute and Johns Hopkins University are running cosmological simulations on Pleiades to model how galaxies and gas change through time.
- Using their Enzo cosmological hydrodynamic code, the "Figuring Out Gas & Galaxies In Enzo (FOGGIE)" project scientists model the co-evolution of galaxies and their gas with a focus on resolving the ultra-diffuse circumgalactic medium (CGM) with unprecedented fidelity.
- The simulations reveal a richly structured CGM full of churning turbulent gas, small clouds, and tenuous hot gas. Results are used to help interpret real observations made by NASA's Hubble Space Telescope and other observatories.
  - Hubble observations show that low-ionization gas, which should be relatively cool and have higher density, often has kinematic structure very similar to more highly ionized gas, which is expected to generically be hotter and lower density.
  - The high-fidelity FOGGIE simulations reveal that sometimes the kinematically coincident gas is co-spatial, but often it is not; rather, the observation is a chance superposition along the line of sight.

** HECC provided supercomputing resources and services in support of this work.*

IMPACT: Coupled with observational data from Hubble and other space-based observatories, simulations run on HECC resources help reveal the physics of the dynamic and structurally complex circumgalactic medium—a crucial step to understanding how galaxies form and evolve.

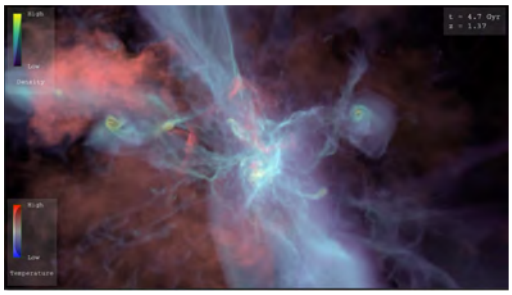

Video from a simulation depicting gas in and around an evolving galaxy across billions of years. Gas density is highest (orange/ yellow) at galaxy's center, but the lower-density gas (blue/black) still has complex structure. *Molly S. Peeples, Space Telescope Science Institute/Johns Hopkins University; Chris Henze, NASA Ames*