

**SC16  
NASA Research Exhibit  
Booth 1511**

**Abstracts are for presentations to be given at the SC16 NASA Exhibit,  
At SC16 in Salt Lake City, UT  
November 14-17, 2016**

**Abstracts will also be posted to NASA@SC16 website  
beginning 11/15/16**

## Aeronautics

### 1) The Fruit of HPC: Experiments in Turbulence

#### Overview

Because it is normally turbulent, the flow of air over a flight vehicle must be approximated using engineering models – models that are in many important cases too inaccurate for industrial purposes. This study addresses this deficiency by using HPC to generate realizations of turbulent flow that serve as benchmark data against which the turbulence models can be tested, and in due course improved. This effort is part of NASA's Transformational Tools & Technologies Project, sponsored by the Fundamental Aerodynamics program, which aims to develop the computational tools necessary to design and certify the next generation of aircraft.

We are particularly interested in model predictions of whether or not the turbulent flow of air immediately adjacent to the vehicle – known as the “boundary layer” flow – remains “attached” to or “separates” from the vehicle surface. This is both because model predictions of separation are too unreliable, and also because, while a separated boundary layer leads to severe degradation of aerodynamic efficiency, the state of incipient separation is often associated with optimal flight conditions. In other words, the task that is arguably the model's most important is the one with which it has the most difficulty.

#### Project Details

HPC is used to perform a series of *virtual experiments* in a *numerical wind-tunnel*, to produce high-fidelity solutions to the equations governing the behavior of turbulent fluid flow involving boundary-layer separation. Although we are not yet able to duplicate exact flight conditions, in terms of the effective size and speed of the aircraft, we can produce cases physically very similar to full-scale behavior. A family of such cases has been considered, and compared to predictions of the models. The discrepancies between the models and the HPC data are currently being assessed, and the first finding is that all models predict separation too early.

#### Results and Impact

The advantage of this HPC-based strategy is that, unlike traditional laboratory experiments, it produces accurate *non-intrusive measurements* of any quantity of interest, anywhere in the flow. These measurements can be used to probe the internal mechanisms of the turbulent separation and the model's attempts to replicate them, as well as generate the statistics the models are called upon to predict. Examples of the latter include the distribution of skin friction over the surface, and the location of the line of mean separation; both have been compared to a range of current state-of-the-art turbulence models, revealing their shortcomings, and quantifying the necessary improvements. These improvements could lead to superior design, in terms of flight characteristics and fuel efficiency, as well as lower risk and cost, of future aircraft.

#### Role of High-End Computing (Why HPC Matters)

Because turbulent flows are unsteady and contain three-dimensional *eddies* that range from very large to very small (and very slow and very fast), this exercise requires the vast amounts of memory and high processor speeds only available via HPC. Without access to HPC (specifically the NAS Pleiades system), the quality and usefulness of

this study, measured both in terms of the number of cases and the values of some of their defining parameters (such as Reynolds number and sweep angle), would have been severely compromised. The progress, compared with a similar study we conducted in 1997, is considerable.

*Presenter at SC16: Gary N Coleman, NASA Langley Research Center,  
g.n.coleman@nasa.gov*

*Technical PI: Philippe R Spalart, Boeing Commercial Airplanes  
philippe.r.spalart@boeing.com*

## **2) An Optimized Multicolor Point-Implicit Solver for Unstructured Grid Applications on GPUs**

### **Overview**

NASA Langley Research Center's FUN3D computational fluid dynamics (CFD) software is used to solve the Navier-Stokes (NS) equations for a broad range of aerodynamics and multidisciplinary applications across the speed range. Accurate and efficient simulations of complex aerodynamic flows are challenging and require significant computational resources. FUN3D uses a fully-implicit strategy based on a finite-volume spatial discretization on mixed-element unstructured grids, which requires the frequent solution of large tightly-coupled systems of block-sparse linear equations of the form  $Ax=b$ . In this work, an efficient implementation of the solver for graphics processing units (GPUs) is proposed. The solver was initially formulated to leverage two existing CUDA library functions. However, numerical experiments showed that the performance of these functions was suboptimal for matrices representative of those encountered in FUN3D simulations. Instead, optimized versions of these functions using varying levels of parallelism have been developed. The new implementation shows performance gains of up to 7x over the approach based on existing CUDA library functions.

### **Project Details**

To develop an efficient GPU implementation of the multicolor point-implicit solver, functions provided by the *cuSPARSE* and *cuBLAS* libraries have first been considered. The function *cusparseSbsrmv* multiplies a block-sparse matrix with a vector, and the function *cublasStrsmBatched* solves block systems of equations by performing forward and backward substitutions using a previously-factored coefficient matrix. Numerical experiments showed that the performance of these library functions is suboptimal for linear systems representative of those encountered in FUN3D simulations.

Optimized implementations of the *cusparseSbsrmv* and *cublasStrsmBatched* functions are developed. To perform a sparse matrix-vector product, an algorithm that allocates a number of warps, or a group of 32 threads, to process a subset of the blocks in one row of the sparse matrix is proposed. A separate kernel is used to perform forward and backward substitutions. Several challenges are encountered, including a variable extent of available parallelism, indirect memory addressing, low arithmetic intensity, and the need to accommodate different block sizes. To address these challenges, particular emphasis is placed on coalesced memory loads, the use of shared memory and pre-fetching, minimal thread divergence within warps, and strategic use of shuffle instructions available on recent hardware.

### **Results and Impact**

All results have been generated using an Intel Xeon E5-2670 dual socket, eight-core host processor with an NVIDIA K40 GPU device. Performance is evaluated using a

series of block-sparse matrices based on tetrahedral meshes. For a given sparsity pattern, matrices are populated using random values over a range of block sizes. Factors critical to the performance of the underlying matrix-vector multiplication and forward and backward substitution operations have been identified, and several challenges specific to the target application have been addressed. Speedups of as much as 7x over the library-based implementation have been observed for a broad range of matrix and block sizes.

### **Role of High-End Computing (Why HPC Matters)**

Accurate and efficient simulations of complex aerodynamic flows require frequent solutions of billions of simultaneous equations. These intensive calculations can be efficiently performed using emerging high performance computing hardware such as GPU accelerators.

*Presenter at SC16: Mohammad Zubair, Old Dominion University, [zubair@cs.odu.edu](mailto:zubair@cs.odu.edu)  
Technical PI: Eric Nielsen, NASA Langley Research Center, [Eric.J.Nielsen@nasa.gov](mailto:Eric.J.Nielsen@nasa.gov)*

## **3) Toward Quieter Skies: New Tools and Technologies to Predict and Mitigate Airframe Noise**

### **Overview**

The projected growth in air travel requires aggressive mitigation of the environmental impact of civil aircraft operations. Aircraft noise adversely affects communities surrounding major airports. Airframe noise is a prominent component of aircraft noise during landing. Under the Flight Demonstrations and Capabilities (FDC) project of the NASA Aeronautics Mission Research Directorate, researchers are working on advancing the state-of-the-art in simulation-based airframe noise prediction methodologies. The effort, which is being performed in collaboration with the Boeing Company, focuses on extending high-fidelity simulations to large civil transports. The complexities of the flow field and the size of the aircraft make achievement of such simulations the “holy grail” of airframe noise prediction.

During approach to landing, noise generated by the airframe is comparable to, and in some instances even louder than, propulsion noise. For large civil transports, the aircraft undercarriage, in particular the main landing gear, is the most dominant noise source. The extreme geometrical complexities associated with the main landing gear, which is composed of a multitude of sub-systems, make this component one of the most challenging configurations for aeroacoustic behavior simulation.

### **Project Details**

As part of the FDC-Boeing collaboration on airframe noise research, we simulated the airflow over a high-fidelity, 26%-scale Boeing 777 main landing gear model both in isolation and as installed on a semi-span model of the same aircraft in landing configuration. Our goal was to accurately capture the noise sources associated with the various components of the main gear. Then, we evaluated the effectiveness of a toboggan-shaped noise-reduction device installed on the gear model. The time-dependent simulations were obtained using the Exa Corporation PowerFLOW® Lattice-Boltzmann flow solver. Approximately three billion volumetric cells were used to discretize the space surrounding the main landing gear.

### **Results and Impact**

Our simulations accurately predicted, for the first time, the noise generated by the main landing gear of a large civil aircraft with all of its extreme geometrical complexities included. Simulating the unsteady flow field for the main landing gear of a large aircraft, which is a complex system-of-systems, was a daunting task. For the isolated gear, the computed farfield noise for flyover and side-line directions is in excellent agreement with wind tunnel measurements of the same model. Adding the toboggan device to the main gear resulted in a 3-4 dB reduction in farfield sound pressure levels over most of the frequency range. Preliminary simulations of the main landing gear installed on the semi-span model indicate that significant interaction occurs between the flows associated with the main gear and high-lift devices.

### **Role of High-End Computing (Why HPC Matters)**

The large-scale calculations needed to conduct any of our simulations require usage of the largest supercomputers. All the simulations were performed on NASA's Pleiades supercomputer using approximately 4,000 processors. Depending on the nature of the information required, the computational runs could last from two to three weeks, resulting in hundreds of terabytes of data. Rendering and effective visualization of such massive datasets were made possible by members of the NAS visualization team at NASA Ames.

### **What's Next**

Our plan is to continue the high-resolution simulations of the 26%-scale semi-span model of a large civil aircraft in landing configuration, with and without the main gear toboggan device installed. Those simulations will pave the way for our future attempt at simulating a full-scale, complete aircraft similar to the Boeing 777 in landing configuration.

*Presenter at SC16 and PI: Mehdi R. Khorrami, NASA Langley Research Center  
mehdi.r.khorrami@nasa.gov*

## **4) Towards the Development of Hypersonic Transportation**

### **Overview**

Supersonic-combustion ramjets are a new class of air-breathing engines that could one day flight you from New York to Tokyo in just two hours! As fast as a rocket, they can run much longer as they do not need to store oxygen to burn: they breath it as conventional plane engines do. Additionally, they do not spend part of the propulsion energy to run a compressor. Instead, the compression naturally results from a complex system of multiple shock-turbulent boundary layer interactions called a "shocktrain" residing at the inlet. The shocktrain sensitivity to external perturbations can be detrimental to the compression ratio, delaying ignition and eventually decreasing thrust. Ensuring efficient burning at supersonic speed is indeed challenging given the very short residence time in the combustor. It is a downward cycle, and understanding both the shocktrain complex response to external perturbations and the mechanics of supersonic combustion are key design considerations.

### **Project Details**

We perform Direct Numerical Simulations (DNS) of dual-mode scramjet isolator to study the impact of frequency-dependent instabilities. The dynamic response to upstream perturbation allows for the identification of resonating frequencies that are of critical

interest in designing a stable compressor. We also work toward developing a better understanding of the underlying physics to improve numerical modeling of such complex systems, by evaluating a posteriori the validity of common Reynolds-Averaged Numerical Simulation (RANS) models with the DNS.

Additionally, it is known that compression shocks leave high-lying vibrational states underpopulated in the post-shock region. This can alter the ignition process because chemical reactions are highly dependent on the internal energy distribution. Numerical simulations of a model scramjet combustor are also performed to understand the impact of such thermodynamic nonequilibrium on ignition.

### **Results and Impact**

The results of this numerical investigation allowed identifying a resonance phenomenon in the shocktrain system, which could potentially provoke a critical engine failure. The DNS also help improve the physical understanding of shock-turbulent-boundary layer interactions. Lastly, a multi-temperature model was implemented to describe species-specific vibrational energies and quantify the amount of nonequilibrium present in the flow. Reaction rates were also derived for such flow from quantum chemistry calculations. The end result is a consequent flame lift-off between the simplistic equilibrium case and the nonequilibrium case.

### **Role of High-End Computing (Why HPC Matters)**

The DNS of the shocktrain are based on a 600 million cells grid, and required several runs on intermediate grids of 50 and 200 million cells to let the structure develop. A typical simulation would take approximately 500,000 processor-hours on the Pleiades supercomputer at NASA Ames Research Center. As we investigate low resonance frequencies, the simulation run time increases as well, sometimes up to 100 flow-through! Likewise, the simulations of the scramjet combustor are based on a 125 million cells grid but require the transport of many species mass fractions and vibrational energies, thereby increasing the computational cost. The storage of TB of data was also facilitated by the dedicated system LOU.

Presenter at SC16: Romain Fiévet, PhD candidate, University of Michigan,  
rfievet@umich.edu

Technical PI: Venkat Raman, Associate Professor, University of Michigan;  
Aaron Auslender, NASA LARC

## **5) Numerical Simulation of Aircraft Noise**

### **Overview**

During takeoff and landing, a combination of jet noise and airframe noise cause large amplitude broadband sound waves. These waves propagate to the ground where communities neighboring an airport are disrupted by the annoying noise levels. Reduction of overall aircraft noise levels is a strong component of the design of next generation subsonic and supersonic aircraft. In order to reduce the noise levels, the underlying physics of the noise generation mechanism must be understood. Numerical simulation is playing a key role, both in understanding the dominant noise generation mechanisms, as well as assessing noise reduction concepts. Validation studies on jet noise and slat noise have been performed to gain confidence in the numerical simulation procedure.

## **Project Details**

The Launch Ascent and Vehicle Aerodynamics (LAVA) solver framework, developed at NASA Ames Research Center, was used to perform high-fidelity time-accurate near-field CFD analysis. Noise generated by either the jet or the slat is then propagated to the far-field using the LAVA Ffowcs Williams-Hawkings acoustic module. Comparisons to available experimental tests performed at NASA Glenn and NASA Langley are used to validate the computed results.

## **Results and Impact**

To simulate both jet noise and slat noise, a higher-order accurate finite-difference scheme for general non-orthogonal curvilinear coordinates was developed in LAVA along with advanced hybrid RANS/LES models. The increased accuracy in both the numerical discretization and the turbulence model allows fine scale structures to be modeled on relatively coarser meshes, reducing the computational resource requirements. In addition, the use of structured overlapping grids allows relatively large grid calculations to be completed in less time due to the high efficiency of the data-structures used in the solver.

Impacts of the project include:

- Validation of the LAVA solver framework for predicting noise generation and sound propagation from both jets and high-lift devices.
- High-fidelity data sets to explore the noise generating mechanisms using advanced post-processing tools also available in LAVA.
- Demonstration of the failure of several popular hybrid RANS/LES models as the mesh is refined to levels much finer than was possible when the models were originally developed.

## **Role of High-End Computing (Why HPC Matters)**

The high-fidelity time-accurate near-field CFD analysis was performed on the NASA Advanced Supercomputing (NAS) computer *Pleiades* using 500-1300 cores for approximately 2 weeks per case. For the acoustic propagation, the Fast-Fourier transform of long duration time series on both solid and permeable surfaces requires large amounts of memory. Large memory nodes on *Pleiades* system containing up to 256 GB of memory are used to efficiently perform this part of the FWH procedure.

## **What's Next (optional)**

As validation of the LAVA solver framework for a large range acoustic applications is coming to completion, the solver is now being applied to next generation aircraft concepts with the Advanced Air Vehicles Program (AAVP) and the Commercial Supersonic Transport (CST) program.

*Presenter at SC16: Cetin Kiris, NASA Ames Research Center, [cetin.c.kiris@nasa.gov](mailto:cetin.c.kiris@nasa.gov)  
Technical PI: Jeffrey Housman, NASA Ames Research Center, [jeffrey.a.housman@nasa.gov](mailto:jeffrey.a.housman@nasa.gov)*

## **6) High-Fidelity Aeroelastic Simulations of Future Airplane Concepts**

### **Overview**

The introduction of lightweight materials and emergence of energy-efficient aircraft designs with flexible wings has spurred the development of aeroelastic models for rapid design optimization and real time shape control. High-fidelity tools play an important role in evaluating these models and new aeroelastic concepts at a low cost and without

damaging existing aircraft or test facilities. In the current project, we investigate viscous effects in a distributed flap system designed to optimize the aerodynamic performance of the wings at cruise, subject to aeroelastic deformations.

### **Project Details**

The drag performance of the Generic Transport Model (GTM) with flexible wings and the Variable Camber Continuous Trailing Edge Flap (VCCTEF) system is studied at fixed lift. The system includes 19 flaps with 3 segments per flap that can be rotated relative to each other. In the baseline configuration the flaps are blended into each other with an elastomer material. An additional configuration, dubbed the “piano finger”, removes the elastomer material leaving a gap between each flap. Different flap deflection settings and fuel loads are compared. The wing deformations are calculated by coupling a 1D beam stick model to the NASA Euler solver Cart3D. The inviscid loads are compared to viscous simulations obtained with the NASA Launch Ascent and Vehicle Aerodynamics (LAVA) solver.

### **Results and Impact**

Accuracy of the inviscid design optimization approach is dependent upon a non-varying offset in neglecting viscous effects across the design space. The difference in pressure drag is the most important since this drives the aeroelastic deformations of the wing. The difference originates from a shifting of the shock location and overestimated pressure gradient through shocks and over the hinge lines of the VCCTEF, as well as flow swirling between the flaps in the case of the piano finger configuration. The drag improvement from the distributed flap system with elastomer is predicted within 25% by inviscid analysis. In contrast, a viscous analysis is recommended for the piano finger, due to the variation of the error on pressure drag with gap spacing. The study showcases how high-fidelity computations can be used for the assessment of modeling approaches. Integration of more elaborate aeroelastic models and associated grid tools is being pursued to support the analysis of increasingly complex problems.

### **Role of High-End Computing (Why HPC Matters)**

Computation of a wing deformation required an average of 5 aeroelastic iterations and 17 Cart3D calls, each taking between 30-90min (grid-dependent) on 72 cores of the NASA Advanced Supercomputing (NAS) shared memory system *Endeavour*. A LAVA viscous simulation typically took 8-12 hours using 320 Ivy-bridge cores of the *Pleiades* at the NAS center.

*Presenter at SC16: Marie Denison, NASA Ames Research Center,  
marie.f.denison@nasa.gov*

*Nhan Nguyen, NASA Ames Research Center, nhan.t.nguyen@nasa.gov*

## **7) New Insight Into Helicopter Dynamic Stall**

### **Overview**

NASA’s Revolutionary Vertical Lift Technology (RVLT) project is working towards increasing a helicopter’s forward-flight speed and performance in an environmentally responsible manner, e.g., to reduce noise levels and fuel consumption. However, these flight conditions can lead to a helicopter’s blade-tip vortices remaining in the path of other rotor blades, causing blade vortex interaction (BVI) and dynamic stall. BVI can generate unacceptable noise levels as rotor blades strike and interact with the blade-tip

vortices. Dynamic stall can also produce very loud noise levels and limit a helicopter's lifting capacity and forward-flight speed.

### **Project Details**

Dynamic stall is especially difficult to accurately predict with computational fluid dynamics (CFD). Poor grid resolution of the vortex wake and turbulent boundary layers on the rotor blades can lead to under-resolving dynamic stall flow phenomena, e.g., BVI and the formation and detachment of a rotor blade's leading edge vortex. Moreover, turbulence models can be less inaccurate when the blade stalls. An important question remains. To what degree does the CFD quantitative accuracy depend on grid-resolving the flow features, and how much does it depend on the turbulence model? Adaptive mesh refinement (AMR) is used to address part of this question by automatically and efficiently refining the vortices in the rotor wake.

### **Results and Impact**

The OVERFLOW Navier-Stokes code is used to evaluate the effects of AMR refinement of the rotor's vortex wake for an isolated Blackhawk helicopter rotor during BVI and dynamic stall events. The rest of the helicopter fuselage and other vehicle components could be included in the computational model, but they are neglected because they do not affect the fundamental nature of the flow.

- The predicted aerodynamic loads did not depend on refining the rotor wake beyond what is normally done for engineering calculations. This surprising result is good news for design engineers who must compute many CFD simulations on coarser grids to reduce the computational time and cost.
- Vortices passing over the top of rotor blades were found to trigger dynamic stall. This new insight shows there is more to dynamic stall than a good turbulence model. This important effect is neglected in two-dimensional turbulence model experiments and CFD validation. A three dimensional CFD flow simulation is required for a complete and accurate representation of dynamic stall.
- Time-dependent flow visualization was crucial in understanding the relationship between blade-tip vortices and dynamic stall.

### **Role of High-End Computing**

These flow simulations were made possible by NASA's Pleiades supercomputer using 5,628 Broadwell cores for 1-2 weeks. Grid sizes varied from 78 million to 1.8 billion grid points using MPI/open-MP with two threads. Solutions typically required 6—36 hours of wall-clock time per rotor revolution, and up to 270 terabytes of disk space to make high-resolution movies of the unsteady flow simulations.

### **What's Next**

Explore the use of AMR on the rotor blades to better resolve the rotor blade's leading edge vortex formation and detachment during dynamic stall. It is anticipated that this will improve the CFD correlation with experiment, and identify to what degree grid resolution and the turbulence model each have on solution accuracy.

Presenter at SC16 and PI: Neal M. Chaderjian, NASA Ames Research Center, [neal.chaderjian@nasa.gov](mailto:neal.chaderjian@nasa.gov)

## **8) Computational Aeroacoustics for Green Aviation Propulsion Technology**

## **Overview**

Over the last few years, there has been renewed interest in pursuing contra-rotating, open-rotor propulsion technology due to its potential to significantly reduce aircraft fuel consumption. A main concern for the design of such systems is that they must meet community and cabin noise standards. Hence, reliable noise prediction capabilities are essential to enable novel and fuel-efficient open-rotor systems. Immersed boundary and structured curvilinear grid approaches have reached a level of maturity to where they are now frequently employed for specific real-world applications within NASA, which makes these noise prediction models more reliable tools for aircraft design.

## **Project Details**

In recent years, NASA initiated several research efforts that have assessed the agency's capabilities for predicting noise source and for analyzing contra-rotating open rotors. This project demonstrates that our higher-order methods provide the ability to perform aeroacoustic analysis of wake-dominated flow fields. One of the key challenges of these models is to conduct efficient high-fidelity simulations of the relatively complex geometry. Our first-of-a-kind aeroacoustic simulation employs an immersed boundary method for an open-rotor application. In addition to applying the immersed boundary method to this moving boundary problem, we conducted a detailed aeroacoustic analysis of the noise-generation mechanisms encountered in the open-rotor flow. In previous work the simulation data—both immersed boundary and curvilinear-based—compared well with available experimental wind tunnel data. The noise-generation mechanisms were analyzed by employing spectral analysis, proper orthogonal decomposition, and the causality method. Several low- to high-fidelity simulation tools were used to provide noise predictions, and their results were compared to the experimental data. In this work, both free and pylon-installed low-speed ( $M=0.2$ ) cases were simulated, revealing pylon effects and allowing for further improvement of the CFD tools.

## **Results and Impact**

We employed higher-order accurate immersed boundary methods to simulate the highly complex flow field around an open rotor. Three key tasks were performed:

- Simulating the open rotor with our novel higher-order immersed boundary method, for low speed conditions.
- Simulating the open rotor installed on a pylon, matching wind tunnel results.
- Analyzing in detail the most dominant noise-generation mechanisms of the open-rotor propulsion system.

Several postprocessing tools (such as Fourier transforms) were used to obtain a detailed understanding of the noise-generation process. For the pylon installed case, we found 1) higher turbulence levels vs no-ptylon case was partially due to an improved wall model, 2) pylon wake chopping enhances blade wake breakup. Through these higher-order accurate simulations, NASA continues to advance the state-of-the-art in computational fluid dynamics (CFD) simulations for green aviation.

## **Role of High-End Computing (Why HPC Matters)**

Using NASA's Pleiades supercomputer, the simulations consumed approximately three weeks of wall-clock time on 500 – 2,000 Intel Haswell cores. These simulations would not be possible without the use of a large cluster like Pleiades. The visualization team at the NASA Advanced Supercomputing (NAS) facility helped generate ultra-high-quality renderings of the open-rotor flow physics.

## **What's Next**

NASA's LAVA team is looking to significantly improve the performance of the solver by investigating a Lattice Boltzmann method (LBM) in the same Cartesian framework, and by optimizing the remainder of the code. Initial tests show great potential due to a demonstrated 50x speedup for flow in an idealized (no complex geometry) setup.

*Presenter at SC16: Michael Barad, NASA Ames Research Center, michael.f.barad@nasa.gov  
Cetin Kiris, NASA Ames Research Center, cetin.c.kiris@nasa.gov*

## **9) CFD Support for Quiet Supersonic Aircraft Design**

### **Overview**

NASA is investigating innovative ways of improving aircraft design to allow for supersonic flight over land. Currently the Federal Aviation Administration (FAA) has a ban on civilian supersonic flights over land due to the sonic boom loudness that such aircraft would generate. To address this issue NASA has been (since 2008) working on developing technologies to enable quiet supersonic aircraft designs. A major feature of these new aircraft designs is the shaping of the aircraft to reduce shock strengths and avoid shocks merging together as they propagate through the atmosphere to the ground. This allows for a much lower noise level when the shocks reach the ground.

### **Project Details**

NASA (Ames, Langley, Glen and Armstrong) have partnered with industry to develop a Low Boom Flight Demonstrator (LBFD) to study a representative low boom aircraft design. The aircraft is designed to be the size of a small business jet but to mimic the shock signature of a larger aircraft to test the feasibility of NASA's current noise emission goals. Part of the design process involves using Computational Fluid Dynamics (CFD) simulations to predict the pressure signature of various design iterations. The CFD simulations help to make informative design choices. This is accomplished by extracting near field pressure signature data from the CFD results, propagating the signature through the atmosphere using propagation codes, and examining the resulting noise signature that the proposed design would generate.

### **Results and Impact**

The computational results enable the designers of the LBFD to assess the sonic boom and vehicle performance prior to a flight test and address any region in the designs where changes are required to the aircraft's design that affect its sonic boom pressure signature (loudness) on the ground. Ensuring that the design has the desired pressure signature is critical because the LBFD X-plane is going to be used to perform various tests, including community fly-overs, to test the validity of a quiet supersonic design. The results of these tests will help inform the designs of future supersonic aircraft designs as well as provide data for the FAA for review of the current ban on supersonic flight over land.

### **Role of High-End Computing (Why HPC Matters)**

CFD simulations of the aircraft in flight require accurate modeling of both the aircraft and the flow field in its vicinity. This is accomplished by generating a computational model of the full aircraft and its nearby airspace; typically this model requires 50 to 100 million grid-cells. A single computational run of a model of this size requires approximately 750 processor-hours on the Pleiades supercomputer at NASA Ames Research Center. By assuming symmetry in our simulations we are able to reduce the computing requirements for each simulation by half by only computing the flow over half of the airplane. The pressure signatures of aircraft designs are extracted from the

simulations obtained at Ames Research Center and provided to Langley Research Center. They are then able to use this data to determine the sonic boom loudness level on the ground of a particular design. The speed of the NASA networks allow the team members from Ames, Langley, and industry to transmit data efficiently as well as to help coordinate efforts.

James Jensen, NASA Ames Research Center  
james.c.jensen@nasa.gov

Susan Cliff, NASA Ames Research Center  
susan.e.cliff@nasa.gov

## Our Planet

### 10) Supporting Petabyte-scale Science on NASA Earth Exchange (NEX)

#### Overview

The NASA Earth Exchange (NEX) is a supercomputing collaboration platform for the Earth science community that provides a mechanism for scientific collaboration and knowledge sharing. NEX combines state-of-the-art supercomputing, Earth system modeling, workflow and knowledge management, and NASA remote sensing data feeds to deliver a complete work environment in which users can explore and analyze large datasets, run modeling codes, collaborate on new or existing projects, and share results among the Earth science communities. NEX supports over 20 large-scale science projects ranging from satellite data processing and analysis to machine learning with volumes of data up to 25PB for the largest projects and combined computing consumption over 9 million CPU-hours in the past year.

#### Project Details

The main goal of NEX is to enable enhanced and more efficient use of Earth observations for NASA Earth science technology, research and applications programs. NEX utilizes the Pleiades supercomputing platform together with over 2PB of satellite, climate and model datasets in order to accelerate research in the Earth sciences. The results and project details can then be shared with the community through the NEX portal – a web-based knowledge network.

NEX supports diverse Earth science communities in the areas of technology development, large-scale data processing systems, climate downscaling, data mining and data analytics, remote sensing, and ecosystem modeling as well as applications in agriculture management. In order to continuously engage the community, we have developed OpenNEX, an outreach platform that provides access to video lectures, tutorials, hands-on labs, and sample applications that anyone can use to access and explore NASA datasets. OpenNEX is structured like a series of workshops, where each focuses on a particular theme.

#### Results and Impact

NEX strives to lower the barrier of entry to data- and compute- intensive science and enable research teams to execute projects at ever-larger scales that would otherwise not be feasible. Example of these projects include:

- Creating a tree-cover map across the United States using 1m resolution data and advanced machine-learning algorithms.
- Development of a set of high resolution, bias-corrected climate change projections that can be used to evaluate climate change impacts on ecosystems, hydrology and human health.
- Global satellite data processing and visualization pipeline for processing of almost 25PB of data.
- Developing a new high-resolution (1ha) global vegetation biomass map by blending a variety of satellite data sets useful for global carbon monitoring.

#### Role of High-End Computing (Why HPC Matters)

The Pleiades supercomputing architecture combined with the massive data store and high-speed network enables NEX to engage scientific communities and provide them with capabilities to execute modeling and data analysis on a grand scale, which was not

previously achievable by most scientists. NEX projects processed and analyzed over 3 petabytes of satellite data over the last year and in doing so utilized over 9 million CPU hours.

*Presenter at SC16: Sangram Ganguly, NASA Ames Research Center, sangram.ganguly@nasa.gov*

*Technical PI's: Petr Votava, NASA Ames Research Center, petr.votava@nasa.gov*

*Ramakrishna Nemani, NASA Ames Research Center, ramakrishna.nemani@nasa.gov*

## **11) Global, High-Resolution Chemistry in GEOS-5**

A start-of-the-art chemistry transport model (CTM), GEOS-Chem, has been coupled to the Goddard Earth Observing System (GEOS-5) general circulation model (GCM). A CTM provides a global 3D picture of atmospheric composition and models the physical and chemical interactions of chemical compounds, while the GCM provides the meteorological inputs needed by the CTM. One of the most important of these compounds is tropospheric ozone, produced when nitric oxides (NO<sub>x</sub>) and hydrocarbon pollutants react in sunlight. Atmospheric chemistry and related processes are highly non-linear, making it crucial for atmospheric chemistry models to adequately capture the spatial and temporal heterogeneity of all involved species. Models running at coarse resolution inherently suffer from a misrepresentation of sharp concentration gradients. To realistically simulate these small-scale concentration patterns, researchers are making a first attempt at a very high-resolution run of GEOS-Chem in GEOS-5.

- GEOS-5 GCM with GEOS-Chem is simulating Jul 2013 to June 2014 at a 12-kilometer (km) resolution.
- To provide a realistic meteorology for GEOS-Chem, the 50-km MERRA-2 re-analysis was downscaled to 12 km to constrain the meteorology.
- Researchers are currently validating the results against the SEAC4RS field campaign and will use the simulation data for observation system observing experiments for the CEO-CAPE mission.

The high-resolution GEOS-Chem run has produced a realistic picture of tropospheric chemistry at this previous untested scale after validating against SEAC4RS data. Some of the parameterizations used by GEOS-Chem will need to be adjusted for the finer resolution, and this work provides a dataset to start that retuning. This work has also highlighted the need for development to reduce the input/output bottleneck imposed by the large amount of emissions data that needs to be ingested into the model.

Presenter at SC16 and Technical PI: Benjamin Auer and Christoph Keller  
Global Modeling and Assimilation Office  
NASA Goddard Space Flight Center

## **12) DeMISTifying Climate Informatics: Visualizing Tomorrow's Forecasts Today with Web World Wind**

Developing a more intuitive means of displaying and observing surface and atmospheric conditions facilitates an improved understanding of our planet's ever-changing climate. To this end, the NCCS has developed an instance of Web World Wind (WWW) that

works in conjunction with Web Map Service (WMS) and the Thematic Real-time Environmental Distributed Data Services (THREDDS) Data Server to provide a more logical display of climate informatics.

- WWW displays an interactive 3D virtual globe or 2D map with high-resolution overlays of climate simulations and previously recorded weather data. It can be easily accessed through any Web Browser and using any Operating System.
- WWW offers high-resolution forecasts through the Forward Processing Model for various metrics on a complete global scale via the querying of the WMS imaging server.
- The THREDDS Data Server (TDS) currently stores a wide array of climate data in various file types and applications, creating a centralized data library that can be easily accessed and interpreted by WMS protocols.

The current state of WWW makes use of the Forward Processing Model that provides a 25 KM resolution global forecast over a 7-day period. On top of this, the NCCS is pleased to be one of the first to make use of 6 KM imagery, permitting a more detailed and accurate 10-day global forecast. Future development will incorporate a vast collection of MERRA/MERRA2 analytics and will allow us to display imagery from the past few decades. Together, these tools enable a more readily accessible and intuitive vehicle for understanding Earth's climate - past, present, and future.

Presenting at SC16 and Technical PI: Ben Bledsoe and Daniel Duffy  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

### **13) Exploring Our Warming World in 5K Ultra HD**

NASA is exploring our warming world in 5K Ultra HD resolution with more than 16 million pixels (horizontal grid cells) in our Goddard Earth Observing System (GEOS) atmospheric model. Using the world class MERRA-2 atmospheric reanalysis from 1980-present as a baseline for our exploration, the GEOS model is used to refine the fidelity of this analysis down to 6km grid cells covering our entire planet from the surface to the top of the mesosphere more than 50 miles above the Earth's surface. The resulting analysis becomes the finest resolution look at the details of our changing planet.

- February of 2016 saw the largest temperature anomaly ever for the globe at 0.84 C warmer than the 1980-2010 climatology.
- In May of 2016 carbon dioxide levels reached a global record high of 404 PPM as concentrations peaked across the Northern Hemisphere in excess of 410 PPM.
- Despite a developing La Nina cooling the Pacific Ocean, July 2016 was not only the warmest July ever at nearly 0.5 C warmer than the 1980-2010 climatology, but the hottest month on record for the globe.

The GEOS model reveals the complexity of carbon dioxide as it is emitted into our atmosphere and transported around the globe, the dispersion of tiny particles from fossil fuel emissions and biomass burning, plumes of ash and sulfate ejected from volcanoes around the world, and a detailed distribution of record warmth across our planet.

Technical PI and Presenter at SC16: William Putman  
Global Modeling and Assimilation Office  
NASA Goddard Space Flight Center

## **14) Improving the Scalability of the GEOS-5 Earth System Model**

The Goddard Earth Observing System (GEOS-5) is a global climate model and assimilation system maintained by the Global Modeling and Assimilation Office (GMAO). GEOS-5 is used as both a near real-time forecast and a comprehensive reanalysis platform. A current GMAO effort is aimed at examining the performance of GEOS-5 at finer resolutions – sub-10-kilometer (km) cell width – that will become the future operational standard. To better utilize high-performance computing (HPC) resources, we are exploring ways to improve the scaling of the model as HPC evolves toward using many-core architectures and MPI+OpenMP hybrid programming.

- In general, for the cubed-sphere dynamics, better scaling is achieved using MPI+OpenMP at high core count, but MPI alone is better for low core count.
- For the dynamics, the core count where MPI+OpenMP becomes beneficial, and the best-performing number of threads per MPI process, is dependent on the model resolution.
- The solar radiation routine appears to scale well using both pure-MPI and MPI+OpenMP. This gives hope that other column-physics routines will be amenable for conversion.

GMAO is determined to continue to move GEOS-5 to higher resolutions to better model and forecast the Earth's atmosphere and climate. Studying GEOS-5's scaling characteristics will allow us to re-architect the model in order to fully utilize all the HPC resources available to us now and in the future.

HPC resources were provided by NASA Center for Climate Simulation (NCCS) and NASA Advanced Supercomputing (NAS).

Technical PI and Presenters at SC16: Matthew Thompson and Hamid Oloso  
Global Modeling and Assimilation Office  
NASA Goddard Space Flight Center

## **15) Immersive Earth Science: Exploring Climate Data with Virtual Reality**

NASA missions and programs create a profusion of science data and information about our planet. Visualization is a key component to effectively analyze and apply Earth science data and derived information to create knowledge. NASA's Science Mission Directorate (SMD) Earth Science Division (ESD) Earth Science Technology Office (ESTO) and Navteca are exploring virtual reality (VR) technology for the next generation of Earth science technology information systems.

- Using a VR headset, the user has an inherently spherical visualization of the Earth instead of any of the standard two-dimensional (2D) projections.
- The three-dimensional (3D) structure of cloud and storm systems from the Global

Precipitation Measurement (GPM) mission, data processed by the Discover supercomputer at the NASA Center for Climate Simulation (NCCS), are visualized using VR.

- The observer interacts with the data, makes selections, and views volumetric data such as the NASA Discover-AQ mission flight paths, in an innovative and intuitive way.

Our results from displaying GPM and Discover-AQ data show that there is great potential for scientific visualization and analysis using VR. Current work includes additional manipulation of 3D and volumetric datasets and an ESTO-sponsored study of scientific applications for VR that will result in the identification of specific use cases. VR increases the potential to gain new insight into the roles that clouds, storm systems, and atmospheric aerosols play in regulating Earth's weather, climate, and air quality.

Presenter at SC16 and Technical PI: Shayna Skolnik and Ramon Ramirez-Linan  
Earth Science Technology Office  
NASA Goddard Space Flight Center

## **16) A Virtual Telescope on the World's Oceans and Sea Ice**

### **Overview**

Researchers from the Estimating the Circulation and Climate of the Ocean (ECCO) consortium are running simulations on Pleiades to produce global "maps" of Earth's ocean and sea-ice system at an unprecedented resolution (~1 km horizontal grid spacing). The simulations are produced with the Massachusetts Institute of Technology general circulation model (MITgcm) across up to 70,000 cores on Pleiades. The global ocean model includes tidal forcing, which results in well-resolved internal tides, and it admits submesoscale ocean dynamics and long internal waves. The resulting model output provides the oceanic equivalent of a telescope that can reveal the state of the entire ocean/sea-ice system in exceptional detail. The model is already providing design guidance in planning NASA satellite missions, for example, Surface Water and Ocean Topography (SWOT) and Winds and Currents Mission (WaCM), and has uses with other agencies, including strategic innovations in next-generation ocean acoustics systems and in academic research.

### **Project Details**

The simulations provide a virtual ocean that admits tidal energetics and submesoscale dynamics not normally represented in global calculations. This extends simulated ocean behavior beyond broadly quasi-geostrophic flows and provides a preliminary example of a next-generation computational approach that will allow us to explicitly probe the interactions between the dominant energetic scales in the ocean and the instabilities that are usually parameterized.

The time-varying fields from these simulations have been recorded at hourly frequency. Comparison with spectra from current meters indicates that the model skillfully represents ocean dynamics down to scales of a few kilometers. These ocean simulations have already proved to be remarkably useful in helping to increase understanding of how large and small scale processes coexist and interact.

### **Results and Impact**

Researchers are using these simulations to:

- Investigate fundamental questions such as how the circulation, chemistry, and biology of the ocean collectively interact with atmospheric carbon.
- Determine how a pollutant plume or debris field might spread from a particular ocean location, or where and when heat is absorbed by or released from the ocean.
- Design satellite-born and in-situ sampling missions that will allow us to better monitor and understand ocean turbulence and its impacts on critical, life-support climate systems of our planet.

### **Role of High-End Computing (Why HPC Matters)**

These high-resolution computational fluid dynamics simulations were performed in support of the NASA Science Mission Directorate. The project used the Pleiades supercomputer to integrate the MITgcm model with 20 billion grid cells for a period of 13 months, resulting in 5,000 terabytes of model output. The I/O layers sustained a throughput of 20 gigabits per second when redistributing and streaming to disk data from 70,000-core experiments. The computations required about 24 million CPU hours on Pleiades.

### **What's Next (optional)**

As computational and storage capabilities continue to increase, this work has the potential to become a transformative strategy for understanding and predicting impact of global ocean circulation on the Earth's climate. A key next step is the coupling of this ocean model with a fully interactive atmosphere of similar resolution.

*Presenter at SC16: Christopher Hill, Massachusetts Institute of Technology  
cnh@mit.edu*

*Technical PI: Dimitris Menemenlis, NASA Jet Propulsion Laboratory  
menemenlis@jpl.nasa.gov*

## Space Exploration

### 17) Improving Fidelity of Launch Vehicle Liftoff Acoustic Simulations

#### Overview

Launch vehicles experience high acoustic loads during ignition and liftoff affected by the interaction of rocket plume generated acoustic waves with launch pad structures. Application of highly parallelized Computational Fluid Dynamics (CFD) analysis tools optimized for application on the NAS computer systems such as the Loci/CHEM program now enable simulation of time-accurate, turbulent, multi-species plume formation and interaction with launch pad geometry and capture the generation of acoustic noise at the source regions in the plume shear layers and impingement regions. These CFD solvers are robust in capturing the acoustic fluctuations, but they are too dissipative to accurately resolve the propagation of the acoustic waves throughout the launch environment domain along the vehicle.

A hybrid Computational Fluid Dynamics and Computational Aero-Acoustics (CFD/CAA) modeling framework has been developed to improve such liftoff acoustic environment predictions. The framework combines the existing highly-scalable NASA production CFD code, Loci/CHEM, with a high-order accurate discontinuous Galerkin (DG) solver, Loci/THRUST, developed in the same computational framework. Loci/THRUST employs a low dissipation, high-order, unstructured DG method to accurately propagate acoustic waves away from the source regions across large distances. The DG solver is currently capable of solving up to 4th order solutions for non-linear, conservative acoustic field propagation. Higher order boundary conditions are implemented to accurately model the reflection and refraction of acoustic waves on launch pad components. The DG solver accepts generalized unstructured meshes, enabling efficient application of common mesh generation tools for CHEM and THRUST simulations. The DG solution is coupled with the CFD solution at interface boundaries placed near the CFD acoustic source regions. Both simulations are executed simultaneously with coordinated boundary condition data exchange.

#### Project Details

Initial application testing and validation of the higher order acoustic propagation technique was carried out against acoustic data from the Ares Scale Model Acoustic Test (ASMAT) experiments. This simulation served to evaluate the capabilities and production readiness of the CFD/CAA framework towards resolving the experimentally observed spectrum of acoustic frequency content. Numerous improvements to the numerical algorithm implementation, boundary conditions and communication process between the simulations were identified and implemented during this project. The initial applications presented here were performed with the 3<sup>rd</sup> order DG solver. Testing with the 4<sup>th</sup> order solver will commence as soon as higher order boundary conditions have been fully implemented.

#### Results and Impact

The new CFD/CAA approach proved highly capable of accurately propagating and conserving the acoustic wave field over the complex launch vehicle and launch pad geometry. Acoustic wave content was preserved to a significantly higher frequency range in the DG solver predictions, especially at sensor locations farther away from the plume source regions in the upper vehicle portions. This improved capability to perform high fidelity computational acoustic field simulations increases the confidence in the specification and understanding of launch acoustic loads design environments through

computational modeling. Understanding of the acoustic environments can be expanded beyond the data available from the limited number of sub-scale tests that could be performed for ASMAT and the current SMAT (SLS Model Acoustic Test). Numerical simulations can be applied in evaluating various sound suppression measures reducing the need for expensive testing.

### **Role of High-End Computing (Why HPC Matters)**

NASA supercomputing resources are instrumental in completing this type of analysis. The computational model requires nearly 300 million mesh cells to simultaneously resolve the launch vehicle and launch pad details and adequately capture the acoustic sources at the rocket plumes. Both simulations are executed simultaneously exchanging transient boundary condition input to the DG solution, utilizing the sophisticated message passing enabled by the high performance NAS architectures. Each of the ASMAT simulations was performed using 2,000 parallel computing processors on NASA's Pleiades supercomputer. The dynamic simulation of the ignition dynamics and acoustic field propagation used nearly 2 million processor-hours over the time frame of two to three weeks, which would have been an intractable problem in the past.

*Presenter at SC16: Peter Liever, NASA Marshall Space Flight Center, peter.a.liever@nasa.gov*

*Technical PI: Jeff West, NASA Marshall Space Flight Center, jeffrey.s.west@nasa.gov*

## **18) CFD-Based Aerodynamic Databases for the Space Launch System**

### **Overview**

NASA's new Space Launch System (SLS) will be the first rocket since the Saturn V to carry astronauts beyond low-earth orbit, carrying 10% more payload than Saturn V and three times the payload of the Space Shuttle. The SLS configuration consists of a center core and two solid rocket boosters, which separate from the core as their fuel is exhausted two minutes from lift-off. During these first two minutes of flight, the vehicle powers its way through strong normal shocks in the transonic regime, through its maximum aerodynamic loads at MaxQ, and is enveloped by gaseous plumes from booster-separation motors. The SLS program is relying on computational fluid dynamic (CFD) simulations to provide much of the data needed to build ascent aerodynamic databases describing the structural load distribution, the surface pressures, and aerodynamic forces on the vehicle.

### **Project Details**

The goal of this work is to run computational simulations to quantify the aerodynamics during the ascent of the SLS vehicle, and provide data to build a number of aerodynamic databases. Previous launch systems relied almost entirely on data from wind tunnel testing for these databases. But with the continued growth of computational resources and capabilities, it has become feasible to use more CFD and rely less on expensive wind-tunnel tests.

A number of different SLS vehicle designs have been and are being analyzed, including: multiple design iterations (SLS-10003, SLS-10005, and SLS-10008) of the first flight vehicle which will launch in the EM-1 mission; multiple designs (28000, 28003, 28004) of the second SLS vehicle, which will be the first to carry astronauts during the EM-2 mission; and the SLS 27000, the cargo version of the rocket.

The project is utilizing multiple CFD tools to perform the analysis, including the NASA-developed flow solvers Cart3D, OVERFLOW, and FUN3D. Recent highlights included a moving-body dynamic simulation of a nominal booster-separation event using the OVERFLOW solver. The CFD run included the simulation of the 22 different plumes that are firing during the separation. These include the 16 BSM plumes, the plumes from the four core-stage main engines, and two plumes from the boosters.

### **Results and Impact**

The project has created a number of different aerodynamic databases for the SLS designs, including distributed line-loads used in structural analysis; surface pressures used in venting and other dynamic analysis; protuberance aerodynamic loads used in the design of the vehicle; and forces and moments on the vehicle during booster separation used to verify the booster-separation system.

### **Role of High-End Computing (Why HPC Matters)**

All of the CFD simulations of the SLS vehicles were run on the Pleiades supercomputer. In each of the past two years, the project has used over 20 million core hours on Pleiades. Depending upon the configuration and flow conditions, each of the Cart3D simulations require approximately 100 to 200 core hours, and the OVERFLOW runs require 10,000 to 30,000 core hours. A total of 12.7 million core hours were used to build the booster separation database, and over 2 million core hours were used to compute the single moving-body dynamic simulation of the booster-separation event. The project typically uses 50 to 100 terabytes of short-term disk storage at any given time.

*Presenter at SC16 and PI: Stuart Rogers, NASA Ames Research Center, stuart.e.rogers@nasa.gov*

*Technical co-PI: Derek Dalle, STC, NASA Ames Research Center, derek.j.dalle@nasa.gov*

## **19) Predictive Modeling for NASA Entry Descent and Landing Missions**

### **Overview**

The Entry Systems Modeling (ESM) technology development project is engaged in maturation of fundamental research developing revolutionary computational models in aerosciences and materials for entry, descent, and landing (EDL) technologies. Examples include:

- Technology first predictive models for shock layer radiation, anchored by detailed quantum mechanics simulations. ESM models have already proved critical for Osiris-Rex, InSight, and Orion.
- Truly predictive models for thermal protection system response during entry. Such a capability is enabling to the end goal of a computational design engine for engineered materials (such as woven).
- First ever demonstration that the dynamic behavior of an entry capsule can be predicted via high-fidelity unsteady analysis.

- High-fidelity simulation of parachute inflation and descent via coupled fluid-structure interaction. The end goal is a predictive capability for parachute inflation loads and dynamics.

EDL Modeling and Simulation (M&S) is critical path for NASA. M&S is used in every mission phase to define mission concepts, select appropriate architectures, design EDL systems and quantify margin and risk, ensure correct system operation, and analyze data returned from the entry. In an environment where it is impossible to fully test EDL concepts on the ground prior to use, accurate M&S capability required to extrapolate ground test results to expected flight performance. The complexity of these high fidelity tools requires the use of the largest supercomputers to provide timely data for mission planners.

*Presenter at SC16: Michael Wright, NASA Ames Research Center,  
michael.j.wright@nasa.gov*

Technical PI: Joseph Brock, NASA Ames Research Center, *joseph.m.brock@nasa.gov*

## **20) How NASA's Astronauts will Survive Earth Re-entry from Deep Space**

### **Overview**

The Orion Multi-Purpose Crew Vehicle (MPCV) will extend NASA's human presence in space beyond low Earth orbit for the first time since the Apollo era. Currently in design, the spacecraft will be launched atop the Space Launch System (SLS), which will be the most powerful rocket ever built, capable of 9.2 million pounds of thrust at lift off. The Orion MPCV will carry the crew for deep space missions, provide emergency abort capability, and safely return the crew during Earth re-entry.

### **Project Details**

Orion's thermal protection system (TPS) protects the vehicle from the extreme heating environment experienced during re-entry. The TPS consists of an Apollo-like heat shield on the forebody and Shuttle-like tiles on the afterbody. The TPS is being designed largely through computational analysis. An aerothermal database was computed to define the convective and radiative heating environments. Together with ground test data and flight data from previous missions, the database is used to select the TPS materials and thicknesses over the vehicle.

### **Results and Impact**

The aerothermal database was used to design the TPS for the Exploration Flight Test-1 (EFT-1) mission which flew December 5, 2014. During the successful flight, the TPS performed perfectly. Post test comparisons of predicted and measured temperatures showed good agreement. The results have been used to update the database which is currently being used to design the TPS for Exploration Mission-1 (EM-1) which will return to Earth after orbiting the Moon.

### **Role of High-End Computing (Why HPC Matters)**

The aerothermal database is composed of 1200 computational fluid dynamics (CFD) solutions. The freestream conditions (velocity, altitude, and angle of attack) spanned the possible flight space of reentry trajectories. The solutions were computed on the Pleiades cluster using grids with 16.4 millions points. Each solution required 48+ hours of wall time using 300 processors. Large data sets (4 GB) were transferred among

various NASA centers throughout the country using high-speed networks, and the solutions are archived using NAS's mass storage system

*Presenter at SC16 and PI: Ryan McDaniel, NASA Ames Research Center,  
ryan.d.mcdaniel@nasa.gov*

## **21) Removing Modeling Assumptions in the Prediction of Buffet Forces**

### **Overview**

NASA is developing the Space Launch System (SLS) as a replacement for the retired Space Transportation System (Space Shuttle). As the system design matures, the current Shuttle-based solid-rocket boosters (SRBs) may be replaced with advanced boosters, either solid or liquid fueled. These advanced boosters may have different geometrical shapes than the SRBs incorporating larger diameters and/or blunter nose caps. As such, Computational Fluid Dynamics (CFD) simulations are used to assess the aerodynamic impacts of alternate booster designs, with a focus on the prediction of the unsteady flowfield that causes buffet. Buffet is a relatively low-frequency (<100 Hz) unsteady pressure loading that impacts the vehicle structural-response as a whole and is known to be a driver in the structural design of launch vehicles. Buffet forces are typically determined using sub-scale wind tunnel tests. The ability to predict these flowfields and resultant loads promise a powerful capability to the SLS and the design of future flight vehicles.

### **Project Details**

The transonic flowfield about SLS cargo configurations was simulated with the Fully Unstructured Three-Dimensional (FUN3D) CFD solver at wind tunnel conditions. Unsteady, time-accurate computations were performed using second-order Delayed-Detached Eddy Simulation (DDES).

Standard practice for predicting structural dynamic response due to aeroacoustic forcing has involved assumptions based on legacy experimental measurement under flat-plate turbulent boundary layers and very approximate model forms of derived quantities representing the behavior of the surface fluctuating pressure cross spectrum. Here too these parameters were initially estimated from the computed cross spectrum. However, it was discovered that the full computed cross spectrum could be used directly, eliminating the model entirely.

### **Results and Impact**

CFD based aeroacoustics and buffet data were found to be in excellent quantitative agreement with ground-based test data. The detailed spatial resolution led to greatly enhanced understanding of loads generated by surface pressure fluctuations. Long duration times were required for statistical convergence due to the considerable variation in pressure signal spectra and levels. Standard Fourier analysis was not adequate, so specialized wavelet methods were developed to reduce the data in order to mitigate effects of the long durations. Use of the full computed surface pressure cross spectrum eliminated a significant amount of the uncertainty in predictions associated with model form assumptions. This translated into being able to reduce unnecessary margins.

### **Role of High-End Computing (Why HPC Matters)**

The pressure time histories for these complex flowfields are often highly non-stationary. This is due to the fact that the pressure signals contain both loud and quiet periods where the spectra and levels change considerably. As a result, the simulations need to be very long in duration in order to approach statistical convergence. Each simulation typically used 12,000 cores of NASA supercomputing resources and required approximately 36 million core-hours.

*Presenter at SC16: William T. Jones, NASA Langley Research Center,  
w.t.jones@nasa.gov*

*Technical PI: Craig L. Street, NASA Langley Research Center, craig.l.streett@nasa.gov*

# Supercomputing

## 22) HECC – Moving to the Future Today

### Overview

NASA's High-End Computing Capability (HECC) Project provides an agency-wide compute environment focused on enabling science and engineering discovery across all technical mission directorates. Meeting that challenge requires an environment that is reliable, stable and expandable. This year we were able to increase the power of NASA's largest High-End Computer and begin the needed process to significantly increase the power and cooling essential for future expansion. With the fielding of the 1.2 PF Electra system in a module, we began the process of evaluating energy efficient expansion options.

### Project Details

HECC meets NASA's competing HPC demands—capacity, capability, and time-critical computing—through a customer-focused team of professionals providing world-class services and hardware to a broad array of technical users. The primary goal of the HECC team is to enable our users to accomplish their mission goals.

This year we migrated to the next-generation LTO-7 tape technology bringing the library capacity to almost half an Exabyte, and increased Pleiades' peak performance to 7.5 petaflops.

We began the critical tasks of increasing the number of systems HECC can support by deploying a Modular HPC system housing the 1.2 PF Electra system. This modular technology promises to have a dramatic impact on our water and electrical demands needed to cool our systems. With a PUE close to 1.05, this promises to significantly reduce our impact on the environment.

The team remains driven to provide the best high-end computing environment to our user community through advances in hardware and improved service delivery.

### Results and Impact

Pleiades moved to its 7<sup>th</sup> generation of Intel Xeon architecture with the installation of the Broadwell racks. The system has been able to advance through seven different processor technologies and three network interfaces while simultaneously delivering the equivalent of over 700,000,000 Westmere node hours. Overall, the HECC project has had a very successful year.

- Pleiades was ranked 15th worldwide in the June Top500 list and 9th on the High Performance Conjugate Gradient (HPCG) list, a benchmark that more closely models the types of work done by our user community.
- LTO-7 tape drives were allowing up to ½ an Exabyte of storage in our library.
- The site was developed and the first module installed with the 1.2 PF Electra system.
- Worked with facilities and headquarters to plot a growth path through 2022.

### What's Next (optional)

HECC will continue to strive to meet the demands of the NASA user community. The facility expansion options under consideration have the potential to have a huge upside in the project's ability to meet these demands.

## **23) Improving Productivity at a Supercomputer Center**

### **Overview**

With key modeling and simulation problems continually growing in size and complexity, optimizing the performance of high-end computing applications is essential. Faster time to solution increases user productivity; it also means the compute resources are more productive, as the time saved can be used for other work. The goal of the Application Performance and Productivity (APP) team at NASA Ames is to maximize both user and system productivity in NASA's High End Computing Capability Project.

### **Project Details**

The APP group has a variety of activities aimed at improving productivity. They assist users in solving compiler/runtime issues with their codes and provide tools to simplify the programming and use of large parallel computing systems. They also develop and maintain benchmark suites to ensure that current systems are working efficiently and that future systems will maximize the science and engineering value to the agency. In the last year, the group expanded its efforts in the area of application performance optimization. Several group members work directly with users to understand and improve codes. In addition to improved time-to-solution, the users are taught how to use performance analysis tools so that they can make further improvements on their own.

### **Results and Impact**

Recently, the APP team has been working in a variety of areas to improve both current and future productivity:

- Combining the strengths of two different tools, the team conducted a study of the I/O characteristics of 12 application codes, including 8 of the 10 with highest usage on the Pleiades system. The I/O in one of the top-10 applications was improved by 36-fold.
- The team has been conducting a study of Knights Landing (KNL), Intel's new Xeon Phi processor. The general impression is that system usability and the application porting process have been improved significantly over KNL's predecessor. For several key applications, KNL performance was better than nodes based on Xeon (Haswell and Broadwell) processors.
- The team has been redefining the benchmark suite used to evaluate system performance. The suite is representative of a typical NASA workload and measures relative performance of different resource types as part of a protocol for establishing charging rates. It is also used in system procurements to make sure that new acquisitions perform well on a typical workload.
- The team has had a number of optimization successes, including a 1.8x speedup of a computational fluid dynamics code used to study a flow separation problem and a 93-fold speedup of an image comparison code used to analyze data from the Solar Dynamics Observatory.

### **What's Next**

In the coming year, the group will characterize performance of a supercomputer deployed in a modular data center environment, including measurements of the electric power requirements of applications. They also plan to explore techniques for using

lightweight instrumentation on a large number of batch jobs to help identify poor resource utilization and potential system issues. They will also investigate how to transform codes to maximize performance on a KNL-based cluster.

*Presenter at SC16: Henry Jin, NASA Ames Research Center, haoqiang.jin@nasa.gov  
Technical PI: Robert Hood, CSRA LLC—NASA Ames Research Center, robert.hood@nasa.gov*

## **24) Testing Code at Exascale with Only a Workstation**

Unfortunately, some critical software defects are only detected/detectable when an application is run at very high resolutions on large numbers of cores. In such situations, traditional methodologies for isolating and debugging the application become extremely expensive:

- Dedicated resources sit idle while developers examine output, probe a debugger, and/or recompile the application to try a fix.
- Sites must license debuggers for possibly millions of cores.
- Developers lose productivity while waiting for dedicated resources.

Therefore, developing new approaches is of vital importance for maximizing the benefits of exascale platforms as they become available.

We are pursuing two technologies that we believe will directly address the costs of testing and debugging exascale applications:

- **MockMPI:** Software mocks are configurable software components that are intended to replace an external dependency for the purposes of testing. MockMPI mocks the standard Message Passing Interface (MPI) and thereby eliminates the dependency on the communication layer. Instead, users specify a sequence of expected MPI calls on a given process, and MockMPI creates the illusion of interactions with an entire cluster. Unit testing and debugging can then proceed on a single node or even a workstation. MockMPI is in an early prototype phase.
- **pFlogger:** The parallel Fortran logging utility allows runtime configuration of application level log messages. Low-level diagnostics can be suppressed in production, but activated for specific components on selected processes for debugging purposes. This software is in the process of being released under an open-source agreement.

Presenter at SC16 and Technical PI: Thomas Clune  
Global Modeling and Assimilation Office  
NASA Goddard Space Flight Center

## **25) Great Scott! 1.21 Jigowatts and 88 Gigabytes per Second: Time Travel with a Time Series Database**

Time Series Databases (TSDBs) are databases optimized for storing colossal volumes

of data indexed by time. Modern TSDBs can ingest hundreds of thousands of data points per second using commodity hardware. This is an ideal fit for monitoring a supercomputer, which may have thousands – if not tens of thousands – of components, each with many distinct metrics.

The NASA Center for Climate Simulation (NCCS) Operations Team is using a TSDB to help monitor the Discover supercomputer with its 3,600+ nodes and additional infrastructure components. Every 10 seconds, the TSDB stores:

- 700,000+ measurements relating to its General Parallel File System (GPFS)
- 11,000+ power-related measurements from rack power distribution units (PDUs)

High-resolution historical data has proven invaluable for both real-time monitoring and post-mortem capabilities, effectively allowing a systems administrator to "time travel" to the moment of a specific event. From a tripped PDU breaker to a user-reported filesystem slowdown, an administrator can examine data from the event in high resolution as if it were occurring in the present moment. Having access to historical data at this resolution has enabled problem diagnosis and root cause analysis that was not possible before.

The collected power usage data has enabled the NCCS to create an aggregated view of power consumption across compute racks, individual Scalable Compute Units (SCUs), room PDUs, and facility uninterruptible power supplies (UPSs). This information has yielded a better understanding of overall power load trends across the facility. Administrators are exploring using TSDBs to store additional information, such as node-level performance data and detailed job resource usage summaries.

Presenter at SC16 and Technical PI: Aaron Knister and Bruce Pfaff  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

## **26) Finding Hurricanes using Machine Learning and High Performance Computing**

Extreme weather is having enormous effects across the globe with major impacts on human life and property. Within a changing climate, there is a vital need to increase the accuracy and timeliness for predicting the onset of extreme weather events. These extreme events happen within a global atmosphere with connections over both space and time that may have significant effects on the development of a potential major storm.

With the recent creation of high-resolution reanalysis data sets, scientists are in a unique position of studying these extreme events in a global climate context. However, the sheer sizes of the data sets pose major problems to exploring known and unknown connections in the atmosphere. The NASA Center for Climate Simulation (NCCS) has been working on the development of a machine learning algorithm to help better predict the onset of tropical cyclones. The NCCS is in a better position than ever given both the availability of the reanalysis data and the infrastructure necessary to perform these calculations.

- The Modern Era Reanalysis for Retrospective Applications version 2 (MERRA2) data set is on the order of 400 TB of size and has global atmospheric and land

data from 1980 to present.

- Data from just under 400 West Pacific tropical cyclones were extracted from the MERRA2 data sets during the months of August and September.
- Using association rule mining, the tropical cyclone data were discretized into a symbolic representation before a clustering algorithm was performed.

The goal of this effort is to not only better understand early conditions in which major tropical cyclones form, but to also understand the process of using a machine learning algorithm on large data sets within a high performance computing environment.

Technical PI and Presenter at SC16: Daniel Duffy  
NASA Center for Climate Simulation  
Goddard Space Flight Center

## **27) Advanced Computing Experiences from Hardware Accelerators to Quantum Computing**

Numerical simulations are essential to NASA's program of Earth observations. Computational demands are ever increasing as computer models become more complex and run at higher and higher resolutions. In the last decade, the advent of hardware accelerators like the Intel Phi and NVIDIA Graphics Processing Units have pushed the fastest computers to tens of petaflops. More recently, application-specific solutions like D-Wave's quantum annealer and IBM's TrueNorth computer are active areas of research. To advance numerical climate simulations we investigate the potential of these architectures, including.

- Implementing OpenMP and structural changes to the finite volume dynamical core of the GEOS-5 climate model to run efficiently on the Intel Phi to determine the potential for performance gains.
- Casting data assimilation and image registration as quadratic unconstrained binary optimization problems suitable for the D-Wave quantum annealer. Researchers ran and tested the performance of the quantum annealer for these applications.
- Investigating the use of Restricted Boltzmann Machine Artificial Neural Networks to carry out image registration, and use of the performance of the D-Wave quantum annealer to train the RBM.

The results of this work have given experience and insight into the efficient use of Intel Phi and D-Wave quantum annealer for carrying out numerical tasks. While the Intel Phi "Knight's Corner" has not led to performance improvements for our applications, we are investigating the newest generation Phi "Knight's Landing" and its potential for numerical climate simulations. The advent of quantum computers is still at its infancy and is an active area of research we continue to pursue.

Presenter and Technical PI: Craig Pelissier and Daniel Duffy  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

## 28) Data Analytics and Storage System (DASS)—Mixing POSIX and RESTful Architectures

The NASA Center for Climate Simulation (NCCS) has deployed the Data Analytics and Storage System (DASS), a combination of scalable storage with compute resources to perform in-situ analytics. This system was designed and deployed to provide a platform for performing large-scale data analytics on complex high-fidelity simulation data.

- Researchers took benchmark performance measurements for both traditional Hadoop benchmarks including TeraSort and Wordcount and large-scale climate analytical operations on NetCDF data.
- The Spatiotemporal Indexing Approach bridges the logical array-based data model and the physical data layout enabling fast data retrieval when executing spatiotemporal queries.
- All testing took place on the Hadoop File System (HDFS) Lustre and the Hadoop Adapter for Luster and GPFS and the GPFS HDFS transparency connector.

As model output data has grown in size, so has the need for storage and analytics. To address this need, the NCCS has deployed the DASS. This system provides large capacity to store commonly used climate simulation datasets for long periods of time while also providing storage proximal computing capabilities for performing data analytics.

Presenter at SC16 and Technical PI: Carrie Spear and Daniel Duffy  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

## 29) Bursting Into Public Clouds For High-Performance Analytics

The NASA Center for Climate Simulation (NCCS) is working to leverage the power of onsite resources coupled with cloud-based assets to expand—on a dynamic basis—the computational power afforded its user population. The objective of this “cloud bursting” is to demonstrate automated overflow processing from the NCCS Advanced Data Analytics Platform (ADAPT) to Amazon Web Services (AWS) using the SubmitOnce component of Cycle Computing’s CycleCloud provisioning tool.

The test case centers on using high-resolution satellite imagery and custom processing algorithms to identify and count trees and shrubs within a coast-to-coast swath of Africa from the Sahara south to the savanna zones. This biomass estimate will reveal how much carbon is stored in the region’s vegetation that could be released as carbon dioxide if those plants burn or die and decompose due to natural or human causes and thus have direct impact on the region’s climate.

- Satellite imagery is organized into 11 Universal Transverse Mercator (UTM) zones, which are roughly 658 by 1,334 kilometers (km) each. UTM zones are further broken down into 100- by 100-km tiles.
- Seamless sub-tiles with image distortions removed are created on ADAPT and then transferred to AWS for the biomass calculations. The batch queueing system in ADAPT communicates with its elastic twin in AWS, enabling Cycle

Computing's tools to move data and initiate processing. AWS then returns results to an ADAPT file system.

The flexibility of the NCCS ADAPT system backed by the relatively infinite AWS resources potentially redefines the traditional science-processing platform. Processing algorithms can be curated on local assets and then hosted in AWS when users want to process large data stores and need many cloud instances to complete processing in a reasonable timeframe.

Presenter at SC16 and Technical PI: Hoot Thompson and Garrison Vaughan  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

### **30) High Performance Computing and OpenStack**

The NASA Center for Climate Simulation (NCCS) at the Goddard Space Flight Center is building a production OpenStack cluster within a physical container, separate from the High Performance Computing (HPC) datacenter. The OpenStack cluster will be a more flexible environment for many use cases: for example, users who need applications not found in the HPC cluster, or who want operating systems not used in the HPC cluster, or whose work patterns are not conducive to a batch scheduler. It will still be possible to run parallel jobs using MPI within an OpenStack tenant, but at a higher latency since the interconnect will be 10 GbE rather than the more traditional Infiniband found in HPC environments. For the NCCS, OpenStack is not being deployed for parallel work; rather it is intended for more inherently parallel, loosely coupled processing.

One very appealing feature of OpenStack is the potential for an enhanced security environment. Whereas the HPC environment is a single security domain, vlan-based OpenStack tenants allow for the creation of different security domains for different groups and to be able to tightly control what those tenants have access to. Large, shared datasets will be available to authorized OpenStack tenants, just as they are accessible from the HPC cluster. Accessing a large parallel filesystem from an OpenStack tenant poses a number of unique challenges in an itinerant cloud environment, which will need to be mitigated in this project.

Presentr at SC16: Jonathan Mills  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

### **31) Climate Analytics as a Service**

Climate model input and output data provide the basis for intellectual work in climate science. As these data sets grow in size, new approaches to data analysis are needed. In efforts to address the big data challenges of climate science, some researchers are moving toward a notion of Climate Analytics-as-a-Service (CAaaS). CAaaS combines high-performance computing and server-side analytics with scalable data management, cloud computing, a notion of adaptive analytics, and domain-specific Application Programming Interfaces (APIs) to improve the accessibility and usability of large collections of climate data. In this section we take a closer look at these concepts and a

specific implementation of CAaaS in NASA's MERRA Analytic Services project. Three components are needed for truly high performance Data Analytics:

- Relevant and interesting data - NASA's MERRA Reanalysis data set is widely used and we use it as a proxy for Climate Model data.
- Convenient and Extensible exposure of the data to users. – Storage-proximal analytics with simple canonical operations; average, maximum, minimum, sum, count, variance and anomaly.
- High Performance Compute and Storage Fabric – We have deployed ADAPT (Advanced Data Analytics Platform) combining purpose built VMs on HPC nodes, 5 PB of data, and a High Performance Hadoop cluster.

Climate Analytics-as-a-Service has reduced tedious data wrangling and analysis from weeks or months to minutes or minutes. There is much more work to be done including extending the canonical operations, introduction of native NetCDF into HDFS and testing other in-memory technologies such as Spark.

Technical PI and Presenter at SC16: Phil Webster  
Computational and Information Sciences and Technology Office  
NASA Goddard Space Flight Center

### ***32) Improving Simulations of Effects of Volcanic Mega-Eruptions on Earth's Climate***

Volcanic mega-eruptions such as 1991's Pinatubo are known to affect climate by cooling much of the Earth. While Pinatubo was the largest eruption of the last century, some eruptions of the last millennium are estimated to have been 10 times larger. In climate model simulations of the past thousand years, mega-eruptions are followed by long cold spells that can dramatically reduce global temperatures. However, following a major international study of climate model performance, completed in 2012, it became widely recognized that proxy measurements of past climate, e.g., tree rings and polar ice cores thousands of years old, do not reveal as drastic, worldwide cooling as did the climate models. A recent NASA study has revealed that to correctly model the climate effects of a volcanic eruption – especially mega-eruptions many times larger than anything in recent history – scientists need to include the atmospheric effects of erupted water vapor, in addition to the effects of sulfur compounds which have previously been the primary focus. Computations such as these were previously infeasible, but the recent availability of more powerful supercomputer resources has enabled scientists to more fully represent their understanding of volcanic emissions' chemistry-climate interactions in computer simulations of Earth's climate.

Technical PI and Presenter at SC16 Ellen Salmon  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center

**33) Bryan O'Gorman - [here](#)**

# The Universe

## 34) Exploring the Sun through High Resolution Modeling

### Overview

At NASA we use cutting edge techniques in supercomputer modeling to gain new insights into the dynamics of the sun. With recent advances in computing power, we are able to create more accurate models that can help us to predict physical phenomena and give us a better understanding of the interior workings of the sun and the stars. Our goal is to create 3D global models which bring together many different physical stellar processes so that we can see how they interact on a global scale. We hope to use these models to create accurate solar magnetic cycles to better understand the effects that they produce and to track their evolution over the star's lifetime.

### Project Details

We use techniques in computational fluid dynamics to create global 3-D models of the sun. We solve for the mass flows of plasma in and around the convective zone, which is where most of the turbulent motion and flow occurs. We simulate how these flows generate and sustain a large global magnetic field and how this field produces the solar cycle.

With the help of the visualization team at the advanced supercomputing division, we use this data to create striking high resolution videos of these models. These visualizations give us a chance to gain a deeper understanding of the underlying physics and help us to improve our understanding of elusive solar phenomena.

### Results and Impact

Our models generate a decade long magnetic solar cycle that corresponds to our observations of the sun. We are gaining new insights into how these magnetic fields are stored and how they propagate throughout the system. These high resolution simulations offer us a greater understanding of how turbulent effects inside of the sun can create magnetic structures at the bottom of the convection zone, which weren't expected from standalone theoretical work. These models offer great learning opportunities for researchers and students and can help explain complicated phenomena that could be difficult to understand.

### Role of High-End Computing (Why HPC Matters)

High-end computing is an indispensable part of exploring the universe today. At NASA we use some of the most advanced computing architecture in the world to help solve problems in helio- and astro-physics. The computing resources on the Pleiades system let us run incredibly detailed simulations with complicated advection algorithms which would take years on regular computers. The data that we generate can also be visualized on the hyperwall at NASA which gives an unprecedented look at the structure of these models.

### What's Next (optional)

We hope to use this data to further improve our models of the sun until we are able to generate complex and realistic models which can give us a predictive power of solar weather similar to that of meteorological models that we use on the earth today.

*Presenter at SC16: Andrey Stejko, NASA Ames Research Center,  
astejko@nas.nasa.gov  
Technical PI: Alexander Kosovichev, NASA Ames Research Center,  
sasha@sun.stanford.edu*

## **35) From Tomography to Material Properties of Thermal Protection Systems**

### **Overview**

Protecting a spacecraft during atmospheric entry is one of highest risk factors that needs to be mitigated during a mission design. NASA is actively developing new material architectures to protect its spacecrafts and probes as they enter planetary atmospheres at hypersonic speeds. At entry speeds from space, air turns into high-temperature plasma, and spacecraft Thermal Protection Systems (TPS) are needed to protect the vehicle payload and manage the heat from the freestream plasma. For extreme speeds such as return from the moon, sacrificial materials such as ablators are the only option as TPS materials.

Modern successful material architectures of spacecraft shields use porous substrate impregnated with phenolic. The substrate is designed to be highly porous to minimize weight and heat conduction, but these properties make it susceptible to penetration of hot gases. To mitigate the rush of gases into the protective substrate, the pores of the substrate are filled with phenolic that decomposes when heated. The decomposition process generates gases that block the incoming plasma.

### **Project Details**

Developing new material architecture involves an extensive design and test cycle. Several material properties are key to a successful design of a spacecraft heatshield. These include conductivity, permeability, material recession rate, etc. To accelerate the design cycle process and reduce the need to extensive testing, NASA is developing modeling and simulation tools that enable characterizing material properties and response to hot plasma.

Realistic digital representation of a substrate can be generated either synthetically or from high-resolution tomography of a built substrate. The digitized representation is then used to compute fundamental properties of the material, and the response of the material to high temperature plasma.

### **Results and Impact**

Direct simulation Monte Carlo (DSMC), a particle based method for approximating the Boltzmann equation, is used to compute the permeability of low-density carbon fiber materials based on their digitized volumes. The method provides an accurate model for the entire boundary layer including the flow within the microstructure, where the size of the fibers may approach the mean-free-path of the flow.

- Numerical computations of permeability were compared with experimental results for FiberForm (used on several NASA missions) show excellent agreement.
- The permeability of other substrates have been computed include those used on the SpaceX vehicle and carbon felt a new substrate in development at NASA

### **Role of High-End Computing (Why HPC Matters)**

The Stochastic PARallel Rarefied-gas Time-accurate Analyzer (SPARTA) developed at Sandia Labs. was used in the current work. Grid is used to group particles by grid cell for purposes of performing collisions and chemistry, as well as efficiently find particle/surface collisions. The datasets used 200x200x200 voxels, and more than 10 millions total cells were used for our simulations. Around 200 millions particles were used in the domain, leading to an average number of particle per cell around 20, needed to reduce statistical scatter.

### **What's Next**

As NASA contemplates large payloads to Mars, missions to Venus and Giant planets, the goal for the next four years is to develop a material design framework that will enable rapid development of mission tailored materials for thermal protection system.

Presenter at SC16: Nagi N. Mansour, NASA Ames Research Center  
[nagi.n.mansour@nasa.gov](mailto:nagi.n.mansour@nasa.gov)

Technical PI: Arnaud Borner, UIUC & NASA Ames Research Center  
[arnaud.p.borner@nasa.gov](mailto:arnaud.p.borner@nasa.gov)

## **36) Data-Constrained Simulation of Coronal Mass Ejection Propagation**

### **Overview**

Coronal Mass Ejections (CMEs) are powerful expulsions of plasma from the surface of the Sun that propagate towards Earth and interact with the Earth's magnetosphere causing aurora, failure in power grids and satellite electronics and GPS systems, and danger to astronauts in space. Hence, there is the need to learn to predict where, when, and how a CME will come about and how it will change as it propagates through the heliosphere on its way through the solar system. This motivates the creation of ever more realistic simulations of CME propagation relevant to NASA goals and objectives. The results of such simulations can be compared to observations from various NASA spacecraft missions.

### **Project Details**

This project utilizes the Space Weather Modeling Framework (SWMF) global magnetohydrodynamics (MHD) code to run simulations of a CME initiated by a realistic non-linear force-free field (NLFFF) model of the 3D magnetic field of solar active region NOAA 11060 that erupted on 08 April 2010. A grid of initial conditions with different amount of free magnetic energy (more and more unstable) are explored. 3D variables of all plasma parameters such as velocity, density, temperature, and magnetic field are produced as a function of position and time as well as simulated white light and extreme ultraviolet (EUV) emission as seen by the SOHO/LASCO, STEREO/COR and EUVI, and SDO/AIA instruments.

### **Results and Impact**

We determined the effect of the different initial conditions (I.C.s) on the morphology and plasma characteristics of the CME ejecta and issued a best fit model to the coronagraph observations. For example increasing the axial flux in the flux rope of the active region increases the speed of the ejecta – our I.C.s span speeds from 300 to 2500 km/s.

- The best-fit model has axial flux of  $11 \times 10^{20} \text{Mx}$  and produces CME velocity of 740 km/s as observed.

- Velocity (t) and height (t) curves are derived and are matched to those from observations.
- The simulated white light images reproduce the deflection of the CME towards the equator as observed by NASA spacecraft.

### **Role of High-End Computing (Why HPC Matters)**

The SWMF code is a massively parallel code, which cannot be run on regular workstations, it can only be run on supercomputers for production runs. This project, in particular, was run on the PLEIADES supercomputer at min 1000 processors for minimum of 16 days per run. 16 days per run produce about 35 min of propagation of the CME and total of 1.5 h f propagation were run for all 8 I.C.s, which makes this a massive supercomputing project. The necessity of such long runs comes from the fact that we have used large resolution to resolve any shocks and diagnose any artifacts since this is the first CME that we analyze with this code and these realistic I.C.s.

### **What's Next**

The next step is to run more events that have interesting signatures at Earth to 1AU with adaptive mesh.

*Presenter/PI Antonia Savcheva, NASA Ames Research Center, asavcheva@cfa.harvard.edu*

## **37) Determining How the Sun's Weather Controls Earth's Space Weather**

### **Overview**

Earth is vulnerable to bursts of X-rays and high-energy protons and electrons produced by energetic eruptions (flares and coronal mass ejections) from the solar surface. These harmful space weather events can harm astronauts, disable satellites, disrupt radio communication and global positioning systems, and interfere with electricity distribution. Earth's space weather is driven by the evolution of the Sun's magnetic field, especially the emergence of new magnetic flux and active regions, commonly known as sunspots. A major NASA science goal is to combine observations and theoretical models of emerging magnetic fields in order to predict the occurrence of flares and coronal mass ejections.

### **Project Details**

This project aims to model the role of solar magneto-convection in the emergence of magnetic fields through the solar surface and the formation of active regions. The outer 10% (two-thirds of the density scale heights) of the solar convection zone is modeled by numerically solving the conservation equations for mass, momentum, and energy plus the induction equation for the magnetic field using a three-dimensional, finite difference, magnetohydrodynamic simulation code.

In our code, ionization energy accounts for two-thirds of the energy transported near the solar surface. Excitation and ionization of the abundant solar elements is included in the equation of state. Radiation from the solar surface produces the low entropy, high-density plasma whose buoyancy work drives the convective motions. Radiative heating and cooling is calculated by solving the radiation transfer equation using a 4 bin multi-group method and assuming local thermodynamic equilibrium. A horizontal magnetic field is specified as a bottom boundary condition.

## **Results and Impact**

Our simulations reveal that convection twists the magnetic field into serpentine loops. When the tops of a loop reach the visible solar surface, a magnetic bipole appears. As more flux emerges, like polarities merge to form strong field concentrations and sunspots. Results from these calculations are used as synthetic data for solar observational techniques. Observational signatures of active regions prior to their emergence is being sought. Add a sentence that ties the title with the results for a lay audience, if possible.

In the next phase of this project, we will use results from solar global dynamo calculations to provide the bottom boundary conditions.

## **Why HPC Matters**

The simulations are being run on 4,032 cores on the Pleiades supercomputer at the NASA Advanced Supercomputing (NAS) facility at NASA's Ames Research Center. Multi-processor computers such as Pleiades are necessary to simulate magneto-convection with large enough dimensions ( $4032^2 \times 500$ ) to include interactions of very disparate spatial and temporal scales, for example, from active regions (100,000 kilometers) to sub-granules (24 kilometers). These deep (20–30 Mm) systems have long turnover times of two to three days, so a long evolutionary time is required for observable structures to develop and emerge. The sophisticated visualization methods necessary to understand simulation results are provided by the NAS visualization team.

*Presenter at SC16: Robert Stein, Michigan State University, [stein@pa.msu.edu](mailto:stein@pa.msu.edu)  
Technical PI: Aake Nordlund, Niels Bohr Institute, [aake@nbi.ku.dk](mailto:aake@nbi.ku.dk)*

## **38) Modeling Nanoscale Light Scattering by the Lunar Regolith**

### **Overview**

When Apollo astronauts brought back soil samples from the Moon, we realized that soils were darker than the rocks that they came from. Additionally, the soils reflect less light at shorter wavelengths than they do at longer wavelengths (referred to as being “redder”). Furthermore, the longer a soil has been exposed on the surface of the Moon, the darker and “redder” it is. This change in reflectance is the result of a number of processes that occur to materials on the surface of airless bodies that we cumulatively refer to as “space weathering.”

Our work is focused on understanding the cause of this change in appearance and the optics and physics behind it.

### **Project Details**

We use a combination of laboratory observations and computer models to characterize the optical effects of space weathering. The idea behind our work is that if we can make a physics based computer model that accurately reproduces the observed spectra of these materials, then we have accounted for the factors responsible for the change in appearance.

We use the Multiple Sphere T-Matrix Model (MSTM) to simulate the reflectance of mixtures of different substances. MSTM allows the user to construct simulated clusters of spheres in any arrangement so long as the surfaces of two spheres don't overlap.

Each sphere is assigned an index of refraction and an extinction coefficient. The model then directly solves Maxwell's Equations at every wave particle interface. The model then outputs information that we can turn into a simulated reflectance profile, allowing us to directly compare our results to laboratory data and remote sensing observations of the lunar surface.

### **Results and Impact**

Space weathering reduces the spectral contrast of the surfaces of airless bodies in the solar system. This can make it extremely difficult to remotely identify the composition of the surface of the Moon, asteroids, and comets. By fully understanding the optical effects of space weathering we can deconvolve the space weathering contribution to the spectrum and better recover compositional information.

Additionally, our research speaks directly to better understanding the role of Mie scattering in natural systems. Mie scattering occurs when the particle scattering the light and the wavelength of the light are approximately the same size. Most models used to determine the composition of planetary surfaces do not account for Mie scattering resulting in poor fits of model data to observations when particle sizes are smaller than about 10 microns.

### **Role of High-End Computing (Why HPC Matters)**

A typical model run for MSTM involves solving several hundred thousand equations as many as 2.5 million times. Due to the nature of the math, MSTM can take advantage of massively parallel computers like NASA's Pleiades Supercomputer to vastly reduce computation times. Runs that take 6-8 hours on Pleiades would take years to run on a regular desktop computer. HPC not only enhances our project, it is the only reason we are able to use this technique in the first place.

### **What's Next (optional)**

We intend to increase the complexity of the particles we are examining with MSTM from spherical clusters of homogeneous particles to more realistic heterogeneous particles with inclusions of different substances and void spaces. This will result in more realistic simulations of complex lunar and asteroid compositions.

*Carey Legett, Stony Brook University, [carey.legett@stonybrook.edu](mailto:carey.legett@stonybrook.edu)*

*Timothy Glotch, Stony Brook University, [timothy.glotch@stonybrook.edu](mailto:timothy.glotch@stonybrook.edu)*

## **39) Highly Parallel Radiative Transfer in Stellar Simulations**

### **Overview**

Stars function the way they do in large part because heat is transferred from point to point within them via light waves. Of course, we also see them, study them, and in fact exist because of the light that leaks out of their interiors into space. In order to simulate stars, this heat transfer must be properly modeled in combination with the other physical processes going on, namely fluid flow, gravitational forces, and magnetic field effects. This whole set of processes is called "radiative magnetohydrodynamics" or "radiative MHD". In our work we do not directly simulate the nuclear reactions that power stars, but their effect is included as a boundary condition specifying the heat flowing upward from the reactions in the stellar core. This work is part of NASA programs that seek to understand stars in terms of their dynamics, evolution, interaction with planetary systems, and role in the large-scale structure of the universe. The results and

implications of stellar simulations interact closely with observations made using NASA satellites and other facilities worldwide.

### **Project Details**

Our group has developed two codes for stellar simulation, StellarBox and StellarSegment. The former is our current workhorse and simulates a small rectangular block of stellar matter that extends a few tens of megameters (Mm) into the interior and a few Mm up into the stellar atmosphere. StellarSegment, which is under development, simulates an orange-segment-like piece of a star, with potential for including the entire star, down to some radius below the surface. Both codes utilize a form of highly accurate radiative transfer calculation that utilizes long “characteristics”, or light rays, that extend in many directions from each point. The computational cost of such a method can be very high, such that it dominates the overall runtime, but by using carefully designed parallel-programming algorithms, this cost can be controlled and made commensurate with that of the rest of the simulation.

### **Results and Impact**

By accurately simulating the radiative heat transfer within a star, combined with advanced turbulence modeling and numerical algorithms, we have achieved simulations of solar and stellar dynamics that exceed the detail that can be obtained observationally, even for the sun. Such high resolution has given us many new insights into the processes that make stars behave as they do and, for the sun especially, has provided new data on how the sun generates its atmospheric dynamics and solar wind. The solar wind interacts strongly with the interplanetary space environment and thereby with spacecraft, human explorers, and the earth’s atmosphere and electrical infrastructure, especially during solar storms, so that understanding its origin at the solar surface is particularly relevant to us all.

### **Role of High-End Computing (Why HPC Matters)**

Of course, none of this work would be possible without high-performance computing. But, as all users of such computers know, high-performance comes with constraints, sometimes severe, on programming techniques and numerical methods. Through many years of experience in using such systems, starting with the Illiac IV, and learning to deal with their limitations and exploit their strengths, it has become possible to use tens of thousands of cores to simulate stars at levels of physical detail that make the results highly credible and therefore of great value to the astrophysical community and ultimately to society in general.

Presenter at SC16: Alan Wray, NASA Ames Research Center  
alan.a.wray@nasa.gov