NASA/SP-20220016341

NASA at SC22 Conference Abstracts

Cetin C. Kiris
NASA Ames Research Center

Aditya S. Ghate
NASA Ames Research Center

Gaetan K. Kenway
NASA Ames Research Center

Leonardo Machado
NASA Ames Research Center

Jared Duensing
NASA Ames Research Center

Timothy Chau
NASA Ames Research Center

Gerrit-Daniel Stich
NASA AMES Research Center

Jeffrey, Housman
NASA AMES Research Center

Oliver Browne
NASA AMES Research Center

Daniel Maldonado
NASA Ames Research Center

Neal M. Chaderjian
NASA Ames Research Center

Patricia Ventura Diaz
NASA Ames Research Center

Seokkwan Yoon
NASA Ames Research Center

David Garcia
NASA Ames Research Center

Scott Neuhoff
NASA Ames Research Center

Marian Nemec
NASA Ames Research Center

Michael Aftosmis
NASA Ames Research Center

Peter Colarco
NASA Goddard Space Flight Center

Feng Li
NASA Goddard Space Flight Center

Sarah Strode
NASA Goddard Space Flight Center

Andrea Molod
NASA Goddard Space Flight Center

Qing Liang
NASA Goddard Space Flight Center

Clara Orbe
NASA Goddard Space Flight Center

Alexander Kosovichev
New Jersey Institute of Technology

Lydia Korre
University of Colorado, Boulder

Megan Damon
NASA Goddard Space Flight Center

Emma Knowland
NASA Goddard Space Flight Center

Christoph Keller
NASA Goddard Space Flight Center

Steven Pawson
NASA Goddard Space Flight Center

Goutam Konapala
NASA Goddard Space Flight Center

Sujay Kumar
NASA Goddard Space Flight Center

Carrie Vuyovich
NASA Goddard Space Flight Center

William M. Putman
NASA Goddard Space Flight Center

Cassi Lochhaas
Space Telescope Science Institute/Johns Hopkins University

Molly Peeples
Space Telescope Science Institute/Johns Hopkins University

Jason Tumlinson
Space Telescope Science Institute/Johns Hopkins University

Ronald Caplan
Predictive Science Inc.

Andres Reyes
Predictive Science Inc.

Tibor Török
Predictive Science Inc.

Cooper Downs
Predictive Science Inc.

John Stefan
New Jersey Institute of Technology

Venkateswara Rao Manga
The University of Arizona

Kevin D. Grossman
NASA Kennedy Space Center

Thomas J. Zega
The University of Arizona

Kyle Augustson
Northwestern University

Daniel Lecoanet
Northwestern University

Michael Way
NASA Goddard Institute for Space Studies

Jordan Angel
NASA Ames Research Center

Michael Barad
NASA Ames Research Center

Josette Bellan
NASA Jet Propulsion Laboratory

Kaushik Balakrishnan
NASA Jet Propulsion Laboratory

Jamie Meeroff
NASA Ames Research Center

Travis Rivord
NASA Marshall Space Flight Center

Peter Liever
NASA Marshall Space Flight Center

Jeff West
NASA Marshall Space Flight Center

William Thigpen
NASA Ames Research Center

Christopher Kung
NASA Goddard Space Flight Center

Purnendu Chakraborty
NASA Goddard Space Flight Center

Hoot Thompson
NASA Goddard Space Flight Center

Daniel Duffy
NASA Goddard Space Flight Center

Ramón Ramírez-Liñán
NASA Goddard Space Flight Center

Mike Little
NASA Goddard Space Flight Center

Tsengdar Lee
NASA High-End Computing

Joseph C. Ferguson
Stanford University

John Thornton
AMA at NASA Ames Research Center

Federico Semeraro
AMA at NASA Ames Research Center

Arnaud Borner
AMA at NASA Ames Research Center

Jordan A. Caraballo-Vega
NASA Goddard Space Flight Center

Jordan A. Caraballo-Vega
NASA Goddard Space Flight Center

Mark L. Carroll
NASA Goddard Space Flight Center

Christopher S.R. Neigh
NASA Goddard Space Flight Center

Caleb S. Spradlin
NASA Goddard Space Flight Center

Thaddeus Norman
NASA Ames Research Center

Jeremy Kravitz
NASA Ames Research Center

Shubha Ranjan
NASA Ames Research Center

James Fischer
Goddard Space Flight Center

International Conference for High Performance Computing, Networking, Storage, and Analysis
Ames Research Center
Toulouse, France, October 2023

October 2023

NASA STI Program ... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- E-mail your question to help@sti.nasa.gov

- Phone the NASA STI Information Desk at 757-864-9658

- Write to:
  NASA STI Information Desk
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681-2199

NASA Research Exhibit at SC22, Dallas, TX
November 14-17
Website: NASA@SC22

Aeronautics

Wall-Modelled Large Eddy Simulations for Aircraft Certification-by-Analysis

Cetin C. Kiris, NASA Ames Research Center
Aditya S. Ghate, NASA Ames Research Center
Gaetan K. Kenway, NASA Ames Research Center

Overview

Accurate computational prediction of aerodynamics for aircraft with swept wings in high-lift configurations (take-off and landing) is notoriously challenging. The flow field for these configurations is dominated by the strong interplay between turbulent boundary layer separation, a variety of off-body vortex tubes, complex wake-boundary layer mergers, and large pressure gradients. This problem has been the focus of the aviation industry's high-lift prediction efforts, and the large variation in predictions made by independent computations underscores the need for a systematic evaluation of current state-of-the-art computational fluid dynamics (CFD) tools, especially those involving scale-resolving turbulence closure strategies, like large-eddy simulation (LES). This is particularly important to enable analysis-based compliance for aircraft certification, one of NASA's future computational aeronautics goals.

Project Details

The Launch, Ascent, and Vehicle Aerodynamics (LAVA) group at NASA's Ames Research Center is attempting to address the CFD challenge posed by high-lift aircraft configurations. The team developed and demonstrated that wall-modeled large-eddy simulations (WMLES) can capture both qualitative and quantitative aspects of high-lift aerodynamics via systematic comparisons against experimental wind tunnel measurements.

The LAVA team utilized two distinct technologies for high-lift WMLES. The first method utilized structured overset curvilinear grids—a type of grid that can be carefully tailored for quality, desired resolution requirements, and are body aligned to represent geometry perfectly. However, due to the intensive manual grid generation effort for this method, the high-quality curvilinear grids required over two months of pre-processing time to complete for the geometry of NASA's High-Lift Common Research model. The second method utilized fully automated structured octree-grids, which reduced the preprocessing time down to roughly two hours of human effort. However, since the resulting grids did not conform with the geometry, the team developed a new technique relying on immersed boundary closure with appropriate near-wall LES closures (such as wall-modeling).

Results and Impact

These simulations were the first to systematically show the qualitative nature of the convergence seen in pressure loading on high-lift geometries with grid refinement. The team demonstrated that grid refinement studies are necessary due to complexities associated with different resolution requirements needed to capture inviscid streamline curvature effects, and also the viscous boundary layer effects responsible for smooth body and corner flow separation that results in aircraft stall. Their study demonstrated that coarser grids can produce accurate results due to "error cancellation," so additional care is needed when evaluating the predictability of CFD results. The team also identified the differences between the aerodynamics occurring in "free-air" versus within a wind tunnel, especially due to the utilization of "semi-span" models mounted on the tunnel wall. They were able to delineate both the inviscid effects from changes to the model geometry, as well as the viscous effects associated with the floor boundary layer, and concluded that proper comparisons between tunnel experiments and CFD requires representation of the tunnel environment in CFD simulations, especially at angles of attack near the aircraft's maximum lift. The LAVA team also observed that the new, automated method increased computational affordability and drastically reduced the human-intensive preprocessing time—from months to just hours—which provided computational predictions at a comparable accuracy-per-unit-cost metric to similar simulations utilizing more complex methods. This successful demonstration of NASA's CFD capability has had a phenomenal impact on a variety of agency projects, including the design of the Transonic Truss-Braced Wing aircraft concept; the study of ice-accretion and resulting loss of aerodynamic performance; and aero-acoustic and vibroacoustic predictions involving high-speed jets, airframe noise, and buffet seen in the forward booster attachment for the NASA's Space Launch System.

Why HPC Matters

The rapid growth seen in scale-resolving technologies in aerospace applications is in large part due to the rapid growth in high-performance computing (HPC) resources. The increase in NASA's HPC capacity, along with the development of new algorithms that leverage the new hardware efficiently, has led to use of LES in predicting aircraft-aerodynamics several decades earlier than scholars in the industry predicted in the early 2010s. The work performed by the LAVA team represents some of the largest LES performed on non-academic geometries, with the largest simulation utilizing over seven billion spatial degrees of freedom. These simulations could represent dynamically relevant turbulent motions as small as two millimeters in length on an aircraft with a 60-meter wing half-span and a seven meter mean aerodynamic chord. The largest simulations the team performed leveraged roughly 120 AMD Rome nodes of NASA's Aitken supercomputer and took about six days to complete for a single angle of attack.

What's Next

The most exciting aspect of the technologies developed and demonstrated for this work is the massive untapped potential for further substantial reduction in simulation wall-time via strong scaling on the latest generation accelerators such as graphics processing units (GPUs). The team anticipates another order of magnitude reduction in wall-times by utilizing more computing resources, and they are looking forward to evaluating these new technologies for other NASA projects and further promoting the inclusion of scale-resolving technologies in everyday CFD applications.

Computational Analysis for the Design of NASA's Side-By-Side Concept Vehicle

Jasim U. Ahmad, NASA Ames Research Center

Overview

With the increasing presence of unmanned aircraft systems (UAS), or drones, in the national airspace, the management for access and operation of these vehicles is required. NASA's management approach is being developed under the unmanned aircraft system traffic management (UTM) program. To determine the aerodynamic characteristics of drones, wind tunnel experiments and computational fluid dynamic (CFD) analysis have been conducted. These experiments and analyses are undertaken to understand the flight capabilities of these vehicles in variable head and cross wind conditions. The focus of this investigation is to model a drone installed in a wind tunnel for varying pitch attitudes and rotor rpm settings. Specifically, the IRIS drone is modeled in the NASA-Ames 7x10-ft wind tunnel. The tunnel mounting hardware and the tunnel enclosure are modeled along with the IRIS drone geometry. The rotors of the drone are modeled using two methodologies: a rotor disk model and full rotating rotors with moving grids. Besides, the disk model is coupled with comprehensive analysis code CAMRAD. The results of the analysis are compared with available experimental data to validate the computational approach.

Project Details

Unmanned Aircraft System (UAS) flow fields are complex and require high-fidelity Navier-Stokes simulations for accurate representation. In general, the flow field of a conventional helicopter as well as multirotor vehicles, whether in forward flight or hover, consists of unsteady aerodynamics with complex three-dimensional wakes dominated by the blade tip vortices. To reduce the complexity of simulating unsteady flows of the rotating blades with moving overset grids, an option is provided in the OVERFLOW computational fluid dynamics analysis to model rotors and propellers as infinitely thin disks. This disk is represented by the momentum source terms in the governing Navier-Stokes equations. Both discrete full rotors and disk models are loosely coupled with the rotorcraft comprehensive analysis CAMRADII.

Results and Impact

The results of these investigations will provide metrics for the safe operation of these UAS vehicles in and around civil populations and in urban settings. In general, the high-fidelity CFD methods are necessary for representing flow physics in detail. The lower-order methods used here are very important tools to accelerate the time requirements for the simulation. It is also important to validate the computation with the experiment. Detailed comparison and validation were performed for both methods. The tunnel geometry is included in the detailed computational effort. The rotor-rotor interactions and interactions with the other components of the tunnel are an important aspect of the CFD validation effort in this work.

Why HPC Matters

High-performance computing resources are required for this large simulation effort. This is even true for low-order methods where one could run multiple simulations efficiently and expediently. The unsteady simulations require tremendous resources to manage the large dataset. In addition to the supercomputer requirement, researchers would be able to assimilate various tools to enhance the capability of the whole process of analysis. The simulations are carried out in NAS systems at NASA Ames Research Center. For this simulation, supercomputers Pleiades, Electra, and Aitkin were heavily utilized.

What’s Next

The CFD tools at the NAS Division will be used as design and simulation tool for next the generation of rotor-based Unmanned Aerial Systems and other concept vehicles. The rotor disk model used in this work will be utilized to accelerate this process.

NASA’s SUSAN Electrofan Aircraft: Designing the Future of Green Aviation

Leonardo Machado, NASA Ames Research Center
Jared Duensing, NASA Ames Research Center
Timothy Chau, NASA Ames Research Center
Cetin Kiris, NASA Ames Research Center

Overview

With the ever-growing concern for the environment, a multidisciplinary team of researchers from NASA’s Ames, Glenn, Langley, and Armstrong Flight Research Centers are investigating the design and performance of a revolutionary hybrid-electric transport aircraft called the SUbsonic Single Aft eNgine (SUSAN) Electrofan, which shows great potential to reduce aviation energy consumption and emissions. This aircraft concept, developed within the the agency’s Convergent Aeronautics Solutions (CAS) project, reduces the number of gas turbofan engines on a conventional single-aisle aircraft from two to one, and leverages distributed electric propulsion (DEP), boundary layer ingestion (BLI), and natural laminar flow (NLF), among other advanced technologies, to achieve an overall reduction in weight, drag, and fuel burn.

These advantages are partly enabled by a hybrid-electric propulsion system powered by a hydrocarbon fuel-consuming tail cone thruster (TCT), which provides 35% of the propulsive force, while also generating electric power via a set of generators to drive wing-mounted electric propulsors responsible for the remaining 65%.

Computational fluid dynamics (CFD) analyses play a critical role in maturing many of the advanced technologies being considered for the SUSAN Electrofan, involving complex multidisciplinary tradeoffs between aerodynamics and propulsion. Accurate simulations of the flow field are required to properly assess the design and potential benefits of these technologies, thereby contributing to a better understanding of the overall vehicle performance.

Project Details

The SUSAN Electrofan can assume a wide range of wing-mounted electric propulsion system configurations. This vast design trade space is being investigated by researchers in the NASA Advanced Supercomputing (NAS) Division at Ames. To accurately assess the potential benefits of DEP and BLI individually, three configurations are considered: Configuration 1 uses two underwing podded propulsors; Configuration 2 features a pylon-mounted DEP system consisting of 16 electric ducted fans in a mail-slot nacelle; and Configuration 3 mounts the mail-slot nacelle DEP system directly onto the lower side of the wing to leverage BLI. Baseline models are developed for each configuration, and successively refined to improve concept performance. The end product is a credible estimate of the relative benefit offered by each configuration.

In parallel, this research also focuses on investigating the aero-propulsive performance of the TCT, which is designed to benefit from BLI. CFD analyses contribute to a comprehensive characterization of the flow ingested by the engine, thereby informing design choices and providing insight into the advantages and design challenges of BLI. One design challenge actively being studied includes inlet flow distortion intensity, which can be addressed through careful reshaping of the fuselage diffuser.

All CFD analyses are performed with the Launch, Ascent, and Vehicle Aerodynamics (LAVA) flow solver developed within the NAS Division’s Computational Aerosciences Branch. The LAVA flow solver considers the compressible, steady Reynolds-Averaged Navier-Stokes (RANS) equations and utilizes an actuator zone model to simulate the propulsive effects of each thrust-producing element.

Results and Impact

By accounting for the interactions between aerodynamics and propulsion, these investigations contribute to further understanding the design and integration of advanced propulsion technologies into next-generation commercial aircraft. These technologies have the potential to significantly reduce aviation energy consumption and emissions—key elements in the future of aviation as the industry pushes for greener aircraft solutions. The relative benefits of DEP and BLI are also quantified on an aircraft systems level, providing valuable insight into the benefits that each technology can provide.

Why HPC Matters

Supercomputing resources available at the NAS facility are paramount for studying these novel propulsion concepts. These resources enable fast and efficient simulation turnaround times that accelerate the exploration of the vast aerodynamic and propulsion trade space while facilitating a rapid exchange of designs and analyses between the various disciplinary teams for a seamless collaborative effort. All simulations were performed on the Pleiades supercomputer, using up to 800 Broadwell cores per simulation.

What’s Next

Design of the SUSAN Electrofan continues, with emphasis on applying high-fidelity aerodynamic shape optimization to improve the aero-propulsive performance of each aircraft component. Future work also includes the generation of an aerodynamic database to support aircraft systems analysis, stability and control, and other efforts. Furthermore, researchers will continue close collaborations on the internal flow path and blade geometry design for each propulsion unit, refinement of the main engine design cycle, structural sizing and analysis, NLF wing technology, and development of scaled model prototypes for flight testing and for exploring control and stability augmentation strategies.

Scale Resolving Jet Noise Simulations to Reduce Airport Noise

Gerrit-Daniel, Stich, NASA AMES Research Center
Jeffrey, Housman, NASA AMES Research Center
Aditya Ghate, NASA AMES Research Center
Cetin, Kiris, NASA AMES Research Center

Overview

NASA's Commercial Supersonic Technology (CST) project has formulated a technical challenge to design a quiet low boom supersonic aircraft that meets Federal Aviation Administration noise regulations. Landing and takeoff noise (LTO) poses a high risk to the commercial supersonic market. However, there is currently no certification noise rule for commercial supersonic aircraft. For this reason, technical committees need reliable noise predictions to assess environmental impacts against the economic benefits of supersonic travel, and it is important that the community agrees on the quality of the underlying data before approving a future regulation.

To address this problem, our team of research scientists at the NASA Advanced Supercomputing (NAS) facility at NASA's Ames Research Center performed wall-modeled large-eddy simulations (WMLES) with the Launch, Ascent, and Vehicle Aerodynamics (LAVA) computational fluid dynamics (CFD) software to predict jet noise for single-stream asymmetric round jets at several flow conditions. The simulations address the new Prediction Uncertainty Reduction technical challenge within the context of the CST project, with the goal to quantify and reduce uncertainties from scale-resolving simulations to assess noise characteristics of next-generation quiet supersonic commercial jets during LTO conditions where the noise from the exhaust jet dominates. Our simulation results show an order of magnitude cost reduction compared to an earlier study of this configuration—attributed to algorithm and software improvements—and demonstrates that WMLES can be used as a cost-competitive approach for jet noise predictions.

Project Details

Under the Transformational Tools and Technologies project, we performed scale-resolving simulations with the LAVA solver to predict jet acoustics for NASA's Low Boom Flight Demonstrator (LBFD). This included using both hybrid Reynolds-Averaged Navier-Stokes/Large-Eddy Simulation and WMLES methods on validation cases with increasing complexity, such as jet-surface interaction noise, multi-stream, and chevron nozzle designs. First, we developed best practices and assessed the uncertainties, strengths, and drawbacks of each simulation method. To better quantify uncertainties, we generated an extensive database comprising several simulations in static and in-flight conditions, and carefully compared them with experimental data. We then applied our lessons learned to a complex configuration, including several jet noise reduction technologies like internally lobed mixers, plugs, and multi-streams.

Results and Impact

By improving WMLES capabilities within the LAVA solver, we successfully predicted jet noise and jet-surface interaction noise. Great performance improvements reduced the overall simulation time by about 65x (since 2019) and the runtime from weeks to hours. The findings from these simulations have resulted in new insights and additional experimental investigations of certain flow phenomena, performed in collaboration with researchers at NASA's Glenn Research Center. The code improvements also enabled generation of a database with several static and in-flight flow configurations. By increasing the output of scale-resolving simulations from a few per year to several per month, we have a better understanding of the uncertainties and strengths of WMLES for jet noise predictions. Gaining confidence in the predictive capabilities of LES for jet noise is another step toward vehicle certification by analysis.

Why HPC Matters

Supercomputers at the NAS facility are instrumental in enabling next-level, large-scale computational aeroacoustics. Due to the nature of the broadband noise generated by jets, it is necessary to run our simulations with a very small timestep over long periods, requiring computational tools that can use current computer hardware very efficiently. Additionally, new load balancing features developed at NAS enabled strong scaling, and we used 128 AMD Rome nodes on the Aitken supercomputer for 250-million-grid-point grid systems—resulting in simulation wall-times of less than two hours to predict broadband noise. Each simulation generates hundreds of terabytes of data that is then analyzed in the frequency domain. For this, special post-processing tools were developed that reduce the cost of accessing data on disk and help go beyond what traditional data analysis tools offer.

What's Next

While great emphasis has been put on making the numerical tools as efficient as possible, gaining speedups of more than 6500%, the manual mesh generation effort is still the biggest bottleneck to increase turnaround times for generating larger databases. The mesh generation for complex configurations utilizing curvilinear meshes is labor intensive and can take an experienced engineer days or weeks of meshing. Ongoing work within the LAVA solver framework is focused on applying WMLES to a mesh paradigm that enables automated meshing, reducing the mesh generation time from weeks to minutes on a modern workstation.

Predicting Buffet Onset on the Transonic Trussed-Based Wing

Oliver Browne, NASA Ames Research Center
Jeffrey Housman, NASA Ames Research Center
Daniel Maldonado, NASA Ames Research Center
Cetin Kiris, NASA Ames Research Center

Overview

The joint NASA/Boeing effort investigating a Transonic Truss-Braced Wing (TTBW) concept design contains advancements in technology that have the potential to improve fuel efficiency of commercial aircraft. The unconventional configuration of the TTBW - which includes a high aspect ratio wing along with and wing and jury struts – leads to complex flow phenomena such as transonic buffet, separated flow and a turbulent wake. Current industry best practices tend to lean towards employing Reynolds-Averaged Navier-Stokes (RANS)-based computational fluid dynamics (CFD) analysis for buffet onset prediction. To accurately predict the onset of buffet and the development of the separated flow, more accurate scale resolving CFD simulations may be required. To provide more insight into the best practices for using scale resolving simulations to predict transonic buffet onset among other challenges, the Advanced Air Transport Technology Project within NASA's Aeronautics Research Mission Directorate launched a collaborative multi-center effort to develop new methods for simulating the TTBW in order to better predict its performance and that of similar truss-braced wing configurations.

Project Details

Researchers in the NASA Advanced Supercomputing (NAS) Division's Computational Aerosciences Branch are simulating the Mach 0.8 cruise configuration of the TTBW for validation purposes and to develop best practices in their scale resolving CFD simulations. A wind tunnel experiment of a half-span model of this TTBW configuration was conducted in the NASA Ames Unitary Plan 11- by 11-foot Transonic Wind Tunnel in January 2022. Large angle-of-attack sweeps at different Mach numbers were ran during the experiment and the data collected is being used to validate the CFD simulations, with a goal to be able to accurately reproduce experimental results using CFD.

The Launch Ascent and Vehicle Aerodynamics (LAVA) team at NAS selected Hybrid RANS/LES (Large Eddy Simulations), HRLES, as the initial choice for the scale resolving simulations approach. As the name suggests, this hybrid method models turbulence inside the boundary layer solving the RANS equations while the turbulence is resolved outside the boundary layer with LES. The dominant transonic buffet phenomena can largely be captured by steady-state RANS or unsteady RANS (URANS). However, RANS/URANS is unsuitable for investigation of the phenomena in the turbulent wake (including deep stall and high-lift configuration) due to excessive dissipation of turbulent motion. Furthermore, RANS has be demonstrated to be an unreliable tool when moving to CLmax and the onset of stall by the LAVA group and others most recently at the High Lift Prediction Workshop 4 (HLPW4). Scale resolving simulations, such as HRLES approaches are better able to resolve the turbulent content and enable more accurate predictions.

Results and Impact
Initial HRLES simulations have been conducted to establish the best practices for the computational grid, an appropriate timestep size and the Spalart-Allmaras (SA) closure to employ with the HRLES approach. It was shown by the LAVA group during HLPW4 that utilizing RANS type grids for HRLES simulations will likely produce a result that is inferior to that of a RANS simulation. Therefore, the grids are purpose built HRLES grids with low aspect ratio grid cells that are appropriate for LES outside the boundary layer. The aerodynamic loads (lift, drag and pitching moment) and surface pressure data from HRLES simulations has been compared with RANS/URANS simulations and the experimental data from the NASA Ames Unitary Plan 11- by 11-foot Transonic Wind Tunnel. We have been considering the in-tunnel configurations with the uncorrected tunnel data and the free air configuration with the corrected tunnel data. Relatively good agreement between the CFD methods and the experiment is observed for the Mach number and angles-of-attacks considered so far. At the higher angles-of-attack considered, the RANS simulations tended to overpredict the location of the shock in the chordwise direction, in both the midboard and outboard regions of the wing, compared to the experiment whereas the HRLES simulations were able to show much better agreement with the experiment in the predicted shock location. These initial simulations have provided important insight into the behavior of transonic buffet – in particular the spanwise development of the unsteady shock motion. The simulations have also given insight into where the computational grids need to be refined further and important information about the shielding function employed with this HRLES approach.

Why HPC Matters
NASA high-performance computing (HPC) resources have already made it possible to run the initial HRLES simulations within a short turnaround time. Scale resolving simulations are more computationally expensive than steady state RANS simulations due to the small timestep size and grid cell sizes needed to resolve the unsteady turbulent content of the flow. Furthermore, to obtain a pseudo-stationary flow field to begin averaging the flow quantities and then a large enough window with which to conduct the time-averaging procedure requires a large number of non-dimensional convective time units (CTUs) or flow-throughs (> 100 CTU's is required at higher angles-of-attack). HRLES approaches are known to have a lot of tuning parameters such as the underlying RANS closure model, the characteristic length scale, the shielding function etc. and being able to run a parameter study to test the sensitivity of these parameters with the efficient throughput provided for by NASA HPC also facilitates quick decision-making and accelerates arriving at the simulation best practices.

What's Next
Once the best practices for scale resolving simulations have been developed for the transonic buffet in the cruise configuration, we will be moving to more complex configurations including studying deep stall and the TTBW in the high lift configurations where high lift devices such as slats and flaps are deployed. Due to T-tail configuration of the empennage that is proposed for the TTBW, it may be prone to deep stall whereby the turbulent wake of the stalled TTBW main wing and strut blankets the tailplane and renders the elevators ineffective and prevents the TTBW from recovering from this type of stall. Accurately capturing the turbulent wake coming from the TTBW main wing and strut and maintaining the turbulent fluctuations until reaching the empennage will be paramount in accurately predicting the deep stall behavior. As mentioned previously, RANS has been shown to be an unreliable tool when moving to higher angles-of-attack and for the high-lift configurations. Chances of success in the abovementioned tasks will be greatly increased by employing scale resolving simulations such HRLES.

A Decade of Discovery in Rotorcraft Simulation with NASA Supercomputing

Neal M. Chaderjian, NASA Ames Research Center

Overview

Rotorcraft vehicles rely on spinning rotor blades to provide the thrust needed for vertical takeoff and landing, and the propulsive force required for forward flight. Examples include the UH60 Blackhawk helicopter, V22 Osprey tiltrotor, and a new class of green aircraft known as Urban Air Mobility (UAM) vehicles. Computational fluid dynamics (CFD) is a powerful tool that, when coupled with supercomputers, can numerically simulate the complex and time-varying flows produced by rotating blades with unrivalled detail.

However, in 2009, the average CFD prediction accuracy for rotorcraft hover performance differed from experiment by 2.4%. The importance of this seemingly small discrepancy can be appreciated knowing that an error of about 0.5% is equivalent to the weight of one passenger! The big question at the time was, why is CFD performing so poorly?

Project Details

NASA's Revolutionary Vertical Lift Technology (RVLT) Project invested significant resources over the last decade to understand these shortcomings and improve the prediction accuracy of CFD methods. High-fidelity CFD simulations use numerical algorithms to solve time-dependent Reynolds-average Navier-Stokes (RANS) equations, which, along with a turbulence model, contain all the relevant physics needed to model rotorcraft flows. A careful study was undertaken to improve simulation accuracy and better understand how numerical choices affect the physics of hover and forward flight.

Results and Impact

Over the past decade, CFD researchers at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center have achieved tremendous improvement in prediction accuracy, including the discovery of previously unknown flow phenomenon in the rotor wake. Research was first directed toward the V22 Osprey rotor in hover, and four key issues were identified and addressed in the CFD solution process: 1) improved the spatial accuracy of the numerical algorithm; 2) improved the standard turbulence model to a hybrid RANS/Large Eddy Simulation (LES) model, known as Detached Eddy Simulation (DES); 3) refined the rotor wake grid resolution; and 4) developed a quick-start procedure to reduce the computer simulation time by fourfold. These advances reduced the CFD error for rotorcraft from 2.4% to 0.2%—an order-of-magnitude improvement at a quarter of the computational cost.

Moreover, researchers efficiently refined the rotor wake below the blades using automated adaptive mesh refinement (AMR), where refined Cartesian meshes were dynamically applied to accurately capture the details of a constantly changing turbulent flow. This resulted in an unexpected discovery: the rotor wake was much more complex than previously thought, as worm-like turbulent structures (see image at right) appeared with grid refinement. The physical process for the formation of these turbulent "worms" was revealed by studying time-dependent flow visualization. Vorticity, a fundamental flow quantity that represents the spinning nature of the flow, is first formed along the rotor blades due to friction (viscosity) and descends downward as a helical sheet. When these sheets pass by a tip vortex, they are pulled into the vortex, much like a tornado pulls in objects along its path. This results in higher vorticity through a process known as vortex stretching, similar to how ice skaters pull in their arms to spin faster. The process that produces the worms stunned the rotorcraft community and resulted in several hover workshops to examine this phenomenon. Was it real or a numerical artifact? Finally, eight years later, an experimental technique known as tomographic particle image velocimetry successfully proved the turbulent worms are in fact real and physical. This demonstrates the power of discovery made through supercomputing.

The dynamic AMR process was also applied to a UH60 Blackhawk helicopter while experiencing dynamic stall in forward flight. It was previously thought that dynamic stall was solely due to a rotor blade pitching up to a high angle of attack, causing a sudden loss of lift and increased blade vibration. However, simulations showed, for the first time, that blade-vortex interaction (BVI) can also cause dynamic stall, as shown in the image at left. As the tip vortex from another blade passed over the blade in the figure, the direction of spin of the vortex caused the outboard angle-of-attack to increase too much, triggering flow separation and stall. On the other hand, the inboard angle-of-attack was decreased, keeping the flow attached and continued to provide lift.

Why HPC Matters

None of this rotorcraft CFD progress would have happened without access to NASA supercomputing resources and advanced time-dependent flow visualization. Rotorcraft flows often require from several hundred million to a few billion grid points to resolve flow details. Typical CFD simulations can take 1,000 – 6,000 processor cores running 24/7 for several days to a couple of weeks on the agency's Pleiades and Aitken supercomputers. Data written out to disk can easily exceed 100 terabytes to generate each animation. This data output provides the most convenient way to explore unknown features in a flow field.

What's Next

One remaining mystery is the premature vortex breakdown of tip vortices in the AMR simulations of the V22 Osprey rotor in hover. This seems to be triggered by vortex pairing, where two separate vortices from different blades wrap around each other like a pretzel. Further research is needed to understand the numerical cause of early vortex breakdown and its remedy.

Simulations for Designing Safe and Efficient Air Taxis

Patricia Ventura Diaz, NASA Ames Research Center
Seokkwan Yoon, NASA Ames Research Center
David Garcia, NASA Ames Research Center

Overview

The advent of Urban Air Mobility (UAM) redefines the concept of transportation in large cities, and envisions a new aviation transportation with highly automated and safe vehicles that will transport passengers and cargo in the urban and suburban areas.

NASA is playing an important role in the development of the urban air mobility ecosystem by identifying key research areas and establishing a fleet of conceptual reference vehicles for UAM applications. Conventional and more unconventional concept vehicles are being developed, including configurations such as tiltwing, quadcopters, ducted rotors, tilt-rotors, ducted-fans, distributed propulsion, and so on.

Project Details

Understanding the complex flow structures in rotary-wing vehicles is key for designing quieter, safer and more efficient UAM vehicles. This is where high-fidelity Computational Fluid Dynamics (CFD) comes into play: using NASA's state-of-the-art supercomputers the flow around different air taxi vehicles is simulated. The vehicles and the air surrounding them are modeled using hundreds of millions of grid points. The flow is solved using NASA's OVERFLOW CFD code, a high-order accurate Navier-Stokes solver. OVERFLOW is loosely-coupled with the rotorcraft comprehensive code CAMRADII. With the two codes, the complete rotorcraft solution is calculated precisely, including aeroelastic effects and trim.

Results and Impact

Among the different vehicle configurations found across the UAM community, tiltwing offers a compelling design. By tilting the wing to 90 degrees, the aircraft can hover, take-off and land vertically as a conventional helicopter. An efficient cruise flight or "airplane-mode" can be achieved tilting the wing back to a horizontal position. In this work, the performance of a tiltwing vehicle concept is analyzed in hover and in cruise.

While most of the vehicle designs considered for UAM have multiple rotors, an often disregarded configuration is the single-main rotor helicopter, despite its importance for its known path for certification. The Quiet Single-Main Rotor (QSMR) helicopter concept has been studied, and different rotor geometries are considered, finding that a drooped tip and reduced tip speed decreases the sound pressure level by 13dB when comparing to the baseline rotor design. However, cruise performance is affected and rotor efficiency drops by 60%.

The tiltwing and QSMR are part of NASA's concept vehicles, intended to focus and guide NASA research activities in support of aircraft development for the emerging UAM market. The advanced CFD tools we are developing will assist in the rotorcraft design process for future UAM vehicles.

Why HPC Matters

The performance of UAM vehicles can only be calculated accurately by solving the fluid equations on overset grids, comprising hundreds of millions of grid points. Simulations of rotorcraft are unsteady as the blades are rotating relative to the airframe. The tiltwing vehicle consists of eight rotors with five blades each, and the QSMR needs a very fine grid near the tip to capture the details of the vortices. Only with the most powerful supercomputers can such complex simulations be solved in just a few days.
The simulations were carried out on NASA's Pleiades, Electra and Aitken supercomputers at the NASA Advanced Supercomputing (NAS) facility, using 1,200 to 3,000 Intel Xeon Broadwell, Ivy Bridge, Skylake and Cascade processors. In addition, NAS's mass storage system is essential to support this work—the data stored for post-processing a single case can take up to 50 terabytes of memory.

What's Next

This work comprises a small part of the efforts conducted over the past few years analyzing the performance of NASA's UAM concept vehicles. Different concepts have been simulated using high-fidelity CFD and NASA's supercomputers, including a side-by-side rotorcraft, a quadcopter, the tilwing and the QSMR vehicles. The next steps will focus on studying the acoustics of the different vehicles, with a clear goal in mind: ensuring the future of UAM by having safer, quieter and more efficient vehicles.

Advancing NASA Software for Sonic Boom Prediction

Jared Duensing, NASA Ames Research Center
Scott Neuhoff, NASA Ames Research Center
Marian Nemec, NASA Ames Research Center
Jeffrey Housman, NASA Ames Research Center
Michael Aftosmis, NASA Ames Research Center
Cetin Kiris, NASA Ames Research Center

Overview

NASA's Quesst (Quiet SuperSonic Technology) mission teams are designing and building the X-59 supersonic aircraft—an advanced low-boom flight demonstrator that will be used to test the noise created by future supersonic transport designs. Community response testing across a wide variety of environments and locations will take place to build a database of noise levels and acceptibility metrics that will be used by the Federal Aviation Administration and International Civil Aviation Organization to set noise level requirements for supersonic overland flight. The most critical feature of the X-59 design is its ability to fly at supersonic speeds without creating an objectionable sonic boom on the ground, as well as its ability to fly at a target loudness level across a wide range of conditions for the flight test campaign. To ensure that the aircraft performs as required, NASA's Commercial Supersonic Technologies (CST) project is developing new flight planning software that will inform pilots how to fly the aircraft to achieve a desired loudness level for testing.

Teams of researchers at NASA's Ames Research Center are using the agency's Cart3D and Launch Ascent and Vehicle Aerodynamics (LAVA) flow solvers to assist in the creation of the flight planning tool, creating a database requiring thousands of high-fidelity computational fluid dynamics (CFD) simulations. To minimize the cost of developing the flight software, the Cart3D and LAVA teams are making various improvements to their analysis tools, increasing the speed of individual calculations and furthering our understanding of low-boom aircraft physics.

Project Details

The creation of the flight planning database requires running thousands of complex CFD simulations, each requiring hundreds of millions of degrees of freedom and requiring thousands of processor hours on NASA high performance computing (HPC) systems. To optimize case efficiency, the LAVA team developed and improved a space-marching technique within the LAVA framework that can be coupled to the CFD, significantly reducing the computational domain and therefore decreasing run times without any loss in accuracy. This method has been utilized by many NASA CFD teams, including LAVA, FUN3D, and USM3D, reducing simulation times by approximately 50%. Additionally, the LAVA team developed an automatic adjoint-based mesh redistribution capability optimized for sonic boom cases and space-marching, achieving the same level of sonic boom prediction accuracy with a quarter of the turnaround time on simpler cases.

Prediction of sonic-boom noise involves a multidisciplinary simulation. The CFD tools are coupled with atmospheric propagation solvers, such as NASA's sBOOM tool, to determine the sonic boom carpet ground effects generated by the aircraft. These simulations are highly susceptible to numerical errors, which may excessively attenuate the pressure waveforms. To address this issue, teams at NASA's Ames and Langley Research Centers developed a reliable error estimate for loudness predictions, which is also used to guide adaptive mesh refinement, controlling the level of discretization error to within a user-specified tolerance. This is an important step in simulation verification of the X-59 sonic boom database and future certification-by-analysis of low-boom aircraft, ensuring that the new designs satisfy noise standards.

Results and Impact

The space-marching and mesh redistribution methods developed by the LAVA team were able to reduce the computational resources required to get a high-fidelity CFD solution by a factor of at least 2. The capability of the space-marching method to be coupled with all of the NASA CFD codes used for the X-59 flight database is leading to significant cost savings over the large number of cases required.

The loudness error estimate for coupled Cart3D-sBOOM simulations developed at NASA Ames by the Cart3D team allows automatic, error-controlled analysis of the X-59. This new tool has provided significant insight into meshing requirements near the aircraft to control numerical errors in the sonic boom predictions on the ground, over 20 kilometers away from the vehicle. In addition, this tool significantly reduces the amount of human intervention the simulation needs, and virtually eliminates the need to handcraft vehicle meshes to minimize signal loss, thereby allowing unprecedented automation on NASA supercomputers.

Why HPC Matters

HPC is critical for simulation-based analysis, design, and uncertainty quantification for the X-59, since these areas require many high-cost simulations to generate high-fidelity databases. The Pleiades, Electra, Aitken, and Endeavour supercomputers at the NASA Advanced Supercomputing facility at Ames have been critical for providing a fast turnaround for this work. For example, a sonic boom carpet database for various X-59 operating conditions and configurations requires several hundred Cart3D-sBOOM simulations, with each simulation requiring 48 cores and approximately three hours of wall clock time. A typical LAVA simulation for the X-59 database requires 1,500 cores and eight hours of wall clock time, with hundreds of cases needing to be generated.

What's Next

With construction of the X-59 well underway at Lockheed Martin's Skunkworks in Palmdale, California, the improvements made to NASA's LAVA and Cart3D frameworks are helping to generate the sonic boom prediction databases needed by Quesst mission engineers to develop advanced flight planning software. Additionally, the CFD database is providing key insights in the relationship between changes in the flight parameters and their effect on ground level noise, improving our understanding of the complex physics of aircraft-formed sonic booms.

Our Planet

Using the NASA Chemistry Climate Model for Future Climate Simulations

Peter Colarco, NASA Goddard Space Flight Center
Feng Li, NASA Goddard Space Flight Center
Sarah Strode, NASA Goddard Space Flight Center
Andrea Molod, NASA Goddard Space Flight Center

Overview

Earth's climate is changing due to the accumulation of human-produced greenhouse gases (GHGs) such as carbon dioxide (CO_2) in the atmosphere. There are many impacts of the increased abundance of GHGs, most notably a long-term increase in the average temperature at the Earth's surface, known as global warming. The effects of global warming are felt in many facets of the climate system, including in the changes of global weather patterns, the retreat of glaciers and sea level rise, and changes in the locations, amounts, and duration of precipitation. These shifts are occurring alongside other changes in the composition and chemistry of Earth's atmosphere, such as the occurrence of the Antarctic Ozone Hole, discovered in the 1980s.

A team of researchers at NASA's Goddard Space Flight Center are using global Earth-system models that include interactive chemistry, called chemistry-climate models (CCMs), to investigate the evolution of the stratospheric ozone layer and its response to the reduction of emissions of ozone depleting substances following the Montreal Protocol treaty instituted in 1989. The emission changes that go into the simulations need to be accounted for in concert with other developments, such as those caused by the rise of GHGs in the atmosphere. In order for the team to have confidence in the~~ir~~ projections generated by their CCM, they must first have confidence in its at simulating the evolving climate state, so they are systematically evaluating their CCM for its ability to simulate basic aspects of GHG-driven climate change.

Project Details

The CCM uses NASA's Goddard Earth Observing System (GEOS) model, running GEOS as a coupled atmosphere-ocean model, building off of previous work producing coupled sub-seasonal to seasonal scale weather forecasts. The CCM includes two parts. The first is a pre-industrial simulation in which the oceanic and atmospheric state is set to conditions expected for the year 1850~~,~~ run for 150 years, assuming fixed pre-industrial CO_2 abundance and pre-industrial emissions of gases and aerosols to develop an approximately equilibrium state of the atmosphere and ocean. The second part of the evaluation starts from the near-equilibrium state of the pre-industrial simulation, but instantaneously quadruples the abdundance of CO_2. This system is then run for 150 years until it also reaches a new quasi-equilibrium state that reflects a climate system warmed by the increased CO_2. These two simulations are used to calculate the equilbrium climate sensitivity (ECS) of our model, a metric that allows us to compare the climate system modeling capabilities of the GEOS CCM to other climate models.

Results and Impact

The ECS of the GEOS CCM is found to be 2.6 °C (nearly 5° F), which is the calculated increase in the Earth's average surface temperature under a doubling of atmospheric CO_2 abundance from pre-industrial levels–a state expected to be met in the real world around the year 2060. This ECS—and other diagnostics computed from these simulations—is well within the range of previously evaluated climate models, which enhances the team's confidence in subsequent projections with the GEOS CCM.

Why HPC Matters

The high performance computing resources at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center, particularly the Pleiades supercomputer, are essential to carrying out this work. These simulations made use of almost twenty 40-core CPU nodes apiece, which allowed the team to run a year of simulated model time in a little over a day. The 300 years of simulated model time here would have taken almost twenty years of real world time to run on a typical desktop computer—assuming it could provide the memory and disk space needed to perform the calculation– while the team achieved results with Pleiades in a little more than six months.

What's Next

This work is foundational to future work with the GEOS CCM by establishing the basic credibility of the coupled atmosphere-ocean model for simulating future climate change. Future work will introduce detailed chemistry mechanisms and impose time varying changes in CO_2 abundance and aerosol and gas emissions to more realistically produce changes in ozone and atmospheric composition in future climate models.

An Innovative Strategy for Ensemble Seasonal Forecasts

Peter Colarco, NASA GSFC
Feng Li, NASA GSFC
Andrea Molod, NASA GSFC
Sarah Strode, NASA GSFC
Qing Liang, NASA GSFC
Clara Orbe, NASA GSFC

Overview

Earth's climate is changing due to the accumulation of human-produced greenhouse gases (GHGs) such as carbon dioxide (CO2) in the atmosphere. There are many impacts of the increased abundance of GHGs, notably a long-term increase in the average temperature at the Earth's surface—global warming. The effects of global warming are felt in many facets of the climate system, including in the changes of global weather patterns, the retreat of glaciers and sea level rise, and changes in the locations, amounts, and duration of precipitation. These changes are occurring alongside other changes in the composition and chemistry of Earth's atmosphere. In particular, the discovery of the Antarctic Ozone Hole in the 1980s led to the institution of the Montreal Protocol, a global treaty that has curbed the production and emission of ozone depleting substances—some of which are themselves potent GHGs—and will over time lead to a recovery of global stratospheric ozone levels.

We use global Earth system models that include interactive chemistry, so-called chemistry-climate models, or CCMs—to investigate the evolution of the stratospheric ozone layer and its response to changing emissions of ozone depleting substances. The emission changes that go into our simulations need to be accounted for in concert with other changes such as those due to the rise of GHGs. In order for us to have confidence in the projections we are making with our CCM we must also have confidence in its basic function at simulating the evolving climate state. In this work we for the first time systematically evaluate our CCM for its ability to simulate basic aspects of GHG-driven climate change.

Project Details

Our CCM uses the NASA Goddard Earth Observing System (GEOS) model. We run GEOS as a coupled atmosphere-ocean model, building off of previous work producing coupled sub-seasonal to seasonal scale weather forecasts. The protocol for our run involves two main simulations. The first is a pre-industrial simulation in which the oceanic and atmospheric state are set to conditions expected for the year 1850. The model is run for 150 years assuming fixed pre-industrial CO2 abundance and pre-industrial emissions of gases and aerosols to develop an approximately equilibrium state of the atmosphere and ocean. Our second simulation starts from the near-equilibrium state of the pre-industrial simulation but instantaneously quadruples the abdundance of CO2. This system is then integrated for 150 years until it also reaches a new quasi-equilibrium state that reflects a climate system warmed by the increased CO2. This pair of simulations is used to calculate the Equilbrium Climate Sensitivity of our model, a metric that allows us to compare the climate system modeling capabilities of the GEOS CCM to other climate models.

Results and Impact

The Equilibrium Climate Sensitivity (ECS) of the GEOS CCM is found to be 2.6 K. This number refers to the increase in the Earth's average surface temperature under a doubling of atmospheric CO2 abundance from pre-industrial levels, a state expected to be met in the real world in around the year 2060 (a change of 2.6 K is a change of almost 5° F). This ECS—and other diagnostics computed from our simulations—is well within the range of previously evaluated climate models, which enhances our confidence in subsequent projections with the GEOS CCM.

Why HPC Matters

NASA HPC resources, particularly the Pleiades supercomputer, are essential to carrying out this work. Our simulations made use of almost twenty 40-core CPU nodes apiece, which allowed us to run a year of simulated model time in a little over a day. The 300 years of simulated model time here would have taken almost twenty years of real world time to run on a typical desktop computer—assuming it could provide the memory and disk space needed to perform the calculation! Our results were generated at Pleiades in a little more than six months.

What's Next

This work is foundational to future work with the GEOS CCM. It established the basic credibility of our coupled atmosphere-ocean model for simulating future climate change. Future work that follows from this will introduce our detailed chemistry mechanisms and impose time varying changes in CO2 abundance and aerosol and gas emissions to more realistically produce changes in ozone and atmospheric composition in future climate.

Investigation of Solar Torsonial Oscillations and their Relation to Activity Cycles

Alexander Kosovichev, New Jersey Institute of Technology

Overview

Torsional oscillations represent bands of fast and slow zonal flows beneath the visible surface of the Sun, which are similar to stream jets in the Earth's atmosphere. Analysis of helioseismology data obtained from two NASA missions in 1996-2022 for two sunspot cycles reveals zones of deceleration of the torsional oscillations inside the Sun. The flow deceleration is caused by magnetic fields generated by the solar dynamo. It originates about 120,000 miles beneath the solar surface in a high latitude region. The zonal flows migrate through the convection zone, revealing patterns of magnetic dynamo waves. The analysis explains the phenomenon of the 'extended solar cycle' observed in the evolving shape of the solar corona and why the polar magnetic field strength predicts the solar maxima.

Project Details

All manifestations of solar activity, from spectral irradiance variations to solar storms and geomagnetic disturbances, are caused by the magnetic fields generated by a dynamo mechanism operating in the convection zone of the Sun. Despite substantial modeling and simulation efforts, our understanding of how the magnetic field is generated, transported to the surface, and forms the solar activity cycles, is poor. The most prominent feature of the solar cycle is the sunspot 'butterfly' diagram: at the beginning of an 11-year sunspot cycle, magnetic sunspot regions emerge at about 30 degrees latitude, and then the sunspot formation zone migrates towards the equator. In addition, during the sunspot maxima, the polarity of the global magnetic field in the Sun's polar regions is reversed.

In 1955, Eugene Parker showed that differential rotation and helical turbulence in the solar convection zone could result in dynamo action in the form of migrating dynamo waves that transport magnetic flux from the deep interior to the solar surface, which can explain the sunspot butterfly diagram. However, in the absence of observational evidence of the dynamo waves, alternative scenarios, called the flux-transport models, were developed. These models suggest that the cyclic evolution of the magnetic field is controlled by meridional circulation. Similarly to a conveyor belt, it transports magnetic flux to the bottom of the convection zone from high to low latitudes, where it emerges to form sunspots.

Results and Impact

The new results provide strong evidence of the dynamo waves and reveal their migration pattern in the form of two branches, migrating towards the poles and the equator. The polar reaches the surface in 1-2 years, while it takes about 8-9 years for the equatorial branch to reach the solar surface and form the sunspot butterfly diagram. Because the polar branch comes to the surface quicker, it explains why the strength of the polar magnetic field predicts the following sunspot maximum. Recent measurements show a significant decrease in the zonal deceleration in the tachocline, indicating that the current sunspot cycle may be weaker than the previous sunspot cycle. Thus, the long-term trend of declining magnetic of the Sun is likely to continue. It may impact the state of Earth's atmosphere, space weather, and climate.

Why HPC Matters

The magnetic field in the convection zone cannot be measured directly, but the magnetic field structure can be tracked indirectly through its effects on the large-scale zonal flows that are similar to stream jets and historically called 'torsional oscillations' because of their cyclic variations synchronized with the magnetic activity cycles. Magnetic forces slow down the flows, and this effect is observed in helioseismology as deceleration of the zonal flows. The flow speed is about 2-4 miles per hour and is measured in frequency shifts of solar oscillations. The measurements are performed for 22 years by two NASA space missions, Solar and Heliospheric Observatory (1996-2010) and Solar Dynamics Observatory (2010-current). The solar acoustic oscillations with a characteristic period of 5 minutes are observed every 45 seconds, and it takes 72 days of uninterrupted observations to obtain one measurement of the subsurface flows in the Sun. The total amount of data processed at the Joint Science Operations Center at Stanford University is about 5 TB. The flow analysis, numerical modeling and visualization were performed at the Supercomputing Division at the NASA Ames Research Center.

What's Next

The next steps are to improve the spatial and temporal resolution of the data analysis techniques and numerical simulation models. It is particularly important to merge models of emerging magnetic flux and sunspot formation with the global models of solar activity. Such synergy will substantially advance our understanding of the dynamo process in the deep solar interior and develop innovative tools for space weather prediction.

Modeling Dynamical Processes in the Solar Interior with 3D Global Simulations

Lydia Korre, University of Colorado, Boulder

Overview

A major challenge in heliophysics is understanding the dynamics driving the almost periodic cyclic magnetic activity observed in the Sun. Within the Sun's outer layers, the combined action of turbulent convection and the helioseismically-observed meridional circulation and differential rotation are thought to give rise to the solar dynamo. In particular, the tachocline, a thin transition shear layer that resides within the convection zone-radiative zone interface, has been presumed to be the seat of the solar dynamo. We employ three-dimensional (3D) global simulations to gain a better understanding of the dynamical processes and balances occurring within the convection zone-radiative zone interface that will help us improve our understanding of the solar tachocline and its role in the solar dynamo process.

Project Details

The convective motions generated in the turbulent convection zone can propagate into the underlying stable radiative layer through a process called overshooting. Overshooting is thought to play a significant role in the Sun and in stars as it can result in mixing of chemical species, and transport angular momentum and magnetic fields. Also, part of the tachocline must coincide with the overshoot region below the base of the convection zone. We have run a series of 3D numerical simulations of overshooting convection in a spherical shell using the open-source convection code Rayleigh. Our setup consists of an outer convective region overlying a stably stratified zone. We have performed a systematic study of overshooting convection both in a non-rotating and in a rotating regime. We have varied the number of density scale-heights in the convection zone, the degree of convective driving and the rotation rate in order to understand the effect of density stratification and rotation on the overshooting dynamics.

Results and Impact

Our results suggest that the amount of overshooting below the base of the convection zone is not a monotonic function of the number of density scale-heights in the convection zone. Instead, it depends on the ratio of the density stratification in the stable zone over the one in the convective region. For the rotating cases, we find that the overshoot lengthscale decreases when rotation becomes stronger and scales as $Ro^{0.23}$ (where Ro is the Rossby number that is defined as the ratio of inertial to Coriolis force). Although the input parameters used in numerical simulations are far from the solar values due to computational constraints arising from the required spatio-temporal resolution, these scaling laws can help us estimate the solar overshoot lengthscale by facilitating its extrapolation to solar values. They also provide us with constraints of these dynamical processes that can be ultimately tested against observations.

Why HPC Matters

Even just beginning to capture the disparate spatio-temporal scales characteristic of solar convection within a simulation requires high spatial resolution and long simulation times. Most of the twenty-two models developed in this study possessed a spatial resolution of 192 × 528 × 1056 (or higher) and required multiple months of walltime to evolve using approximately 4,100 Broadwell compute cores at NASA Pleiades. The resources required to evolve these models and to store/analyze their output are far beyond that available on a modern desktop or a typical campus-sized cluster. Moreover, the large compute allocation (~1,400,000 SBUs) required for this effort would have been extremely difficult to obtain through peer-level, national-scale supercomputing initiatives. Put simply, these state-of-the art global simulations are necessary for our gaining a better understanding of the perplexing dynamics occurring in the solar interior but would have been impossible without the support of the NASA HEC program.

What’s Next

We are currently assessing which parameter regimes lead to differential rotation profiles consistent with those observed in the Sun. Within those solar-like regimes, we plan to investigate mean flows and angular momentum transport in the convection zone and the overshoot region as well as their dependence on different input parameters. The next step is to add magnetism and study the different dynamical processes associated with the interaction of both overshooting convective motions and large-scale mean flows with magnetic fields. These efforts will ultimately shed new light on the dynamo process and facilitate the creation of a self-consistent, fully-nonlinear model of the tachocline dynamics.

High-Resolution Air Pollution Forecasting Using the NASA GEOS Model

Megan Damon, NASA Goddard Space Flight Center
Emma Knowland, NASA Goddard Space Flight Center
Christoph Keller, NASA Goddard Space Flight Center
Steven Pawson, NASA Goddard Space Flight Center

Overview

Particulate matter, ozone, nitrogen dioxide, and other pollutants all have significant health impacts on humans and vegetation. Currently, surface air pollutant observations are sparse, leaving vast areas of the globe not well understood. NASA global models with state-of-the-science atmospheric chemistry equations help fill in the gaps left by these observations and provide insight on future conditions. Due to advances in computing and software engineering, the same computer models used for weather forecasting can be used as a foundation for better predicting air pollution.

Project Details

Forecasting air quality has typically been performed at the city, state, or country scale. To simulate the evolution of air pollutants at the global scale at high spatial resolution requires powerful supercomputing resources such as those provided by the NASA Center for Climate Simulation (NCCS). NASA Goddard Space Flight Center's Global Modeling and Assimilation Office (GMAO) is using its Goddard Earth Observing System (GEOS) weather forecasting model to run global, high-resolution atmospheric composition simulations capable of providing air quality forecasts in near-real time. The GEOS Composition Forecast (GEOS-CF) model system's resolution is 25 kilometers (16 miles) from Earth's surface to the lower mesosphere, achieving 10 times higher resolution than conventional global simulations.

Results and Impact

Air quality can vary dramatically by both region and time of day. With a 5-day global air quality forecast, GEOS-CF provides high-quality information designed to support NASA flight campaigns, instrument teams, and satellite missions, as well as aid in decision making regarding human health and agriculture.

Recent and ongoing NASA field campaigns and instrument teams supported by GEOS-CF include the TRacking Aerosol Convection Interactions ExpeRiment – Air Quality (TRACER-AQ), the Arctic-Boreal Vulnerability Experiment (ABoVE), and the Asian Summer Monsoon Chemical and Climate Impact Project (ACCLIP). The GEOS-CF team provides custom visualizations for each field campaign, which has revolutionized campaign operators' ability to target scientifically significant events with observations from aircraft, satellites, ships, and other platforms.

Why HPC Matters

On the NCCS Discover supercomputer, GEOS-CF uses more than 3,500 compute processes across nodes containing 28 CPUs with 4.5 gigabytes of memory per CPU. The model takes advantage of PFIO (Parallel Fortran I/O server), which allows I/O offloading to dedicated nodes that are reserved for I/O operations rather than model execution. With I/O being performed asynchronously, each day GEOS-CF can produce 5-day forecasts within 9 hours.

The forecasting system produces approximately 1 terabyte of NetCDF data daily, which is then compressed to 450 gigabytes before being saved to online storage. Data visualizations are primarily provided via the NCCS-hosted FLUID system in the form of datagrams, surface maps, and total column plots. For our R&D efforts, we are integrating eViz, a Python Command Line Interface (CLI) visualization package that has a highly configurable interface.

What's Next

NASA's GEOS-CF team is participating with the Environment and Climate Change Canada (ECCC) agency in a multi-model evaluation of air quality forecasts for North America. In terms of workflow, GEOS-CF is moving towards a Python-based workflow system that will perform experiment setup and manage simulation execution, handle data post-processing, and enable the creation of plots and movies for analysis. Through the workflow effort, we have identified small, interchangeable tasks to aid GEOS-CF researchers in accomplishing both reanalysis and forecasts more effectively. We are also using both the Google Earth Engine and Amazon Web Services (AWS) for disseminating two-dimensional hourly-averaged air quality fields for local government applications. We are planning to move to a higher vertical resolution in the future — from 72 to 91 model levels — and to add chemical constituent data assimilation.

Improving Satellite-Based Snow Estimates with GPU-Accelerated Deep Learning

Goutam Konapala, NASA Goddard Space Flight Center
Sujay Kumar, NASA Goddard Space Flight Center

Carrie Vuyovich, NASA Goddard Space Flight Center

Overview

Snow conditions on the Earth's land surface – vital components of the global hydrological cycle – play a critical role in determining local and regional climate and seasonal water availability. Despite its well-recognized importance, the accurate characterization of snowpack states is challenging due to the highly variable nature of snow formation, accumulation, and melt processes, and there is a growing need for snow estimates that are both reliable and timely. To support this need, a team of researchers in the Biospheric Sciences and Hydrological Sciences Laboratories of the Earth Sciences Division at the NASA Goddard Space Flight Center are developing an improved snow dynamics system using a novel deep learning (DL) approach. This DL-based data system estimates daily snow depth and snow water equivalent for the continental United States using far fewer computing resources but with much greater accuracy.

Project Details

The core algorithm of the developed simulation system is a GPU-accelerated version of a DL approach known as Convolutional Neural Networks (CNN). This particular CNN uses passive microwave information from the Advanced Microwave Scanning Radiometer 2 (AMSR2) aboard the GCOM-W satellite as a key data source to train the model against *in-situ* snow measurements collected from various sources. The trained DL model then generates spatially continuous snow estimates across the continental U.S. at 10-kilometer resolution. The researchers are currently expanding this DL system to provide estimates for global usage.

Results and Impact

The GPU-accelerated CNN algorithm was able to provide snow depth simulations with a 60% improvement in performance compared to the currently operational snow depth estimates using only satellite data. Moreover, both the inference and training can be done using a single GPU for the entire continental U.S. compared to numerous nodes using a traditional physics-based snow model. Due to the DL model's ability to simulate snow depth far more accurately using less resources, it will be a valuable tool for understanding the role of the snow dynamics in the spatiotemporal changes in hydrological cycle.

Why HPC Matters

High-end computing (HEC) platforms at the NASA Center for Climate Simulation (NCCS) are critical for providing essential resources for preparing our modeling systems for emerging HEC systems and exploring new science developments. These DL-based snow simulations were made possible due to unrestricted access to the Tesla V100 GPUs on the NCCS Explore/ADAPT Science Cloud.

What's Next

Given the success of GPU-accelerated snow dynamics, the team at NASA Goddard is now preparing to expand these simulations to a global scale. The team is also exploring the possibility of applying DL to Global Precipitation Measurement (GPM) mission data and building a long-term climate data record by merging multiple satellite products.

The Evolution of GEOS as a Digital Twin for Earth System Science

William M. Putman, NASA Goddard Space Flight Center

Overview

The NASA Global Earth Observing System (GEOS) model and assimilation system is NASA's flagship system for enhancing the use of NASA's extensive Earth observations. GEOS enables forefront modeling studies that support the planning of future observation missions and advanced data assimilation studies, maximizing the impacts of NASA data for use in weather prediction, climate research, oceanic and atmospheric circulation, and other fields of study. NASA Goddard Space Flight Center's Global Modeling and Assimilation Office (GMAO) is using historical observations to produce digital models of the Earth system's climate, which has long been a goal at the center of GEOS development. With the vast expansion of machine learning capabilities and improved programming paradigms for graphics processing units (GPUs), GEOS is now poised to provide an experimental framework within NASA for the creation of true Earth system digital twins for weather and climate.

Project Details

The GEOS framework is ideal for Earth digital twin creation due to its flexibility: it supports both high temporal and 3D physical resolution as well as the exploration of increasingly complex Earth systems such as advanced carbon cycles, global aerosols, ocean dynamics, and biogeochemistry. As such, a digital twin for carbon studies represents a unique configuration of GEOS as compared to a digital twin for severe storms and tornados. The balance between complexity, resolution, and computing capability drives the application of GEOS to the production of Earth digital twins, which can play roles in improving the use of NASA observations or defining new Earth observing system missions.

Artificial intelligence (AI) is being integrated into GEOS chemistry components to enable improved and faster simulations of these complex Earth system processes. GMAO is exploring use of AI for replacing complex physics parameterizations for clouds, aerosols, and radiation processes. The entire GEOS code base is being refactored to introduce a domain-specific language to optimize the use of GPUs on computationally intense portions of GEOS, including the dynamics and transport of aerosols. These advances within GEOS demonstrate the continued evolution of GMAO's Earth digital twin capability as NASA explores new computing platforms and tackles the challenges of producing and analyzing petabytes of data for research and development.

Results and Impact

GMAO will highlight the advances in data assimilation capabilities and the creation of Earth digital twins with GEOS across a range of capabilities including: coupled ocean-atmosphere Earth system modeling, advanced studies of carbon emission and transport at ultra-high resolutions, and pushing the edge of compute capability in the creation of a 1.5-kilometer-resolution, global digital twin.

Why HPC Matters

The NASA Center for Climate Simulation (NCCS) and the NASA Advanced Supercomputing (NAS) Division provide essential high-end computing (HEC) resources for preparing GEOS modeling systems for emerging HEC systems and exploring new science developments.

What's Next

While current GEOS configurations require thousands of computational cores, configurations over the coming decade will require scalability approaching millions of cores. NASA's GMAO will continue to prepare its Earth system models for future HEC platforms by exposing new depths of parallelism within existing GEOS application components.

The Universe

Tracking the Fuel for Galaxy Growth

Cassi Lochhaas, Space Telescope Science Institute
Molly Peeples, Space Telescope Science Institute/Johns Hopkins University
Jason Tumlinson, Space Telescope Science Institute/Johns

Overview

Thanks in large part to NASA's Hubble Space Telescope, we now know that the parts of galaxies we can see – the glowing stars and gas – are actually surrounded by a giant reservoir of diffuse and nearly invisible gas that we call the circumgalactic medium (CGM). Galaxies form new stars out of gaseous material that is accreted from the CGM, so understanding the properties and origin of this material is fundamental to understanding galaxy evolution. However, the processes that affect accreting gas and how exactly it makes its journey through the CGM to the galaxy are open questions. The Hubble Space Telescope can only observe this gas if there happens to be a distant bright object behind it that can illuminate it like a lighthouse shining through fog. There are only a handful of bright lighthouses behind any given galaxy, so we turn to simulations to help us fill in the gaps.

Project Details

The FOGGIE project (Figuring Out Gas & Galaxies In Enzo) is a suite of galaxy evolution simulations. We start by simulating a large chunk of the Universe, then pick a handful of galaxies within this chunk on which to "zoom in." These galaxies are then re-simulated at a much higher resolution. In particular, FOGGIE specializes in resolving the diffuse gas surrounding galaxies, the circumgalactic medium. By forcing a strict resolution requirement not just within the galaxies themselves, but also in the space surrounding them, FOGGIE can achieve equisite resolution in the diffuse circumgalactic gas. This high resolution is necessary when it comes to understanding how galaxies get their gas: by resolving the small-scale interactions between inflowing gas, outflowing gas, and gas at a variety of temperatures and densities, we can test different theoretical models of galaxy evolution, compare to observations, and ultimately understand galaxy feeding and growth.

Results and Impact

The FOGGIE simulations reveal that most of the gas falling through the CGM onto the galaxy is predominantly cooler than the rest of the CGM gas that is not flowing inward. They also show that the accreting gas hasn't been affected by the heavy elements that are formed in stars, suggesting that it is nearly pristine gas that has never been located near another galaxy. Importantly, the high resolution of FOGGIE reveals interface regions, where the accreting gas slides by and interacts with the other, warmer gas in the CGM. These interface regions are subject to instabilities that help the cool and pristine accreting gas heat up and mix into the warmer gas surrounding it. Ultimately, the FOGGIE simulations are helping us understand how this fresh gas reaches the galaxy and what happens to it along the way through the CGM.

Why HPC Matters

The FOGGIE simulations would not be possible without significant computing power. Each of the six galaxies simulated in the FOGGIE suite is run on 512 Intel Xeon Haswell nodes on NASA Pleiades' supercomputer and takes 12-18 months of wallclock time, or several million node-hours, to simulate the galaxies through 13.7 billion years of cosmic time. Each node has 128 GB of memory, necessary to simulate the nearly 200 million resolution elements in the gas and stars within and around the simulated galaxies. Such a large number of resolution elements produces large output files as well, up to 15-30 GB per simulation output. FOGGIE also has exquisite time resolution, with a simulation output every 5 million years of cosmic time, leading to thousands of simulation outputs per galaxy and hundreds of terrabytes of data in total. Once the simulations are run, analyzing the outputs to produce visualizations and answer scientific questions takes additional computational resources, requiring hundreds of gigabytes of memory, fast I/O speeds, and several hundred thousand more node-hours of code running on Pleiades.

What's Next

The next generation of the FOGGIE simulations will include larger cosmic volumes and higher resolution in the regions outside of galaxies. This requires moving to the new Enzo-E code, which is scalable to thousands of processors and will enable a truly revolutionary number of resolution elements to be simulated at once. Until then, the FOGGIE team is continuing to analyze the current generation of FOGGIE simulations to understand how galaxies form and evolve.

Make Your Own Solar Storms!

Ronald Caplan, Predictive Science Inc.
Andres Reyes, Predictive Science Inc.
Tibor Török, Predictive Science Inc.
Cooper Downs, Predictive Science Inc.

Overview

Solar storms, including coronal mass ejections (CMEs), are large explosive events on the Sun that are capable of ejecting billions of tons of magnetized, million-degree plasma into space. When these CMEs reach Earth, they can disrupt radio transmissions, damage satellites, and severely impact power transmission grids, leading to extended large-scale power outages. Predicting such eruptions and their properties is therefore of great importance, and a major goal of NASA's Heliophysics Research Program.

Due to their inherent complexity, the mechanisms that control CME properties are poorly understood and challenging to model empirically. Therefore, data-constrained magnetohydrodynamic (MHD) simulations of CMEs are a promising path forward for improving our understanding and advancing space weather prediction capabilities.

To facilitate such simulations for community use, we have developed an interactive modeling framework called CORHEL-AMCG that allows non-expert users to routinely model multiple CMEs in a realistic coronal and solar-wind environment and propagate them to Earth.

The CORHEL-AMCG framework is designed to run on high performance computing platforms as well as in-house multi-GPU servers. This is a step towards future operational space-weather forecasting and rapid turn-around for research-focused CME simulations.

Project Details

The CORHEL-AMCG framework consists of two main components. The first component is a web-based interface called the Automated Multiple CME Generator (AMCG), which is used to design one or more pre-eruptive magnetic configurations, full MHD simulations that generate the background state of the solar atmosphere out to Earth, and high-fidelity CME simulations from Sun to Earth. All simulations are performed with the second framework component - a heavily updated version of the CORHEL modeling suite. CORHEL consists of a collection of tools and simulation codes (including the Magnetohydrodynamic Algorithm outside a Sphere (MAS) code) linked together for ease of use. It also includes auto-generated web-based visualization reports for each stage of the CME design and simulation.

Results and Impact

We have developed the CORHEL-AMCG framework and have used it to successfully model real solar storm events involving multiple eruptions. These event simulations where then used to model and study solar energetic particle events, which can pose a danger to Astronauts in interplanetary space.

Why HPC Matters

The direct simulations of solar storms in COREHL-AMCG require running moderate to large simulations with grids of up to 50 million cells and integrating time-dependent equations for millions of cycles. We therefore require the use of high performance parallel computations.

The run times and resources used for a typical sequence of simulations needed by CORHEL-AMCG using the dual Intel Skylake Xeon Gold 6148 nodes on NASA NAS's Electra supercomputer are:

- Twisted field design eruption: ~1.5 hours using ***8*** CPU nodes (320 CPU cores)
- Solar atmosphere background: ~4 hours using ***8*** CPU nodes (320 CPU cores)
- Solar storm eruption and propagation: ~20 hours using **32** CPU nodes (1,280 CPU cores)

The first two simulations above are small enough to be able to run on a single V100 GPU node on Pleiades, with run times:

- Twisted field design eruption: ~0.75 hours using ***1*** GPU node (4 GPUs)
- Solar atmosphere background: ~6 hours using ***1*** GPU node (4 GPUs)

What's Next

When the CORHEL-AMCG framework is complete, it will be delivered to NASA's Community Coordinated Modeling Center. There, it will be integrated into a public run-on-demand system, allowing researchers to routinely generate solar storm simulations from Sun to Earth. This routine modeling of real events will allow us to learn much more about them, leading to improvements in space weather predictive capabilities.

Simulating Sunquakes: The Impact of Flares on the Sun

John Stefan, New Jersey Institute of Technology

Overview

Solar flares occur when magnetic energy is released on the Sun in the form of heat and light. Most people are familiar with the effects of solar flares on Earth: extra static on the radio, poor GPS navigation, and stunning auroras if the conditions are right. An often overlooked effect of solar flares is their impact on the Sun itself. Some solar flares, particularly the strong ones, can excite the Sun-equivalent of an earthquake: a sunquake. Where earthquakes are generated by tectonic plates shifting inside the Earth, sunquakes are generated when fast-moving particles accelerated by solar flares hit the Sun's surface and transfer their energy in the form of pressure waves. Our work looks at how the Sun responds to the energy released by solar flares, and how this corresponds to what we see in observations. The results of this work give us a deeper understanding of how energy is transported during a solar flare.

Project Details

In our model, we consider how the pressure, density, and velocity of plasma changes as sunquakes travel through the Sun's interior. Because the density and pressure in the Sun increase with depth, a portion of the waves that travel downward is reflected back up to the surface. Some data about the flare, such as the location and quantity of energy released, can be recovered by examining how long these waves take to return to the surface, as well as how strong the reflected wave is. We use data from modeled flares as an input for our model, where we can adjust the total energy of the flare and the minimum energy of the particles accelerated by the flare. Each of these parameters strongly affects the types and amplitude of waves that are excited.

Results and Impact

Initially, we expected the flares that had more high-energy particles to generate the strongest sunquakes. Since these high-energy particles are moving so quickly, they're more likely to travel through the Sun's atmosphere without collisions and can deposit their energy into denser regions more effectively. However, we found the flares with more lower-energy particles generated the strongest sunquakes. These particles can't penetrate as deeply into the Sun, yet are able to accelerate the thin layers of the solar atmosphere to very high speeds—efficiently transferring their energy into sunquakes that travel throughout the Sun. Additionally, flares with lower-energy particles tended to generate strong acoustic-gravity waves (compared to the normal pressure waves), which haven't been measured in observed sunquakes. These findings indicate that a "sweet spot" exists in terms of the particle energies in flares that cause some flares to generate sunquakes while others do not.

Why HPC Matters

Compared to the size of the Sun, the flare energy is released in a relatively small area of a few hundred kilometers (km) and up to 1,000 km across. The sunquake that we observed travels tens to hundreds of thousands of km. Resolving these scales requires many grid points in the simulation, which in turn generate very large data files. The high performance computing (HPC) resources available at the NASA Advanced Supercomputing (NAS) facility are key to working with these large files, which require more memory to open than several laptops combined. The Lou data storage system provides a quick and reliable means of storing and accessing the generated data. Additionally, our model requires several hundred processor-cores of the Pleiades supercomputer to perform our simulations in a reasonable timeframe.

What's Next

To date, we have only considered one type of particle—protons—and looked at how effective certain energies are at generating sunquakes. We know that electrons are also accelerated by solar flares and deposit their energy differently than protons. In future work, we will examine the effectiveness of electrons in generating sunquakes and whether observations can be explained by a combination of the two.

A Quantum-Mechanics Driven Computational Framework for Cosmochemistry and Materials Modeling

Venkateswara Rao Manga, The University of Arizona
Kevin D. Grossman, NASA, Kennedy Space Center
Thomas J. Zega, The University of Arizona

Overview

Computational paradigm that simulates materials starting from quantum-mechanical descriptions of atoms has been playing a crucial role in different NASA missions ranging from cosmochemistry to novel alloy design for oxygen production on moon. Crucial to this *ab initio* framework of materials computation and modeling, capturing the physics across different length scales, is supercomputing resources such as NASA/NAS. With such powerful resources, computational materials research is leading to atomic-scale insights into various thermophysical and thermochemical behaviors, in addition to the improvement of the existing models of materials behaviour. One of the achievements of this computational materials science is that it has been driving technological advances at an accelerated speed. Pertaining to the field of cosmochemistry, it has been our goal to describe thermal and chemcial landscape of our early solar nebula from exprimental investigation of meteorites in conjunction with the materials modeling. In the case of alloy develoment, NASA's goal to produce oxygen on moon from lunar soil via an electrolysis route demands ultra-refractory materials that can work in the extreme temperatures and oxidation/reducing conditions. In this pursuit, NASA/NAS supercomputing resources, have enabled us; 1. to a nebular thermochronometer, 2. to predict condensation phase daigrams and 3. to conduct rapid screening of ultra-refractory electrodes.

Project Details

A NASA emerging world grant aims to deduce the thermodynamic landscape of the high temperature region of our early solar nebula. Calcium-aluminum-rich inclusions (CAIs) that are found within meteorites, are known to be the first solids to have condensed in the high-temperature region of the protoplanetary disk and offer clues into the early landscape. These mineral assemblages with a myriad of microstructures, crystal chemistries and defect structures when analyzed in conjunction with the thermodynamic models, enable us to regenerate the early nebula, especially, in the high temperature region of the disk. In another NASA project which aims to extract oxygen out of lunar soil, molten regolith electrolysis (MRE) runs at temperature as high as 1873 K and availability ultra-refractory electrode materials have been a limiting aspect in realization of the technology. Motivated by these research interest, we have undertaken thermodynamic and kinetic modeling of materials within the *ab initio* computational framework.

Results and Impact

The first principles driven thermodynamic modeling of CAI phases has led to predictions of the condensation sequences of refractory and ultra-refractory minerals within the high-temperature region of the early solar nebula. The computations of crystal chemistries in phases such as spinel ($MgAl_2O_4$) have led to the development of a nebular thermochronometer to determine the temperatures of certain physical processes in the protoplanetary disk. Similarly, an on-going modeling of crystal chemistry of another CAI phase (hibonite) is enabling development of an oxybarometer, which when applied to the CAI assemblages help us to gauge the thermal and chemical conditions that prevailed during their formation. Our recently published condensation phase diagrams have also revealed the deviations of the mixing gas-and-dust system from equilibrium condensation. One of the most important impacts of this work is to show the interplay between thermochemistry and the gas-and-dust dynamics in the nebula.
Pertaining to MRE electrode development, the materials modeling effort has predicted superior alloy compositions which can work at the targeted temperatures in the molten regolith. The computational research also progressing to help optimize the electrolysis process by predicting an efficient melt pretreatment step wherein the regolith is compositionally tamed to produce oxygen in the targeted temperature-window of operation.

Why HPC Matters

The quantum mehcanical based first-principles calculations which are crucial to our thermodynamic modeling of meteoretic and UR material systems are made possible by the NASA/NAS supercomputing power. A density functional theory (DFT) based code, Vienna ab inito simulation package (VASP) is employed to this end. VASP simulations demand 100 – 600 CPUs/job with a run duration that can go upto 100 hours. Calculations include different properties and materials systems and allocations in the range of 50 – 100 millions cpu-hrs, at least, are annually requested. The material systems include oxides, phosphide, platinum group alloys, transitions metal alloys etc. For example, phonon calculations of Allendeite, $Sc_4Zr_3O_{12}$, which is named after the Allende meteorite, were run with 600 cpus per job with a walltime of ~60 hrs. The phonon calculations of platinum group alloys are also equally intensive.

What’s Next

Recent investigations of primitive solids in meteorites have revealed many ultra-refractory (UR) phases such as Sc- and Zr-bearing $Sc_4Zr_3O_{12}$. Measuring the high-temperature thermochemistry of these mineral solutions in a huge composition space of interest is a difficult task. An additional challenge in modeling the microstructure evolution of these UR phases is to incorporate the effects of kinetic processes with underlying solid-state diffusion. To this end, we have proposed to employ the first-principles quantum-mechanics driven computational framework to predict the thermochemistry and atomic transport in the entire temperature and composition space of all the UR phases pertaining to meteorites and MRE technology.

New Findings in Solar Magnetic Cycles and Waves in Stellar Interiors

Kyle Augustson, Northwestern University
Daniel Lecoanet, Northwestern University

Overview

The Sun and other stars possess striking features at their surface like magnetic starspots and in their interiors such as differential rotation. What remains a mystery is how these magnetic fields arise and what internal processes drive their variability. Is it instabilities arising from a near-surface shear, the convection itself, or other processes that lead to their creation and destruction during a magnetic cycle? And, how do these magnetic fields influence the Sun’s atmosphere and Earth’s delicate magnetic shield? As part of the effort to understand the Sun, it is informative to examine what magnetic structures are stable inside stars. And, does evidence for those deep interior magnetic fields appear in the surface wave patterns of other stars? Can other important processes such as chemical mixing be assessed through changes in the wave patterns?

Project Details

This project aims to answer these questions through the study of magnetic dynamos, convection, and mixing processes, including how they interact to produce waves that can be detected through asteroseismology -- the study of surface waves that can be used to probe a star's interior like seismology can probe the Earth's interior.

Results and Impact

Using NASA Advanced Supercomputing (NAS) resources, we found that certain magnetic field configurations that were previously believed to be stable are unstable, meaning that magnetic equilibria may be more difficult to obtain. In finding a new description of convective penetration for models of stellar evolution, we simulated its influence on the transition between convectively stable and unstable regions in both rotating and non-rotating systems. We ascertained that some solar and stellar dynamos are linked to the magneto-rotational instability (MRI) if there is sufficient radial differential rotation, showing an alternative path for stellar dynamo theory. Recent 3D simulations capture both the convection, the wave excitation region, and the wave propagation to estimate their surface amplitudes, allowing us to contact stellar observations.

Why HPC Matters

NASA high-end computing (HEC) resources enable our work with the flexible and spectral framework that is the PDE Dedalus code much of which requires large-scale computation. We solve several classes of problems with Dedalus on these resources: eigenvalue, initial value, and boundary value problems. Each class uses different HEC resources. Eigenvalue problems require substantial amounts of memory, few cores per node, and solve a given problem over a large set of parameters in parallel. Initial value problems are usually nonlinear, requiring many cores, little memory, and many iterations to complete a run at a single set of parameters. Boundary value problems, like eigenvalue problems, are parallelized only over the parameter set. Recently, we have added support for PETSc and SLEPc allowing us to add a layer of parallelism the BVP and EVP problems, over both parameters and for each parameter. We worked with experts at the NAS facility at Ames Research Center to optimize parts of the Dedalus code on NAS hardware resources used for these the simulations and to build software stacks that support many users.

What's Next

We aim to better understand the nonlinear MRI in the solar near-surface shear layer, provide new magnetic equilibrium solutions for rotating stars, and estimate surface wave amplitudes to guide the interpretation of stellar observations.

Modeling Atmospheres with ROCKE-3D: Solar and Extrasolar Worlds

Michael Way, NASA Goddard Institute for Space Studies

Overview

Understanding the climates of Venus, Earth, and Mars through time is critical to understanding the conditions under which life may begin and flourish. Earth has been habitable for nearly 4 billion years, yet our neighboring worlds Venus and Mars are cautionary tales, where no signs of present or previous life have been found thus far. Our planetary research group at the NASA Goddard Institute for Space Studies (GISS) develops and utilizes a 3D General Circulation Model (GCM) known as ROCKE-3D to model the atmospheres and/or oceans of terrestrial worlds. These simulations assist NASA in targeting neighboring and extrasolar worlds with specific missions (e.g., Mars Curiosity rover, James Webb Space Telescope) and can act as benchmarks to constrain our understanding of newly discovered exoplanets in the habitable zone.

Project Details

Recent ROCKE-3D climate simulations include a) an exoplanetary world around Proxima Centauri, b) a temperate ancient Venus, c) a deep future-time Earth with two different supercontinent configurations, and d) an ancient hypothetical ancient lunar atmosphere.

We use the best available data to set up our climate simulations. For Proxima Centauri b (a super Earth exoplanet that orbits an M-type star) we know nothing more than its orbital characteristics, estimates of mass/density, and how much sunlight it receives, so we model a simple water-covered "aquaplanet." For ancient Venus we use present-day topography and supply it with five different water reservoirs that fit within the error bars of past water estimates. For future Earth, we use sophisticated geodynamics models to construct two alternative supercontinent outcomes 200–250 million years into the future alongside the expected increase in solar luminosity. For the Moon, we use estimates of volcanic water outgassing to model where such water would be deposited long-term – at the poles in regions devoid of sunlight.

Results and Impact

Proxima Centauri b: Using the first fully coupled atmosphere+ocean model of this world, we discovered that the transport of heat by the ocean is a key factor in its likely habitability. Compared to a simulation with a shallow simple ocean, our simulation with a fully coupled dynamic ocean produced a global mean surface temperature 16°C warmer!

Ancient Venus: We were the first to use a forward modeling approach to piece together an entire climate history for Venus that fits within observational constraints. We demonstrated that Venus could have had temperate conditions for much of its history and may have been the first habitable world of this solar system.

Earth's Future Supercontinents: The Aurica low-latitude supercontinent has a global mean surface temperature 7°C warmer than today. This is expected given its lack of high-latitude continents like on present-day Earth and its relative 2.5% increase in solar luminosity. The Amasia supercontinent (with the main continents pushed north of the equator while Antarctica remains in place) is 3–7°C warmer. The +3°C result is surprising but shows that the mean altitudes of mountain ranges play a key role in ice sheet coverage driving the temperature lower.

Ancient Lunar Atmosphere: Ours was the first time a 3D GCM was applied to study the effects of volcanic outgassing at lower latitudes from the maria plains and transport of water to the permanently shadowed regions of the poles. We will soon be able to put constraints on the amount of water delivered, which may be measured by Artemis program astronauts in the near future.

Why HPC Matters

The NASA Center for Climate Simulation (NCCS) is vital to the success of our projects. The constancy of CPU type and increasing speed and storage enable our relatively small development group to put resources into science development and less into reconfiguring code for the "latest and greatest" computing architecture. The four projects detailed used on the order of 7 million core-hours on the NCCS Discover supercomputer. Compiler support by the highly capable NCCS staff is also important to our success. In addition, we rely on the NCCS DataPortal to distribute our results simply and quickly to the community.

What's Next

The ROCKE-3D modeling group will continue to apply their expertise to constraining the habitability of new exoplanetary worlds being discovered by NASA's Transiting Exoplanet Survey Satellite and Webb missions, more accurately modeling the long-term climate history of Venus guided by results from new missions sent by NASA (DAVINCI and VERITAS) and the European Space Agency (EnVision), and helping to discern the origin of water in the Moon’s permanently shadowed polar regions in conjunction with expected revolutionary discoveries from NASA’s Artemis program.

Human Exploration

Multiphase Simulations of the Launch Environment at NASA’s Kennedy Space Center

Jordan Angel, NASA Ames Research Center
Michael Barad, NASA Ames Research Center
Scott Neuhoff, NASA Ames Research Center
Cetin Kiris, NASA Ames Research Center

Overview

To support NASA’s Artemis missions launching aboard the new Space Launch System (SLS), as well as launches by the agency’s commercial partners, the Launch Ascent and Vehicle Aerodynamics (LAVA) team in the NASA Advanced Supercomputing (NAS) Division are pushing the state of the art for numerically predicting ignition overpressure (IOP) waves in launch environments. These waves are formed during the ignition of rocket engines and pose a safety hazard by creating unwanted vibrations that may damage the vehicle or support structures. NASA launch complexes have long used water injection in the flame trench and at exhaust openings to absorb some of the powerful acoustic wave energy. Despite the widespread use of water-based sound suppression systems, engineers are still working to gain a full understanding of their effectiveness, due to the extreme complexity of the launch environment. Vital information about the IOP intensity is needed to ensure mission safety, yet achieving accurate numerical prediction of IOP intensity has been impeded by a lack of capable analysis tools.

Project Details

To address this problem, the LAVA team developed a new computational approach for simulating multiphase (gas-liquid) flows that can robustly solve the extreme flow conditions in the launch environment and still accurately capture the acoustic IOP wave propagation. In collaboration with the Launch Services Program (LSP) at NASA’s Kennedy Space Center (KSC), the LAVA team used this new capability to simulate the Space Test Program (STP)-2 mission launch, which flew on SpaceX's Falcon Heavy rocket from KSC’s Launch Complex (LC) 39A. The simulation showed excellent agreement with flight measurements and was a significant breakthrough for computational IOP prediction. The LAVA team is applying this multiphase capability to carry out challenging realistic launch simulations, in order to study acoustic wave propagation from the complex interactions of rocket engine plumes with the flame trench and water-based sound suppression systems.

Simulations of the launch environment provide rich details of the complex interactions between supersonic rocket plumes and liquid water. The LAVA multiphase solver has now been applied to other launch configurations, including the SLS Scale Model Acoustic Test (SMAT) and the full-scale SLS launch environment. As more cases are run and comparisons are made to measured data, the LAVA team continues to modify and improve their best practices to achieve the best predictions possible.

Results and Impact

The impact of this new capability is also being felt in the Artemis missions. To help ensure the safety of the upcoming Artemis flights, NASA Exploration Ground Systems (EGS) engineers needed estimates of the acoustic loads on the Mobile Launcher (ML) surface during the SLS ignition. The LAVA team created full-scale simulations for the SLS at LC-39B, including the sound suppression system—providing EGS engineers with SLS acoustic IOP predictions, along with the predicted pressure field on support structures.

The impacts of this work are twofold: first, the new capability represents an achievement in fundamental research in applied mathematics; and second, these methods were able to be quickly adopted into the LAVA framework, which opened the possibility of applying these new methods to practical engineering problems. With the increasing number of NASA and commercial launch missions, the availability of a well-validated and robust multiphase CFD tool will serve the agency now and in the future.

Why HPC Matters

Each of the LAVA team's high-resolution simulations of the launch environment at LC-39B used approximately 400–500 million grid cells and ran several weeks on 8,000 cores of either the Electra supercomputer's Skylake processors or Aitken's Cascade Lake processors. Because these simulations are time dependent, they generate roughly 400 terabytes of data for each run. Visualization of these results provides insight into the complicated flow conditions of the launch environment, and required reliable data storage, high-performance data I/O, and the powerful visualization tools provided by the NAS Division's Visualization and Data Analysis team—who always strive to achieve the greatest possible utilization of complex datasets generated by NAS computational resources.

What's Next

With the Artemis I flight, new flight data will be available to further validate the LAVA solver's predictive capability—serving as an effective ''blind'' validation case for the code. Furthermore, the code is currently being modified to run on GPU accelerators to enhance turnaround time and leverage new compute architectures. The goal is to provide crucial CFD information for designs in a reasonable turnaround time to accommodate launch pad modifications, changes to the water-based sound suppression system, and new vehicle designs.

Analyzing Crater Formation During Descent on a Planetary Body with a Soil Surface

Josette Bellan, NASA Jet Propulsion Laboratory
Kaushik Balakrishnan, NASA Jet Propulsion Laboratory

Overview

Controlled landing of a spacecraft on a soil-surface planetary body, such as on the Moon, Mars, Titan, asteroids, etc., is achieved by using thrust from positioning rockets on the vehicle to direct the lander to the desired location on the surface. The supersonic jets from the rockets invariably interact with the planetary soil, forming craters and ejecting bursts of a mixture of gas and particles that could impact landing and crew safety. During NASA's 1969 Apollo 12 mission to the Moon, lunar soil particles, called "regolith," were blown from the surface at speeds up to 70 meters per second as the spacecraft descended. Regolith particles traveling at such high velocities can damage the landing vehicle, particularly if they strike vulnerable locations, so understanding the morphology of the crater formation from supersonic jets is critical for future space mission design.

The cratering and whirlwind of particles caused by the positioning rocket's plume can also cause a problem in interpretation of radar altimetry, which is used to detect the planetary surface and ensure that the descending lander does not crash. Regolith particles displaced by the rocket jets could strike the radar sensors and affect the accuracy of radar measurements, obscure cameras positioned on landers, or damage the coating of the lander antenna, thereby changing the antenna's thermal properties and impeding transmission of information from the lander.

Project Details

To support future missions to the Moon and other planetary bodies, researchers at NASA's Jet Propulsion Laboratory in Pasadena, California are using modeling and simulation tools on agency supercomputers to better understand the interaction between a supersonic plume and the soil it disturbs. Their three-dimensional and time dependent model consists of the conservation equations for mass, momentum, species, and energy for both a jet plume and the interacting soil medium. These equations include phase interaction terms—which account for a variety of dependances and variables involved in the interactions—for each set of equations. The simulations will be used to explain the intricate details of jet-induced cratering to not only augment our physical understanding of the problem, but also help the design of future planetary landing missions.

Results and Impact

The researchers obtained results from several simulations at different initial ratios of jet-to-ambient gas densities. Comparisons of results with photos from the Mars Curiosity rover showed that the model is able to accurately capture the shape of craters during their formation. The team also studied cratering on Titan, Saturn's largest moon, which is known to have a dense atmosphere.

Their results showed that when the jet-to-ambient gas density ratio is approximately 1, craters have a conical cross-section with a relatively small diameter, are relatively deep, have ripples on their internal wall, and exhibit larger bursts of particles expelled from the crater's rim. In contrast, craters formed when the jet-to-ambient gas density ratio is larger than 1 have parabolic cross-sections, relatively large diameters, are shallow, have two concentric craters with the outside crater being shallower that the central crater, and the particle bursts are subdued but have higher velocities, which means that they will travel larger distances.

At the same jet-to-ambient gas density ratio, the crater depth is determined by the jet density: the larger the jet density, the deeper the crater. Examination of the soil material forced out of the crater showed that the particles accumulated outside the crater rim had generally two distinct behaviors according to the regions of the crater from which they originated. Most of the crater-depth-produced bursts of soil (a mixture of particles and gas) had particles with a volumetric fraction (i.e., volume of particles in the gas/particle mixture) similar to that of the undisturbed soil and were uniformly radially distributed, whereas the ejected soil from closer to the surface had a much smaller particle volume fraction and displayed an oscillatory pattern of radial distribution from the crater rim.

Why HPC Matters

NASA's high-performance computing resources have been crucial in conducting the high-fidelity numerical simulations required for this project. The computational grid consists of more than 150 million cells, with each cell solving 13 partial differential equations at each time step. The researchers' massively-parallel code typically used 1,024 – 2,048 processors of the Pleiades supercomputer at the NASA Advanced Supercomputing facility over approximately six weeks for each run, and are the only jet-induced cratering simulations performed at this scale in terms of compute and memory requirements to date.

What's Next

The current model may be used to simulate upcoming NASA missions, such as the Mars Sample Return landing scenario for scheduled for 2028. These efforts also provide the springboard to develop a model for predicting cratering on the Moon or other planetary bodies where the jet-to-atmosphere density ratio is extremely small. The new model is under development, with the first step aiming to predict the evolution of a supersonic plume in an extremely low density gas. This model will help with mission design for future lunar landing missions, such as the Artemis III mission planned for launch in 2024.

Building Aerodynamic Databases for NASA's Space Launch System

Jamie Meeroff, NASA Ames Research Center

Overview

NASA's Artemis program is the agency's mission to return to the Moon by 2025 with the first woman and the first person of color, with the goal to eventually reach Mars. Artemis uses NASA's Space Launch System (SLS), one of the most powerful rockets ever constructed, to send astronauts and cargo to lunar orbit. It is critical to understand the aerodynamic forces on the SLS and Orion crew vehicle during launch and ascent through Earth's atmosphere, so a team of computational fluid dynamics (CFD) researchers at NASA's Ames Research Center are developing detailed aerodynamic databases to ensure that the payload survives the ascent through the Earth's atmosphere. These CFD databases simulate a plethora of flight conditions and complex interactions of the various vehicle components that are difficult to model in an experimental setting. High- performance computing (HPC) resources are critical in providing the necessary computational power to generate enough accurate CFD simulations to produce the quality of database needed to assure mission safety.

Project Details

The team used CFD simulations with various NASA-developed solvers, including OVERFLOW, FUN3D, and Cart3D, to construct aerodynamic databases for both the ascent and booster separation—when the two solid rocket boosters separate from the main vehicle—portions of an SLS launch. Tens of thousands of individual CFD simulations have been run supporting various SLS configurations and different Artemis missions. Both the ascent and separation phases of flight require researchers to model complex flow physics, including the interaction of multiple rocket plumes.

While computationally-derived databases were used in the later days of the Space Shuttle program, CFD models have greatly improved as the power and efficiency of NASA HPC as grown, adding increased complexity, finer details down to the smallest protuberances, increased volume of data, and faster turnaround times. Furthermore, modern software practices, including revision control and in-house case management software (programmed in Python), have been paramount for effective data control and the minimization of human error.

Results and Impact

The databases developed by the CFD team at Ames are used by various groups in the SLS program, including NASA's Guidance, Navigation, and Control team and agency partners like Boeing, Northrop Grumman, and the United Launch Alliance. They have been used for the development of large-scale items such as building ascent trajectories and verifying the structural integrity for vehicle protuberances during ascent. Ultimately, these databases are critical in ensuring the success of all Artemis missions and any future missions using the SLS.

Why HPC Matters

All of the CFD simulations needed to build these aerodynamic databases were run on the Pleiades, Electra, and Aitken supercomputers at the NASA Advanced Supercomputing facility over the duration of the SLS program. Millions of processor hours were needed to generate enough data to provide adequate coverage for each of the simulated phases of flight. It would be impossible to reproduce the data generated for these databases without significant HPC resources. Reduced-order methods would require numerous simplifying assumptions that would increase the uncertainty of the data provided and put the mission—and astronauts—at risk.

What's Next

As the Space Launch System transitions from a developmental program to an active one, flight data will be used to verify computational methods and results, and to identify potential improvements to the databases themselves. SLS development continues as new vehicle configurations are introduced and new missions are planned. The ever-expanding NASA HPC resources will allow for new and improved databases to be generated using lessons learned from earlier SLS flights. Ultimately these newer databases will continue to increase future SLS mission safety.

Simulating Supersonic Parachute Inflation for Mars Landers

Michael Barad, NASA Ames Research Center

Overview

NASA's Perseverance rover completed its journey from Earth to Mars using an Entry, Descent, and Landing (EDL) profile similar to many of its predecessors, including the Mars Science Laboratory (MSL). This process is infamously called the "seven minutes of terror," because hundreds of critical events need to happen successfully in a very precise sequence—without intervention from Earth—due to the 11-minute signal lag. Roughly four minutes into descent, the capsule deploys a parachute to slow it down until it's close enough to the surface (about 1.3 miles) to begin its powered descent. The parachute must inflate as evenly as possible despite the turbulent wake of the capsule, and without any rips or tears to the incredibly thin woven fabric—at about 80 microns, not quite the width of a human hair. This is one of the most dangerous portions of EDL and is notoriously challenging to predict. The planning time and cost of flight tests to certify any changes to the current process are pacing items for the development of next-generation parachute systems. To accelerate this process, and reduce costs and risk, NASA's Launch, Ascent, and Vehicle Aerodynamics (LAVA) team—funded by the agency's Entry Systems Modeling Project—is developing the capability to simulate supersonic parachute inflation.

Project Details

Supersonic parachute inflation is challenging to predict accurately because it requires modeling the two-way interaction (known as fluid-structure interaction, FSI), between the very flexible parachute fabric (the structure) and the air (the fluid) as it rushes by at supersonic speed. The LAVA team is developing the capability to perform such challenging FSI simulations by leveraging its Cartesian adaptive mesh refinement computational fluid dynamics (CFD) solver and coupling it to its recently developed finite element structural dynamics solver.

Results and Impact

As a first step to gain confidence in the new tool's ability to accurately predict the complex process of supersonic parachute inflation, the LAVA team simulated a portion of the Advanced Supersonic Parachute Inflation Research Experiments (ASPIRE) SR01 flight test, where a version of the MSL's parachute was inflated at supersonic conditions in Earth's upper atmosphere. The simulations predicted the peak opening load to within 10% of that measured during the flight test, but showed discrepancies in the inflation rate and the dynamics immediately following the peak, under-predicting the partial collapse and rebound of the parachute. The team continues to investigate potential sources of error, such as the effect of the rapid deceleration of the vehicle when the parachute inflates, to improve the tool and associated best practices.

With further validation and verification, the ability to simulate supersonic parachute inflation can be used in the design process and pre-flight qualification of new parachute systems. It can also be used to reduce risk by simulating many different scenarios that would be too costly to do in flight tests—for example, side winds of varying strengths. Moreover, certain fine-scale details that are too difficult to record during flight tests can be extracted from simulations, such as the timing, location, and magnitude of the maximum fabric stress—information that can be used to determine when and where a parachute may fail during a real-world mission.

Why HPC Matters

Simulating supersonic parachute inflation is more computationally demanding than typical CFD or structural dynamics simulations, as it requires the two different methodologies and disparate parallelization strategies to work in unison. Despite these challenges, the LAVA team was able to very quickly optimize the FSI tool to solve for hundreds of millions of degrees-of-freedom in the fluid and structural domains. A typical simulation uses 2,560 Cascade Lake cores on NASA's Aitken supercomputer and runs for just over a day, corresponding to about 0.5 seconds of time simulated. It generates dozens of terabytes of data, which can be visualized and analyzed in parallel. Hundreds of such simulations must be completed to develop best practices and thoroughly validate this new capability for various parachutes operating at different flow conditions.

What's Next

The LAVA team will continue to optimize the FSI tool to reduce the turnaround time of supersonic parachute inflation simulations by improving its parallel scalability, and its solution-based AMR algorithms. The team is also working to include more physical phenomena in the simulations to better represent real-world parachute inflation by considering the effect of the deceleration of the capsule and canopy, initial parachute shape and flow conditions, and more representative parachute structural properties.

Multiphase Simulations of the SLS Launch Environment

Travis Rivord, NASA Marshall Space Flight Center

Brandon Williams, NASA Marshall Space Flight Center

Overview

NASA's Space Launch System (SLS), which will send astronauts back to the Moon in the next few years with the Artemis Program, is powered by four RS-25 engines and two solid rocket boosters (SRBs). During launch, the SLS propulsion system generates intense acoustics and other powerful waves, such as ignition overpressure (IOP) which, if unmitigated, have the potential to damage the vehicle and possibly cause loss of mission or crew. To protect the vehicle from these powerful waves, the SLS launch pad design includes an ignition overpressure/sound suppression (IOP/SS) system that sprays 270,000 gallons per minute of water very close to the SRB and RS-25 nozzles. The SRB and RS-25 engine plumes, and the proximity of the IOP/SS water, create a complex multiphase (gas and liquid) environment during the ignition sequence. The interplay among these systems creates challenges related to water spray into/onto engine nozzles, potential debris transport, and additional transient loads due to strong plume-water interactions—all of which the SLS vehicle must be able to withstand.

Prior to the first uncrewed Artemis launch this year, the SLS multiphase liftoff environment was largely unknown due to differences from the Space Shuttle and other programs. Some data was available from tests of individual systems, but no integrated testing or analysis was available. Even post-launch analysis of Artemis I cannot provide a full understanding of the complex physics involved due to limited (or obstructed) camera views and instrumentation. Computational fluid dynamics (CFD) is being used to investigate the details of the multiphase environment that could not be measured, help interpret the data gathered from the launch, and ultimately identify phenomena that are a concern for future flights.

Project Details

Engineers at NASA's Marshall Space Flight Center performed simulations of the SLS ignition sequence using the Loci/STREAM-Volume of Fluid (VoF) multiphase CFD solver. Prior efforts successfully validated the CFD solver against various tests, giving analysts confidence in their ability to simulate the SLS multiphase liftoff environment prior to the Artemis I launch. The CFD simulation of the SLS ignition sequence was conducted in three phases. First, the IOP/SS water system was simulated for approximately six seconds to reach a quasi-steady state. Next, the RS-25 engine plumes were activated and held at full power for one second. Finally, the SRB booster was activated and the simulation was carried out until the time just prior to vehicle motion. This simulation process mimics the conditions that exist at launch.

Results and Impact

The SLS ignition sequence simulation results provide deep understanding of the underlying physics that occur during launch. Observations from the simulation include reduction of water splashing into/onto the engine nozzles, change in angling of the dense water sheets, and the origin of the powerful ignition overpressure (IOP) wave. These observations directly inform the SLS program on subjects including plume-water-induced side loads, debris transport, and the launch acoustics environment. Additionally, with post-launch comparison of CFD observations to flight data, these tools can be applied to launch vehicles and environments other than SLS with confidence.

Why HPC Matters

The SLS ignition sequence CFD simulations are conducted on meshes of up to hundreds of millions of cells, on thousands of processors, for weeks at a time. These simulations generate terabytes of data that must also be stored and archived for future use on HPC systems. Simply put, the CFD simulations would not be possible without NASA's HPC resources.

What's Next

Continued comparisons between the Artemis I flight data and the CFD simulations will both improve confidence in the CFD results and provide deeper understanding into the SLS multiphase launch environment. This will be used to provide insight for decision making for the first crewed SLS flight, Artemis II. Future simulations will target new configurations of the SLS launch environment, including the

IOP/SS water required to support the more powerful variants of the SLS vehicle, such as Block 1B crewed and cargo configurations. Additionally, this capability provides NASA the ability to investigate launch environments for vehicles other than the SLS to support other agency space missions.

Rocket Plume-Surface Interaction Simulations for Moon and Mars Landings

Peter Liever, NASA Marshall Space Flight Center
Jeff West, NASA Marshall Space Flight Center

Overview

Plume-surface interaction (PSI) between a lander's engine plumes and regolith–the fine soil found on solid planetary bodies–creates many hazards during landing: obscuration, contamination by dust particle clouds, high-energy ejecta streams, and landing area cratering. The Fluid Dynamics Branch at NASA's Marshall Space Flight Center (MSFC) has developed simulation tools to provide predictive PSI environments to teams across the agency, including those working on the Human Lander System, Commercial Lunar Payload Services, and future crewed Mars landers. Development required the extension of proven computational fluid dynamics (CFD) tools for plume flow applications in low pressure atmosphere as found on Mars and vacuum surroundings found on the Moon. The team at Marshall performed gas-particle multi-phase interaction modeling of surface scouring and cratering with the Gas-Granular Flow Solver (Loci/GGFS), which features a Eulerian-Eulerian modeling approach–treating both the gas and granular material efficiently as continuum phases–to directly predict erosion, crater formation, and particle streams. Alternative modeling approaches that represent the soil as discrete Lagrangian particles and track hundreds of billions of particles would result in prohibitive computational cost for realistic landing site scales.

Project Details

The team validated the Loci/CHEM CFD application for low-pressure and near-vacuum plume impingement against NASA experiments at MSFC and Langley Research Center. For lunar mixed continuum-rarefied flow simulations, the team used the Loci/CHEM-Boltzmann extension, which switches cell-by-cell to a rarefied solver only where needed. The experiments indicate that rarefaction¬–the lessening of density to levels where gas cannot be represented as a continuum fluid anymore–has a significant amplification effect on surface shear stresses where the rocket plume and regolith meet, the dominant influence in PSI viscous erosion and scouring.

To better understand the importance of jagged lunar regolith particle mixture effects, the team at the Fluid Dynamics Branch created novel multi-phase model implementation within GGFS, which enabled them to model the roughness and interlocking of irregular particle shapes and dense packing from a wide particle size distribution–features that resist penetration of gas flow into the particle bed and reduce erosion and fluidization. To generate closure model databases for drag, shear, kinetic energy dissipation, and other properties needed in the Eulerian particle phase algorithms, they preprocessed particle-particle and gas-particle interaction simulations using the Discrete Element Method (DEM). The CFD-based predictive tools were then validated against the results from NASA experiments of plume impingement on surrogate regolith material surfaces in low pressure and near-vacuum test chambers.

Results and Impact

These simulations provided insight into the plume structure dynamics and the significance of rarefied flow effects in the surface erosion mechanisms. The ability to perform full-scale lander PSI simulations with GGFS in lunar conditions was demonstrated for a model of the Apollo Lunar Module–the only crewed vehicle to land anywhere outside of Earth to date–at an altitude of five meters above the regolith surface. The Lunar Module simulations were also used to progressively test and improve modeling of important effects of irregular particle shape and multi-size mixture compositions on erosion levels. Application testing and validation of GGFS is now underway using sub-scale PSI experiments recently performed in vacuum chambers at MSFC using a variety of materials to simulate regolith from spheres, mono-disperse sand, sieved mono-disperse lunar simulant jagged particles, to full size range BP-1 lunar simulant. The experiments were designed to validate and guide how the particle phase modeling approach would need to be refined.

Why HPC Matters

Simulation of gas-particle surface erosion is an inherently time-dependent process requiring time-accurate simulations with long run times and large datasets. The simulations performed by the team used over a hundred million cells, and required 4,000 to 8,000 processors on the agency's Aitken, Pleiades, and Electra systems across two to four weeks of runtime. The NASA Advanced Supercomputing (NAS) Division at Ames Research Center provided efficient and reliable high-performance computing (HPC) resources, including storing and archiving terabytes of data for analysis and visualization postprocessing.

What's Next

GGFS maturation towards production readiness continues with improvements in two-phase flow algorithm implementation, validation of the particle physics models, and simulation optimization on agency HPC assets. Model validation will be performed against the vacuum chamber PSI experimental data generated with a variety of particle mixture compositions up to full range lunar simulants. The team's current goal is to improve turnaround time of simulations to rapidly support customer projects, since current PSI 3D simulation run times are four weeks or more. They have identified the presence of numerous particle physics submodels and computational cost disparity between rarefied and continuum algorithms that cause simulation latency and processor load imbalances, contributing to long run times. Besides algorithm and simulation process optimization, porting submodels to graphics processing units (GPUs) is an attractive option that would allow the CPUs to perform the mainline solver work more efficiently without wait for sub-models. NAS's CPU/GPU offerings will be a key enabler for the PSI development team in pursuing this approach.

Supercomputing

HECC: Evolving to Meet Tomorrow's Requirements

William Thigpen, NASA Ames Research Center

Overview

NASA's challenging mission to explore space and understand the universe and the Earth within it requires the Agency to innovatively apply and extend humankind's most advanced capabilities, technologies, and knowledge. The High-End Computing Capability (HECC) Portfolio provide a suite of such powerful leading-edge tools. The HECC Portfolio is for the benefit of all programmatic Mission Directorates—currently, the Aeronautics Research Mission Directorate (ARMD), the Exploration Systems Development Mission Directorate (ESDMD), the Science Mission Directorate (SMD), the Space Operations Mission Directorate (SOMD) and the Space Technology Mission Directorate (STMD). Additionally HECC supports the NASA Engineering and Safety Center (NESC).

Project Details

At its most fundamental level, the HECC Portfolio accelerates scientific and engineering discoveries by providing services that address NASA's current computing, computational and data analytic needs and engages in research and development that will create the productive environments of the future.

HECC meets NASA's current HPC requirements through the provision of world-class computing architectures both on the cloud and at NASA Ames, designed to maximize return on investment and services that multiply the impacts of those assets. Each year, HECC determines the amount of compute time to be provided to its customers and tracks usage against the target's allocations to ARMD, ESDMD, SMD, SOMD, STMD and NESC.

Computer time is measured in wall-clock hours and is described in terms of Standard Billing Units (SBUs). One SBU is currently defined as the work accomplished during a wall-clock hour on one Intel "Broadwell" node, which has two 14-core E5-2680v4 processors with a clock speed of 2.4 GHz. Broadwell nodes are equipped with 2,400-MHz DDR4 memory to provide higher memory bandwidth. Since HECC systems comprise many model types, this standard allows work to be normalized for all users across all the systems.

SBU allocations can be used on local or cloud assets. The allocation also provides users access to a suite of comprehensive, integrated services that include:

- Supercomputing Systems Services
- Application Optimization Services
- Data Publication and Discovery Services
- Visualization Services
- Data Analytics Services
- Networking Services

While meeting the current needs of NASA, HECC will partner with academia, industry, and other government agencies to prepare for a new generation of computers and ensure that the future generations of advanced computing technologies provide suitable platforms to support NASA's engineering and research challenges. The state of the art in physics-based simulations is changing, largely due to the Department of Energy's (DOE) Exascale Computing Project (ECP), which is spending $1.7 billion over 7 years to "deliver exascale simulation and data science innovations and solutions to national problems." For ECP, the path to exascale will employ GPU-accelerated nodes; preparing applications for those new architectures is difficult and time consuming. Data analytics is also exploding driven by a large increase in the amount and sophistication of the data produced by next-generation satellites, large-scale modeling applications and experimental environments such as wind tunnels along with the power of data analytics tools now emerging from the marketplace.

NASA’s ability to maintain its premiere world position in science and engineering and provide the nation with optimal solutions to our complex problems will be heavily impacted by how we can exploit and drive the emerging technologies in hardware and software (physics-based, data analytics, machine learning, and tools) that will increase the effectiveness of our scientists and engineers.

In addition to machines like ECP’s, the HPC market will continue its focus on machines without GPUs. Many of these may use high-bandwidth memory (HBM) on the processor chips to improve performance or specialized processors for special domains such as artificial intelligence and machine learning. For example, there are machines currently in design that will couple ARM processors with HBM. Much of NASA’s HPC workload, such as the OVERFLOW computational fluid dynamics code, are memory bandwidth bound and would benefit greatly from such an architecture. In this case, preparing applications is likely to be much more straightforward, but changes in the code and the underlying algorithms are likely still needed to take advantage of HBM and to reduce data traffic across the deep memory hierarchy.

To stay on the forefront of physics-based simulations and data analytic workload, HECC will prepare now for the paradigm changes that will be coming in the exascale era and beyond. That preparation involves identifying areas that are limiting the accuracy and/or the time to solution. If research into solutions for these limitations is underway by other projects such as DOE’s ECP, HECC will partner with the ongoing activity. In other situations, HECC will award contracts to conduct research in the areas of concern.

HECC will organize its activities along the three main technical axes of that project’s organization:

- Application Development: work with the Mission Directorates to identify NASA’s most important applications for the future, and partner with subject matter experts to upgrade them now to ensure they will be ready for the exascale era.
- Software Technology: investigate approaches and packages that facilitate the porting of existing applications or the development of new ones. Work with academia and industry to exploit the power of data analytics to provide better insight into complex data sets.

Hardware and Integration: provide access to exascale-like hardware and programming systems to support the first two areas. HECC will explore arrangements with the other government agencies to get access to their systems, or access to cloud-based resources, or in-house pathfinding systems acquired through the New Technology area of HECC’s budget. This area will also focus on evaluating how different exascale computing architectures suit NASA’s HPC requirements.

Results and Impact

This presentation focuses on the advances that HECC has achieved or is achieving to address the computing challenges of the agency. A new cluster, Cabeus, will provide scientists and engineers with 3 times more General-Purpose Graphics Processing Unit (GPGPU) capacity than is currently available. In December, we will begin installation of 128-node GPGPU expansion. This expansion contains 512 Nvidia A-100 GPUs. When delivered this will be joined with the existing GPGPUs to build Cabeus. The new GPGPUs will provide approximately 5 PetaFLOPS (10^{15} FLoating point Operations Per Second) resulting in a new system exceeding 7.6 PetaFLOPS.

We continually learn as we enhance our environment. Our rapid growth to multiple facilities with multiple systems resulted in an InfiniBand network that was sub optimal for our expanding environment. A second focus is the evolution from a single system focused multi-dimensional hypercube with storage internal to the hypercube to a data island connected by a fat tree and accessed through Lustre routers by all the compute clusters. Aitken, Pleiades and Electra maintain their expandable hypercubes while the hyperwall and new GPU cluster, Cabeus, utilize Fat Trees. This maintains the expandable nature of our compute systems while takin advantage of the better performance and simplicity of the fat tree.

The cloud efforts this year have had 2 focus areas. We were able to transition a traditional CFD team from on-premises resources to the cloud. Performance on the cloud was excellent and the team was able to make significant progress without the delay often seen locally. Additionally, we were able to work with some science teams to burst in from the cloud for compute intensive components of their workflows.

The experiment of housing supercomputers in modular facilities exceeded expectations providing growth opportunities that would not have been possible in our older facilities. The modular facilities increased power available for computing from 5.8MW to 10.9MW with the potential to grow to 37MW. The PUE for the 3 module is 1.05. We contracted with AZZ to build a new module with an additional 3.7MW of power. This will house the continuing growth components of Aitken. The initial expansion is 4 racks, 512 Milan-based nodes.

The hyperwall display is growing from 245 million to 1 billion pixels growing from a 23' x 8'11" to a 32' x 9'3" 128-display wall. The growth of the displays is forcing a redesign of the display room resulting in a room focused on the viewing environment. The new facility will be available 1st quarter of 2023.

What's Next

- HECC will continue to optimize the environment on site as well as facilitate access to cloud and other external vendors buildin coupled environments that take advantage of what they all have to offer. Additionally, HECC is addressing reliabillity and robustness by deploying a data module by the compute modules. Currently, all file systems are in our fixed facility (N258). The Data Module will allow systems that are not in N258 to continue operations during a planned or unplanned outage in N258. We are working toward having the home file systems, the archive and critical infrastructure systems duplicated, large file systems will be shared between N258 and the data module. Note that critical data can be recovered from the archive during prolonged outages allowing all users to continue operations. We are also looking at improviong our infrastructure in N258. One of our current chillers is irreparable. Rather than replace it with a similar design, we are replacing all chillers with newer, more energy efficient models. This will improve the PUE from 1.34 to 1.28.

Domain-Specific Language (DSL) Adoption into NASA's Goddard Earth Observing System (GEOS) Model Code

Christopher Kung, NASA Goddard Space Flight Center
Purnendu Chakraborty, NASA Goddard Space Flight Center

Overview

NASA's Global Modeling and Assimilation Office (GMAO) is undertaking an effort to develop a next-generation Goddard Earth Observing System (GEOS) model. Recognizing that accelerator-based high-performance computing (HPC) systems provide a promising platform to meet future GEOS requirements, the GMAO is looking to leverage such systems by incorporating a domain-specific language (DSL) into GEOS. DSL adoption brings an opportunity for code portability of GEOS across multiple architectures. It also enables developers to use a higher-level language, which improves development productivity by abstracting the details of the computing architecture. To jump-start DSL adoption, the GMAO leveraged a GridTools for Python (GT4Py) DSL port of the Finite-Volume Cubed-Sphere Dynamical Core (FV3) from the Allen Institute for AI (AI2). The GMAO is developing an approach to incorporate the GT4Py-port of FV3 into GEOS and determine whether incorporating the DSL-ported FV3 makes it tractable to achieve performance gains on accelerators over CPUs.

Project Details

The GMAO is taking incremental steps to create an accelerator-based GEOS codebase. The first step involved the integration of the GT4Py port of FV3 (gtFV3) into GEOS by creating a code interface layer that connects the Fortran-based GEOS codebase with the Python-based gtFV3. Since Fortran can interoperate with C-based code, the interface was implemented using the C Foreign Function Interface (CFFI) for Python, which allows C code to interface with Python. This interface passed relevant data between gtFV3 and the GEOS framework and also executed gtFV3. Developers successfully integrated the interface into the GEOS framework, demonstrating that gtFV3 can accurately execute an atmospheric model by comparing averaged zonal wind values from a CPU-based GEOS run.

Results and Impact

The GMAO verified that gtFV3 was successfully integrated into the GEOS framework by comparing 30- and 90-day U-zonal wind averages between gtFV3 running on NVIDIA graphical processor units (GPUs) and a CPU-based FV3 GEOS run. Currently, the performance of executing gtFV3 in GEOS is comparable to the current CPU-based FV3 implementation. AI2 has shown that gtFV3 executes about 3.5 times faster on GPUs versus CPUs, so the GMAO will try to achieve the same performance with gtFV3 within GEOS.

The project is forward-looking for the GMAO, which develops the GEOS model used to simulate and predict actions of the Earth's atmosphere and oceans. Since the DSL hides the details of the computing architecture from the developer, it provides GEOS with the capability to utilize current architectures – such as general-purpose GPUs (GPGPUs) and potential future architectures such as vector engines and ARM processors – without having to rewrite the codebase.

Why HPC Matters

Global atmospheric and ocean simulations such as those running with GEOS require signficant HPC resources to accurately resolve global climate patterns. Currently, advancements in HPC compute capabilities are being provided by accelerator-based computing such as GPUs. Enabling GEOS to utilize these architectures will provide the capability to tackle higher-resolution and more complex global simulations accurately.

What's Next

With heterogeneous architectures, properly managing data movement is key to achieving optimal performance, and this project will explore how to expand the GEOS software framework to support accelerator data transfers and refactor GEOS to minimize data movement. Creating a true global climate model will require porting additional physics schemes to GPU architectures. The GMAO will aim to create an aquaplanet simulation as a next verification model.

Making the Cloud Do What NASA Scientists Want

Hoot Thompson, NASA Goddard Space Flight Center
Daniel Duffy, NASA Goddard Space Flight Center
Ramón Ramírez-Liñán, NASA Goddard Space Flight Center

Overview

The NASA Science Managed Cloud Environment (SMCE) team has been working directly with multidisciplinary scientists to generate a cloud-native, collaborative, scientific research, and development environment. The SMCE team has assembled JupyterHub, Python, Dask, high-performance computing (HPC) parallel clusters, and other commonly used tools into a rapidly deployable environment used for large-scale analytics, data processing, and artificial intelligence and machine learning (AI/ML) applications. This initial work lays the foundation for creating an open science analytics platform – the Elastic Science Cloud Analytics Platform (ESCAP) – deployable across NASA and non-NASA research institutions and fueled by Infrastructure as Code best practices.

Project Details

The collaboration between the SMCE team and NASA scientists has yielded numerous requirements and desired features that have guided the design and implementation of ESCAP. Through a close integration between JupyterHub and Parallel Cluster, users can submit and monitor HPC jobs while analyzing data as soon as it is created. The SMCE team is seeking to achieve an even closer integration between Jupyter notebooks and HPC resources through shared user and group identities and JupyterLab extensions to enable browsing and downloading both public (as offered by Amazon Web Services, Microsoft Azure, and Google Cloud) and private datasets accessible from the notebooks. Technologies being explored for integration include Eclipse Che and easier methods of introducing user-defined Dask containers. Ever in the picture is a secure environment that promotes user flexibility, collaboration, and rapid research and development while minimizing risk.

Results and Impact

The outcome of this initiative will be a fully open-source, cloud-based solution that is easily and rapidly deployable and can be dynamically provisioned and shut down, as needed, to curtail unnecessary costs. This would create a reference cloud-based architecture for commercial cloud computing and on-premises computing capabilities for NASA scientists and collaborators. In addition, scientists could publish analyses, tools, and even web-end points side-by-side with the data for users to quickly enable a wide range of use cases.

Why HPC Matters

The end game of the SMCE is to enable and promote reduced time to science. HPC is vital to meeting the computational demands of NASA's environmental research objectives. Yet, since NASA has several disconnected data repositories and models, fully integrating them requires a new approach. Due to participation by a broader science community, this game-changing open framework will spawn rapid growth and benefit a wide range of research efforts.

What's Next

Future work by the SMCE team will include easing access to NASA's vast data repositories for NASA scientists and other affiliated researchers and promoting a seamless processing experience. For single sign-on, the team will focus on facilitating data transit between computing platforms (HPC, Dask, Notebooks, etc.) by integrating authentication and authorization across all platform components. Computationally, new testing will include additional processor options such as ARM- and AMD-based architectures, along with GPU systems for both HPC and AI/ML.

Will AI/ML Take Over HPC and the World?

Daniel Duffy, NASA Goddard Space Flight Center
Mike Little, NASA Goddard Space Flight Center

Overview

The applications and use of artificial intelligence and machine learning (AI/ML) have exploded in recent years due to a convergence of compute hardware acceleration, software frameworks, and data availability. From autonomous vehicles to ML-informed suggestions of which movie to watch next on streaming platforms based on past selection patterns, to deep fakes, AI/ML is inundating our lives as well as high-performance computing (HPC). Should we be concerned? What should we do about this trend, and how will it affect HPC at NASA? Join us to explore current research, the potential pitfalls and dangers of using AI/ML for science, and how HPC centers can be better prepared.

Project Details

The NASA High-End Computing (HEC) Program supports science and engineering research across all Agency mission directorates, and there has been a major increase in the requirements to support AI/ML in the last few years. Across the HEC Program, systems specifically designed to support AI/ML are being deployed for training scientific data, for inference, and for the creation of surrogates to augment physics-based models.

In addition to these requirements, NASA is also embarking upon the Transform to Open Science (TOPS) mission that is focused on ensuring accessibility, reproducibility, transparency, and inclusivity in science. NASA's HEC centers must be able to balance these rapidly emerging science requirements with initiatives like TOPS to meet the future needs of NASA scientists and engineers.

Results and Impact

The initial deployments of AI/ML systems specifically designed for HPC have already yielded amazing science results at NASA:

- Data scientist Brian Powell's research explores the Transiting Exoplanet Survey Satellite (TESS) Full Frame Images (FFIs) using ML to discover never-seen quadruple and sextuple star systems.
- Research scientist Jim Tucker is mapping carbon biomass by measuring the canopy area and height of trees from space over areas larger than the continental United States.
- Research scientist Mark Carroll and collaborators are using NASA's Ice, Cloud, and land Elevation Satellite (ICESat-2) measurements to create bioclimatic predictors and apply to future climate scenarios.
- Application scientists are training models to augment physics-based components to increase the performance for large-scale, high-resolution simulations.

From these scientific studies, NASA's HEC centers have learned critical lessons for how to best support AI/ML applications.

Why HPC Matters

Given NASA's mission, NASA's data centers are uniquely positioned to have accces to and analyze extremely large datasets. Observations of the Earth have been at petascales for decades and are growing at exponential rates in upcoming missions, with solar observations already in the multi-petabyte range. To adequately train and perform inference across these large datasets, HPC capabilities are needed.

As NASA's scientific models move toward exascale computing capabilities, there is a push toward creating models that couple physical-based systems with AI/ML trained surrogates. An example of this is NASA's high-resolution Goddard Earth Observing System (GEOS) weather model: GEOS uses traditional physics-based calculations but is coupled with an inference model that can perform spatiotemporal downscaling with sufficient accuracy for science applications. These types of applications represent the future of NASA research as we move toward creating a digital twin of the universe, and that can only be performed using HPC.

What's Next

This unique combination of requirements – continuing traditional HPC applications, supporting AI/ML, and providing capabilities to support the transformation to open science – is a significant challenge for NASA. In addition to the rapid changes in the technology landscape (e.g., processors, file systems, and software stacks), NASA HEC centers are faced with the challenges of supporting ethical and open scientific research. Through a partnership with the American Geophysical Union (AGU), new standards are emerging for ethical AI/ML research to be announced in the December meeting. The need for openness and reproducibility will affect the HPC environments and will become important components for future system architectures

Future Directions for NASA High-End Computing – Illustrated by an Earth Science Project

Tsengdar Lee, NASA High-End Computing
William Putman, NASA Goddard Space Flight Center
Dan Duffy, NASA Goddard Space Flight Center

Overview

NASA's Weather and Atmospheric Dynamics Focus Area (WADFA) is continuing our ambition to better understand the Earth system and to improve the capability to predict weather and extreme weather events. In the next couple of years, NASA's Global Modeling and Assimilation Office (GMAO) will develop and mature a coupled modeling system for sub-seasonal to seasonal (S2S) prediction and to produce

Modern-Era Retrospective analysis for Research and Applications version 3 (MERRA-3) data product based on coupled data assimilation.

Two external forces are impacting the modeling efforts at NASA. First is the availability of exascale computing systems. Already, the Frontier system at Oak Ridge National Laboratory (ORNL) has four AMD Graphics Processing Units (GPUs) on a node. In order to take advantage of the exascale computing system architecture, a modernization and refactoring of our modeling systems is a must. At almost the same time, NASA's Science Mission Directorate (SMD) is working to update and implement a new data management and computing system strategy with an open-source science initiative (OSSI) to make scientific data and software open in order to encourage collaboration. Under OSSI, the Core Services for Science Discovery (CSSD) project will develop core data and computing services and Transform to OPen Science (TOPS) project will build a community to accelerate the adoption of open source.

In this presentation, we discuss how NASA's HPC program may port our codes to the exascale systems and to align with OSSI to accomplish the ambitious goals of WADFA.

Project Details

NASA's High-End-Computing program is supporting NASA Center for Climate Simulations for a GEOS code modernization and porting effort to better prepare the Earth System Modeling (ESM) efforts for the exascale systems. After two years of painstaking effort, we begin to see the end of the tunnel. In another year or two, NASA will have an ESM that would be capable of running at extremely high resolution on an exascale computing system. The goal of the project is to reach the threshold set by DOE's exascale computing project and be able to run on 1/5 of an exascale computing system.

Results and Impact

HEC program's effort to modernize and port codes to exascale computing systems will enable NASA to reach the 1/5 system threshold set by DOE to run codes on their exascale computing. Without this capability, some of NASA's most innovative, ground-breaking science and engineering projects would not be possible.

Why HPC Matters

NASA High-End Computing supports an array of computationally taxing projects to enable the Agency to confront grand challenges and drive advances in science and engineering for the benefit of humanity. Leadership-class projects like recent Human-scale Mars Landers simulations, Artemis II booster separation simulations, and the next-generation Goddard Earth Observing System (GEOS) model are enabled by these HPC resources.

What's Next

In order to tackle the Agency's most taxing projects, we must be able to port our codes to exascale systems. This work has already begun through the Domain Specific Language (DSL) project launched three years ago. In order to meet our goals and successfully complete this project by 2024, we must recruit more experienced software engineers with particular expertise in GPU architecture.

Microscale Modeling of NASA Thermal Protection Materials

Joseph C. Ferguson, Stanford University
John Thornton, AMA at NASA Ames Research Center
Federico Semeraro, AMA at NASA Ames Research Center
Arnaud Borner, AMA at NASA Ames Research Center

Overview

To better understand and design thermal protection systems (TPS) for spacecraft, NASA scientists use numerical modeling to simulate the structure and behavior of heat shield materials at multiple levels, ranging from full-scale TPS modeling on a vehicle down to the atomic scale. For microscale modeling, where the material microstructure is simulated, an imaging technique called X-ray microtomography is used to generate high-resolution 3D images of a material's internal structure. Over the past decade, X-ray microtomography has altered the materials science field with its ability to non-destructively analyze a material microstructure. NASA's Porous Microstructure Analysis (PuMA) software, developed at Ames Research Center, provides a robust and efficient framework for computing material properties based on these 3D microstructures, and it has provided unprecedented insights into materials relevant for different missions, including heat-shields, parachute fabrics, meteorites, and other advanced composites.

Project Details

The PuMA software was originally developed as a tool to accurately predict material properties for TPS based on 3D microstructural models provided by microtomography. From a materials microstructure, PuMA can compute a material's porosity, surface area, pore diameters, effective thermal or electrical conductivity, anisotropic conductivity, material orientation, elasticity, tortuosity, and permeability. PuMA also includes solvers to simulate transient thermal transport and the decomposition of materials via oxidation. The software has grown extensively to include the ability to generate advanced synthetic microstructures with the aim of tailoring and optimizing their performance. Since PuMA was written with modularity and flexibility as software principles, the framework has provided the ability to easily add new physical models as the agency tests new TPS materials.

PuMA gives researchers the capability to model a comprehensive spectrum of material properties, from the most fundamental geometric features of a microstructure to advanced anisotropic thermoelastic and chemical properties, and can generate artificial microstructures ranging from simple analytical shapes to complex fibrous woven and non-woven geometries. Researchers can combine material generation and material characterization to enable parametric studies and sensitivity analyses to optimize the microstructural performance, which can inform mission design decisions and reliability assessments for TPS materials.

NASA has released PuMA as an open-source software with the goal of providing the scientific community access to both the specific physical solvers and a generic framework to build custom functionality for materials science.

Results and Impact

The governing equations solved by PuMA are applicable across a wide range of physical problems and length scales, from nano- to meter-scale multi-material systems. Because of this breadth in physical models, PuMA has been adopted in a uniquely diverse range of research topics, including advanced composites, concrete, 3D printed material analysis, geological flows, porous media combustion, meteorite analysis, parachute fabrics, batteries, and even the analysis of flow within teeth cavities. Users of PuMA are spread across NASA, government institutions, domestic and international universities, and private companies.

Why HPC Matters

The microscale modeling effort at NASA Ames heavily utilizes the high-performance computing (HPC) resources at the NASA Advanced Supercomputing facility, including the petascale Pleiades and Electra systems. The microscale flow modeling research uses both PuMA and the SPARTA direct simulation Monte Carlo tool, developed at Sandia National Laboratories, and relies on NASA HPC for large scale computations.

What's Next

Ultimately the PuMA development team plans to continue to improve and build upon the numerical solvers in the framework. As an open source code, they expect continued development and involvement from academic users, and believe that the field of microscale analysis will open the door for optimized material design that the capabilities in PuMA will help pave the way toward.

Advancing Geospatial Data Structures on GPUs: Mapping the Earth at Fine Scale

Jordan A. Caraballo-Vega, NASA Goddard Space Flight Center

Mark L. Carroll, NASA Goddard Space Flight Center
Christopher S.R. Neigh, NASA Goddard Space Flight Center
Caleb S. Spradlin, NASA Goddard Space Flight Center

Overview

The growing number of Earth observation satellite constellations is producing exponential increases in the amount of data produced. Petabytes of Earth imagery are coming down from hundreds of Earth-observing sensors in orbit. Planet's satellite constellation alone generates more than 11 terabytes of data per day. The scientific community is dealing with a fast-moving Big Data problem that necessitates multi-disciplinary engineering, high-performance computing (HPC), and software intervention to succeed. We present an end-to-end, GPU-accelerated pipeline for inference of machine learning (ML) and deep learning (DL) models on very high-resolution geospatial data.

Project Details

The Data Science Group (DSG) at NASA's Computational and Information Sciences and Technology Office (CISTO) is working with scientists and industry to accelerate the development of innovative science powered by artificial intelligence and machine learning (AI/ML). The goal of this project is to accelerate the inference of individual very high-resolution (< 3 meter) scenes by removing the CPU-dependency of traditional software and by efficiently leveraging HPC hardware.

The DSG builds GPU-accelerated Earth science data structures powered by NVIDIA RAPIDS and distributed by Xarray + Dask. With a complete support for cuML, PyTorch, and TensorFlow ML and DL frameworks, we perform all calculations and processing directly on the GPUs, avoiding bottlenecks from in-memory offloading. Multi-node and multi-GPU strategies are implemented for greater acceleration and parallelization. Once inference is performed, we can use NVIDIA's cuCIM library for additional GPU-accelerated postprocessing. Since executing this software uses a large amount of compute resources, we containerize it for greater portability.

Results and Impact

The DSG has processed thousands of visible and near-infrared very high-resolution images for the segmentation of clouds and land cover and individual object detection across the tropics and Arctic regions. Performing Earth observations at more local scales is important to better understand and monitor Earth dynamics. However, previous efforts using traditional computing methods have greatly underperformed in the presence of diverse landscapes and study areas smaller than regional scales. By using ML and DL models combined with HPC resources, we can monitor Earth at a level of detail that will accelerate the scientific study of regional dynamics using NASA's massive Earth data holdings and overcome existing Big Data challenges.

Why HPC Matters

Development occurs in the Explore/Advanced Data Analytics PlaTForm (ADAPT) computing environment at the NASA Center for Climate Simulation (NCCS). All of the GPU compute is done in the Promoting Research in Science using Machine Learning (PRISM) cluster and the new NCCS Discover Scalable Compute Unit 16 (SCU16) expansion. The PRISM cluster has 22 nodes each containing 40 Intel cores, 4 NVIDIA V100s, and 768 gigabytes (GB) of RAM; it includes an additional NVIDIA DGX system with 8 NVIDIA A100s and 1 terabyte (TB) of RAM. The Discover SCU expansion includes an additional 12 GPU nodes comprised of 2 x 24 AMD Epyc Rome cores, 4 NVIDIA A100 cores, and 512 GB of RAM.

What's Next

NASA Goddard's Data Science Group plans to continue scaling up the inference of large extents of Earth at very-high resolution. This includes exploring integration of NVIDIA's Triton system into the inference framework to better support cloud-native applications, while allowing the team to perform cloud bursting when needed.

Powering Synthetic Data Generation with HPC and Dask

Thaddeus Norman, NASA Ames Research Center
Jeremy Kravitz, NASA Ames Research Center
Shubha Ranjan, NASA Ames Research Center

Overview

Environmental issues concerning water quality and climate change have increased interest in the study of phytoplankton–microscopic marine algae. These organisms make up the foundation of the aquatic food web across the globe and serve several critical roles in recycling carbon and nitrogen. Their responsiveness to nutrient availability and water temperature by changing their distribution and color make them key indicators of change within their ecosystem, and these changes can be observed from Earth-observing satellites.

In order to better understand the observational data, the Biospheric Science team at NASA's Ames Research Center is developing the Spectral Water Inversion Processor and Emulator (SWIPE) tool, which uses synthetic data to help in the modeling and analysis of the optical properties of water containing algae. However, the process of generating synthetic data is time consuming, so the Biospheric Science team reached out to the Data Science team at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center to help speed up the process.

Project Details

There are four interrelated components to the SWIPE tool: modeling, image processing, big data, and deep learning. The modeling component includes the Equivalent Algal Population (EAP) code that uses a two-layer Mie scattering method (which relates the diameters and optical properties of microscopic particles to the wavelengths of scattered light) to calculate the optical properties of different plankton and phytoplankton species. This synthetic dataset is then used by the image processing component to determine the amount of algae particles in the imaged waterway. EAP data is stored in SWIPE's big data library, which is used by the deep learning component to generate analytical and predictive models of the water quality.

Generating SWIPE's synthetic datasets is a time-intensive process, taking over two hours to complete on a laptop computer for just one species of phytoplankton. At the request of the Ames Biospheric Science Team, the NAS Data Science team took a close look at the EAP component with the goal of reducing this bottleneck in their production pipeline. The Data Science team embedded Dask–a flexible open-source Python library for parallel computing–into the SWIPE tool to allow the distribution of the EAP code across multiple processors and nodes on agency high-performance computing systems, like the Pleiades supercomputer at the NAS facility.

Results and Impact

By using the Dask framework, the NAS Data Science team was able to make use of modules such as "dask.delayed" to run SWIPE's EAP code in parallel, increasing the number of processors allocated to the job. This method reduced the dataset generation time to approximately 18 minutes using Pleiades, compared to two hours on a laptop. Since each Pleiades Broadwell node has 28 processors, approximately 160 datasets can be generated in the same time as a laptop could generate one. With over 70 phytoplankton species being modeled–and more candidate species to be identified–this speed-up significantly reduced the Biospheric Science team's time-to-solution.

Why HPC Matters

The NAS Data Science team's efforts to integrate frameworks such as Dask into the SWIPE tool supports the Biospheric Science team's ultimate goal of building a neural network model that will use observational data to determine water quality from satellite and in-situ imagery. By leveraging the computational power of the Pleiades supercomputer and code parallelization, the two teams were able to increase the rate of dataset generation sixfold, which would not have been possible without access to agency high-performance computing resources and knowledgeable support staff.

What’s Next

With the dramatic decrease in the amount of time needed to generate synthetic data, the Ames Biospheric Science team is examining more phytoplankton species as potential candidates to be added to the SWIPE database. Adding the characteristics of additional species will provide researchers with a more accurate determination of water quality for a larger range of waterways around the world.

Beowulf Cluster Computing Inducted into Space Technology Hall of Fame

James Fischer, Goddard Space Flight Center

Overview

In April 2022, Beowulf Cluster Computing was inducted into the Space Technology Hall of Fame during the 37th Space Symposium hosted by Space Foundation in Colorado Springs, Colorado. This international award recognizes the impact of an idea first conceived at NASA Goddard Space Flight Center in 1993.

Project Details

The Beowulf approach for constructing powerful computing clusters uses only commodity parts: personal computer boards, Ethernet interconnects, and open-source Linux. This idea was born in NASA Goddard’s Space Data and Computing Division, which provided supercomputer access to NASA scientists for analyzing satellite data and running their computer models. Providing top-of-the-line Cray supercomputer cycles to all of the NASA scientists who needed them was too costly. Beowulf clusters were 10 times cheaper.

Under the Federal High-Performance Computing and Communications (HPCC) Program, most of NASA’s Earth and Space Sciences Project effort was focused on gaining access to the biggest scalable parallel computers in the U.S. to put Grand Challenge scientists on a path to achieving teraflops computing. A less visible project milestone was to develop a workstation made of parallel computers that could sustain one gigaflops (billion floating-point operations per second) at a cost of $50,000. This milestone inspired the ideas that became Beowulf.

The Beowulf concept using all commodity hardware and open-source software came from then-NASA research scientist Thomas Sterling. Using commodity hardware meant that each system would have the lowest possible cost and allow multiple Beowulf systems, which in turn would enable many software developers. Teaming with Sterling on Beowulf development was Don Becker, whom Sterling recruited to NASA for the project; a small staff, often students of existing university collaborators; and NASA Goddard’s in-house team of computational scientists.

Results and Impact

The NASA Beowulf team’s "How to Build a Beowulf" tutorials and the 1997 RedHat release of multi-platform CDs of Beowulf Linux brought commodity tools for parallel computing to the mainstream for the first time, and the team went on to publish a popular book on this topic in 1999. NASA’s investment in Beowulf was around four people per year for four years. Following that, Beowulf took a life of its own, and the NASA team became a contributor to the project, rather than its coordinator.

Why HPC Matters

After having proved successful at NASA, Beowulf Cluster Computing rapidly spread into supercomputing facilities globally. This inaugural method of building high-performance computing clusters has continued to be used for the past 20 years because it delivers superior price/performance. Currently, most of the world's TOP500 computer systems use the Beowulf approach.

What’s Next

Space Foundation's recognition of Beowulf's lasting contribution to high-end computing, along with NASA Goddard's role in its initiation, has cast a public spotlight on the creative work of a group of technical collaborators who took Beowulf from concept in 1993 to an easily and widely replicable kit form by 1997.

This new and welcome interest has re-energized this author to organize memories and records from that timeframe to construct a narrative that better brings into focus the breadth of work done during those four years and that identifies and honors the individual contributors. This narrative of innovation in computing technology will include its origins in the earliest days of space flight and will likely take the form of papers and book chapters.