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.

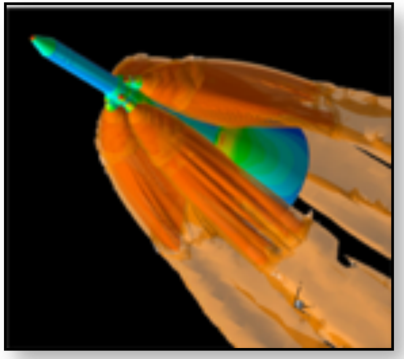

Over 1500 simulations utilized ~2 million processor hours to study launch abort systems on the next-generation crew transport vehicle.

The formation of vortex filaments and their roll-up into a single, prominent vortex at each tip on a Gulfstream aircraft.

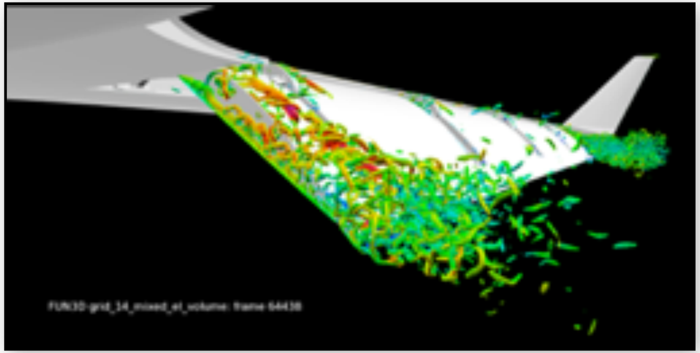

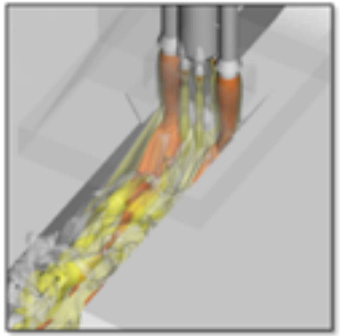

Over 4 million hours were used over a four-month project to evaluate future designed of the next generation launch complex at Kennedy Space Center.

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.

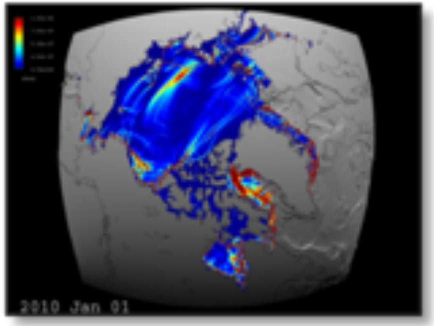

NASA is looking at the oceans, running 100’s of jobs on Pleiades using up to 10,000 processors. Looking at the role of the oceans in the global carbon cycle is enabled by access to large processing and storage assets.

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.

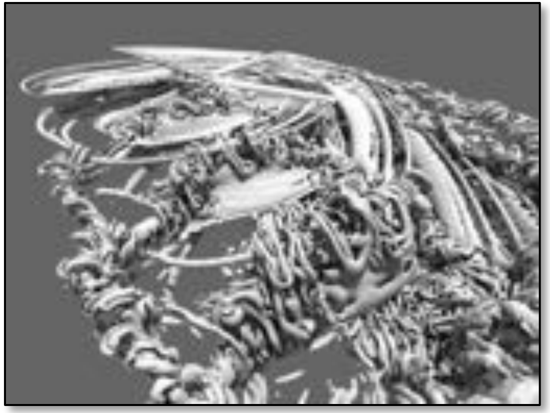

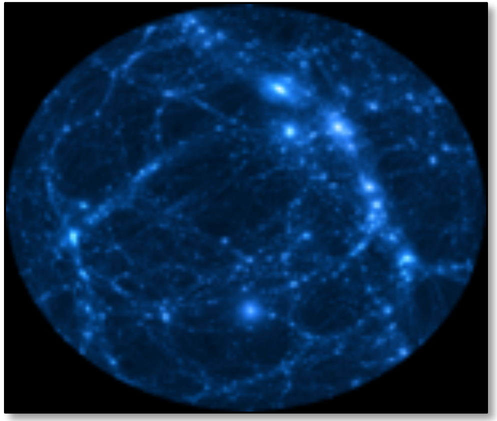

To complete the Bolshoi simulation, which traces how the largest galaxies and galaxy structures in the universe were formed billions of years ago, astrophysicists ran their code for 18 straight days, consuming millions of hours of computer time, and generating massive amounts of data

NASA's HEC Requirements: Time Critical

NASA also needs HEC resources that can handle time-sensitive, mission-critical applications on demand (maintain readiness).

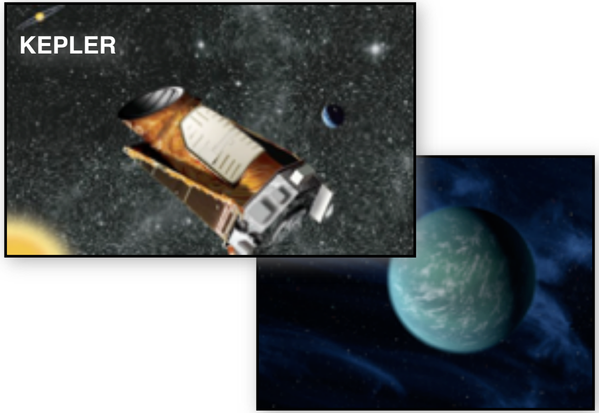

HECC enables enormous planetary transit searches to be completed in less than a day, as opposed to more than a month on the Kepler SOC systems, with significantly improved accuracy and effectiveness of the software pipeline.

REENTRY

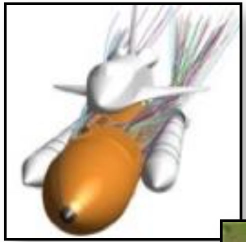

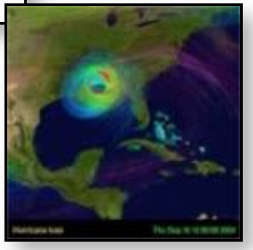

STORM PREDICTION

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.

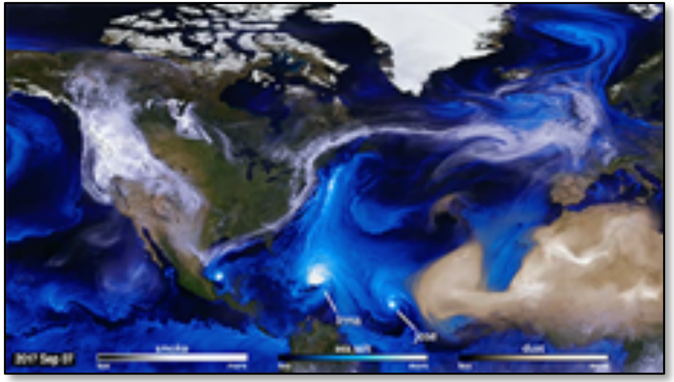

Sea salt particles get caught up in the swirling winds of hurricanes Irma and Jose on September 7, 2017. NASA captured the interaction of hurricanes and aerosols during the 2017 Atlantic hurricane season by combining satellite data with sophisticated models that describe the underlying physical processes.
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.

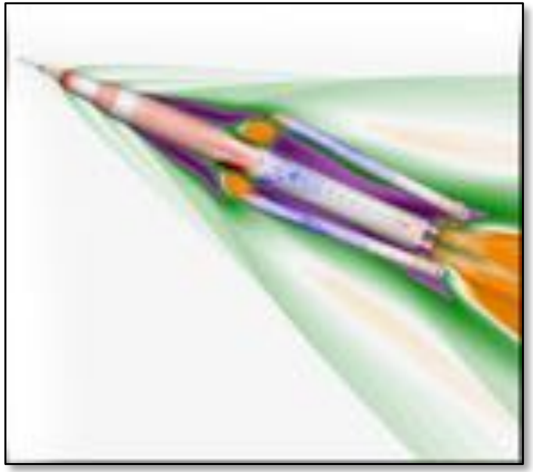

Space Launch System Block IB booster flow field simulated using NAseparationSA's FUN3D code. The crewed vehicle features a different, more powerful second stage. *Jamie Meeroff, Henry Lee, NASA/Ames.*

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.

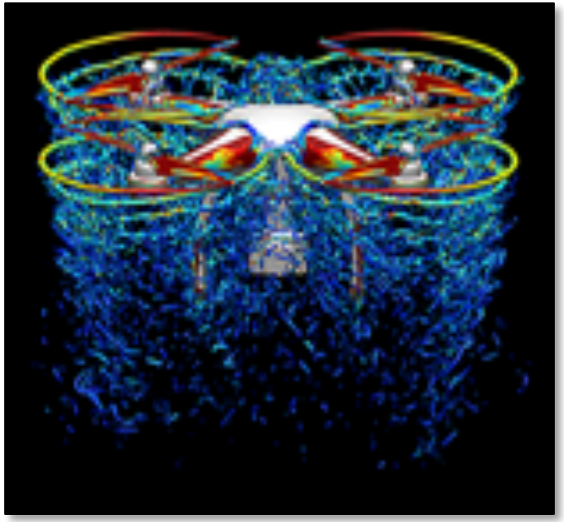

Visualization of the flow of NASA's modified design of a complete DJI Phantom 3 quadcopter configuration in hover. Simulations revealed the complex motions of air due to interactions between the rotors and the airframe. *Patricia Ventura Diaz, Tim Sandstrom, NASA/Ames*

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

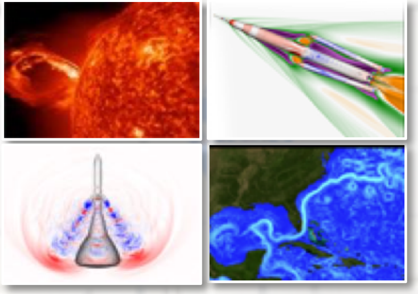

D-Wave 2000Q quantum system

– 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.

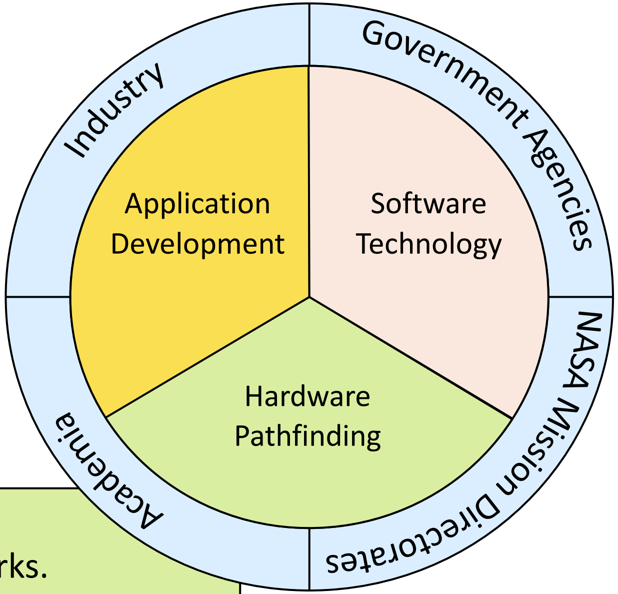

- Research and develop programming approaches, for accelerated and non-accelerated architectures.
- Optimize and port techniques/tools.
- Optimize and port techniques/tools for visualization, data analytics, AI, and machine learning.
- Research and develop efficient math libraries.
- New file system 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

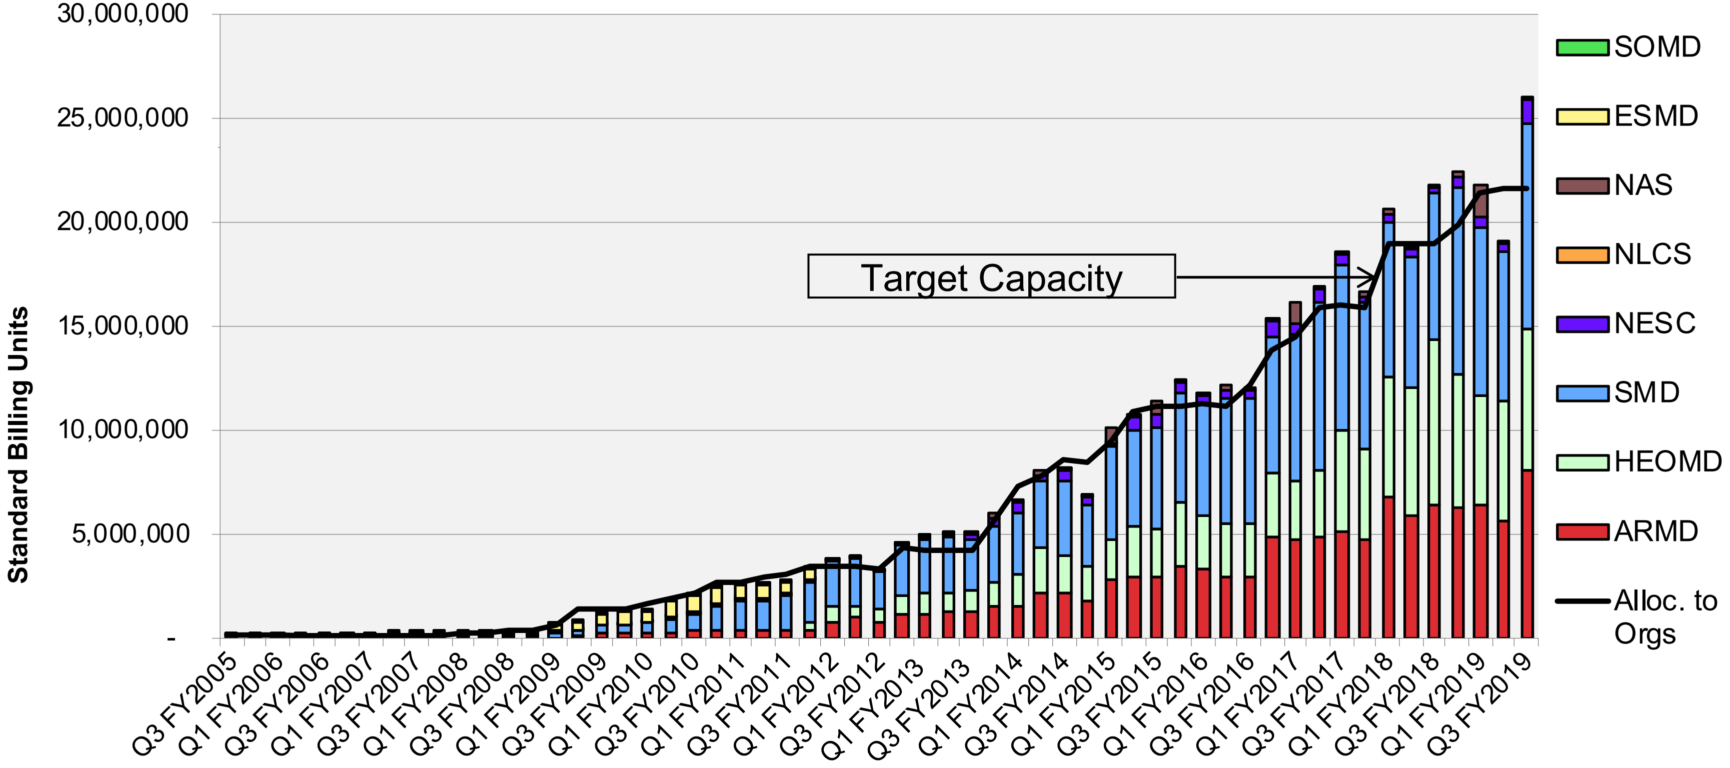

* 1 SBU (Standard Billing Unit) represents the work that can be performed on a dual-socket Broadwell node in one hour.

Aeronautics Support (25,490,311 SBUs)

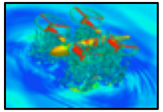

Advanced Air Vehicles

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.

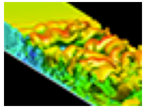

Transformative Aeronautics Concepts

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.

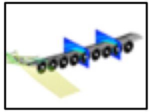

Integrated Aviation Systems

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)

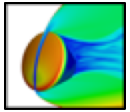

Multi-Purpose Crew Vehicle

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.

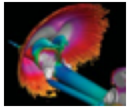

Space Launch Systems

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.

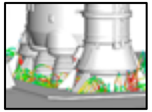

HEOMD - Space Flight Operation & General

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.

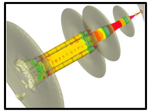

NASA Engineering & Safety Center

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)

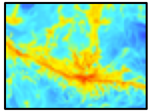

Astrophysics

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.

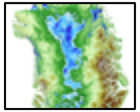

Earth Science

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.

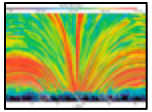

Heliophysics

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.

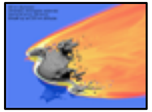

Planetary Science

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

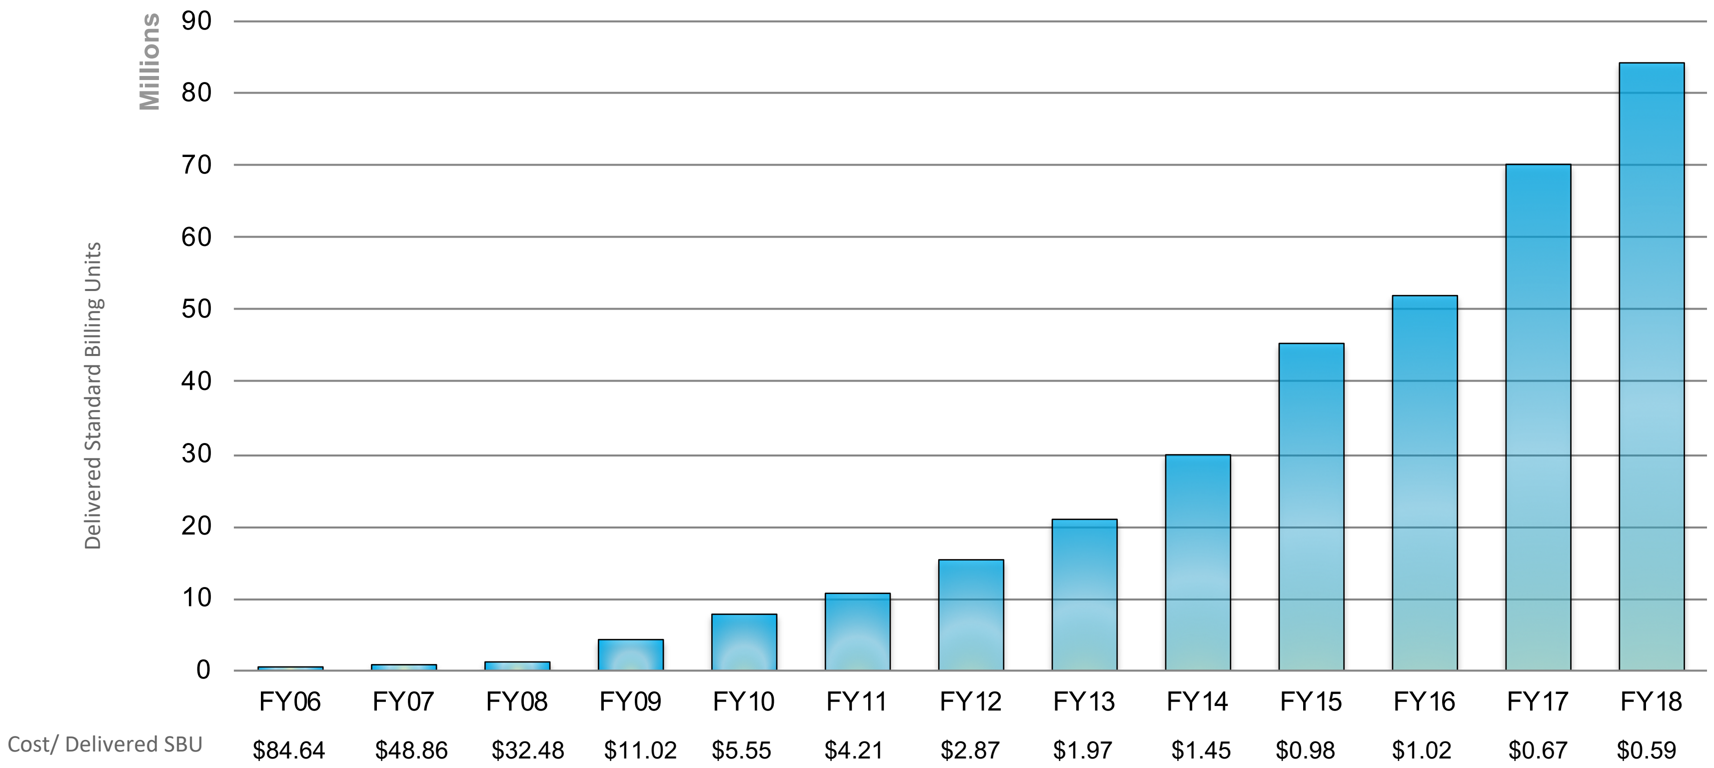

Modular Computing

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

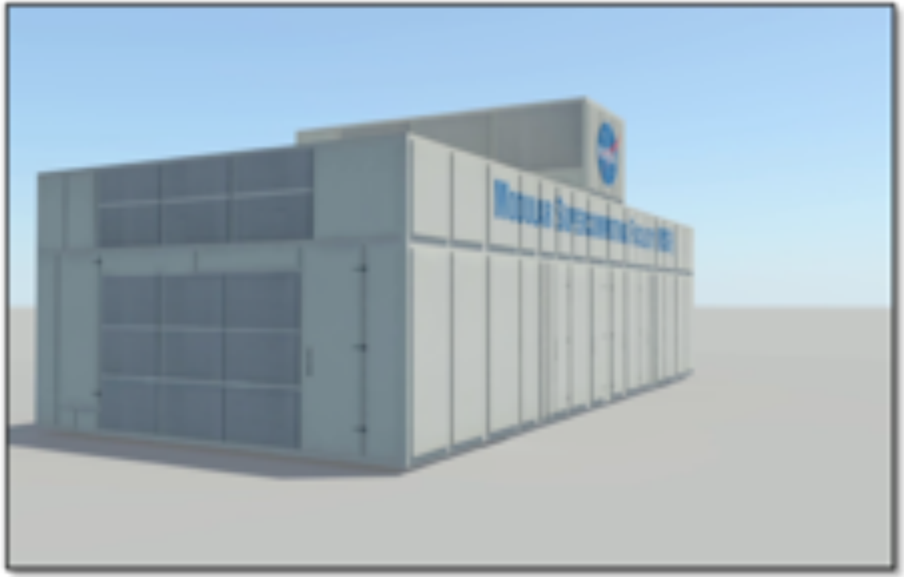

Module 2 Assembly

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)

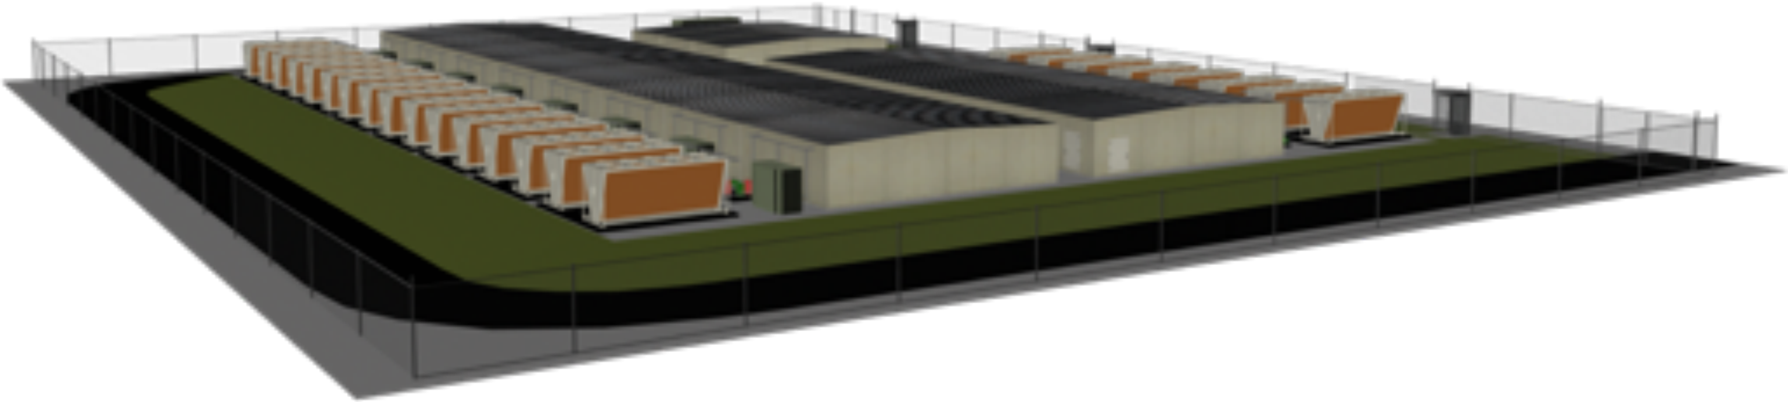

NAS Facility Expansion Accomplishments

- **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.

Pleaides Upgrades 2008 – 2016

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.

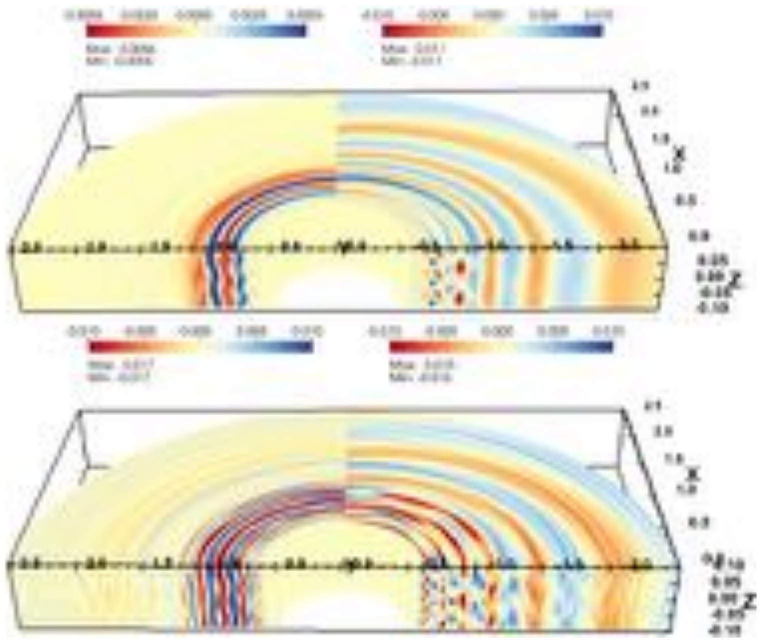

Top: Results with standard resolution (512x128x128). Bottom: 4x higher resolution (2048x512x512). While the spiral wave instability (SWI) develops with both resolutions, the SWI-induced turbulence is better resolved with a larger number of grid cells. *Jaehan Bae, Carnegie Institute of Washington.*

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.

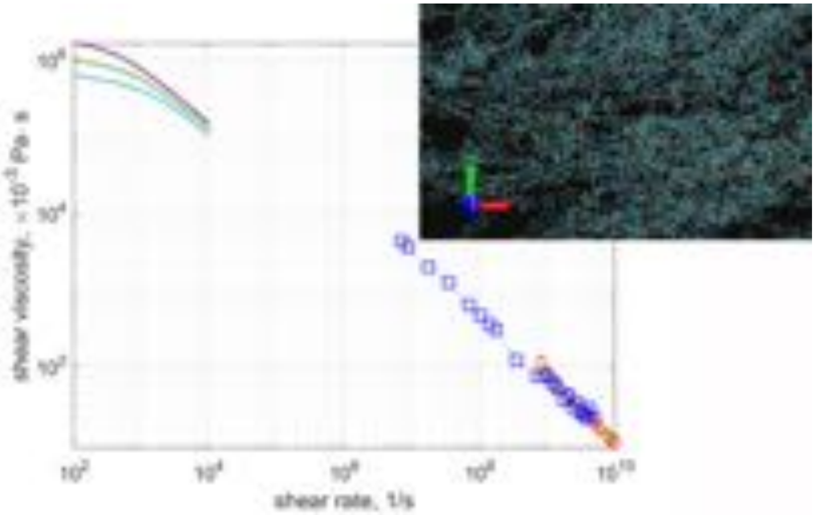

This graph shows atomistically resolved molecular dynamics (MD) predictions of the shear viscosity for flat polymer-polymer interfaces (squares at lower right). These can be compared with experimental data on bulk samples (lines at upper left). The inset shows MD predictions obtained on the Pleiades 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.

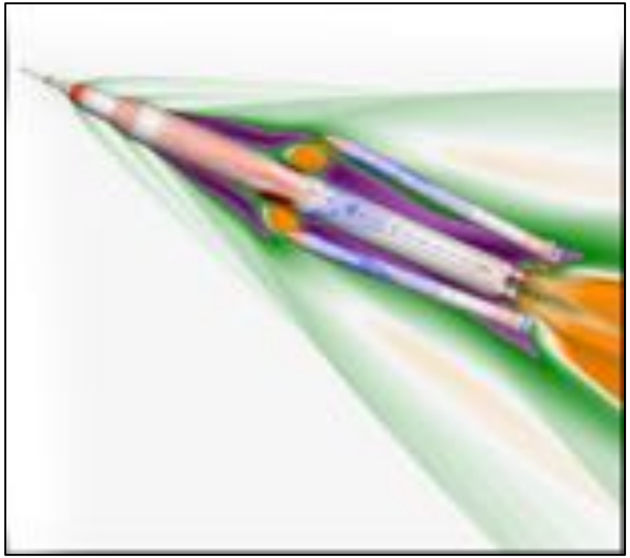

SLS Block 1 vehicle flight and wind tunnel geometries, showing both the flight flow field (left) and the wind tunnel flow field (right), simulated using NASA's FUN3D code. This vehicle will be used for Exploration Mission 1, the first SLS flight. The vehicle surface is colored by pressure contours, where 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`

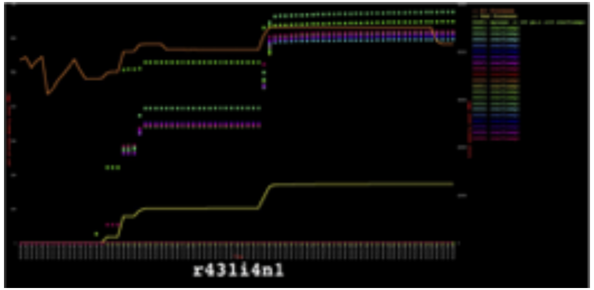

Screenshot from vnuma memory tool.

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:

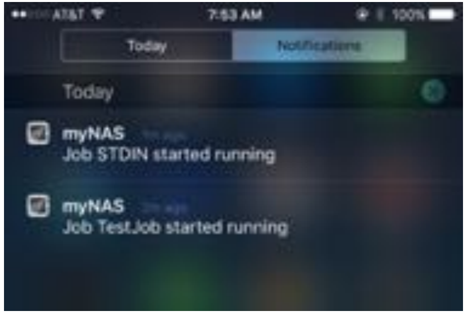

Did You Know …

- **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.

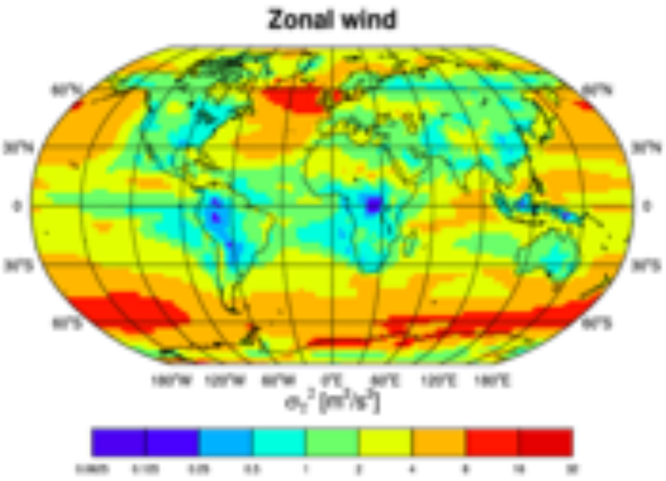

Climate models are the main tools used to simulate future climate conditions. However, the predictions made by these models are still subject to large uncertainties, and many computational runs with varying initial conditions are required to gauge the sensitivity of climate predictions to various model 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.

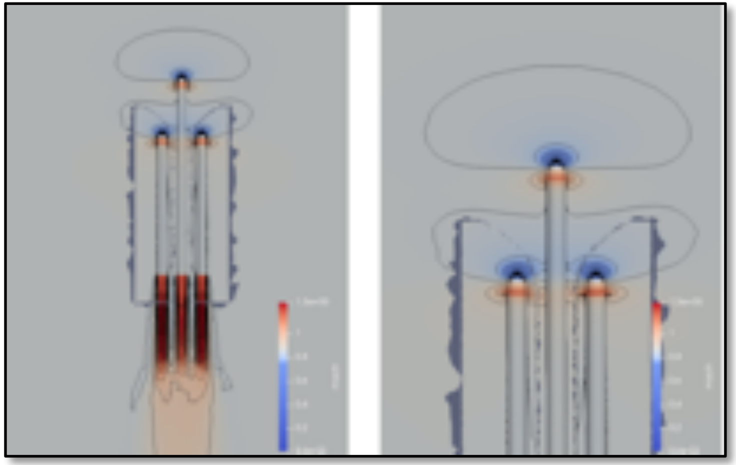

HyperSolve steady overset simulation of notional heavy-lift vehicle. The Mach contours on 3 domains are displayed on the left-hand side. Close up Mach contours are shown on the right-hand side. *Matthew O'Connell, Cameron Druyor, and Kyle Thompson, NASA/Langley*

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.

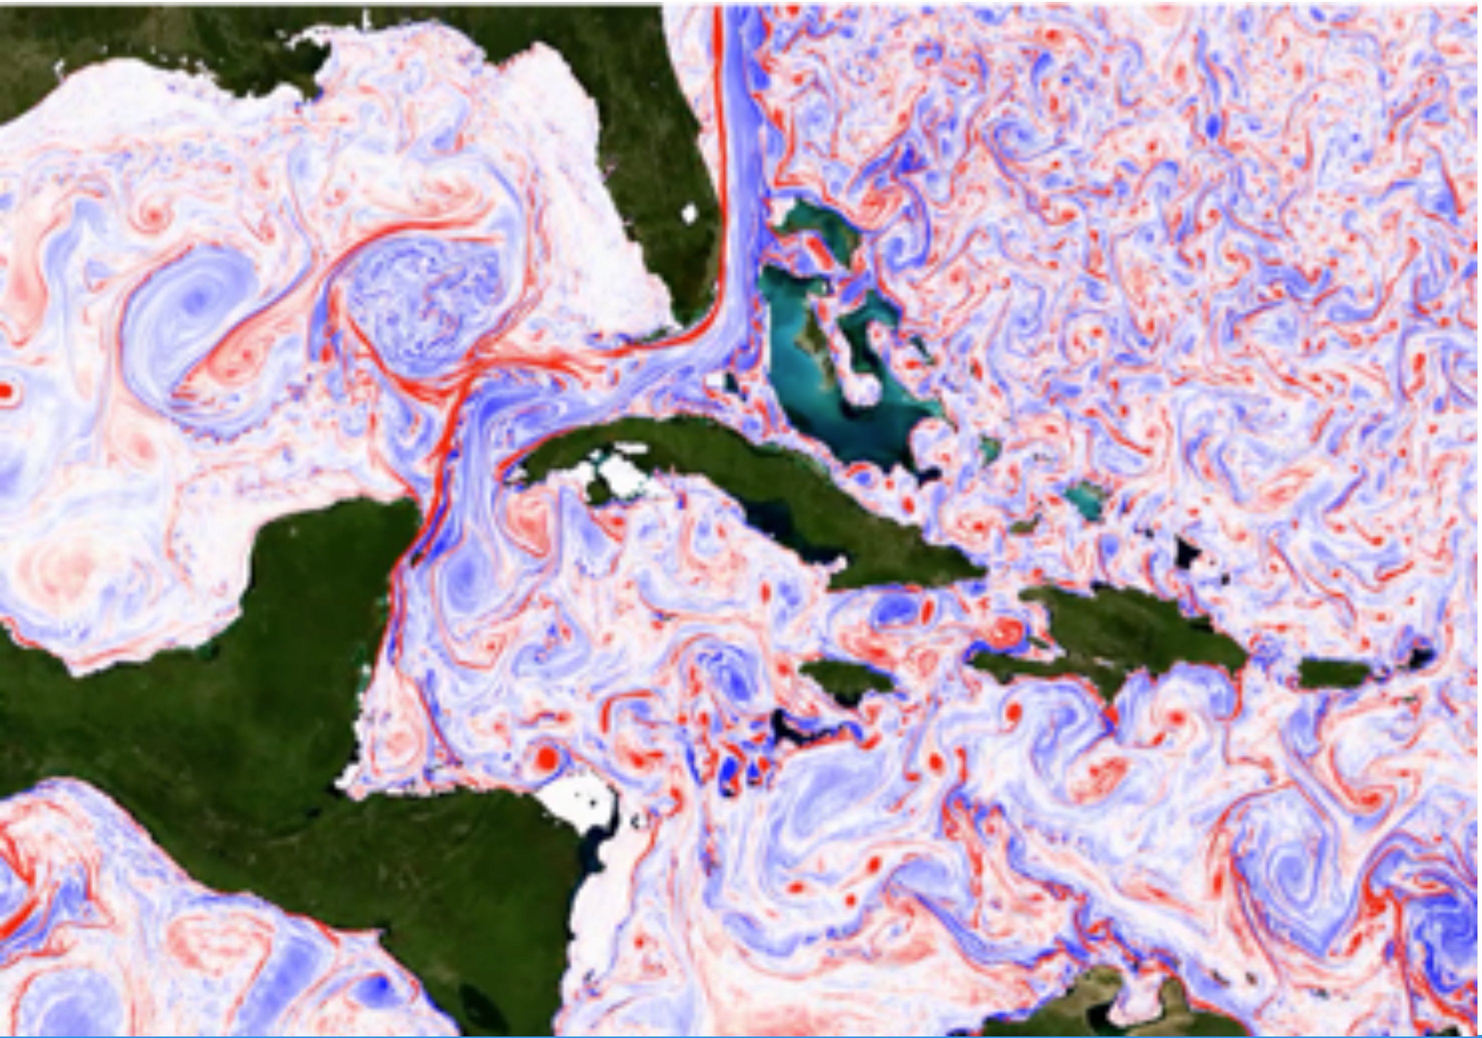

MITgcm Ocean Model:
22 billion gridpoints

Tools Developed for ECCO Data Analysis

The “multimovie” player:

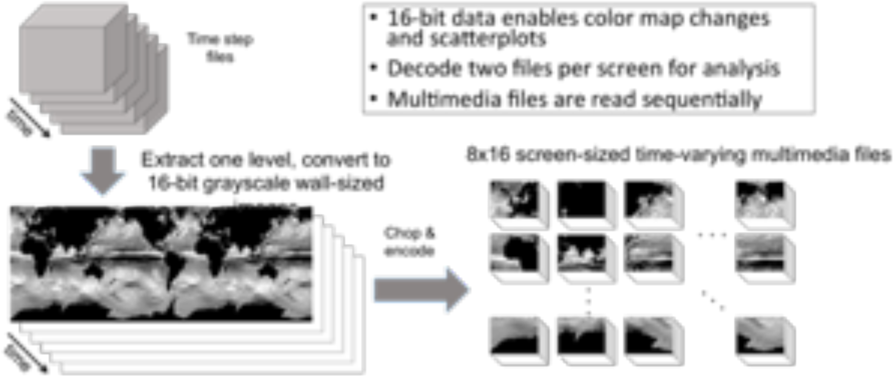

Combining 2 screens for comparisons:

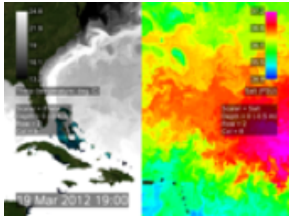

Showing 2 scalars in split screen

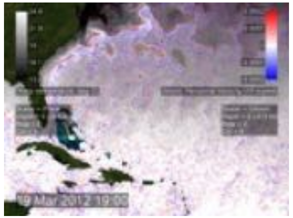

Overlaying extremes of one scalar on top of another

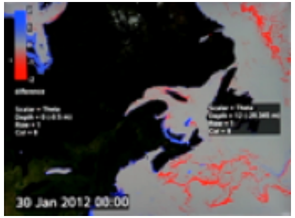

Subtracting two values, approximating a gradient

Scatterplots:

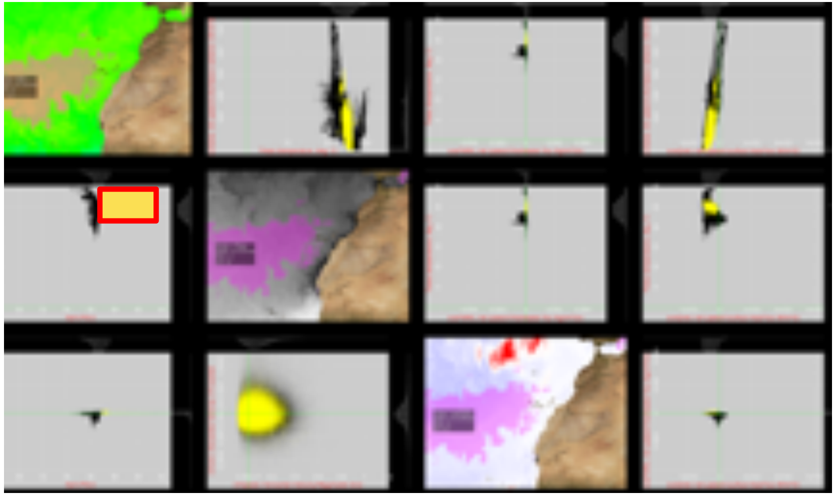

A 3 by 4 array of scatterplots with some map view displays. All of the displays show data for the same tile (geographical location), the eastern Atlantic ocean off North Africa. The red selection box is on the screen in the second row, first column. The selected points are shown in yellow on the scatterplot screens, and in light magenta on the map view screens.

These tools have also been used for analyzing Mars climate model data.

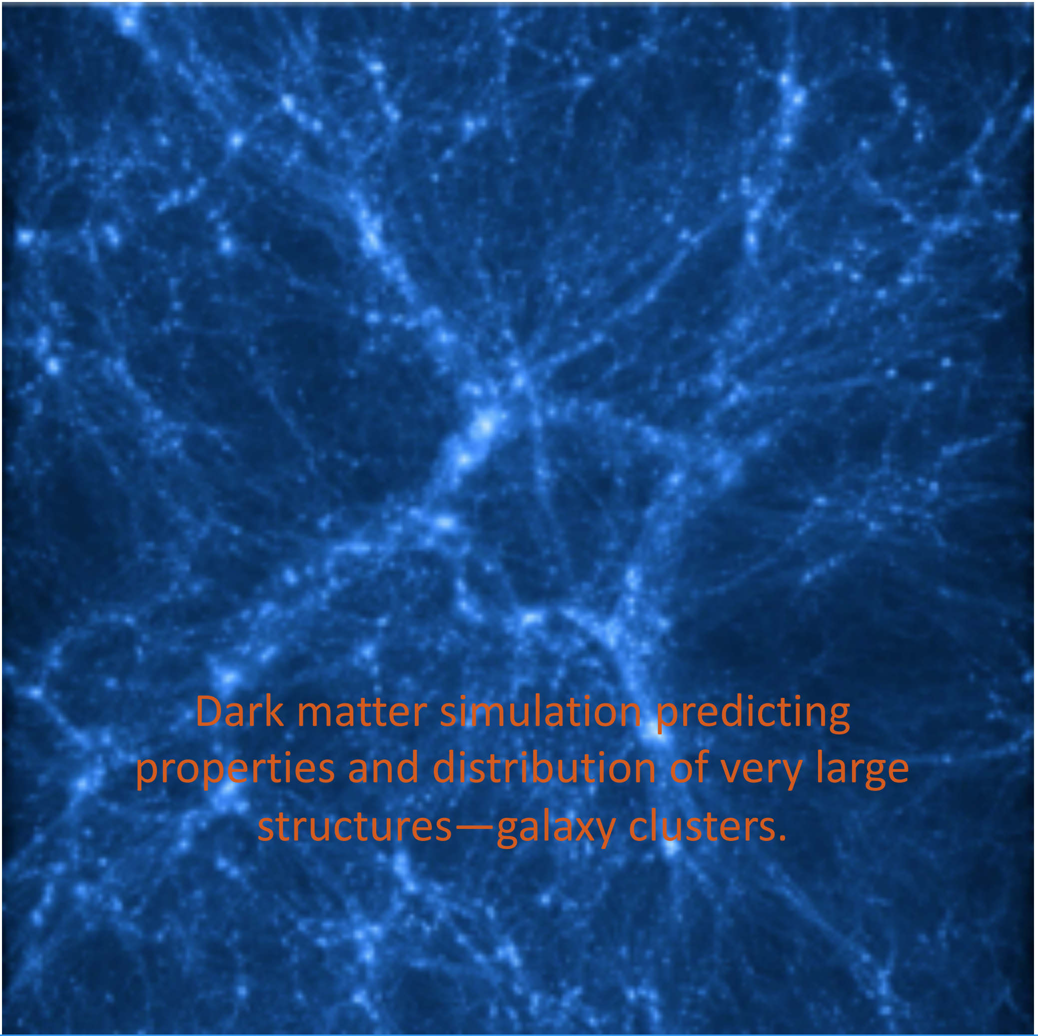

Bolshoi Cosmological Simulation: 10 billion points

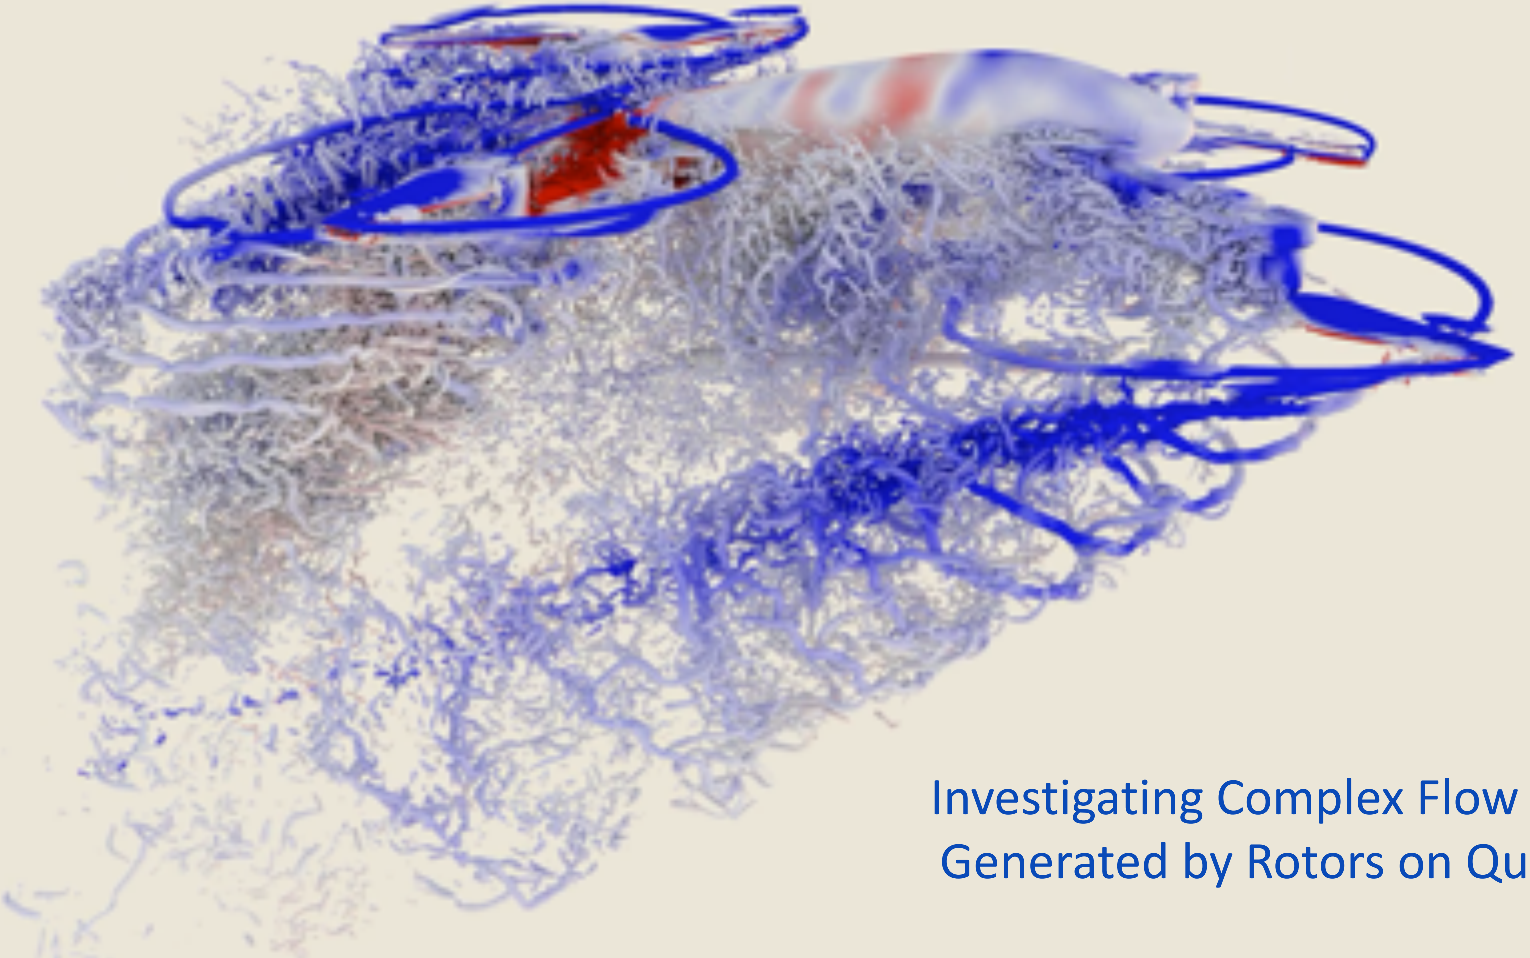

Investigating Complex Flow Structures
Generated by Rotors on Quadcoptors

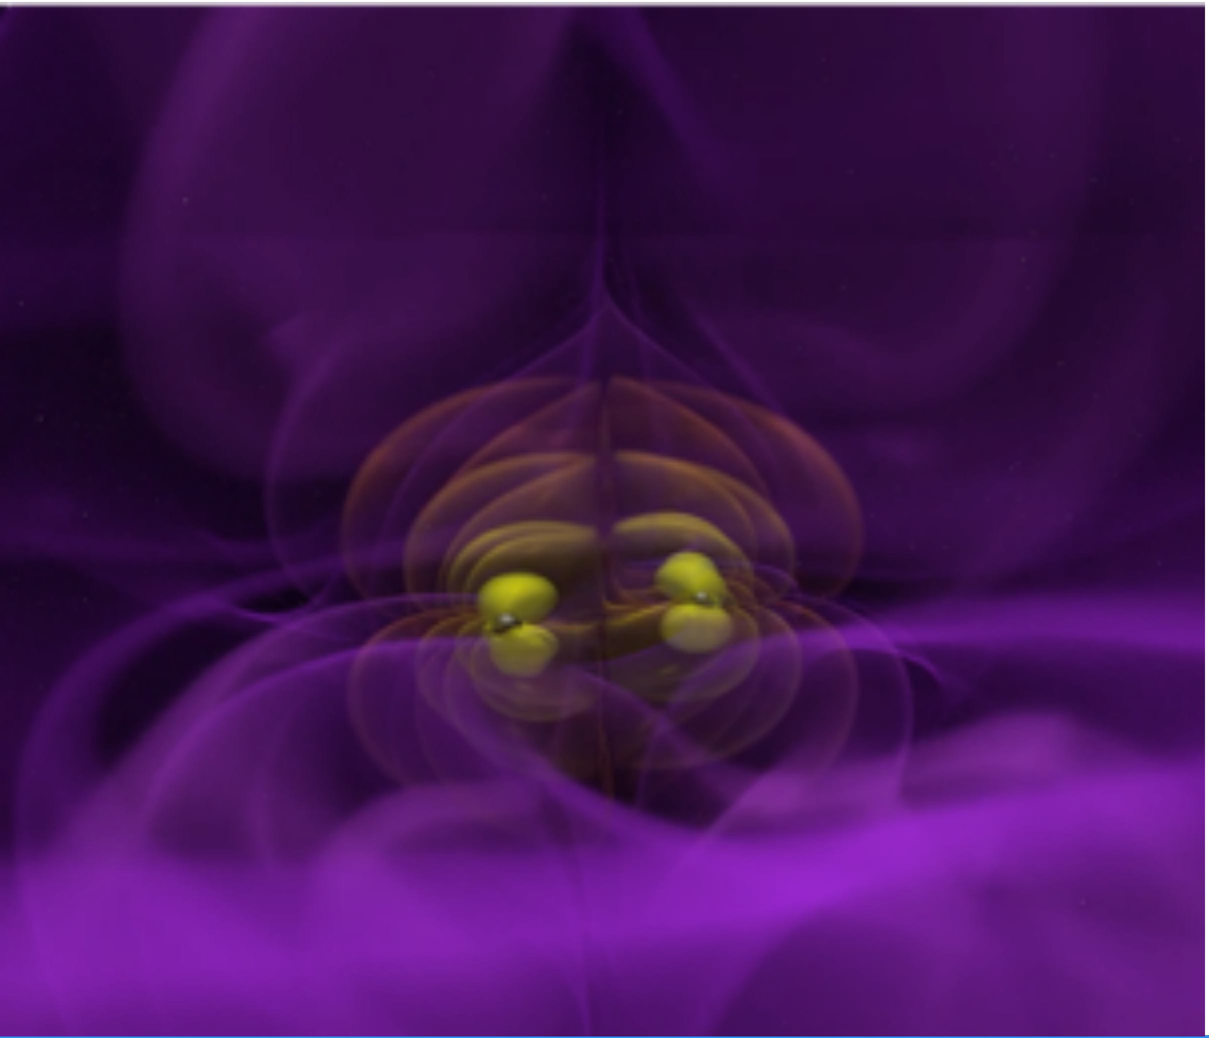

Gravitational Radiation From Spiraling Black Holes: LIGO First Detection

Particle Trace of Solid Rocket Booster Separation on Space Launch System

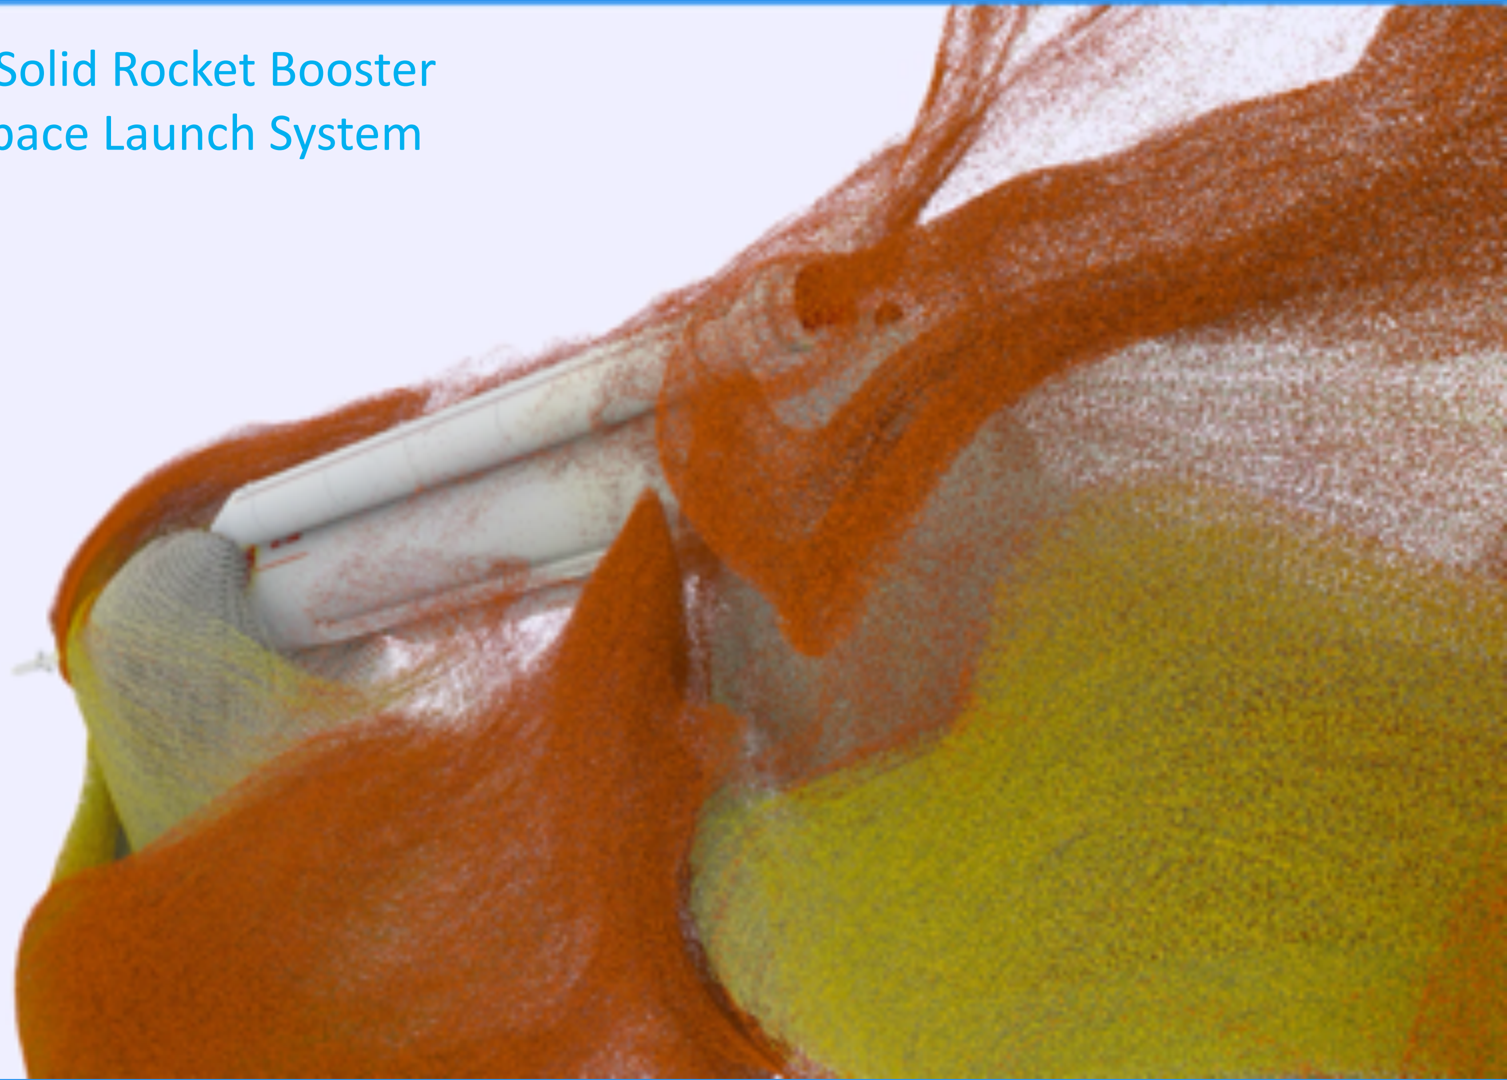

0 ' 0 "

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.

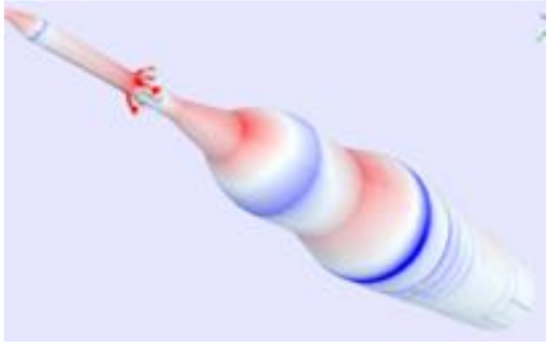

Rendering of the uPSP results, showing local pressure differences on the surface of an SLS model.

We're Listening!

Ideas

Suggestions

Questions

Contact us any time

Chris Henze
chris.henze@nasa.gov
1-650-604-3959

Data Analytics/Publication and Discovery

Shubha Ranjan
Lead, Data Analytics, Publication, and Discovery
shubha.ranjan@nasa.gov

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.

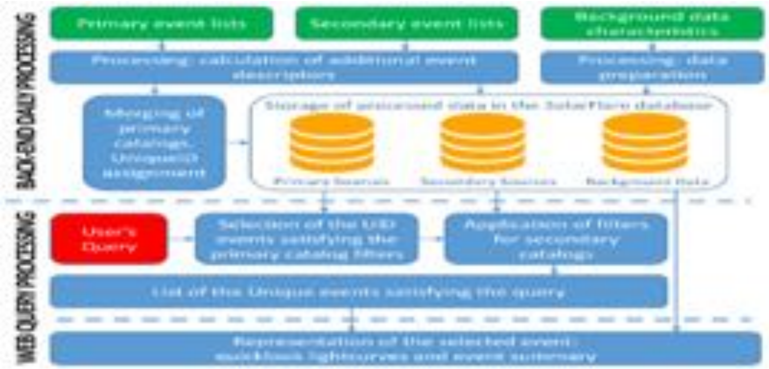

More Successes: Metrics for Data Portal Usage (Peak Downloads over 100 K)

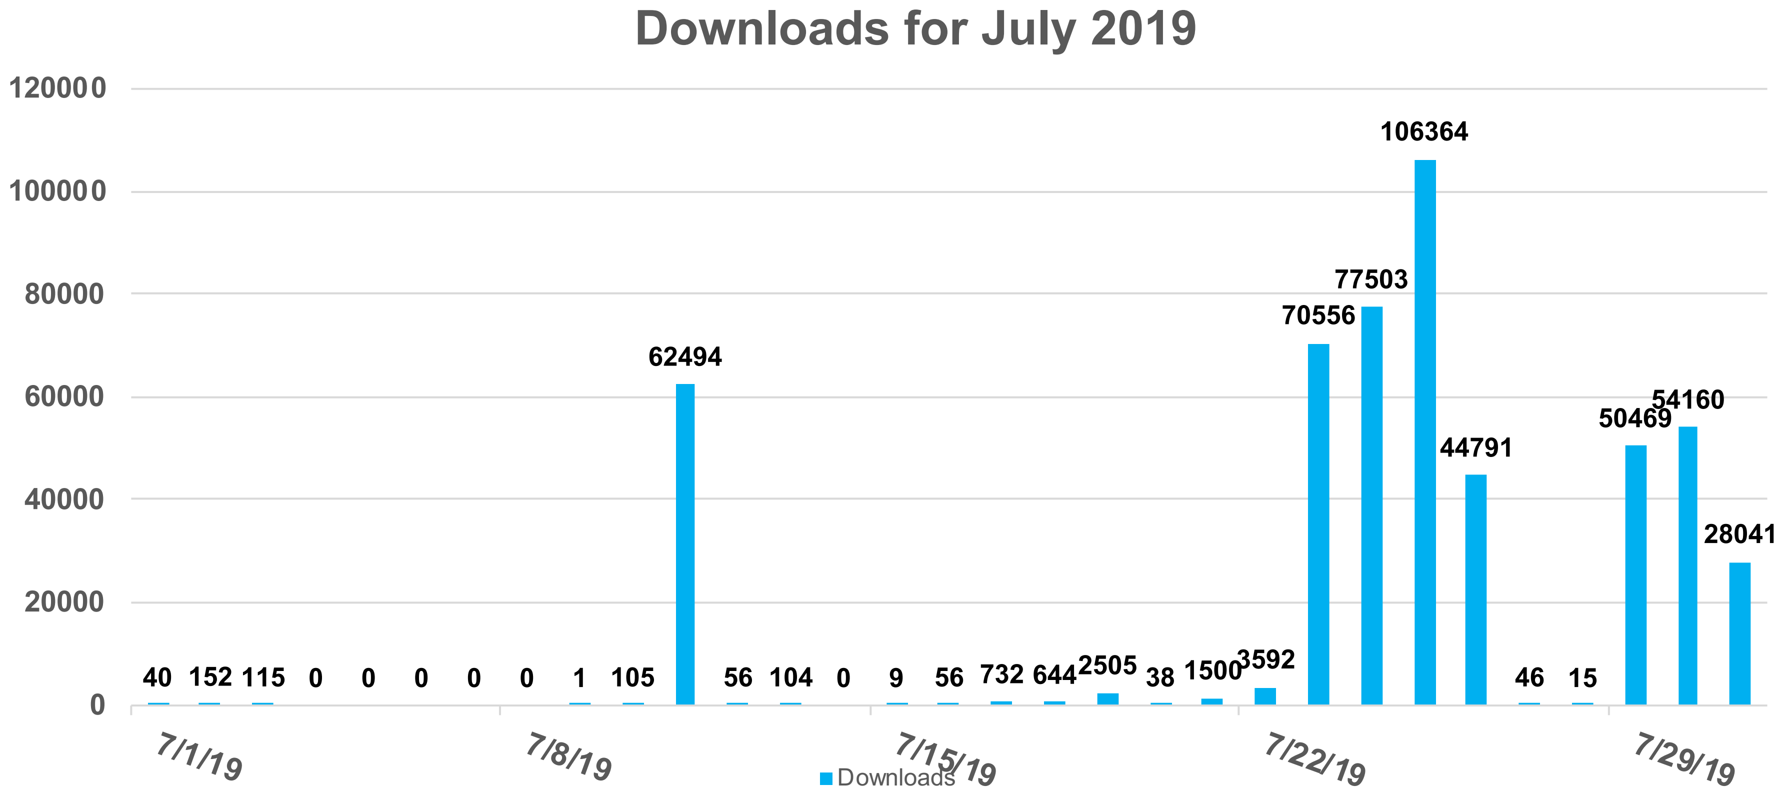

Yet More Successes: Machine Learning

- 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

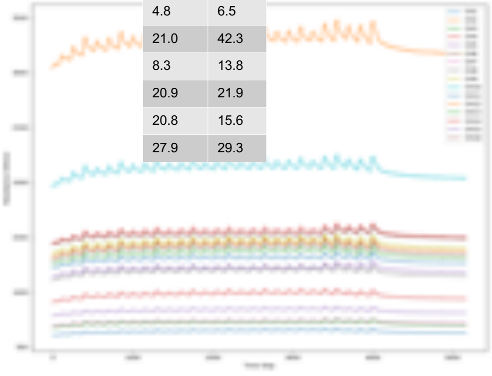

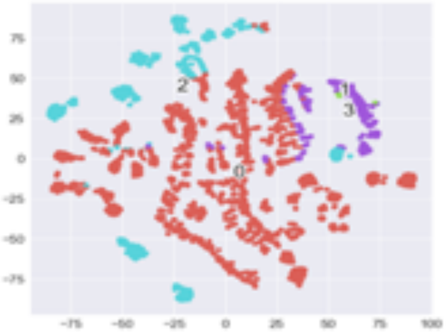

We Listened: Your Feedback Created Changes

- **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.

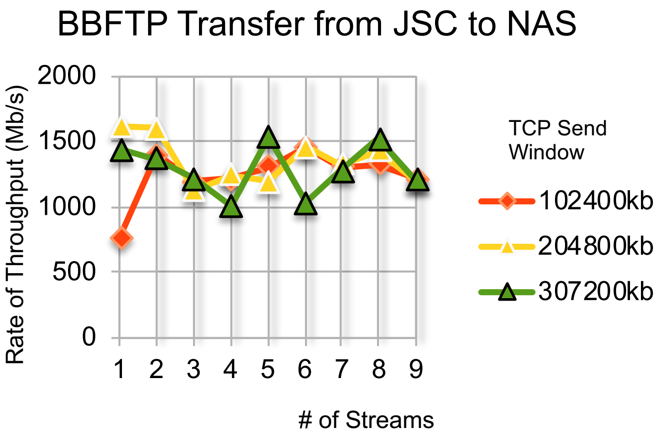

This chart shows data transfer rates (throughput) and the number of parallel streams for a single transfer between Johnson Space Center and NASA Ames using the bbFTP utility. HECC engineers achieved end-to-end performance tuning by adjusting the TCP window size (represented by each line on the graph.) 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