





National Aeronautics and Space Administration



# NASA SC19




NP-2019-10-02-ARC  
www.nasa.gov

November 18 – 21, 2019 • Denver, CO • Booth #655



## NASA BOOTH #655 DIRECTORY



RESEARCH AREA	DEMO TITLE	PRESENTER	LOCATION	PAGE
<b>AERONAUTICS</b>	Predicting Quadcopter Drone Noise Using the Lattice Boltzmann Method	Barad, Michael	H	4
	Using CFD to Develop NASA's X-57 Maxwell Flight Simulator	Duensing, Jared	D1	5
	Airframe Noise Prediction Grand Challenge: Simulating Large Airliner Landings	Khorrami, Mehdi R.	H	6
	Minimizing Sonic Boom through Simulation-Based Design: The X-59 Airplane	Nemec, Marian	S	7
	Predicting Jet Noise for Full-Scale Low-Boom Aircraft	Stich, Gerrit	H	8
	HPC Framework for Predicting High Altitude Relight in Aircraft Engines	Tang, Yihao	D3	9
	Towards Urban Air Mobility: NASA's Quadcopter Air Taxi Concept	Ventura Diaz, Patricia	H	10
<b>HUMAN SPACE FLIGHT</b>	Towards Multiphysics Prediction Capability for the KSC Launch Environment	Angel, Jordan	S	11
	Predicting Orion Pad Abort Vibrations to Keep Artemis Astronauts Safe	Cadieux, Francois	H	12
	Simulating Dream Chaser® Spacecraft Aerodynamics: Subsonic through Hypersonic	Jeffries, Michael	D2	13
	Building Booster Separation Aerodynamic Databases for Artemis II	Meeroff, Jamie	D4	14
	Revolutionizing Assessment of Unsteady Flows Using Pressure-Sensitive Paint	Roozeboom, Nettie	H	15
	Microscale Analysis of Spacecraft Heat Shields	Thornton, John	H	16

<b>OUR PLANET</b>	Fine-Scale Forest Modeling Across the North American Boreal Zone	Foster, Adrianna	D2	17
	Realistic Simulations of the Coupled Atmosphere-Ocean-Ice System	Hill, Christopher	S	18
	HPC and the Topography of the Earth's Surface	Morin, Paul	H	19
	The Future of Global Numerical Weather Prediction with GEOS	Putman, William	H	20
	Simulating Air Quality and Volcanic Plumes with GEOS CCM	Sherman, Elliot	D3	21
	Virtual World, Real Understanding: Using Virtual Reality to Visualize Earth Data	Skolnik, Shayna	D3	22
	Tree Census from Space: Quantifying Woody Biomass Using Machine Learning	Tucker, Compton	H	23
<b>SUPERCOMPUTING</b>	Teaching Your Legacy Code New (GPU) Tricks	Caplan, Ronald	D1	24
	Facilitating Science Through Centralized Storage	Carriere, Laura	D1	25
	Accelerating Science with AI and Machine Learning	Carroll, Mark	S	26
	Jumping the Queue: From NASA to the Commercial Cloud	Hood, Robert	S	27
	Exploring Bespoke Supercomputer Options for Weather and Climate Workloads	Lee, Tsengdar/ Bob Sorensen	H	28
	Achieving Quantum Supremacy with Noisy Intermediate-Scale Quantum Devices	Mandra, Salvatore	H	29
	High-Performance Computing in the Azure Cloud	Ormseth, Reid	D4	30
	What Could Possibly Go Wrong? – The Myth of the Smooth Install	Pfaff, Bruce	S	31
	The Evolution of NASA's High-End Computing Capabilities	Thigpen, William	H	32
<b>THE UNIVERSE</b>	Exploring the Origins of Extreme Magnetism in Red Dwarf Stars	Bice, Connor	S	33
	How Does Space Weather Affect Mercury, Earth, and Uranus?	Dong, Chuanfei	S	34
	Simulating Rocky Exoplanet Atmospheres with General Circulation Models	Faucher, Thomas	H	35
	Simulating the Water Cycle and Cloud Formation on Mars	Kahre, Melinda	H	36
	Modeling the Solar Corona to Study Sources of Space Weather Disturbances	Kitiashvili, Irina	S	37
	Discovering Distinct Cycles in the Global Solar Dynamo	Matilsky, Loren	D2	38
	Simulating the Diffuse Gas Surrounding Galaxies	Peeples, Molly	H	39
	Numerically Modeling the Weather and Climate of Mars	Richardson, Mark	D4	40
	The Milky Way and Its Neighbors: Simulating Ultra-Faint Dwarf Galaxies	Van Nest, Jordan	S	41

## AERONAUTICS

### Predicting Quadcopter Drone Noise Using the Lattice Boltzmann Method


The market for new vertical takeoff and landing vehicles, including autonomous urban air taxis and drones for applications such as package delivery, imaging, and surveillance, continues to grow rapidly. While electrification of the propulsion system and complete autonomy for operational efficiency are key pacing items, aerodynamic noise continues to be the biggest roadblock to community acceptance and widespread adoption. Novel computational fluid dynamics (CFD) approaches like the Lattice Boltzmann Method (LBM) promise to be ten times faster than traditional methods—and extremely accurate—for analyzing these types of flows. We used the Lattice Boltzmann solver within NASA’s Launch Ascent and Vehicle Aerodynamics (LAVA) solver framework to predict the aerodynamic noise generated by an isolated Straight Up Imaging (SUI) quadcopter drone in forward flight, derived from first principles. 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.

In detailed comparisons with data from a dedicated experiment, our predicted mean thrust is within 1% of the value measured experimentally. The predicted far-field sound pressure levels also compare very well with measurements. Sensitivity to grid resolution of both the predicted performance and the aerodynamic noise generated by the propeller blades are carefully documented. To the best of our knowledge, the present study represents the first successful prediction of drone noise from first principles using the Lattice Boltzmann approach. The computational efficiency of the solver and the complete absence of labor-intensive manual volume mesh generation in the workflow are key factors that make routine aeroacoustic analysis of urban air taxis and drones from first principles possible. Assessment of the LAVA LBM for quadcopter noise predictions showed the following results: (1) 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; (2) no manual CFD mesh generation was required—the model used adaptive mesh refinement (AMR) based on running statistics of pressure fluctuations and turbulent kinetic energy; (3) simulated 30-rotor rotations with 568 million cells were completed in under 125 hours, despite LBM using only Cartesian isotropic cells, which make it very expensive to resolve thin propellers; and (4) LBM is attractive because it is the only tool to demonstrate it can capture complex rotorcraft broadband wake turbulence noise so far (however, it is rather expensive for aerodynamic performance and tonal noise).

The simulations were run on hundreds of Intel Xeon Skylake cores on the Electra supercomputer at the NASA Advanced Supercomputing (NAS) facility, and several terabytes of data were generated for each calculation. Data analysis and visualization were conducted using the facility’s high-performance systems and custom in-house visualization software. High-performance computing is the key enabler that makes performing high-fidelity quadcopter noise simulations possible within useful turn-around times.

*Michael Barad, Cetin Kiris, NASA Ames Research Center  
michael.f.barad@nasa.gov, cetin.c.kiris@nasa.gov*

*Passive particles tracing the complex flow structures generated by four sets of blades spinning at 133.1, 134.9, 163.2, and 164.6 Hz. This 5% tip chord simulation of a Straight Up Imaging (SUI) quadcopter in forward flight was conducted using NASA’s high-performance LAVA Lattice-Boltzmann flow solver.  
Francois Cadieux, Timothy Sandstrom, NASA/Ames*



## AERONAUTICS


### Using CFD to Develop NASA's X-57 Maxwell Flight Simulator

In recent years, NASA has developed the X-57 “Maxwell” electric aircraft, which aims to achieve a fivefold reduction in energy consumption compared to conventional aircraft propulsion. Aeronautics research teams from NASA's Ames Research Center, Armstrong Flight Research Center, and Langley Research Center are collaborating to construct multiple aerodynamic databases to develop a stability and control model for the X-57's flight simulator. Populating this database requires simulation and analysis of multiple flight conditions, aircraft configurations, and propulsion settings. To date, 800 cases have been completed and more than 2,000 additional CFD simulations will be run using the Launch Ascent and Vehicle Aerodynamics (LAVA) flow solver developed at Ames.

Aerodynamic databases for the X-57 program are populated by performing steady-state Reynolds-Averaged Navier-Stokes (RANS) computational fluid dynamics (CFD) simulations for a variety of conditions across the flight envelope. Integrated aerodynamic loads and moments over the aircraft surface—the main quantities of interest—will help to assess stability and control characteristics of the aircraft. Multiple CFD codes have been used for this study to gain confidence and quantify uncertainty, including NASA's LAVA flow solver and Star-CCM+ commercial solver.

The key goal of this effort is to develop an accurate model of the X-57 aerodynamic performance and stability characteristics to implement in the aircraft flight simulator. The X-57 contains 12 high-lift propellers across each wing, providing distributed propulsion during takeoff and landing, and two cruise propellers at the wing tips, which are the primary source of thrust during cruise. We are producing databases for each of the propeller settings and comparing the effects to the corresponding no-power case. Accurately predicting the impacts of propulsion ensures the flight simulator will perform in a manner consistent with actual flight.

Preliminary validation studies performed prior to building the initial database demonstrated that all flow solvers used for this analysis could predict aerodynamic forces and moments within 8% of the experimental data, and within 5% across codes. Studies that addressed best practices for CFD mesh generation, turbulence model usage, and numerical scheme selection were a crucial step to success.



*Modeling aircraft propulsion efficiently and accurately is key to understanding the aerodynamic effects of NASA's X-57 “Maxwell” aircraft's high-lift and cruise propellers. Using actuator zones—a method to replicate the propulsion effects of moving propellers—balances computational cost and accuracy in the power-on CFD simulation case. In this simulation image, pressure is shown on the aircraft surface and streamwise velocity is shown on the planar slice for selected cruise propeller settings. Jared Duensing, NASA/Ames*

Many takeaways and lessons learned from this analysis will be valuable to future NASA aeronautics research. For example, a process for pre-database preparation and validation was established to reduce the risk of inaccurate CFD predictions that would waste computing resources. We anticipate that this process can be applied to many future aeronautics applications that require the simulation of a large number of conditions. This procedure, as well as recently developed pre- and post-processing tools, will likely reduce the amount of overhead and computational cost required for future NASA programs that require database development.

Large-scale CFD simulations like these cannot be performed without the use of high-performance parallel computing resources. All of the database simulations completed so far for this project were run across 1,200 Intel Ivy Bridge cores on NASA's Pleiades supercomputer, with each case requiring 16 hours of wall time—a net usage of 16 million core hours for all simulations run to date.

*Jared Duensing, Cetin Kiris, NASA Ames Research Center  
jared.c.duensing@nasa.gov, cetin.c.kiris@nasa.gov*

## AERONAUTICS

### Airframe Noise Prediction Grand Challenge: Simulating Large Airliner Landings

Reduction of aircraft noise is a high-priority goal of NASA's Aeronautics Research Mission Directorate. Airframe noise contributes significantly to the total noise generated by modern aircraft during approach and landing. 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.

Physically accurate prediction of airframe noise for a full-scale, large airliner with high-lift devices and undercarriage deployed is a grand challenge for the aerospace community. Such simulations were deemed unattainable just a few years ago. A simulation methodology we used previously to generate high-fidelity, aeroacoustic data for a full-scale, isolated Boeing 777 nose landing configuration.

In partnership with the Boeing Company to study airframe noise, we are using Dassault Systemes's Lattice Boltzmann code PowerFLOW to investigate the complex, unsteady flow field around a full-scale, large civil transport to obtain a greater understanding of the prominent noise-generation mechanisms and locate the major sources.

For large, twin-aisle civil airliners, the undercarriage and wing high-lift devices are the dominant sources of airframe noise. The extreme geometrical complexity of these aircraft during landing generates noise-producing flow features that contain a broad range of spatio-temporal scales. Physics-based prediction of landing-gear noise requires a fundamental understanding of the time-dependent flow field produced by each component and the resulting interactions among these flow fields. To gain insight into these phenomena, we are running simulations on the Pleiades supercomputer at the NASA Advanced Supercomputing (NAS) facility. Our goals are: to accurately compute the far-field noise signature of a full-scale, twin-aisle airliner in landing configuration; and to validate the predicted results against measured acoustic data obtained from flight tests of the same aircraft.

These simulations were run on low-resolution meshes to establish best practices for achieving the numerical, spatial, and temporal resolution required to properly capture the key noise-producing flow features. We are conducting ongoing, preliminary comparisons of the predicted levels and frequency content of the noise spectrum on the ground with flight test data, which so far indicate that good agreement will be achieved. These encouraging results bode well for the suitability of our current simulation approach to solving this airframe noise grand challenge.

High-fidelity airframe noise simulations of the complete airliner require a vast number of calculations that can only be performed on supercomputers. Our current low-resolution simulations required 14 billion cells to discretize the volume surrounding the complete aircraft with landing gear and wing high-lift devices deployed. The simulation was run on 6,000 Pleiades cores and used over 1.4 million processor hours. To compute the far-field noise spectrum, the simulation generated over 50 terabytes of data. The combined capabilities of Pleiades and the post-processing expertise available at the NAS facility have enabled high-resolution definition and visualization of the prominent flow features over a broad range of spatial-temporal scales.

To capture the high-frequency segment of the far-field noise spectrum, we will perform the simulations on a finer spatial grid. We anticipate that these fine-resolution simulations will comprise 28 billion cells and require 7.5 million processor hours, distributed over 10,000 to 12,000 Pleiades cores.

*Mehdi Khorrami, NASA Langley Research Center  
mehdi.r.khorrami@nasa.gov*



*The fluctuating surface pressure resulting from the flow field around a large airliner produces broadband sound that propagates to the ground. Red denotes positive pressure fluctuations and blue denotes negative pressure fluctuations. Benedikt Koenig, Dassault Systemes; Patrick Moran, NASA/Ames*


## AERONAUTICS

### Minimizing Sonic Boom through Simulation-Based Design: The X-59 Airplane

One of NASA's six Strategic Thrusts for the agency's Aeronautics Research Mission Directorate is "Innovation in Commercial Supersonic Aircraft," with a vision of fast air travel available to a broad spectrum of the traveling public. Unlike the Concorde, future supersonic aircraft will be efficient, affordable, and environmentally responsible, generating an acceptable level of en route noise (sonic booms). The first major step toward this goal is the ongoing construction of the new X-59 Quiet SuperSonic Technology X-plane to demonstrate technologies that reduce the loudness of sonic booms to gentle thumps. The X-59 will fly over select communities to gather data on human responses to low-boom flights. This data will be delivered to U.S. and international regulators to help set certification standards for supersonic commercial aircraft and enable supersonic overland flights.

To achieve low-boom performance, we used high-fidelity computational fluid dynamics (CFD) simulations extensively to shape the X-59. While performance optimization of traditional aircraft primarily involves maximizing the lift-to-drag ratio, prediction and minimization of sonic boom noise is an additional requirement that is unique to commercial supersonic flight. In a typical simulation, a nearfield pressure signature is extracted from the simulated flowfield and propagated from a flight altitude of around 55,000 feet through the atmosphere to the ground to obtain the sonic boom noise carpet. NASA's production-level Cart3D simulation package for CFD has been used extensively to determine the pressure field near the aircraft and evaluate the ground noise carpet of each major design evolution of the X-59. Cart3D is coupled with both an atmospheric propagation solver to estimate the noise level on the ground, and uncertainty quantification tools to provide uncertainty estimates in the pressure signatures due to variations in the aircraft's operating conditions and configuration.

These Cart3D simulations have contributed to various design improvements throughout the project, 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; most recently we used it to verify the design of a new wind-tunnel strut that holds the aircraft model in position but does not interfere with its pressure signature. A unique capability of this CFD simulation package is its automatic mesh refinement capability, which minimizes numerical errors in the simulations. This not only ensures automatic verification of Cart3D simulations, but is also used to guide meshing requirements for other CFD solvers used in NASA's X-59 and Commercial Supersonic Technology projects.



*Computational schlieren image (similar to those produced by the schlieren technique in experiments) from a high-resolution Cart3D simulation. The dark and bright regions represent shockwaves and expansions, respectively. The complex shock system of the X-59 is shown, including the conical shocks behind the aircraft. The simulation includes propulsion effects, air probes, and secondary air system ducts. Significantly 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*

Our simulations are run on the Pleiades, Electra, and Endeavour supercomputers at the NASA Advanced Supercomputing facility. The large capacities of Pleiades and Electra are critical for handling the thousands of cases required for aerodynamic performance databases, shape optimization, and uncertainty quantification of the X-plane. We use the Electra system for most production runs, where a typical case turnaround is roughly an hour on several hundred cores. On the shared-memory Endeavour system, our fine-mesh simulations complete in about 8 hours on 512 cores.

*Marian Nemec, Michael Aftosmis, NASA Ames Research Center  
marian.nemec@nasa.gov, michael.aftosmis@nasa.gov*

## AERONAUTICS

### Predicting Jet Noise for Full-Scale Low-Boom Aircraft

NASA's supersonic vehicle research addresses the development of tools, technologies, and knowledge to help eliminate today's technical barriers to practical commercial supersonic flight. Although the major focus of this development has been on reducing the ground signature at supersonic cruise, in order for an airplane to be certified it must also satisfy noise constraints during takeoff and landing when it is traveling at subsonic speeds. The dominant source of sound during this phase is jet noise. Computational aeroacoustic simulation tools can be used to assess new designs for NASA's Low-Boom Flight Demonstrator aircraft at these lower speeds. We are further scrutinizing both modeling and simulation practices and the use of hybrid Reynolds-averaged Navier Stokes/Large-Eddy Simulations (RANS/LES) methods, and are developing new methods to predict jet noise—including surface interaction noise—on a full flight configuration airplane.

While great advancements have been made toward predicting jet noise with computational fluid dynamics (CFD) solvers—including high-fidelity simulations at flight Reynolds numbers and some modeling of simple but important geometric features, such as jet surface interaction noise—most simulations rely on high-fidelity, wall-modeled LES methods and are still limited to simple geometries due to the complexity of the problem. However, the scale-resolving simulations needed to predict accurate far-field noise are still too computationally demanding to investigate full airframe interaction noise using wall-modeled LES.

A possible strategy to achieve full-scale airframe simulations including jet-noise prediction is the use of a hybrid RANS/LES approach, in which most parts of the aircraft are simulated using unsteady RANS, while critical areas where most of the noise is created are simulated using LES. As a step toward predicting jet acoustics for the Low-Boom Flight Demonstrator, scientists at NASA's Ames Research Center performed scale-resolving simulations with the Launch, Ascent and Vehicle Aerodynamics (LAVA) solver scrutinizing hybrid RANS/LES methods on validation cases with increasing complexity, including jet surface interaction noise and chevron nozzle designs.

By improving zonal hybrid RANS/LES capabilities (funded by NASA's Transformational Tools and Technologies (TTT) project) within the LAVA solver framework, we achieved a successful prediction of jet noise as well as jet surface interaction noise within NASA's Commercial Supersonic Transport (CST) program. Findings from these simulations have already resulted in new insights and additional experimental investigations of certain flow phenomena. Gaining confidence in the predictive capabilities of hybrid RANS/LES for jet noise is another step toward vehicle certification by analysis.



*Illustration of NASA's planned Low-Boom Flight Demonstrator aircraft. The close proximity of the tail cone surfaces and jet motivated the work of jet surface interaction noise prediction (for example, jet scrubbing noise) with the use of hybrid Reynolds-averaged Navier Stokes/Large-Eddy Simulation methods. NASA, Lockheed Martin*

The Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility are instrumental in enabling these next-level, large-scale computational aeroacoustics simulations over hundreds of nodes and thousands of cores. Each simulation generates hundreds of terabytes of data that is then analyzed in the frequency domain, which requires parallelized post-processing tools that reduce the cost of accessing data on disk. Raw computational power is not enough—large and efficient distributed filesystems are also critical. In addition, NAS's visualization services help us go beyond what traditional data analysis tools offer.

*Gerrit-Daniel Stich, Cetin Kiris, NASA Ames Research Center  
gerrit-daniel.stich@nasa.gov, cetin.c.kiris@nasa.gov*

## AERONAUTICS

### HPC Framework for Predicting High-Altitude Relight in Aircraft Engines


When aircraft operate at high-altitude conditions, low ambient air pressures and temperatures can lead to engine flame-out situations and hamper relight attempts. In high-altitude relight, an aircraft engine must ignite within a certain period after its initial flameout. To accomplish this, one or more ignition kernels—pockets of high-energy flow and free radicals—are introduced into the combustion chamber by firing the engine ignitor. The kernel develops into a flame front, and eventually reaches flame stabilization or extinction, depending on the kernel's initial state and the turbulent flow evolution. There are two main challenges for predicting such events: conflicts between simulation fidelity and computational capacity, particularly with strong turbulent flow and realistic jet fuels; and the underlying stochasticity of such a problem due to turbulence and the initial spark discharge.

We have developed a comprehensive modeling framework to handle these challenges. The framework includes: a tabulated chemistry model that models non-premixed forced ignition of realistic jet fuels, incorporated into a turbulent modeling framework of large-eddy simulation (LES), using a presumed probability density function approach; and an uncertainty quantification approach, based on Monte Carlo simulations using the tabulated chemistry model, to estimate the statistical property of engine relight problems—the ignition probability.

Using this framework, we performed a series of LES simulations to generate enough samples to represent the statistical effects of aleatoric parameters (e.g., turbulent flow field and the initial kernel energy) on the ignition outcome. We then used the simulation outputs to reconstruct the ignition probability using our uncertainty quantification strategy. In a typical simulation of a model combustor representing high-altitude relight, two main processes were analyzed: upstream fuel mixing, and downstream non-premixed spark ignition. Upstream, the inflow is divided into the upper and lower flow by a splitter plate; only the upper flow is fuel-mixed, while the bottom flow remains pure air. The two flows start to mix after entering the downstream chamber, where a commercial aircraft ignitor is mounted on the floor plane. With each spark discharge, a kernel is ejected from the ignitor top surface into the combustor chamber, like a pulse-jet-in-crossflow problem.

The purpose of this layout is to mimic the non-premixed, turbulent, stochastic nature of altitude relight problems. In particular, the kernel needs to transit through the non-reacting bottom flow before mixing with the fuel-seeded upper flow to ignite, while two important physical effects are competing with each other: the entrainment of fuel and the dissipation of kernel energy. Eventually, ignition succeeds or fails depending on the initial kernel status (such as size and ejection speed) and the turbulent flow conditions (such as fuel mixing and strain rate).

Results show that our model can qualitatively and quantitatively predict non-premixed turbulent forced-ignition processes for a range of different operating conditions. In addition, the LES tabulation-based modeling framework provides the capability to efficiently generate abundant high-fidelity data of turbulent forced-ignition processes. This data can be applied as input for data-driven science techniques used to study patterns of altitude relight problems, benefitting future engine design.



*The pulse-jet-in-crossflow effect of the ignition kernel is visualized from a cross-section of the combustor. The shape of the flame front (yellow) follows those flow streak lines (red) integrated from the initial kernel ejection boundary. To capture such complex flow dynamics requires high-fidelity simulation of the turbulent flow field, which is indispensable to accurately predict the ignition process. Yihao Tang, University of Michigan; Malik Hassanaly, National Renewable Energy Laboratory*

More than 800 simulations were performed for this project, using approximately one million core hours on NASA's Pleiades supercomputer. High-performance computing is essential for such applications, for both the compute capability and for storing and transferring large numbers of datasets.

*Yihao Tang, Venkat Raman, University of Michigan  
yhtang@umich.edu, ramanvr@umich.edu*

## AERONAUTICS

### Towards Urban Air Mobility: NASA's Quadcopter Air Taxi Concept

With many cities around the U.S. growing rapidly, their transportation systems are being pushed to the limits. Urban Air Mobility (UAM) is envisioned to be the future air transportation system over populated areas, where everything from small package-delivery drones to passenger-carrying air taxis are able to interact safely and efficiently. UAM is the new era in air travel that NASA is working on right now.

The capacity of multi-rotor vehicles to perform vertical take-off and landing (VTOL), together with their great maneuverability, make them an excellent choice for UAM aircraft. The accurate prediction of multi-rotor vehicles' performance and acoustics is very challenging, due to the unsteady and complex flows, as well as the aerodynamic interactions of the rotors. High-fidelity computational fluid dynamics (CFD) simulations play a key role in the development of next-generation air vehicles.

In this work, we simulate NASA's quadcopter UAM concept vehicle using high-fidelity CFD, within the scope of NASA's Revolutionary Vertical Lift Technology (RVLT) project. One of the goals of RVLT is to develop and validate advanced high-fidelity CFD tools with improved physical realism and numerical solution accuracy that can be used in the rotorcraft design process.


The quadcopter and surrounding air are modeled using hundreds of millions of grid points with overset volume grids. The flow is solved using NASA's OVERFLOW CFD code, a high-order accurate Navier-Stokes solver. Then, the rotorcraft trim and rotor blade motions are predicted with the comprehensive rotorcraft code, CAMRAD II. The final solution is obtained by loosely coupling OVERFLOW and CAMRAD II.

The quadcopter air taxi is one of NASA's concept vehicles, intended to focus and guide NASA research activities in support of aircraft development for the emerging UAM market. Advanced CFD tools are being developed to simulate various UAM concepts from NASA. Past work showed that side-by-side rotors increase the cruise performance compared to non-overlapped twin rotors. In this work, we study the performance for a six-passenger quadcopter vehicle. The rotor-rotor and rotor-airframe interactions were analyzed, resulting in a new design where the rear rotors are placed above the front rotors, reducing the interactions while keeping a compact configuration. The advanced CFD tools we are developing will assist in the rotorcraft design process for future multi-rotor UAM vehicles. Coupling these tools with the comprehensive CAMRAD II code provides accurate rotor trim and motion for UAM vehicles.

The performance of the quadcopter vehicle can only be calculated accurately by solving the fluid equations on overset grids, comprising hundreds of millions of grid points. The rotation of the blades makes the problem unsteady, and the quadcopter consists of four three-bladed rotors—making the problem even more complex. In addition, the blades design allows for flapping (adding relative in-plane motion). Only with the most powerful supercomputers can such a complex simulation be solved in just a few days.

The simulations were carried out on NASA's Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility, using 1,200 to 3,000 Intel Xeon Broadwell, Ivy Bridge, and Skylake processors. In addition, NAS's mass storage system is essential to support this work—the data stored for post-processing a single case can take up to 40 terabytes of memory.

*Patricia Ventura Diaz, Seokkwan Yoon, NASA Ames Research Center  
patricia.venturadiaz@nasa.gov, s.yoon@nasa.gov*



*Visualization of the flow of NASA's six-passenger quadcopter concept for urban air mobility (UAM), in edgewise flight (front-oblique view). The quadcopter vehicle design reduces rotor-rotor interactions between the front and rear rotors, while keeping an efficient and compact configuration. The background shows the pressure oscillations, where white is high and black is low. This image reveals the complexity of the flow for a multi-rotor configuration, where many rotors interact with each other and the different components.  
Patricia Ventura Diaz, NASA/Ames*

## HUMAN SPACE FLIGHT

### Towards Multiphysics Prediction Capability for the KSC Launch Environment

When engineers at NASA's Kennedy Space Center (KSC) redesigned and rebuilt the main flame deflector (MFD) at KSC's Launch Complex 39B, a team of experts from the NASA Advanced Supercomputing (NAS) Division was called in to apply high-resolution computational fluid dynamics (CFD) to help identify thermal, pressure, and flow environments on and around the geometrically complex MFD. The NAS Division's Launch Ascent and Vehicle Analysis (LAVA) team investigated both ignition overpressure waves and thermal conjugate heat transfer simulations of the MFD, the mobile launcher, and the tower, including umbilical hardware and other key components of the launch pad.

The LAVA team is currently working to improve the fidelity of these launch environment simulations by including the Ignition Overpressure Protection and Sound Suppression (IOP/SS) system. This system delivers almost a half-million gallons of water to the mobile launcher and MFD to suppress strong acoustic waves that could damage the structure or launch vehicle. Incorporating water into the simulations poses difficulties, in part due to massive variations in density and the presence of powerful shock waves. The LAVA team is in the early stages of developing multiphase capabilities that will be utilized in future launch simulations.


For the MFD redesign project, the LAVA team completed large-scale CFD simulations of the highly detailed, as-built pad geometry, which required the development of appropriate boundary conditions, meshing, and lift-off information. The simulations were run with the LAVA code and post-processed using statistical tools from the LAVA framework and high-end custom visualization techniques that were also developed at NAS. The extremely complex geometric topologies were easily handled using the high-order accurate LAVA immersed boundary Cartesian Navier-Stokes approach. The simulations are some of the largest conducted at NAS—the LAVA Cartesian simulation contained over 555 million cells and over 500,000 timesteps, and utilized more than 4,000 cores on NASA's Electra supercomputer.

For the current project, the driving force is the need to incorporate relevant physics into the simulations to improve confidence in safety and mission success. The results of these simulations, obtained with advanced post-processing, will be used to evaluate the ability of the IOP/SS system to prevent violent acoustic waves that could damage the launch structure or vehicle. Incorporating the water deluge system into the new simulations requires modifications to the governing equations and implementation of new numerical techniques. Analysis of both thermal and pressure environments will be conducted and compared with the previous simulations that do not include the IOP/SS system.

The results of the launch environment simulations were used to guide project engineers in the MFD redesign. The new multiphase results will give a more complete picture of the extreme conditions the launch system experiences and will help NASA reduce mission risk—averting issues and potentially saving significant amounts of time and money. This collaboration also represents a great opportunity to advance NASA Ames Research Center's key expertise in the launch acoustics area, and to further improve the LAVA software.

The high-resolution simulations of the launch environment required approximately five weeks of simulation time on Electra using up to 4,000 Intel Xeon Skylake cores. With the added complexity of modeling the IOP/SS system, the need for powerful computational resources also grows. Because these simulations are time-dependent, they generate hundreds of terabytes of data for each calculation.

*Jordan B. Angel, Cetin Kiris, NASA Ames Research Center  
jordan.b.angel@nasa.gov, cetin.c.kiris@nasa.gov*



*Snapshot from a simulation of launch ignition for NASA's next-generation Space Launch System. Surfaces are colored by pressure (red is high; blue is low) while particles are colored by temperature (orange is hot; black is cooler).  
Michael F. Barad, Timothy Sandstrom, NASA/Ames*

## HUMAN SPACE FLIGHT


### Predicting Orion Pad Abort Vibrations to Keep Artemis Astronauts Safe

NASA's Artemis Program aims to transport humans to interplanetary destinations such as the Moon—and eventually Mars—and return them safely back to Earth via the Orion crew module. During launch, Orion will sit atop the Space Launch System (SLS), a rocket more powerful than the Saturn V that propelled the Apollo capsules into space. In the event of a problem during launch, Orion's Launch Abort System (LAS) is designed to carry the spacecraft and its crew to safety, away from the rocket. The LAS is positioned on top of Orion during takeoff, and jettisoned once it has safely reached low-Earth orbit.

In order to propel the crew module away to safety, the LAS uses an abort motor that produces four large high-speed exhaust plumes that flow along the sides of Orion, generating extremely strong acoustic vibrations. These vibrations need to be carefully analyzed to ensure the safety of the astronauts. To better understand the effects of acceleration on the strength of the vibrations generated by the plumes on the surface of the LAS, we produced a detailed turbulence-resolving simulation of the Pad Abort 1 (PA-1) flight test conducted in 2011, where the Orion LAS accelerates to over ten times the gravity of Earth.

The Launch, Ascent, and Vehicle Aerodynamics (LAVA) team at NASA's Ames Research Center has been closely collaborating with the Orion Loads and Dynamics team at Johnson Space Center since 2017 to help characterize the vibrations imparted by the LAS abort motor plumes onto the LAS structure. During this time, the team has run multiple simulations for various launch abort scenarios using the cutting-edge LAVA computational fluid dynamics (CFD) software. To investigate the relationship between vehicle acceleration and unsteady loads on the LAS, and to validate the LAVA predictions with flight test data, we simulated the PA-1 flight test and compared it to a simulation where the vehicle is at rest during the abort sequence.

The simulation of the accelerating PA-1 flight test is unprecedented in scale and scope, and it demonstrates a new capability for NASA. It is the first turbulence-resolving simulation of an accelerating pad abort scenario with enough fidelity to accurately predict the acoustic vibrations on the surface of the LAS when compared to the flight test data—and it was completed in a short-enough turnaround time to affect engineering decisions. Results have already contributed to a better understanding of the effects of acceleration and angle of attack on the acoustic environment. The success of this simulation also opens the door to more analysis that can help reduce uncertainty for abort scenarios that are difficult or too expensive to test, further helping to reduce risk and ensure the safety of astronauts when using the LAS.



*Snapshot of the Pad Abort 1 flight test simulation showing the detailed turbulent exhaust plume physics captured. Density (red is low, blue is high) is shown on the vehicle surface and a plane cutting through two abort motor nozzles. The hot, high-velocity exhaust gas has a lower density than air when it exits the nozzle. Its difference in speed with respect to the slower moving air around the vehicle creates turbulent eddies that result in pressure fluctuations on the vehicle surface. Michael Barad, Tim Sandstrom, NASA/Ames*

The PA-1 simulation was performed with an adapted Cartesian mesh with 600 million cells, and completed for nearly 572,000 time steps corresponding to 1.25 seconds of real time. It made full use of 400 Skylake nodes, corresponding to 16,000 cores, via hybrid MPI-OpenMP parallelism on NASA's Electra supercomputer for 1.5 months, generating nearly 200 terabytes of data. The NASA Advanced Supercomputing Division's visualization experts provided passive particle renderings that helped enrich our understanding of the fluid dynamics of the launch abort motor plumes.

*Francois Cadieux, Cetin Kiris, NASA Ames Research Center  
francois.cadieux@nasa.gov, cetin.c.kiris@nasa.gov*

## HUMAN SPACE FLIGHT

### Simulating Dream Chaser® Spacecraft Aerodynamics: Subsonic through Hypersonic

Sierra Nevada Corporation's (SNC) reusable Dream Chaser spacecraft will transport cargo to the International Space Station starting in fall 2021, as part of NASA's Commercial Resupply Services missions. During entry, descent, and landing, Dream Chaser will fly through Mach numbers ranging from above 20, about 15,224 miles per hour to less than 0.4—roughly 304 mph—as it lands on a runway. Vital to the success of Dream Chaser is the accurate characterization of aerodynamic forces on the airframe and control surfaces during atmospheric flight. Additional aerodynamic interactions occur at hypersonic velocities due to Reaction Control System (RCS) firings. Predictions of aerodynamic forces and RCS aerodynamic increments were all modelled using extensive CFD.


Wind tunnels and computational fluid dynamics (CFD) methods are now complementary tools in the development of aerodynamic and aerothermal databases. Using wind tunnel data as anchoring data, we developed and expanded the Dream Chaser databases using thousands of aerodynamic and aerothermal predictions; a host of NASA-developed CFD packages, including: FUN3D; CART3D; DPLR (Data Parallel Line Relaxation); and the Direct Simulation Monte Carlo (DSMC) Analysis Code, DAC97. These simulations were run using NASA's Pleiades and Endeavour supercomputers at NASA Ames Research Center.

These CFD predictions were used heavily in the development of aerodynamic and aerothermal preflight databases. These databases are used by SNC's Guidance, Navigation and Control (GN&C) for preflight predictions and control analysis and by thermal designers in the development of the Thermal Protection Systems (TPS) used to protect the vehicle during re-entry.

The partnership with NASA allowed the SNC team access to vast computing resources and resulted in CFD simulations of thousands of unique cases along the Dream Chaser entry, descent, and landing (EDL) trajectory. Using the Pleiades and Endeavour supercomputers at the NASA Advanced Supercomputing (NAS) facility, the SNC aerodynamic team performed approximately 20,000 high-fidelity CFD simulations to calculate the aerodynamic forces and moments acting on the Dream Chaser vehicle, control surface interactions, and aerodynamic increments by firing RCS thrusters throughout EDL. These simulations have required tens of millions of core-hours, and store multiple terabytes of data on NAS systems. None of this would have been possible without the help and resources of the NAS Division.

In fall 2021, Dream Chaser will launch on its first voyage to the International Space Station (ISS) from Kennedy Space Center (KSC), the first of six missions. Each mission will conclude with the return of valuable experiments from the ISS, returning to the Space Shuttle Landing Facility at KSC. NASA is planning to capture images of the Dream Chaser during its first return flight using Scientifically Calibrated In-Flight Imagery technology. It is expected the heating observations will be coupled with extensive CFD, along the Best Estimated Trajectory, to calibrate thermal predictions and update thermal databases after the flight. The long term benefit to NASA would be increased payload capability to and from the ISS through improvements to the Dream Chaser spacecraft.

*Matt Opgenorth, Michael Jeffries, Sierra Nevada Corporation  
matt.opgenorth@sncorp.com, michael.jeffries@sncorp.com*



*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 plume of the thruster impinges on the wing, changing the pressure distribution and therefore changing the vehicle's aerodynamics. Matt Opgenorth, Sierra Nevada Corporation*


## HUMAN SPACE FLIGHT

### Building Booster Separation Aerodynamic Databases for Artemis II

The Artemis program is NASA's mission to land the first woman and next man on the Moon by 2024. Artemis II, the first crewed flight of the program, will send humans to near-lunar space for the first time since the end of the Apollo program. This mission consists of an Orion crew capsule launched atop NASA's Space Launch System (SLS), one of the most powerful rockets ever constructed. The SLS configuration consists of two solid rocket boosters attached to a center core. Two minutes after liftoff, the solid rocket boosters separate from the center core by means of 16 booster separation motors (BSMs) firing simultaneously. Aerodynamic forces play a critical role during this separation, and accurate databases are necessary to ensure that boosters move safely from the core stage without recontacting it. These databases are constructed using computational fluid dynamics (CFD), as the flight conditions and complex interactions of the various components are difficult to model in an experimental setting.

CFD simulations made using NASA's FUN3D flow solver were used 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, and all 16 BSMs. The aerodynamic data itself is a function of 13 independent variables: six describing position (three for translation and three for rotation); four describing free-stream conditions (angles of attack/sideslip, Mach number, and air density); and three describing thrust conditions (core, booster, and BSMs). In addition, separate sets of computations were run to simulate the possible effects that could be caused by one of the core stage engines failing.

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.



*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. Slices of the flow are taken on vehicle centerline and through the left booster's separation motors. 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, Henry Lee, NASA/Ames*

All CFD simulations were run on the Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility over the course of a three-month period. The nominal database required 5,780 FUN3D simulations, while the core stage engine failure databases required over 8,000 total simulations. Each case used approximately 2,500 core hours to complete, and the entire project needed anywhere from 50 to 100 terabytes of short-term storage at any given time.

*Jamie Meeroff, Derek Dalle, NASA Ames Research Center  
jamie.g.meeroff@nasa.gov, derek.j.dalle@nasa.gov*

## HUMAN SPACE FLIGHT

### Revolutionizing Assessment of Unsteady Flows Using Pressure-Sensitive Paint

Many important aerospace problems involve unsteady, separated aerodynamic flows, and the key measurement necessary for understanding these flows is surface pressure. Traditionally, hundreds of unsteady pressure sensors are used in wind tunnel tests to assess flows around aerospace vehicles and components. However, a new state-of-the-art approach is needed to decrease the overly conservative loads produced by traditional methods. An emerging optical technology, Unsteady Pressure-Sensitive Paint (uPSP), is now being used to measure global fluctuating surface pressures. Recently, uPSP was used in tests run on a model of NASA's Space Launch System in the Unitary Plan Wind Tunnel (UPWT), located at NASA's Ames Research Center, to determine the aeroacoustics environment and assist in developing the buffet forcing functions. The high spatial resolution of the uPSP data collected during the test will help engineers develop a more efficient, higher-performing vehicle.

The goal of the recent test campaign was to demonstrate the power of uPSP to assess low-frequency forces, while also addressing one of the main challenges this method faces when deployed in a production wind tunnel: access to sufficient data storage and compute resources to store and process the enormous amount of data produced during the tests. Working with high-performance computing (HPC) experts at the NASA Advanced Supercomputing (NAS) facility, we established the Red Rover project to develop a secure, uni-directional, high-bandwidth connection from the UPWT to the compute and storage resources at the NAS facility, also located at Ames—about a half-mile away from the wind tunnel.


Previously, transferring uPSP data from the UPWT to the NAS supercomputers for processing required physically transporting hard drives from one facility to the other. Just uploading the data to multiple drives took over a month; clearly, this is not a sustainable approach for expanding the amount of data collected during wind tunnel campaigns. The Red Rover project aimed to transfer, process, and visualize the uPSP data, recorded by four high-speed cameras, within one hour of data collection, allowing researchers to make informed decisions about test matrix prioritization.

During the test campaign, over 600 data points and more than 150 terabytes of data were acquired and securely transmitted to the NAS facility. In addition, NAS experts parallelized portions of the code used to process the uPSP data, dramatically reducing processing time from one day per test point to one minute. By connecting the power of this high spatio-temporal measurement to the power of the Pleiades supercomputer, along with the ability to quickly process and visualize the results on the NAS facility's hyperwall, near-real-time design decisions are made possible—establishing a new path toward decreasing design cycle time and mitigating over-design of aerospace vehicles due to insufficient tools.

Traditionally, uPSP has been used in small facilities, with the amount of data acquired relatively small and manageable. Since NASA researchers have recently matured uPSP technology, it can be deployed in large production wind tunnels. High-performance computing (HPC) resources are required to store and process the enormous amount of raw data produced in uPSP tests in these larger facilities. Combined with HPC resources, uPSP provides a new tool to efficiently design and assess wind tunnel models in order to develop next-generation aerospace vehicles.

The project's success is a clear demonstration of the reward when experimental and computational facilities and experts are integrated with the goal of solving the challenging problems within aerospace.

*Nettie Roozeboom, Ji Lie, NASA Ames Research Center  
nettie.roozeboom@nasa.gov, ji.lie-1@nasa.gov*



*Pressure distribution as measured by the uPSP on NASA's SLS SM-1 cargo vehicle (yellow is higher pressure; blue is lower). The data was processed on NASA's Pleiades supercomputer and visualized by the NAS Division's Visualization and Data Analysis team. Jennifer Baerny, Paul Stremel, Timothy Sandstrom, NASA/Ames*

## HUMAN SPACE FLIGHT

### Microscale Analysis of Spacecraft Heat Shields

Imagine entering Earth's atmosphere after returning from the outer solar system. A heat shield less than two inches thick protects you from temperatures up to 2,900 degrees Celsius (5,252 degrees Fahrenheit)—the actual conditions experienced by NASA's Stardust capsule during reentry in 2006. The capsule was protected by a heat shield composed of Phenolic Impregnated Carbon Ablator (PICA), a lightweight material invented by NASA to withstand these high temperatures.


PICA and most other ablative heat shield materials consist of multiscale composites with complex microstructures. Evaluating these materials is difficult, requiring precise knowledge of their properties. To support this effort, NASA scientists are developing research codes to compute material properties and simulate ablation at the microscale. Simulations run on the agency's supercomputers are leading to a better understanding of current heat shields, which could lead to improved materials that would allow spacecraft to survive higher temperatures. Utilizing these tools along with experiments, researchers are working to push the limits of spaceflight, allowing for greater flexibility in future missions.

To analyze the materials in ablative heat shields at the microscale, the Predictive Material Modeling team in the NASA Advanced Supercomputing (NAS) Division at Ames Research Center has developed the Porous Microstructure Analysis (PuMA) software. This software is capable of importing sets of grayscale images representing a material's microstructure to create a geometrically accurate computational domain. Once the microstructure is imported, we use finite volume and particle-based methods to compute properties for the material, including thermal conductivity, porosity, and tortuosity. Ablation of the microstructure due to oxidation can also be simulated using a random walk method. We also use Direct Simulation Monte Carlo (DSMC) methods to analyze heat shield microstructures. These high-fidelity, massively parallel simulations generate millions of molecules that interact with each other and the material microstructure. DSMC allows for more complex environments with many chemical reactions.

All of these capabilities have been extensively tested and verified using numerous test cases obtained from literature or developed in-house. Our findings show the computed material properties are accurate for many materials with known properties. The ablation simulations are found to be qualitatively accurate, however, a lack of experimental data has prevented true validation. Both PuMA and DSMC have been used to analyze materials in current heat shields, as well as those being considered for future missions. Our massively parallel DSMC simulations require around 2,000 processors and run in a few hours. The simulations are typically run using Intel Xeon Broadwell nodes on the Pleiades supercomputer at the NAS facility.

PuMA was originally developed using OpenMP for shared memory systems. While a large multi-processor workstation is generally sufficient for computations, we have also performed parametric studies requiring many PuMA simulations to run in parallel. As with the DSMC simulations, these studies are generally run on Pleiades with each parallel simulation on a Broadwell node with 28 processors. Work has also been done to add support for GPUs in PuMA using OpenACC. Testing is done on Electra Skylake nodes using NVIDIA Tesla V100 GPUs. NASA is currently investigating woven carbon fiber composites for use as heat shields due in part to their customizability and desirable thermal properties. The added complexity of woven materials brings new challenges and opportunities for these codes going forward.

*John M. Thornton, Nagi N. Mansour, NASA Ames Research Center  
john.m.thornton@nasa.gov, nagi.n.mansour@nasa.gov*



*Frame from a simulation showing microscale ablation of a heat shield due to oxidation. Particles enter from the top and react as they collide with the microstructure of the material. As the simulation progresses, the surface material is eaten away, allowing the particles to diffuse further downward. The material represented here is the carbon fiber precursor to the Phenolic Impregnated Carbon Ablator (PICA), invented by NASA for use in spacecraft heat shields. Joseph C. Ferguson, NASA/Ames*

## OUR PLANET

### Fine-Scale Forest Modeling Across the North American Boreal Zone

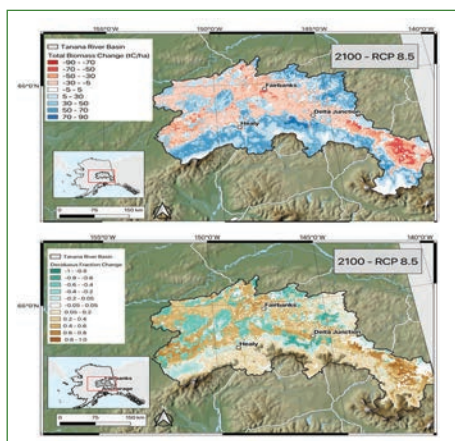
The North American boreal forest is highly influenced by complex interactions between vegetation, wildfire, soil conditions, and climate. Climate change is occurring faster in this region than anywhere else on Earth and is likely to impact all of these interactions, ultimately leading to changes in vegetation composition and structure. High-resolution forest modeling can aid in simulating these potential changes by considering these important vegetation drivers at a fine scale.

Scientists updated an individual tree-based forest model (University of Virginia Forest Model Enhanced - UVAFME) for better performance within the boreal zone. Specifically, the updates improved simulation of permafrost dynamics, litter and organic matter decay, and wildfire dynamics, and their interactions with vegetation and soil characteristics. Scientists tested these updates against 244 field observation sites across Alaska. They compared model output to observations of soil moisture, permafrost depth, forest biomass, species composition, and other biotic and abiotic characteristics. Results showed that the model could reliably produce these observations across a wide variety of site conditions. Next, the model was applied on a 2-kilometer (km) by 2-km grid across interior Alaska's Tanana River Basin (~115,000 km<sup>2</sup> area), equaling about 28,500 sites. The model was used to simulate forest dynamics under two climate change scenarios: Representative Concentration Pathway (RCP) 4.5—a more conservative climate change prediction—and RCP 8.5—a more extreme prediction.

UVAFME predicted that decreasing biomass and increasing deciduous fraction is likely under those climate change scenarios. These changes have the capacity to feed back to climate through changes in forest structure, which impacts surface roughness, albedo (which in turn impacts energy fluxes), and evaporative flux. The model results increase understanding of the potential futures of the Alaskan boreal forests and demonstrate the importance of fine-scale interactions between trees and their environment. Importantly, accurate simulation of forest dynamics and characteristics was not possible until all model updates were finalized and all linkages between vegetation and soil characteristics were implemented. This finding demonstrates the importance of a fine-scale model like UVAFME for simulating future forest state. Scientists used the NASA Center for Climate Simulation's ADAPT Science Cloud to conduct code refinements and testing as well as the model simulations for this work. They tested the model on ADAPT's Arctic-Boreal Vulnerability Experiment (ABoVE) nodes, running jobs interactively. Once ready to run the large simulation across the Tanana, they used a SLURM script to submit batches of jobs to 10 nodes at a time. Each job consisted of ~2,850 sites to simulate and only used 1 core and minimal RAM. The major sink for this type of simulation is time. It takes about 40 seconds to run one site. If run in sequence on a personal computer, running all of the Tanana would take 1,140,000 seconds (i.e., 317 hours)! However, distributing the job across 10 nodes reduces runtime to only about 32 hours. Thus, ADAPT was an invaluable resource for timely testing and production of this work.

In the future, UVAFME will be applied across the entire North American boreal region, providing further estimates of future vegetation trajectories across a much larger region. Additional model updates will include the incorporation of insect disturbance.

*Adrianna Foster, Scott Goetz, Northern Arizona University  
adrianna.foster@nau.edu, scott.goetz@nau.edu*



*These maps show model predictions of climate change's effects on Alaska Tanana River Basin vegetation in the year 2100. The top map depicts forest biomass (in tonnes of carbon per hectare), with dark red indicating large decreases and dark blue indicating large increases. The bottom map depicts deciduous fraction (percent of biomass that is deciduous species), with dark brown indicating strong increases and dark teal green indicating strong decreases. Adrianna Foster, Scott Goetz, Northern Arizona University*

## OUR PLANET

### Realistic Simulations of the Coupled Atmosphere-Ocean-Ice System

Recent studies have emphasized the importance of small-scale atmospheric and oceanic processes for driving air-sea exchange processes—for example, individual clouds and oceanic storms, called mesoscale eddies, which are not resolved by current weather and climate models. In this project, we study these processes by coupling two flagship data-assimilating NASA models: 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); and 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. What is missing in both the atmospheric and oceanic high-resolution simulations is the ability to accurately represent coupled air-sea exchange processes. We are addressing this deficiency by running coupled ocean-atmosphere simulations, where the GEOS atmospheric model and the ECCO ocean model are both run at km-scale resolution.

The high-resolution GEOS and ECCO simulations enable the study of how swirls of ocean water, called eddies, affect circulation of both the ocean and the atmosphere. Eddies represent a major pathway for transferring heat from the ocean to the atmosphere. Growing evidence shows eddies have remote impacts on wintertime European blocking (causing anomalously cold and dry weather), monsoon rainfall, and rainfall along the U.S. North Pacific Coast. To date, we have developed simulations with horizontal grid spacing as small as 7-km for the GEOS atmosphere and 2–4-km for the ECCO ocean. The first diagnostic of 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.


We expect that the coupled simulations will be used by scientists in the ocean research community for many of the same purposes as the separate oceanic and atmospheric simulations have already been used, but with much increased realism in the description of air-sea interactions and oceanic and atmospheric planetary boundary layers. These simulations are guiding our understanding of coupled air-sea exchange processes and their parameterization in coarser-resolution weather and climate models. They are also being used to guide the utilization of existing satellite observations and the development of new instruments. In particular, a realistic coupled nature simulation will be an invaluable resource for carrying out NASA's Observation System Simulation Experiments for use in existing, planned, and proposed missions.

NASA supercomputing resources are revolutionizing Earth science studies by enabling increasingly realistic simulations of the coupled ice-ocean-atmosphere system. The simulations were carried out on up to 70,000 cores on the Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility, and the Discover system at the NASA Center for Climate Simulation (NCCS), and are being made available to the scientific community on NAS and NCCS data portals.

In setting up the coupled GEOS/ECCO simulations, we encountered both computational and scientific issues, such as the efficiency of input/output (I/O), a lack of compatibility between the high-resolution atmospheric and oceanic land masks, the representation of sea ice in the coupled system, and the proper method to distribute runoff into the ocean to avoid instability. Our ongoing work aims to address these difficulties and to decrease the grid spacing to 3 km in the atmosphere and 1–2 km in the ocean.

*Chris Hill, Massachusetts Institute of Technology  
cnh@mit.edu*

*Dimitris Menemenlis, Jet Propulsion Laboratory  
menemenlis@jpl.nasa.gov*



*Arctic sea ice concentration and thickness from an internal-wave-admitting 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 called "leads." Since many interactive processes between ocean, atmosphere, and ice take place along leads in sea ice, resolving these thin lines improves the representation of air-sea interactions in polar regions. Dimitris Menemenlis, NASA/JPL; Nils Hutter, Alfred-Wegener Institute*

## OUR PLANET

### HPC and the Topography of the Earth's Surface

The University of Minnesota's Polar Geospatial Center (PGC) and NASA Goddard Space Flight center are constructing Digital Elevation Models (DEMs) of the Earth from DigitalGlobe satellite imagery through the NextView License for scientific research. This project seeks to generate DEMs over wide areas of interest to NASA and National Science Foundation researchers, including the Arctic, Antarctic, and semi-arid regions around the globe.


To create the DEMs, scientists are using imagery from multiple satellites (WorldView-1, WorldView-2, WorldView-3, and GeoEye-1) with resolutions down to 2 meters (m) x 2 m per pixel. Overlapping pairs of the high-resolution images are auto-correlated and run through the Surface Extraction from TIN-based Search-space Minimization (SETSM) software package to compute the DEMs.

The project is using hundreds of thousands of high-resolution visible and near-infrared images, resulting in the analysis of many petabytes of data. The DEMs can be computed to differing resolutions. Each increase in resolution raises the amount of compute time required for SETSM to converge, and the size of the resulting DEM grows accordingly.

The end result will be a DEM baseline for significant portions of the Earth applicable across a wide variety of scientific inquiries. Heights can be used to establish current measurements of ice and glaciers around the world. In addition, scientists can create high-resolution maps of existing vegetation, specifically tree canopy heights across areas where other measurements do not have sufficient resolution to even resolve individual trees. Once a baseline is created, changes in these elevations can be used to study how specific environments are evolving in a changing climate.

For example, ice height changes of the Vavilov Ice Cap in the high Russian Arctic were measured from 2013 through 2016. Rapid changes were easily seen from the satellite imagery, with the ice moving fast enough to create a fan-shaped edge of the glacier and spread across the Kara Sea. Similarly, the DEMs can be used to look at changes in vegetation over time, either due to natural events (such as volcanic eruptions) or anthropomorphic effects (such as clear cutting). Ultimately, these DEMs will provide scientists with a consistent, high-resolution, and constantly updating picture of how the Earth is changing over time.

High-performance computing and storage are critical to generating the DEMs. In addition to collaborating with NASA, PGC has an agreement with the National Center for Supercomputing Applications (NCSA) and The Ohio State University to produce DEMs for the entire Earth using the NCSA Blue Waters supercomputer. At NASA, DEMs for the Arctic boreal regions and high mountain terrains in Asia are being generated using the Discover supercomputer at the NASA Center for Climate Simulation. These DEMs are stored for NASA scientists to use for fundamental research into how climate change is affecting these regions and to also help define the next generation of remote sensing platforms needed to continue to study the Earth.



*This computed Digital Elevation Model (DEM) captures the three-dimensional structure of glaciers, ice sheets, and other topographical features in northern Canada, Greenland, and Iceland. By incorporating high-resolution satellite imagery taken over multiple years, DEMs become a powerful tool for tracking how the Earth is changing over time. Paul Morin, University of Minnesota*

In addition to continuing to support the creation and automation of the DEM processing and storing of data, NASA and PGC are looking to collaborate on algorithms to increase the resolutions of the resulting height measurements. Employing a similar method developed for the Multi-Angle Implementation of Atmospheric Correction (MAIAC), researchers can improve the measurements by removing the effects of certain atmospheric conditions, such as aerosols.

*Paul Morin, University of Minnesota  
lpaul@umn.edu*

*Compton Tucker, NASA Goddard Space Flight Center  
compton.j.tucker@nasa.gov*

## OUR PLANET

### The Future of Global Numerical Weather Prediction with GEOS


The prediction capabilities of global models have continuously evolved from traditional, medium-range, global weather prediction applications to span scales from hourly prediction of convective-scale storms to seasonal Earth system prediction. This evolution has increased the demands on the system infrastructure design and workflow to achieve the required performance on modern high-performance computing (HPC) platforms. The planned evolution of NASA's Goddard Earth Observing System (GEOS) modeling and assimilation system will stress the capabilities of conventional HPC, overwhelming the available compute cycles at the NASA Center for Climate Simulation (NCCS) in the coming 5–10 years.

Such demands have led Global Modeling and Assimilation Office (GMAO) developers to re-design the key elements of GEOS assimilation and modeling systems to achieve significant performance gains on anticipated exascale platforms. Transitioning the assimilation system to a new, unified, interagency framework—the Joint Effort for Data assimilation Integration (JEDI) system—positions GEOS to exploit new and efficient data assimilation algorithms within an Earth system context that fully integrates atmospheric and oceanic models. GMAO researchers are exploring the suitability of the GEOS model, leveraging a domain-specific language approach and artificial intelligence to accelerate computational performance and data exchange efficiency of the integrated Earth system model. We will highlight recent efforts in these areas and emphasize the demand for further development and re-design to achieve the science objectives in support of NASA's Earth system modeling and assimilation missions.

The future of global, numerical weather prediction (NWP) can be explored today with the recent expansion of the NCCS Discover supercomputer to 6.7 petaflops. Using up to 80,000 Intel Skylake processors, an advanced version of the GEOS model produced near-real-time global weather forecasts at resolutions of 6 kilometers that will be typical of operational NWP systems in the year 2025 and beyond. These forecasts include the extreme tornado season across the United States in May–June 2019 and the extreme Category 5 Hurricane Dorian in August–September 2019. This advanced version of GEOS produces significant improvements in the representation and prediction of these severe events.

High-end computing (HEC) at NASA—specifically the NCCS and the NASA Advanced Supercomputing (NAS) Division—provides an essential resource for preparing our modeling systems for emerging HEC systems and exploring new science developments. While current GEOS configurations require thousands of computational cores, configurations over the coming decade will require scalability approaching millions of cores. Exposing additional parallelism within GEOS becomes essential for reaching these scalability targets. Partnerships with the NCCS and NAS support our exploration of traditional hybrid parallelism approaches within GEOS today and prepare our models for future HEC platforms by exposing new depths of parallelism within our application components.

*William Putman, NASA Goddard Space Flight Center  
william.m.putman@nasa.gov*



*A satellite observation (left) and a 48-hour Goddard Earth Observing System (GEOS) model forecast (right) show the cloud structure of Hurricane Dorian centered over the northern Bahamas on September 2, 2019. At this time, Dorian was an intense Category 5 hurricane with a large well-defined eye. The bright red and pink colors circling the cleared eye represent the coldest and most intense convection/storminess in the eye wall and circulation of Dorian.  
William Putman, NASA/Goddard*

## OUR PLANET

### Simulating Air Quality and Volcanic Plumes with GEOS CCM

The Goddard Earth Observing System (GEOS), developed at NASA's Global Modeling and Assimilation Office (GMAO), is an Earth-system model featuring numerous connected components that work together to facilitate the simulation of Earth System processes, produce near-real-time forecasts, and generate long-term reanalysis products. Using the GEOS Chemistry-Climate Model configuration (CCM), which includes fully interactive aerosols and gas-phase chemistry, we examined how changing the resolution of an atmosphere model affects simulation results, with the goal of improving predictions of air quality and the impacts of volcanic eruptions on surface weather. We specifically examined the impacts of stratospheric ozone intrusions on air quality and of volcanic eruptions on surface meteorology.

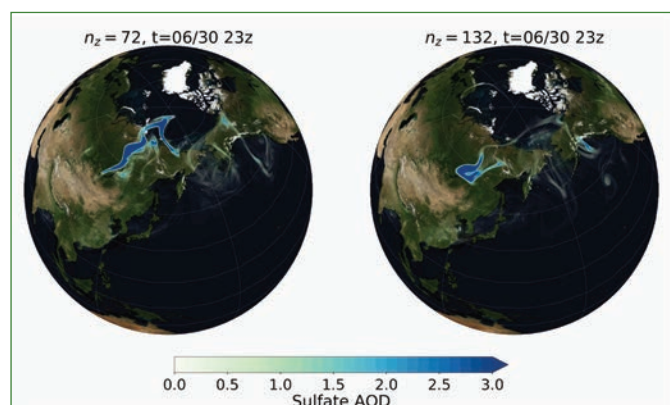
To investigate how increasing the resolution of the model's vertical mesh would affect the simulated transfer of constituents between the stratosphere and the troposphere, we conducted two simulations, one with 72 vertical model levels (the standard configuration for GEOS) and another with 132 levels. The additional levels allowed for more detailed representation of the atmosphere's vertical structure, particularly at the interface region between the stratosphere and troposphere. Our simulated stratospheric intrusion and volcanic plume events were both characterized by the exchange of material (such as ozone and volcanic aerosols) between the stratosphere and troposphere, which made the simulations especially sensitive to the model's vertical resolution.

Our experiments showed that increasing the vertical resolution of the model improves simulation accuracy of both the stratospheric intrusion events and the volcanic eruptions. These improvements will allow us to better interpret NASA satellite observations of real events in Earth's atmosphere, and make better air quality forecasts.

NASA's HPC resources are essential for running and further developing the GEOS model. A typical model configuration can require hundreds to thousands of processors and terabytes of memory, and analyzing its output requires archiving terabytes to petabytes of disk space. For this project, we ran our simulations across 200 nodes on the Pleiades supercomputer at the NASA Advanced Supercomputing (NAS) facility, and stored 63 TB of data on NAS storage resources. As model development continues, and the need for higher resolution simulations grows, so will the need for additional HPC resources.

Simulations like these lay the groundwork for increasing the complexity and resolution of GEOS models.

*Elliot Sherman, Peter Colarco, K. Emma Knowland, NASA Goddard Space Flight Center  
elliot.m.sherman@nasa.gov, peter.r.colarco@nasa.gov, k.e.knowland@nasa.gov*



*Two different realizations from the last day of a GEOS model simulation of the volcanic aerosol plume produced by the Raikoke volcano on the Kuril Islands, Russia. The plume loading is represented by the sulfate aerosol optical depth (AOD), a measure of how much material is in the atmospheric column. The left image is from a simulation with 72 vertical levels, while the image at right is from a simulation with 132 levels, showing the different transport of the plume in each simulation. Peter Colarco, NASA/Goddard; Parker Case, University of Colorado*

## OUR PLANET

### Virtual World, Real Understanding: Using Virtual Reality to Visualize Earth Data

Have you ever seen storm surge warnings during a hurricane? What does a storm surge of 6 feet imply, versus 12 feet? What impact does that have on buildings, people, roads, and bridges? Sometimes we talk about data, but it is difficult to visualize the reality of what the data actually means. NASA collects a massive amount of scientific data that is useful for real-world, practical applications, including disasters, if it can be presented in a digestible and relatable way. Navteca, the NASA Applied Sciences Disaster Applications Group (within NASA's Science Mission Directorate), and the NASA Information Technology Communications Directorate have been researching virtual reality (VR) technology for the next generation of Earth science technology information systems and advanced data visualization. The Disaster Applications Group promotes the use of Earth observations to improve disaster prediction, preparation, response, and recovery.

This project began with an evaluation of the ArcGIS portal and data residing at the NASA Center for Climate Simulation (NCCS). Tests showed that through an application programming interface (API) it was possible to query and serve the disaster application data. The data was then overlaid atop accurate 3D terrain and structures to create visualizations.

The team developed an interface to the 3D virtual city models where users can order and view data from relevant disaster datasets in an immersive, interactive VR environment. In one custom VR interface, a 3D model of San Juan, Puerto Rico displays an interactive water level tool to demonstrate the severe flooding after Hurricane Maria and simulated flooding for different category hurricanes. The 3D topography and buildings allow users to quickly understand flooding impacts based on complex numerical output from sources like the NOAA Sea, Lake and Overland Surges from Hurricanes (SLOSH) Model.


An important achievement has been importing new datasets from the NCCS-hosted NASA portal in near-real time. As new datasets are made available, the project team has been successful ingesting them into the VR environment without additional coding or development. For example, after the Ridgecrest, CA, 7.1 earthquake in July 2019, the Disaster Applications Group and Navteca simultaneously presented new InSAR data in both the web portal and in the VR demo. Similarly, VR visualizations of soil moisture for at-risk coastal areas were created during Hurricane Dorian as it moved up the Eastern seaboard.

This project contributes to growing research demonstrating VR technology as a significant new approach to enhancing the way that scientists and researchers view and communicate results and data in an engaging, interactive format. It furthers the NASA Applied Sciences mission by leveraging innovation and technology to find practical use for NASA Earth observations.

By focusing on disaster scenarios, we are investigating whether VR technology can help develop enhanced tools that inform data-driven decision making. This type of VR visualization could significantly impact understanding of GIS and Earth observation data and also improve how users and stakeholders collaborate and communicate about disaster impacts and changes over time.

Earth observation and remote sensing data product generation uses high-end computing resources at the NCCS. Additionally, the NCCS-hosted NASA Disaster Applications data portal was an integral part of this project's success. Consolidating many datasets in one place creates a "one-stop shop" for users in search of data related to specific disaster events and allows ingesting different data overlays in near-real time.

*Shayna Skolnik, Adam Silcott, NASA Headquarters/Navteca  
shayna.s.skolnik@nasa.gov, adam.d.silcott@nasa.gov*



*The project team created an interface to the NASA Disaster Applications GIS portal hosted at the NASA Center for Climate Simulation (NCCS). In an intuitive way for users to visualize flooding dynamically, simulated flooding is overlaid on land surfaces using realistic water textures. This interactive tool shows the potential impacts of severe flooding, as well as the capacity to visualize other Earth data such as hurricanes, fires, landslides, floods, and volcanic eruptions. Adam Silcott, NASA HQ/Navteca*

## OUR PLANET

### Tree Census from Space: Quantifying Woody Biomass Using Machine Learning

NASA scientists are determining the woody carbon content of discrete trees in the arid and semi-arid zones south of the Sahara Desert in Africa, regions that provide early indicators of climate change's full effects. They are mapping individual-tree crown area and determining heights for any tree or bush with an area of two square meters or more.

The goal of this project is to quantify the woody biomass (amount of carbon) in the trees and bushes in the hot dry tropical, arid, and semi-arid regions of Sub-Saharan Africa, an area of 10 million square kilometers.

We used hundreds of thousands of visible and near-infrared high-resolution images from four different DigitalGlobe satellites. These images are accessed through the NextView License for NASA's scientific research and have a resolution of approximately 50 centimeters per pixel. To reduce any multiple counting of tree pixels, the images are first orthorectified and then mosaicked using Python and GDAL to create panchromatic and NDVI mosaics.


Once the mosaics are completed, training data are compiled from multiple sites across Africa. Machine learning (ML) then delineates trees by outlining the tree canopies and producing shapefiles of the tree canopies in question. The execution of the ML approach uses a large amount of compute resources.

Initial results for the first UTM zone on the Sub-Saharan West Coast have shown that a highly accurate canopy mapping of trees is possible, with hundreds of millions of individual trees identified and their crown areas determined. Previous tree-mapping efforts with lower resolution have greatly underestimated the number of individual trees.

It is important to understand how biomass measurements of the Earth change over time, because this information can be related to changing climatic conditions. The world's arid and semi-arid regions respond quickly to changes in climate. Because of the low density of trees in arid and semi-arid regions, changes in these areas are more easily quantified.

Creating the mosaics requires substantial compute and storage resources: this is why the NASA Center for Climate Simulation's Advanced Data Analytics Platform (ADAPT) was ideal for this portion of the work. Virtual machines with specialized software stacks connected to a high-performance computing (HPC) file system (IBM Spectrum Scale cluster) take massive amounts of raw images as input and create output data for analysis. The output data are a three-dimensional array of 2.3x1,013 elements.

Analysis at such a massive scale requires substantial HPC resources. This project leverages an innovative combination of using ADAPT to build the mosaics and train the model, while using the Discover supercomputer to run the trained model that, in turn, identifies tree canopies.



*This figure illustrates our mapping of individual tree crowns and stereographic retrieval of their height. This is done at the 50x50 centimeter x-y scale with a vertical accuracy of  $\pm 1$  meter. This figure is from the Sahel Zone in northern Senegal, on the south side of the Sahara Desert.  
Compton Tucker, NASA/Goddard*

Researchers will produce a census of all trees across the entirety of Sub-Saharan Africa—from 12 degrees north latitude to 24 degrees north latitude, from the Atlantic Ocean to the Red Sea—using satellite data and ML. Stereo data will be used to estimate tree heights. These two measurements will provide an estimate of the total amount of carbon in woody biomass over this large region of the African continent. By analyzing data from multiple years, it will be possible to calculate changes in biomass over the past decade and link these to the changing, abiotic aspects of climate.

*Compton Tucker, Katherine Melocik, NASA Goddard Space Flight Center  
compton.j.tucker@nasa.gov, katherine.a.melocick@nasa.gov*

## SUPERCOMPUTING

### Teaching Your Legacy Code New (GPU) Tricks

Accelerated computing using graphics processing units (GPUs) has become an industry standard for high-performance, energy-efficient computing. However, adapting the technology is challenging, especially for large, long-developed legacy codes maintained by domain scientists. The OpenACC parallel programming model allows GPU acceleration through the use of special code comments. This greatly eases the porting process and produces a single source code that's able to run on both CPUs and GPUs. CORona-HELiosphere (CORHEL), a state-of-the-art modeling suite developed by Predictive Science Inc., is used extensively to generate solar wind and solar eruption simulations for research and space weather applications. It is also used as the backbone of the Solar Particle Event Threat Assessment Tool (STAT). Adding GPU acceleration to CORHEL would greatly benefit these applications, paving the way towards the eventual goal of real-time forecasting.

CORHEL is a great example of a long-standing legacy code package. Developed by domain scientists for over 15 years, it contains multiple model applications, totaling over 100,000 lines of code. Rewriting CORHEL with GPU-specific APIs is simply not feasible due to time, money, and developer learning curve considerations. By using OpenACC, we can rapidly add GPU acceleration with minimal impact on the original code base, and maintain full portability—the code still works as before on older CPU systems.

We approach adding OpenACC to CORHEL on a model-by-model incremental basis. Occasionally, it is necessary (or advantageous) to refactor and/or modify the original CPU code to better position it for OpenACC acceleration. Almost always, this refactoring is minimal and improves the CPU code in both performance and clarity. After each implementation, the resulting code is extensively validated on CPUs and GPUs, and then optimizations are explored.

Using the above approach, we successfully added OpenACC to several models within CORHEL. Timing results on various GPU systems, including the NVIDIA V100 GPU nodes on NASA's Pleiades supercomputer, show very good performance compared to CPUs. For example, a simplified solar eruption simulation that takes over eight hours on a Sandy Bridge 16-core CPU node, and over 3 hours on a Skylake 40-core CPU node, takes only about 12 minutes to complete on a single 8xV100 GPU node. These results demonstrate the low-cost, high-performance advantages of using GPUs on legacy applications. In addition to the value of running CORHEL on GPUs, this project also demonstrates the potential of using OpenACC to port other applications run at NASA—both legacy and new—to GPUs.

In order to develop and test our OpenACC GPU implementation of CORHEL, we utilized many high-performance computing and local resources including the recently deployed NVIDIA 32GB V100 GPU nodes on NASA's Pleiades system. Running on Pleiades' four-GPU and eight-GPU nodes resulted in large speedups over using the CPU nodes. We plan to continue implementing OpenACC into the remaining models in CORHEL, most notably its thermodynamic magnetohydrodynamic model used for large realistic coronal background and solar eruption simulations. This will greatly improve the time-to-science in studying these solutions in detail, as well as improve time-to-solution for advanced space weather applications.

*Ronald M. Caplan, Jon A. Linker, Predictive Science Inc.  
caplanr@predsci.com, linkerj@predsci.com*

*OpenACC is an application programming interface that allows the addition of GPU acceleration to legacy codes through the insertion of special code comments. We have used OpenACC to port selected models of the CORona-HELiosphere software suite to GPUs. Magnetic field lines are taken from a snapshot of a simplified solar eruption simulation. Ronald Caplan, Jon Linker, Predictive Science Inc*



## SUPERCOMPUTING

### Facilitating Science Through Centralized Storage

Imagine having access to interrelated NASA data from the cryosphere, atmosphere, hydrosphere, and geosphere all stored in the same place. Imagine being able to run your climate analysis on model output from the Global Modeling and Assimilation Office (GMAO) GEOS code or on exoplanet atmosphere output from the Goddard Institute for Space Studies (GISS). Imagine being able to do Machine Learning on the latest GMAO Nature Run output or six different reanalyses. The NASA Center for Climate Simulation (NCCS) Centralized Storage System (CSS) provides the means to do all of these things: a location for curated NASA data products to be shared among researchers and NCCS compute environments, including Discover, our traditional high-performance computing (HPC) environment; on-premise Cloud; Data Services; and, in the near future, Machine and Deep Learning compute.

The NCCS provides 50 petabytes (PB) of online storage to Discover, 95 PB of tape storage, and 3.5 PB of online storage to Cloud. CSS provides an additional 17 PB of shared storage for curated data, with another 17 PB coming online in the new year. The NCCS uses IBM's Spectrum Scale General Parallel File System (GPFS) to store the data. High-speed internal InfiniBand and 40 Gigabit Ethernet networks connect CSS to the compute and Data Services environments. Data must be covered by a Data Management Plan, have an identified curator, and be archived at a separate facility. It is then made available read-only to all our environments.

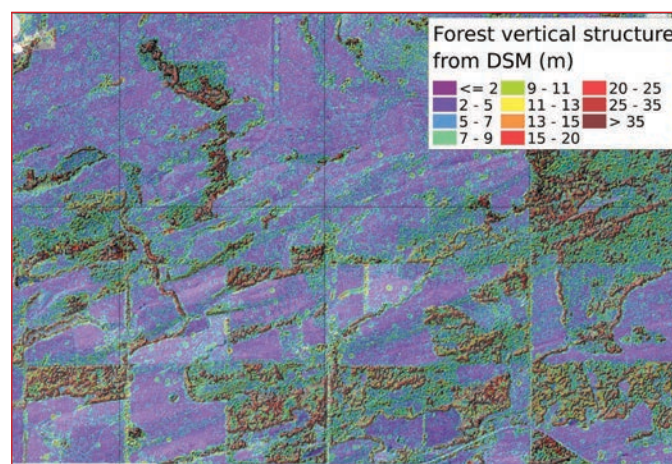
Access to datasets such as Landsat, MODIS, GMAO forecast products, National Geospatial-Intelligence Agency (NGA) high-resolution imagery, ICESat-2, and reanalyses allows researchers to focus on science rather than data transfer, storage, and management. In addition, scientists can apply new techniques such as Machine and Deep Learning on these massive data holdings. Research collaborations such as NASA's Arctic Boreal Vulnerability Experiment (ABoVE) use large volumes of Earth observations, including high-resolution visible-wavelength imagery, on Discover to create digital elevation models (DEMs) and store them in CSS. This data is mounted on virtual hosts in the ADAPT Science Cloud for investigations using customized analytics environments to examine how the Arctic boreal environment (e.g., permafrost) is changing. High-resolution ABoVE analysis data generated on ADAPT can in turn be used on Discover, for example, in investigating high-mountain snowmelt and ice sheet-ocean interactions.

The Head in the Clouds workflow—a unique NASA effort to count trees and shrubs in sub-Saharan Africa conducted on both Discover and ADAPT—has been greatly simplified with the availability of the NGA data in both environments. HPC provides the compute resources to process, analyze, and visualize these large datasets. Access to tens of petabytes of complementary Earth observations and analysis data—coupled with Discover's powerful 6.7-petaflop HPC environment, ADAPT's hundreds of customizable virtualized hosts, and GPUs for Machine Learning—is proving valuable to the NASA science community. This proximity of related data, processing, and services provides a unique opportunity to take advantage of interdisciplinary datasets and new research paradigms (e.g., Earth scientists studying cryospheric processes, the carbon cycle, and biomass).

NCCS will increase available storage to 34 PB, add an 80-node GPU cluster to enable more Machine and Deep Learning opportunities, and provide search and discovery functionality as well as usage metrics to inform data sharing and eventual disposition.

*Laura Carriere, NASA Goddard Space Flight Center  
laura.carriere@nasa.gov*

*Bennett Samowich, NASA Goddard Space Flight Center  
bennett.r.samowich@nasa.gov*



*Research on tree canopy structures in a managed boreal landscape in central British Columbia will benefit from centralized storage shared between the NASA Center for Climate Simulation (NCCS) Discover and ADAPT research environments. Shown here are vertical structure patterns in that forest at 1-meter resolution. Tree heights are represented in colors ranging from purple (less than 2 meters) to rust (over 35 meters). Paul Montesano, Chris Neigh, NASA/Goddard*

## SUPERCOMPUTING

### Accelerating Science with AI and Machine Learning

Data holdings at NASA are growing at a geometric rate, with an estimated 100 petabytes of Earth science imagery alone. Traditional methods of analyzing these data are insufficient to produce answers in a reasonable time frame. Scientists have turned to Artificial Intelligence and machine learning (AI/ML) methods to facilitate the analysis of the large volumes of data. Here we present several projects that use various AI/ML methods in concert with high-performance computing (HPC) to derive products and inferences from image products including satellite imagery and model outputs. In addition, we will outline our plan to develop a unified approach to data science that incorporates AI/ML.

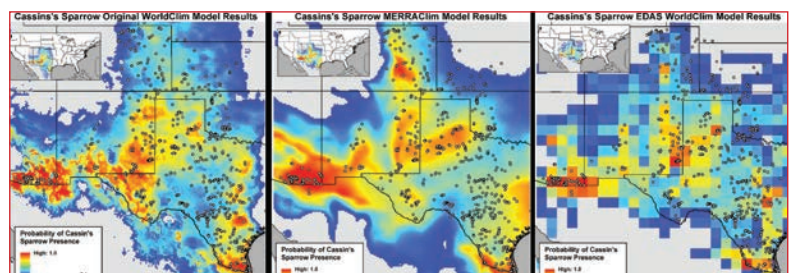
The data science and analytics group in the Computational and Information Sciences and Technology Office (CISTO) is working with scientists and industry to take advantage of all available resources to accelerate science. This group has previously built data analytic services for climate reanalysis data, including the MERRA Analytic Service, Earth Data Analytics Service, and CREATE-V.

Currently, we are working with scientists to support landslide detection with commercial, high-resolution data, automatic detection of lake depth from Landsat data, routine mapping of surface water extent to support reservoir management, and using maximum entropy modeling to predict habitat suitability for species of particular interest. Our goal is to generate a software code base that uses ML techniques and is optimized for HPC systems and is reusable for many different projects.

Physics-based models for weather and climate prediction are the standard because these models are trained to incorporate Earth system information and model the natural world as closely as possible. As the length of the observational record increases, it becomes possible to use ML to predict conditions based on past real-world events. The figures show a prediction of monsoon strength using Artificial Neural Network and climate data. This ML model identifies teleconnections between the climate variables surface temperature and atmospheric pressure with the strength of the ensuing monsoon season. The figures also show early results of habitat suitability mapping, which demonstrate that using MERRA inputs can yield comparable results to the standard WorldClim data inputs. WorldClim provides 19 climate-related variables (temperature, precipitation, etc.) that can be used as discriminatory variables in a diverse study. MERRA is a climate reanalysis dataset that provides hundreds of climate variables produced by the GEOS-5 climate model. The next step is to include the full suite of MERRA variables in the model in a Monte Carlo simulation to identify any key variables not in the WorldClim data that may contribute to improvements in the habitat suitability map.

Development occurs in the Advanced Data Analytics PlaTform (ADAPT) computing environment at the NASA Center for Climate Simulation, but the design process includes consideration for running in the commercial cloud. All code is developed with standard computer science principles in mind to maximize the portability of the code to different processing environments and is parallelized to optimize performance.

*Mark Carroll, NASA Goddard Space Flight Center  
mark.carroll@nasa.gov*



*Species Distribution Models (SDM) use climate information (temperature, precipitation, etc.) to identify habitat that is suitable for a species of interest. The goal of this research is to use MaxEnt software with three different input datasets (WorldClim, WorldClim augmented with MERRA, and MERRA) to create habitat suitability maps for Cassin's Sparrow in North America. The results indicate that MERRA can be used to produce comparable results to WorldClim (standard dataset used for SDM). Jian Li, Roger Gill, NASA/Goddard*

## SUPERCOMPUTING

### Jumping the Queue: From NASA to the Commercial Cloud

NASA's High-End Computing Capability (HECC) Project provides high-performance computing (HPC) resources and services to more than 1,500 scientists and engineers across the agency. Recently, we made it possible for users to run their jobs seamlessly on commercial cloud resources.

Last year, HECC conducted a trade study to determine how commercial cloud resources could be used in a cost-effective way. The study found that while replacing all in-house resources with those in the cloud would be cost prohibitive, some portions of users' HPC workloads might be cost effective to run on commercial clouds. The study also recommended starting an effort to provide a commercial cloud capability. In response to these recommendations, we have undertaken a three-phased approach to making commercial cloud resources available to NASA's science and engineering projects.

In the first phase, we implemented a pilot project that allowed a few select users to burst jobs from HECC front-end nodes to Amazon Web Services (AWS) resources. From the users' perspective, they only needed to annotate their batch job scripts—we developed a cloudbursting mechanism that took care of allocating the AWS resources and staging files between HECC systems, located at the NASA Advanced Supercomputing (NAS) facility, and AWS systems. By using GPU-accelerated nodes at AWS, which were not available on HECC systems at the time, the users were able to make significant advances in their research.

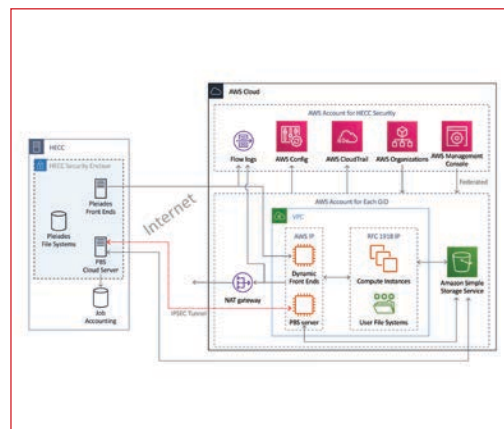
The second phase is built on the pilot project and makes AWS access available to all HECC users. In contrast to the first phase, where HECC paid the cloud charges, users are now responsible for them. This pay-per-use system makes cloud access sustainable for HECC's budget, and lets the user decide whether the opportunity cost of waiting in a batch queue for HECC resources exceeds the cost of "jumping the queue" by running the job at AWS. We are also enabling export-controlled work through the use of AWS GovCloud.

In the third phase, we are adding mechanisms for non-HECC users to access cloud resources to meet their science and engineering needs. We are currently analyzing user workflows to gather requirements, and will provide a user environment either by using web-based software available from cloud vendors or by building a suitable in-house environment.

In the pilot project, two users were able to make significant advances in their research through the use of GPU-accelerated nodes at AWS. One researcher who studies spiral waves in protoplanetary disks increased the resolution of his modeling fourfold—from 512-by-128-by-128, to 2,048-by-512-by-512. The second improved the accuracy of atomistically informed predictions for 3D manufacturing, narrowing the gap between the simulation-derived predictions and experimental data by more than two orders of magnitude.

Overall, by providing researchers with a seamless way to use AWS systems to run their jobs instead of waiting to access HECC resources, the users can determine for themselves whether the benefits of jumping the queue are worth the cost of running on commercial cloud resources. This is an important factor in helping NASA stay cost effective while meeting its current and future computing needs for its science and engineering projects, missions, and programs.

Robert Hood, Steve Heistand, NASA Ames Research Center  
 robert.hood@nasa.gov, steve.heistand@nasa.gov



*Two considerations are driving the process of integrating AWS resources into the HECC environment: convenience for users and security for NASA. Users can run AWS jobs either through cloudbursting via the HECC Portable Batch System (PBS) server or by logging directly into an AWS front end. Files can be staged between HECC systems and AWS or stored at AWS. Security is achieved by keeping all AWS and system logs in a separate AWS account, which also acts as a gateway for HECC staff to access the AWS console when needed. Steve Heistand, Sherry Chang, NASA/Ames*

## SUPERCOMPUTING

### Exploring Bespoke Supercomputer Options for Weather and Climate Workloads

In this study, we researched the options available to NASA and others to develop a high-performance computing (HPC) system specifically tailored for weather and climate research. We surveyed weather and climate users and potential HPC suppliers to gather key insights. Their key concerns centered on limitations—of memory, storage latency, and bandwidth—and the current reliance on GPUs not well suited to today’s weather/climate workloads. Our recommendations include a NASA initiative to: 1) more accurately assess the range of existing and planned workloads; 2) better provide specific hardware and software requirements; and 3) help determine the various strengths and opportunities of any potential HPC design.

NASA set out to explore alternative architecture solutions inspired by the 2016 release of ARM licenses, the formation of the Gen-Z Consortium, the availability of High Bandwidth Memory, and user community difficulties in adopting GPU-based solutions. One possibility is to specify system requirements for a bespoke supercomputer tuned for agency weather and climate workloads. NASA contracted with Hyperion Research to study the options and provide follow-on recommendations. The study was divided into two phases: Phase 1 involved interviewing weather and climate modeling centers around the world for requirements, and Phase 2 entailed surveying HPC vendors for their responses to the requirements.

Survey respondents provided a number of suggestions for future development, noting that any project to build such a system would need to be a community effort. Currently, such projects are seen as Center/Lab specific or based on “Grand Challenge” programs with relatively short-term horizons. Another major theme centered on bringing Artificial Intelligence—particularly machine learning capabilities—into the programming mix, but that opinion was not universally supported.

All of the vendors surveyed showed interest in exploring development of a bespoke HPC system for NASA. However, most indicated that the option for building a special-purpose, one-off system or even a series of such systems over time could not be economically justified—regardless of the amount of non-recurring engineering (NRE)—unless that particular design had value within the wider weather community or even across a number of complementary verticals. Other key points included the desire for NASA to: 1) conduct a robust, pre-competitive, codesign effort with potential vendors; and 2) reexamine and modernize NASA’s existing codebase to take advantage of existing and anticipated progress in high-end hardware and software.

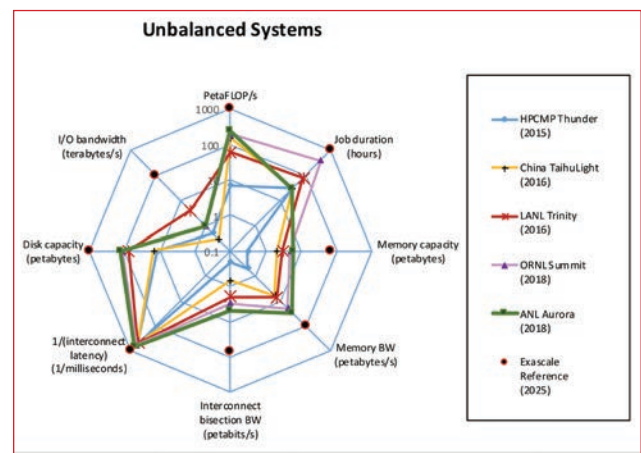
NASA scientists and engineers need the power and scale of HPC resources to understand the behavior of extremely complex systems such as weather, climate, heliophysics, astrophysics, aerospace vehicles, and launch systems.

Recommendations for next steps include: 1) mounting a NASA initiative to more accurately assess the range of existing and planned workloads that better meet specific hardware and software requirements; 2) composing benchmarks of mini-application suites, test cases, or even full applications to help determine the various strengths and opportunities of any potential HPC design; and 3) organizing a pre-competitive co-design conference, or series of conferences, that brings together interested commercial vendors to discuss various options and opportunities. A key goal of such co-design efforts would be to discuss projections of advances in HPC hardware and software that would better align with NASA requirements but are also in keeping with larger commercial trends.

*Tsengdar Lee, NASA Headquarters  
tsengdar.j.lee@nasa.gov*

*Robert Sorensen, Hyperion Research  
bsorensen@hyperionres.com*

*This figure, reproduced from the U.S. Department of Defense High Performance Computing Modernization Program, illustrates recent high-performance computing systems compared to a desired reference system. Areas of concern include memory capacity, I/O bandwidth, interconnect bisection bandwidth (BW), and memory BW. The NASA-Hyperion Research study reconfirms these concerns and suggests follow-on activities to better characterize the desired weather and climate workloads as compared to the reference system. Tsengdar Lee, NASA/Headquarters*



## SUPERCOMPUTING

### Achieving Quantum Supremacy with Noisy Intermediate-Scale Quantum Devices

Quantum computing promises a great deal of compute power, with the potential to revolutionize current computational capabilities and provide large speedups for solving problems in optimization, chemistry, materials science, and machine learning. “Quantum supremacy” is the ability of a programmable quantum processor to perform a well-defined task beyond the capability of the most powerful supercomputers on Earth. However, a demonstration of quantum supremacy on today’s noisy, intermediate-scale quantum (NISQ) devices needs a careful comparison with state-of-the-art classical simulators. Such a demonstration is on the verge of taking place, and will likely involve sampling from a random quantum circuit (RQC). It is well known that sampling from RQCs using classical computers is exponentially hard with the size and depth of RQCs. Achieving quantum supremacy experimentally will represent a historic milestone in quantum computing, and unlock a wide range of possibilities, including the better understanding of the performance of quantum algorithms on real quantum hardware.

Researchers in the Quantum Artificial Intelligence Lab (QuAIL) at NASA’s Ames Research Center in Silicon Valley, in collaboration with Google and Oak Ridge National Laboratory (ORNL), have developed a quantum circuit simulator. The goal is to provide the quantum computing community with a high-performing simulator to benchmark the performance of classical supercomputers against quantum processors, assess the advantage of noisy quantum devices in performing sampling tasks, and aid in the calibration of quantum hardware. Our simulator, called qFlex, is based on tensor network contractions, which provide a large amount of arithmetic intensity to the simulation procedure, and show high performance on current supercomputers. We implement a back-end algorithm for the contractions on CPUs, and use a highly optimized tensor contraction library for graphics processing unit (GPU)-based runs. In addition, our goal is to perform highly efficient (memory and time-wise) contractions of hard quantum circuits.

We have implemented a quantum circuit simulator that pushes classical supercomputers to performances above 60% of their theoretical peak, providing a fair comparison of quantum hardware with the world’s top supercomputers. We are able to benchmark square lattices of qubits, as well as Google’s Bristlecone quantum processor. Our simulator raises the bar for a quantum supremacy demonstration by benchmarking real hardware, solidifying a potential supremacy claim. It also provides further insight on the hardness of the problem and expands hardware verification possibilities. In addition, it provides a real-world application for supercomputers with efficiencies close to those of the LINPACK benchmark.



A dilution refrigerator for Google’s quantum computer. Image courtesy of Erik Lucero, Google

Visualization and data analysis experts at the NASA Advanced Supercomputing facility at Ames helped with code optimization and parallelization on both the Pleiades and Electra supercomputers. We used more than 90% of the nodes on these systems over about 12 hours, achieving 64% of the peak performance of both clusters together. Such high efficiencies on top HPC systems provide a comparison for quantum processors’ computational power and a threshold for them to surpass. In addition, verifying the quantum hardware requires large amounts of classical computational resources, which are only provided by large clusters.

*Salvatore Mandra, NASA Ames Research Center  
salvatore.mandra@nasa.gov*

*Benjamin Villalonga, University of Illinois at Urbana-Champaign  
vlllnc2@illinois.edu*

## SUPERCOMPUTING

### High-Performance Computing in the Azure Cloud

Cloud-based computing provides the possibility of highly dynamic resources that can be spun up on demand to meet computational needs. The NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center (GSFC) is testing cloud resources as a method for providing computational resources to meet a wide variety of science use cases. One example includes the creation of a digital twin of the Goddard Earth Observing System (GEOS) model, which provides research forecasts for NASA applications. In the event of on-premises system downtimes, the cloud-computing digital twin can provide continuous operations for these types of applications. In addition, many science applications and data analytics require a collaborative environment for NASA and non-NASA contributors. Cloud computing resources could readily be provisioned for these types of requirements.

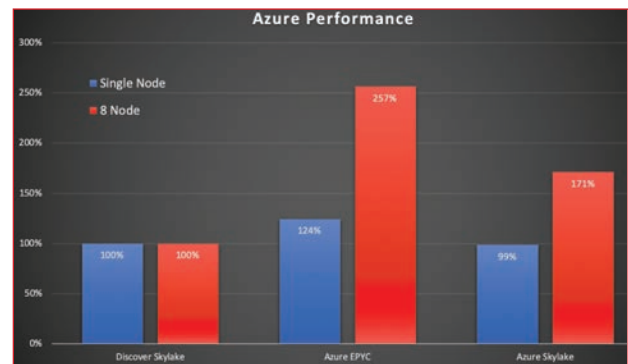
Microsoft has recently deployed Azure offerings targeted directly at high-performance computing (HPC) resources, with top-of-the-line CPUs and InfiniBand-based, remote direct memory access (RDMA) networking—critical for multi-node HPC workloads. These resources provide the NCCS with an ideal prototype for testing both operational HPC capabilities and application performance. This project prototyped the dynamic creation and rapid deployment of HPC clusters in Azure. These clusters were then used to benchmark highly relevant applications to the NCCS and create a cost comparison between on-site resources and the Azure cloud. As an added benefit, the Azure cloud has multiple processor types that can be used for vendor comparison.

We showed that, while cloud computing can provide commensurate performance for HPC applications, there are many challenges when employing cloud resources for HPC applications. Many basic assumptions of typical data center technologies clash with typical cloud architectures and best practices. We evaluated multiple best practices for rapid deployment of dynamic HPC clusters in the cloud and developed methodologies for quickly creating dynamic HPC clusters.

For static workloads, cloud-based resources are not yet competitively priced versus on-premises compute clusters. However, the rapid ability to stand up or expand resources provides a solution for production workloads and meets a wide variety of requirements, including on-premises downtimes. Alternately, cloud-based resources can provide workspaces that are isolated from strict security environments, allowing a location for collaboration between government and non-government personnel for low-sensitive, scientific data exploration. The ability to leverage these capabilities for an operational center similar to the NCCS provides users with more options to accelerate their scientific research.

Recent offerings from Amazon that include Elastic Fabric for higher-performance networking and managed Lustre environments offer interesting capabilities. In addition, the deployment frameworks recently developed on Azure will be evaluated for their cross-platform capabilities and directly compared against the features and performance of alternate cloud vendors.

*Reid Ormseth, NASA Goddard Space Flight Center  
reid.g.ormseth@nasa.gov*



*This graph represents the performance of the NASA GEOS climate model in the cloud using both Intel Skylake and AMD Naples versus Intel Skylake on premises. Single-node run, shown in blue, uses 2x3x6 grid decomposition and 36 cores; multi-node run, shown in red, uses 4x6x8 grid decomposition and 288 cores. Performance is relative to the on-premises Discover Skylake cluster. Lower is better! Mark Potts, RedLine Performance; Reid Ormseth, NASA/Goddard*

## SUPERCOMPUTING

### What Could Possibly Go Wrong? – The Myth of the Smooth Install

The Discover supercomputer cluster is the heart of the high-performance computing (HPC) services offered by the NASA Center for Climate Simulation (NCCS). The NCCS provides HPC and data storage services for NASA's Science Mission Directorate. Discover users include hundreds of NASA and university researchers who have come to rely on the regular technology upgrades of the cluster to maximize and accelerate their production of scientific results over the long term.

The NCCS has steadily increased the capacity of the Discover cluster since it was first installed in 2006 through the addition of Scalable Compute Units (SCUs). The latest expansion, SCU15, was integrated into Discover in the summer of 2019. This 2-petaflop upgrade increased the cluster's size to over 4,100 compute nodes containing over 129,000 CPUs and brought its total computing capacity to 6.7 petaflops. That translates into the ability to process over 6 quadrillion operations per second.

SCU15 is the first expansion to the portion of Discover using an Omni-Path high-speed network, while the rest of the cluster operates over multiple InfiniBand high-speed networks. This expansion also increased the number of Intel Skylake CPUs within Discover to over 46,000, allowing users to run higher-resolution models on the system.

SCU15 represents the 18th upgrade to the Discover supercomputer since its inception. Expanding from a 3.3-teraflop system with a mere five racks of hardware to a 6.7-petaflop system with over 80 racks of hardware has led to growing pains along the way. Everything from physical space to power, cooling, processors, interconnects, and cables provided challenges that were overcome for every successful deployment.

Each SCU installation had its own unique challenges. The NCCS has learned many lessons over the years that were successfully applied to subsequent installations. The installation of SCU15 was no exception. Installed in July 2019, it consists of many components that were similar or identical to the existing hardware deployed throughout the cluster. While the SCU15 vendor was a new partner for the NCCS, the challenge remained the same: to integrate the latest technology upgrade into the cluster as quickly and seamlessly as possible in order to provide additional computing capacity to the NCCS user community as fast as possible. Relying on 13 years of experience performing system upgrades, what could possibly go wrong?

With the latest upgrade to the Discover supercomputer, the NCCS continues to provide state-of-the-art HPC resources to the NASA High-End Computing (HEC) Program users in the form of more of the latest, most efficient processors with one of the fastest high-speed interconnects available today. This upgrade, as with all previous upgrades, will enable advances in NASA science and engineering research that would not be possible without computational resources like Discover.

*Bruce Pfaff, NASA Goddard Space Flight Center  
bruce.pfaff@nasa.gov*



*Scalable Compute Unit 15 (SCU15) is part of the Discover supercomputer cluster at the NASA Center for Climate Simulation (NCCS). SCU15 was installed in July 2019. It is an expansion to the Skylake and Omni-Path portion of the Discover cluster.  
Chris Meaney, NASA/Goddard*

## SUPERCOMPUTING

### The Evolution of NASA's High-End Computing Capabilities

For over 35 years, the NASA Advanced Supercomputing (NAS) Division at NASA's Ames Research Center has housed and managed the U.S. space agency's most powerful supercomputing assets. Focused on high-performance computing (HPC) technologies, efficient operations, and user success, the NAS Division has worked with industry to deploy a series of highly successful systems that have enabled scientific and engineering achievements across the agency.

The complementary role of the High-End Computing Capability (HECC) project is evolving to meet NASA's future challenges in returning to the Moon as a pathway to Mars, while continuing exciting research in aeronautics, space exploration, and Earth science.

In Fiscal Year 2019, HECC established NASA's newest Modular Supercomputing Facility (MSF), which houses an initial modular system named Aitken. This facility is the successful culmination of work accomplished over the past three years, through initial deployment of a prototype modular system and the subsequent construction of support infrastructure on a one-acre site that connects the system to the main NAS facility. Taking advantage of the temperate climate of the San Francisco Bay Area, the new facility utilizes outside air to cool NASA's next-generation supercomputers.

With the commissioning of this site, the NAS facility can potentially support 12 compute modules that will allow HECC to run more large jobs, helping to meet the agency's demand for supercomputing resources. Aitken is an HPE system with 1,152 nodes; if fully populated with similar nodes, the MSF would have over 40,000 nodes and 1.6 million cores in 12 interconnected modules.

In Fiscal Year 2019, the modular approach saved approximately 2 million kilowatt-hours of power and over 3 million gallons of water, when compared to what would be used if the hardware had been installed in the primary NAS facility. The module's power usage effectiveness (PUE) rating is an impressive 1.05—the same PUE as the prototype system—which far surpasses industry standards. The technology facilitates an operational approach to new system deployments that are not constrained by facility limitations; each expansion will be housed in a module that is customized to best support the systems deployed. The MSF will be able to provide up to 30 megawatts of power for future expansion.



*NASA's Modular Supercomputing Facility at Ames Research Center. The facility houses the agency's newest petascale supercomputer, named Aitken, to help solve NASA's most challenging problems in an environmentally conscious way that provides flexibility, power efficiency, and cost savings. Aitken's name comes from a large impact crater on the surface of the far side of the Moon's south pole, which was named for Robert Grant Aitken, an American astronomer who specialized in binary star systems. Derek Shaw, NASA/Ames*

This evolving technology allows NASA to continue providing essential HPC resources to accelerate and enhance the agency's mission of space exploration, scientific discovery, and aeronautics research by continually creating and ensuring optimal use of one of the most productive HPC environments in the world.

In the next year, the HECC project will concentrate on working with current and potential users across NASA to identify a comprehensive set of requirements necessary to achieve the agency's science and engineering goals. This will be used to design NASA's HPC roadmap and identify capabilities that will facilitate the successful transition to emerging technologies. The project will then identify risks that could impede these advances and work with industry to determine if the risks will be mitigated by planned technology development. When risk areas are not mitigated, HECC will partner with other NASA projects and programs, other national and international agencies, industry, and academia to conduct promising research that could lead to solutions such for risk areas.

*William Thigpen, NASA Ames Research Center  
william.w.thigpen@nasa.gov*

## THE UNIVERSE

### Exploring the Origins of Extreme Magnetism in Red Dwarf Stars

With NASA missions such as Kepler, K2—and now TESS—yielding ever more success in detecting exoplanets, it is clear that the question of habitability goes beyond simply being a terrestrial planet in a star’s “Goldilocks Zone.” As we have learned as residents of the solar system, stellar magnetic activity in the form of wind, flares, and coronal mass ejections can have substantial impacts on planetary atmospheres and surfaces. On and around Earth, the X-rays and charged particles released by the Sun’s activity can disrupt our communications networks, threaten power grids, and pose serious health risks to astronauts. For planets with less effective magnetic shielding, such as Mars or Venus, this activity can lead to dramatic atmospheric losses and alterations of atmospheric chemistry.


The majority of stars in the universe, including the solar system’s nearest neighbor, Proxima-Centauri, are red dwarfs—also known as M-dwarfs—but not one of them is visible to the naked eye. Compared to the Sun, M-dwarfs are small, cold, and very dim, but this has not diminished their magnetic activity. Instead, many M-dwarfs appear to have surfaces carpeted by strong magnetic fields, and give off frequent flares that may be a thousand times more energetic than those typical of the Sun. The radiation environment around such stars could have devastating consequences for their planets’ potential habitability. By employing high-resolution magnetohydrodynamic (MHD) simulations of these unassuming stars’ interiors, using NASA supercomputers, we are building toward an understanding of how M-dwarfs generate such intense magnetism, how that magnetism shapes their multi-billion-year evolution, and what it might mean for the exoplanets orbiting them.

We use the open source Rayleigh code to study the magnetized convection occurring within M-dwarf stars. So far, early (more massive) M-dwarfs have been completely untouched as the subject of thorough computational studies, so we have been undertaking extensive studies of the types of flows and dynamos that can be achieved within such stars. We have taken great interest in the role that a layer of rotational shear, known as the tachocline, may play in shaping the global fields of these stars, especially as they relate to the process of stellar spin-down.

The project is still young, but has so far yielded a number of results. Our models have demonstrated that the tachocline can act as a reservoir for magnetic fields in these stars, stabilizing and prolonging their global cycles even when it is not the primary generation site. When the tachocline participates directly in field generation, it enables certain dynamo configurations that are unstable without it. We found that transitions between dynamo states in these stars are capable of releasing sufficient magnetic energy in the form of buoyantly rising flux ropes to power observed superflares, whose deep origins are yet to be well explained. Finally, we found that the presence of a tachocline in these stars leads to their surface fields being stronger and organized on larger spatial scales, both of which allow these stars to shed angular momentum through their magnetized winds more quickly than stars without tachoclines.

This project could not exist without NASA high-performance computing resources. We make extensive use of the Pleiades supercomputer: each model requires anywhere from 1,024 to 4,096 cores and runs for dozens to hundreds of hours, depending on what result we hope to get out of it. In the process, each run generates multiple terabytes of output, which quickly balloons beyond what we are able to store locally.

*Connor Bice, Juri Toomre, University of Colorado Boulder and JILA  
connor.bice@colorado.edu, jtoomre@lcd.colorado.edu*



*Visualization of the near-surface convection achieved in one model of an M-type star, with its northern pole tipped toward the viewer to reveal structure. Broad, red upflows are separated by narrow, blue downflows. Near the equator, the influence of the fast stellar rotation causes the convection to organize into rolling, banana-shaped cells. A band of diminished convection separates these structures from the more isotropic flows at high latitudes.  
Connor Bice, Juri Toomre, University of Colorado Boulder and JILA*

## THE UNIVERSE

### How Does Space Weather Affect Mercury, Earth, and Uranus?

Space weather can have devastating effects on human health and technology. Improving our understanding of its impact on Earth can potentially lead to the mitigation of space weather damage through forecasting. Moreover, investigations of the solar wind's interactions with other planets will enhance the science returns of NASA's missions such as MESSENGER to Mercury, Magnetospheric Multiscale around Earth, and the Voyager flyby past Uranus. While significant discoveries are made by NASA's direct spacecraft observations, in situ measurements are often taken at limited points along the trajectories of planetary orbits or flybys. Such limitations can be alleviated by global numerical simulations, which allow interpretation of in situ measurements in a three-dimensional context and distinguish temporal from spatial fluctuations.

The primary science focus of our work is to study the solar wind interaction with Mercury, Earth, and Uranus by employing a 3D magnetosphere model containing the key electron physics that are essential to reproduce and interpret the observations. Many global models are based on either a fluid or hybrid (kinetic ion particles and massless electron fluid) approach for planets in our solar system and beyond. However, due to a lack of detailed electron physics, none of these models accurately treat collisionless magnetic reconnection. To solve this issue with affordable computational costs, we developed a novel multi-moment multifluid model capable of reproducing certain critical aspects of the reconnection physics from fully kinetic simulations.

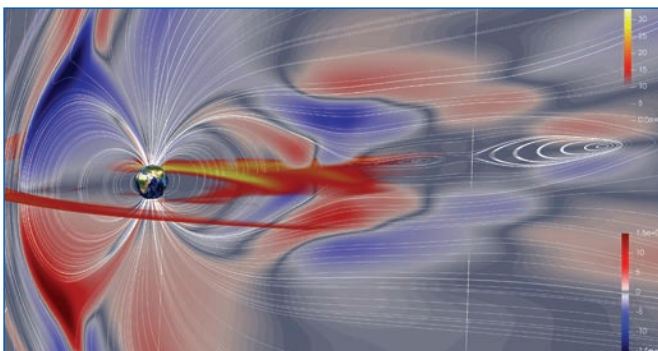
For the first time, we performed global (on the scale of many tens of planetary radii) simulations of interactions between the solar wind and planetary magnetospheres, including Mercury, Earth, and Uranus, using the ten-moment multifluid (multi-ion and single-electron) model. This model solves the continuity, momentum, and pressure tensor equations of each charged species, as well as the full Maxwell equations. It is noteworthy that the new model supports all kinds of electromagnetic waves and contains lower-order kinetic physics. Given that a kinetic particle model needs hundreds of particles in each cell for eliminating particle noise, while a fluid model shows no noise with only one fluid parcel per cell, one simulation can get at least 10 times faster by employing this kinetic fluid code.

The simulation results reveal highly dynamic responses of the planetary magnetospheres to varying external conditions. The model can reproduce observations beyond the traditional magnetohydrodynamic approach, including dawn-dusk asymmetries in a planet's magnetotail and field-aligned currents. The new model is essential for capturing the electron physics associated with collisionless magnetic reconnection in a planet's magnetosphere.

Given the significant impact of stellar winds on exoplanets residing within close-in habitable zones of M dwarfs, the investigation of the solar wind interaction with planets under extreme space weather conditions may have important implications for studying (exo)planetary habitability.

The Pleiades and Electra supercomputers at the NASA Advanced Supercomputing facility make it possible to perform our 3D high-resolution magnetospheric simulations and study the responses of different planets to varying space weather conditions. We used an unprecedented fine-grid resolution—the smallest grid size is 0.01 planetary radii—for this type of simulation, which require up to 1,300 nodes (that is, 52,000 cores).

*Chuanfei Dong, Liang Wang, Princeton University  
dcfy@princeton.edu, lw6@princeton.edu*



*Perspective view of the Earth's magnetosphere from a ten-moment multifluid simulation. The solar wind impinges from the left and carries interplanetary magnetic field lines. White streamlines are magnetic field lines in the meridional plane. Two magnetic bubbles on the night side (right side) are magnetic flux ropes. Blue-red contours in the meridional plane represent the out-of-plane component of the magnetic field in nanotesla (nT). Yellow-scale contours in the equatorial plane represent ion number density in unit per cubic centimeter (cm<sup>-3</sup>). The sharp boundary in ion density at the dayside (left side) marks the bow shock due to the fast-moving solar wind. Chuanfei Dong, Liang Wang, Princeton University*

## THE UNIVERSE

### Simulating Rocky Exoplanet Atmospheres with General Circulation Models

Are we alone in the Universe? This question has been asked since the first humans observed the night sky. Do stars other than Sun also have planets orbiting them? This second question was answered in 1992 with the discovery of the first exoplanet. Thousands of exoplanets have now been discovered, and some of the largest and hottest (the so-called “hot Jupiters” or “hot Neptunes”) have already been characterized by current telescopes such as the Hubble Space Telescope. To search for life-as-we-know-it, astronomers target rocky (and thus Earth-like) planets orbiting within the “habitable zone.” This zone corresponds to the planet-to-star distance where liquid water can be sustained on the planetary surface at atmospheric conditions similar to Earth. While dozens of such planets have already been discovered, their atmospheric characterization requires more advanced telescopes, such as NASA’s James Webb Space Telescope (JWST), that are going to be operational in the 2020 decade. Among the prime targets for such characterization are Proxima Centauri b—the closest rocky exoplanet to Earth—and TRAPPIST-1—a seven-rocky-planet system.


Upstream of JWST observations, scientists are using 3D general circulation models (GCMs) to evaluate potential climate states and their associated temporal- and spatial-dependent observable signals. NASA researchers have performed GCM simulations on the NASA Center for Climate Simulation (NCCS) Discover supercomputer. These simulations of Proxima Centauri b and TRAPPIST-1e assumed a modern Earth-like atmospheric composition and used two different GCMs: ROCKE3D, developed at the NASA Goddard Institute for Space Studies, and LMD-G, developed at the Laboratoire de Météorologie Dynamique in Paris, France. Presentations will discuss the habitability of these planets and their potential for atmospheric characterization.

Simulations show that Proxima Centauri b and TRAPPIST-1e are both capable of retaining liquid water on their surfaces for a variety of atmospheric compositions, making them potentially habitable. However, planets orbiting M dwarf stars may face tremendous challenges to preserving their atmospheres against stellar flares. JWST observations will be needed to answer the question “Can planets around M dwarf stars have atmospheres?”

GCM simulations aim to resolve—in 3D—the physics and dynamics of a planet’s atmosphere and therefore require significant computational resources. The two GCMs used in this work, ROCKE3D and LMD-G, have been run on the NCCS Discover supercomputer. Depending on spatial resolution, approximately 50 cores are simultaneously used per simulation, which can last from a few days to a few weeks.

*Thomas J. Fauchez, NASA Goddard Space Flight Center  
thomas.j.fauchez@nasa.gov*

*Anthony D. Del Genio, NASA Goddard Institute for Space Studies  
anthony.d.delgenio@nasa.gov*



*Surface temperature for the Proxima Centauri b exoplanet, assumed to be fully covered by an ocean with a modern Earth-like atmosphere, as simulated with the ROCKE3D global circulation model developed at the NASA Goddard Institute for Space Studies. The plot is centered at the substellar point, the point on the planet’s surface at which the star is at zenith. Temperature is above freezing near the substellar point, leading to liquid water and therefore habitable conditions at the surface. Anthony D. Del Genio, Michael J. Way, NASA Goddard Institute for Space Studies*

## THE UNIVERSE

### Simulating the Water Cycle and Cloud Formation on Mars

The water cycle is one of three coupled seasonal cycles that control the climate of Mars. A full understanding of how water is exchanged between the atmosphere and the surface, moves through the atmosphere, and forms clouds requires the use of sophisticated global climate models (GCMs) that have been adapted to Mars. One such GCM, under development for several decades at NASA's Mars Climate Modeling Group at Ames Research Center, has recently been updated to include the cubed-sphere finite volume dynamical core developed at the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory (NOAA/GFDL). This modern dynamical core is highly parallelized, scalable, and flexible, which enables global, multi-year, high-resolution modeling of Mars' atmosphere and the ability to investigate a multitude of open questions about water, clouds, and the Martian climate.


Modern GCMs adapted to Mars yield critical insight into the importance of the water cycle and water ice clouds in the Martian climate system. Water ice clouds form when atmospheric dust is present to provide ice nuclei and when cloud nucleation and growth are thermodynamically favored. Like on Earth, clouds can either cool or warm the atmosphere and/or the planet's surface depending on their location, altitude, and optical properties. We use the GCM developed at NASA's Mars Climate Modeling Center to investigate the water cycle and how water ice clouds affect the Martian climate.

We have carefully compared the GCM results with spacecraft observations and find that radiative effects of Martian clouds, even though they are thin compared to Earth's clouds, significantly change the thermal structure of the atmosphere and intensity of its global-scale wind systems. In this way, they help control the movement of water around the planet. In the distant past when orbital properties were different, we have found that clouds could provide a strong greenhouse effect and warm the surface by tens of Kelvin. Our studies directly address NASA's goals to understand how bodies in the solar system evolved and they assist the agency's mission planning activities.

The high-performance computing resources provided by the NASA Advanced Supercomputing (NAS) Division enable three necessary requirements for the success of this project: visualization of model results, high spatial resolution capability, and fast throughput. Visualization of the four-dimensional output from GCMs is absolutely critical for interpreting the results, and running at high spatial resolution on hundreds or thousands of computational cores allows us to simulate important processes not resolvable at common grid spacings. Since we need to simulate many years of the Mars climate system to address the science questions of interest, the computational power of NASA's Pleiades supercomputer allows us to study the impact of many different processes on practical time scales.

We will continue to improve our representation of physical processes, increase the model spatial resolution, and conduct longer and longer simulations to study interannual variability.

*Melinda Kahre, NASA Ames Research Center  
melinda.a.kahre@nasa.gov*



*Simulated water ice clouds over the Tharsis Montes volcanoes during a summer afternoon on Mars. The simulation was run on NASA's Pleiades supercomputer.  
Alex Kling, David Ellsworth, NASA/Ames*

## THE UNIVERSE

### Modeling the Solar Corona to Study Sources of Space Weather Disturbances

Our everyday lives are closely linked to solar radiation and dynamics. The Sun is an active star with quasi-periodic, 11-year cycles of magnetic activity. During periods of high solar activity, strong plasma eruptions such as flares and coronal mass ejections—accompanied by high-energy radiation and energetic particles—disturb the near-Earth space environment and cause magnetic storms. Such activity can affect communication, power grids, and navigation and space technologies, and can also affect human health.


Detailed modeling of the complex dynamics of the magnetized turbulent solar plasma allows us to understand the origins of solar activity and the triggering mechanisms of plasma eruptions. Such studies provide critical support for interpreting observational data from NASA's space missions, such as the Solar Dynamics Observatory, IRIS, Hinode, and the Parker Solar Probe, and contribute to preparations for the Artemis Program and human space exploration in general. We performed 3D radiative magnetohydrodynamics (MHD) simulations to study the processes of plasma dynamics and the formation of magnetic and thermodynamic structures, which lead to energy accumulation in the solar corona and its subsequent release. These studies help us understand the complex evolution of the solar atmosphere and corona and provide insight into the physical processes that are important for improving space weather forecasts.

Our goal is to perform 3D radiative MHD simulations of solar dynamics, in which the computational domain includes sub-surface layers of the Sun as well as the solar atmosphere and corona, to study processes of plasma heating and plasma eruptions. Such studies are critical for building reliable dynamical forecasts of coronal activity. The simulations have been performed for different solar conditions using the StellarBox code developed at NASA's Ames Research Center. The code takes into account the effects of turbulence, ionized plasma, magnetic fields, chemical composition, and radiation in the highly stratified realistic conditions on the Sun.

Our simulation results reveal that initially uniform weak magnetic fields are greatly amplified due to small-scale dynamo action below the visible surface of the Sun, and cause the spontaneous formation of tornado-like plasma structures and eruptions. The strongest helical flows originate in strong, 1-kiloGauss magnetic field patches formed on the solar surface. The helical magnetized structures extend into the corona and produce Alfvén waves that drive plasma eruptions. The simulations also reveal many important details that are unresolved in observational data, such as the Kelvin-Helmholtz instability of the magnetic structures and plasma downflows in the corona. The current study is primarily focused on high-resolution simulations that model conditions in local areas of the Sun in order to understand the origin of various types of wave and magnetic self-organization phenomena and the triggering mechanisms of small-scale eruptions observed by NASA satellites.

3D modeling of turbulent ionized plasma under realistic solar conditions requires large-scale computational resources. We use the Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility to perform our high-fidelity simulations. Multidimensional massive data flows in such simulations require a very high-speed network for interprocessor communication and for postprocessing and visualization. In this work we primarily used Intel Ivy Bridge and Skylake nodes. The NAS teams provide several critical support needs for this project: 24/7 user support, visualization, and computer science expertise.

*Irina Kitiashvili, Alan Wray, NASA Ames Research Center  
irina.n.kitiashvili@nasa.gov, alan.a.wray@nasa.gov*



*View of a coronal structure from the top of the computational domain, revealed by tracking particle motions. Red shows regions where the solar plasma is heated above 1 million degrees Kelvin. Visualizations are performed by advecting particles seeded at 1,400 miles above the solar surface. Irina Kitiashvili, Timothy Sandstrom, NASA/Ames*

## THE UNIVERSE

### Discovering Distinct Cycles in the Global Solar Dynamo

Modern society is intimately dependent on electronic communications, satellite observations, and large-scale power grids, all of which are frequently affected by space weather phenomena such as solar flares and coronal mass ejections. All space weather fundamentally has its roots in the solar magnetic field, but the global dynamo producing this magnetism remains difficult to constrain. Global simulations of the Sun's interior magnetism, run on NASA's massively parallel supercomputers, provide the most effective route toward understanding the global solar dynamo. Using these simulations in concert with extensive NASA observations of the Sun's surface magnetic field will ultimately help scientists predict space weather, allowing society to be better prepared for a possibly catastrophic space weather event

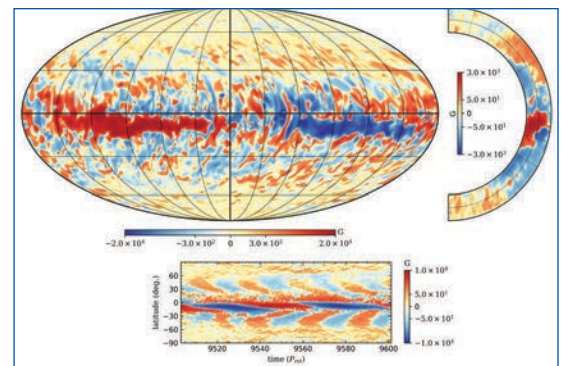
Scientists don't know where sunspots originate in the highly turbulent solar convection zone (CZ), or the physical causes of sunspot emergence. The goal of this project is to characterize the nonlinear dynamo responsible for the cycles in solar magnetic activity (the most prominent of which is the 22-year sunspot cycle). We use the sophisticated, open source, nonlinear magnetohydrodynamic (MHD) code Rayleigh to simulate a magnetized fluid in a rotating, convecting spherical shell. Our simulations reproduce regular cycles during which globally connected "wreaths" of magnetism reverse polarity and migrate toward the equator, two features highly reminiscent of the sunspot cycle. Additionally, there is an asymmetric cycle, during which strong magnetism preferentially occupies one hemisphere at certain longitudes, reminiscent of the "active longitudes" and "active hemispheres" observed on the Sun. This new class of simulations allows us to better synthesize our understanding of distinct features of the solar cycle.

Our new class of simulations builds on previous numerical work aimed at understanding equatorward propagation in the sunspot cycle, regular polarity reversals, and active longitudes. We offer for the first time a systematic, integrated perspective on these disparate phenomena and present a detailed description of the MHD processes at work in our simulations. We will produce publicly available data products that will allow modelers of the photosphere to cleanly compare our results with observations from NASA's Solar and Heliospheric Observatory (SOHO) and Solar Dynamics Observatory (SDO) missions.

A global simulation of the solar dynamo that hopes to capture much of the small-scale, turbulent structures associated with convective flows requires several million core-hours of computer time. Thus, our work would not be possible without the infrastructure of the Pleiades supercomputer at the NASA Advanced Supercomputing (NAS) facility. Over the last year, we used about 15 million core-hours for this project and generated over 20 terabytes of data. We frequently relied on the excellent support provided by NAS staff to allow the Rayleigh code to scale up to run on 8,192 Broadwell cores, or 293 Broadwell nodes. The videos for this work, as well as much of the MHD analysis, were all made on the high-memory (756 GB) data analysis nodes on Pleiades.

Historically, Parker's "interface dynamo" has been used to explain the regular 22-year sunspot cycle. This theory attributes a central role to the tachocline of shear at the base of the CZ, where the differential rotation (fast equator and slow poles) transitions sharply to solid body rotation. The next step in the current study is to include a stably stratified region below the CZ and a tachocline of shear at the transition between these two regions.

*Loren Matilsky, Juri Toomre, University of Colorado Boulder and JILA  
loren.matilsky@colorado.edu, jtoomre@lcd.colorado.edu*



*Snapshot of the azimuthal magnetic field achieved in one of the global solar dynamo simulations produced at the University of Colorado Boulder Department of Astrophysical and Planetary Sciences. Cuts through the field are taken on a spherical surface at mid-depth (upper left), along a meridian (upper left), and in time-latitude space (bottom). Loren Matilsky, University of Colorado Boulder and JILA*

## THE UNIVERSE

### Simulating the Diffuse Gas Surrounding Galaxies


Galaxies grow by converting gas into stars. They acquire this gas from their diffuse environs, known as the circumgalactic medium (CGM). As stars evolve and die as explosive supernovas, they eject freshly produced heavy elements such as oxygen and iron back into the CGM, along with prodigious amounts of energy and momentum. Some of this gas will then recycle back into galaxies to fuel future generations of star formation. Thus, understanding the physics of this dynamic and structurally complex medium is crucial to understanding how galaxies form and evolve. Thanks in part to thousands of hours of observations made with NASA's Hubble Space Telescope and other observatories, we now know the CGM to be a massive and extensive reservoir of galactic gas and heavy elements. Circumgalactic gas is generally too diffuse to be seen in emission, so rather than taking pictures of it, we observe it in absorption towards bright background light sources, akin to using cosmic lighthouses to observe this cosmic fog.

The Figuring Out Gas & Galaxies In Enzo (FOGGIE) project uses the Enzo cosmological hydrodynamic code to model the co-evolution of galaxies and their gas, with a focus on resolving the ultra-diffuse CGM with unprecedented fidelity. We "observe" these simulations in a manner similar to how the real CGM is observed; these synthetic data help us interpret the real observations as, unlike in the real universe, we know what the underlying gas physics is in the simulations. These simulations reveal a richly structured circumgalactic medium full of churning turbulent gas, small clouds, and tenuous hot gas. Observations made with Hubble have shown that low-ionization gas, which should be relatively cool and higher-density, often has kinematic structure very similar to more highly ionized gas, which is expected to generically be hotter and lower-density. With FOGGIE's exquisite resolution, we have started to disentangle this puzzle: sometimes kinematically coincident gas is actually co-spatial, but often it is not; rather, the observation is a chance superposition along the line of sight. By directly relating the observational signatures in the synthetic data back to their underlying physical conditions, we are learning how to better glean information about the real CGM and its relation to galaxies.

The FOGGIE simulations could not be run without high-performance computing resources. Evolving a single galaxy and its gaseous halo at high resolution across cosmic time typically takes a few million core-hours on the Pleiades supercomputer at the NASA Advanced Supercomputing (NAS) facility; this corresponds to wall-clock times of over a year and hundreds of thousands of node-hours on the system's 128-GB-memory Intel Xeon Broadwell and Haswell nodes. The large memory per core is critical because these kinds of simulations are particularly memory intensive: they are complex, multiphysics calculations following many fields (density, temperature, chemical composition, fluxes, and so on), combined with the tens of millions of resolution elements needed to resolve the CGM. In addition, the long-term mass storage system at the NAS facility is important for enabling analysis: each output can be ~15 gigabytes, and with 2,500 outputs per run, it's crucial to have access to hundreds of terabytes of space with fast I/O in order to analyze the detailed time evolution of the galaxies and their gas. The next steps for FOGGIE are to extend the small-scale resolution of the CGM to larger volumes for more galaxies and over more of cosmic time. To do this, we are transitioning to the Enzo-E code, an updated version of Enzo that is scalable to thousands of processors.

*Molly S. Peeples, Space Telescope Science Institute/Johns Hopkins University  
molly@stsci.edu*

*Snapshot of a visualization showing gas (left) and stars (right) in and around a simulated galaxy. While the starlight is easy to observe, it does not tell the whole story of what the galaxy is doing. The gas density is highest (orange/yellow) at the center of the galaxy, but the lower-density gas (blue/black) still exhibits complex structure. The simulation was run by the Figuring Out Gas & Galaxies In Enzo (FOGGIE) project using the Enzo cosmological hydrodynamic code. Raymond Simons, Space Telescope Science Institute (STScI); Molly S. Peeples, STScI/Johns Hopkins University*



## THE UNIVERSE

### Numerically Modeling the Weather and Climate of Mars

Martian atmospheric and climate modeling is of both scientific and engineering importance. As a global system that is in many ways analogous to the Earth's climate system, better study of Mars allows us to expand our understanding of fundamental planetary climate system behaviors, and allows us to test terrestrial meteorological and climate models in novel regimes. These models are then used in practical application to assist in the design and operation of NASA spacecraft bound for Mars.

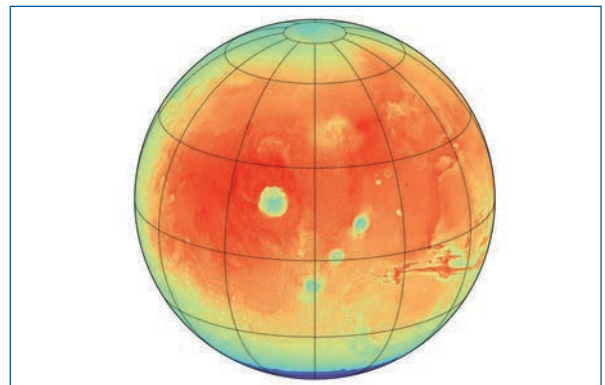
In our research, we have developed Martian adaptations of the widely used National Center for Atmospheric Research (NCAR) Weather Research and Forecasting (WRF) and Model for Prediction Across Scales (MPAS) terrestrial micro-, meso-, and global scale models. The models are used on NASA supercomputers to study processes on scales from the decimeter-scale boundary layer winds that will influence the flight of the 2020 Mars Helicopter, to the long-term evolution of Mars' global dust and water cycles, and the role of trace gases and carbon dioxide clouds in warming the ancient climate system.

Most of the core problems in Martian meteorology and climate dynamics require study via detailed fluid dynamical simulation. These problems are computationally intense and require efficient numerical approaches to take maximum advantage of parallel computation. WRF is designed to use such systems, and can be efficiently run across hundreds of processes. WRF is sufficiently flexible that the same system can be configured to run as a limited-area boundary layer model (grid spacing of about 10-100 meters), as a full global model (grid spacing of about 10-300 kilometers), or as a nested model with multiple domains at varied resolution embedded within each other. However, the structured nature of the WRF mesh limits the linearity of scaling as the number of processes exceeds several hundred. The latest NCAR model, MPAS, addresses this with an unstructured mesh that scales much more effectively. The ability to run higher-resolution, larger-domain simulations enables several advances, including: much more detailed examination of processes involving the interaction between organized convective systems and global scale circulations (such as dust storm growth); analysis of the interaction of boundary layer thermal and shear shows with mesoscale flow systems; and study of the coupled lower atmosphere and thermosphere.

Numerical modeling allows us to assess disparate spacecraft and telescope observations within a consistent theoretical framework. Our results enable a wide range of observations to be conceptually understood. For example, they allowed Curiosity rover observations of surface pressure to be understood in terms of a locally induced hydrostatic rebalancing of air over the widely varying local topographic relief, in addition to the more commonly understood large-scale thermal tides. Our modeling has also led to better understanding of the roles of surface dust availability and putative orbital angular momentum coupling in the occurrence of global dust storms. The wide variety of studies supported by our modeling helps NASA achieve its goal of understanding how planetary systems work, and assists the agency in designing and operating spacecraft.

Our research program would be impossible without NASA's Pleiades supercomputer. Typically, a single low-resolution, decadal global simulation requires from tens to hundreds of processes for about 10 days. Our highest resolution eddy-resolving simulations can require days across thousands of processes. For some of these cases, intermediate data volumes in the range of tens of terabytes are required in order to generate final archived data files of up to hundreds of gigabytes.

*Mark Richardson, Aeolis Research  
mir@aeolisresearch.com*



*Daytime near-surface (8 meters) air temperature simulated by the Mars Weather Research and Forecasting (MarsWRF) model projected on a globe. Deep blue: 150 Kelvin (K); deep red: 230 K. Major Martian volcanoes are identifiable due to their cooler air temperatures, and appear as roughly oval shapes. Upper left: Olympus Mons; Arsia, Pavonis, and Ascraeus Mons aligned center right. Lower right: the warmer Valles Marineris. Color pixels represent individual model grid cells at less than 12 km spacing. Mark Richardson, Aeolis Research*

## THE UNIVERSE

### The Milky Way and Its Neighbors: Simulating Ultra-Faint Dwarf Galaxies

Recent surveys made with observatories such as NASA’s Hubble Space Telescope have greatly increased our knowledge of nearby dwarf and ultra-faint dwarf (UFD) galaxies. However, many uncertainties remain—including galaxy star formation histories, chemical compositions, kinematics, and distributions around the Milky Way. Modern cosmological hydrodynamic simulations have not been able to probe UFDs in the distant field, and because our best observational samples are near the Milky Way, the ability to interpret current and future observations requires Milky Way simulations that are capable of resolving down to the UFD range.

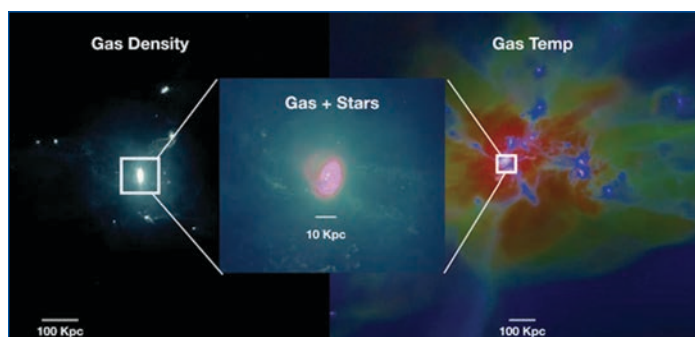
To meet this challenge, we developed our “DC Justice League” suite of Milky Way cosmological zoom-in simulations, which we leverage to study the small galaxies neighboring these Milky Way analogs and make predictions for future surveys by space observatories now in development, such as NASA’s Wide Field Infrared Survey Telescope (WFIRST) and James Webb Space Telescope (JWST). The DC Justice League simulations were produced using the fully cosmological, hydrodynamic code ChaNGa (Charm N-body GrAvity solver), massively parallelized to take advantage of HPC resources using a novel Charm++ scheme. ChaNGa contains the physics to model supernova feedback, star formation, and gas outflows in order to create realistic simulated galaxies easily comparable to observations from modern telescopes.

The DC Justice League is the first suite of cosmological simulations capable of resolving down to the ultra-faint regime. The simulated dwarf and UFD galaxies reproduce a variety of observations across many orders of magnitude in luminosity, with a wide range of shapes, star formation histories, and gas contents. This large sample of simulated dwarfs and UFDs around the Milky Way will provide an ideal tool for analyzing and understanding upcoming observations by WFIRST, JWST, and other telescopes.

In one notable result, our simulations showed several UFD and faint dwarf galaxies that have appreciable gas mass, despite not being star-forming. This prediction comes as a surprise to observers hunting for these objects.

We used the significant HPC resources at the NASA Advanced Supercomputing (NAS) facility in order to produce the DC Justice League simulations. This year alone, our flagship “Sandra” simulation of the Milky Way used 1.4 million node hours (28 million processor hours) on the Pleiades supercomputer, each job running on 400 Ivy Bridge nodes. Terabytes of data output from our simulations are being stored on the NAS facility’s mass storage system. Without these HPC resources, it would not be possible to achieve our groundbreaking simulations.

*Jordan Van Nest, Ferah Munshi, University of Oklahoma  
jdvannest@ou.edu, fdm@ou.edu*



*A rendering of the Milky Way-like galaxy simulation “Sandra” in gas density (left), gas and stars (middle) and gas temperature (right). In the left image, the brightest regions show the densest gas; in the middle image, the red and blue represent old and young stars, respectively; and in the right image, hot gas is red, and cold gas is blue. Notice that around Sandra there are a plethora of smaller, dwarf galaxies, which are the focus of this study. Elaad Applebaum, Rutgers University; Ferah Munshi, University of Oklahoma*

**Notes**

**For the latest schedule of events taking place in the booth,  
check our website at:**

**[www.nas.nasa.gov/SC19](http://www.nas.nasa.gov/SC19)**

**Or follow us on Twitter:**



**[@NASA\\_Supercomp](https://twitter.com/NASA_Supercomp)**

*For more information on any of NASA's SC19 demos, you can either talk to a researcher or visit <http://www.nas.nasa.gov/SC19/demonstrations.html>*