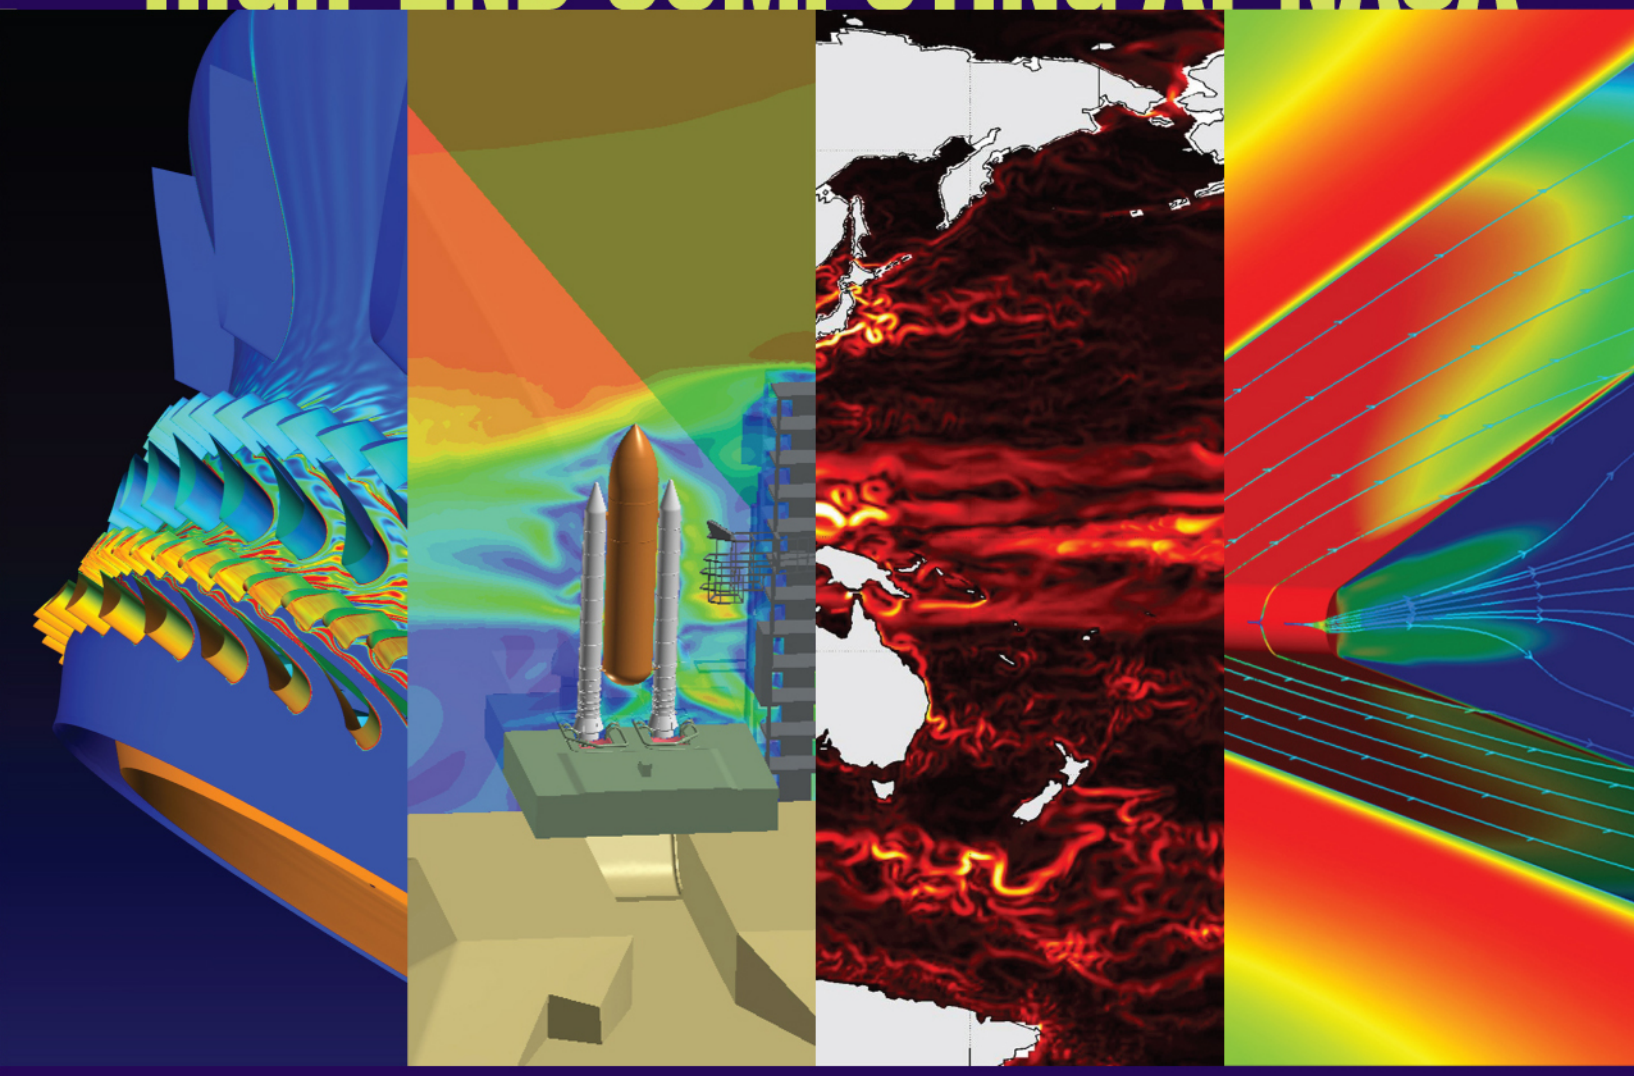
National Aeronautics and Space Administration

HIGH-END COMPUTING AT NASA



Highlights of science and engineering achievements on NASA's supercomputers

2007–2008

HIGH-END COMPUTING

VISION STATEMENT

For NASA's science, engineering, and technology community to consider high-end computing resources and services as essential tools that broadly enable rapid advances in knowledge and insight, and dramatically enhance mission achievements for the Agency.



Images shown on front cover (from left to right):

- Instantaneous view of the 5 blade-row high-pressure aircraft engine turbine coarse mesh simulation with the flow colored by vorticity to show velocity non-uniformities. The blade surfaces are colored by static pressure.
- Velocity magnitude color contours on a vertical plane passing through the centerlines of both Space Shuttle solid rocket motors.
- A global, coupled ocean-ice model simulates variables such as ocean current speed (shown here) and ice concentration.
- Sample wedge in an arc jet flow on the Orion crew exploration vehicle thermal protection system.

PROGRAM LETTER

June 1, 2009

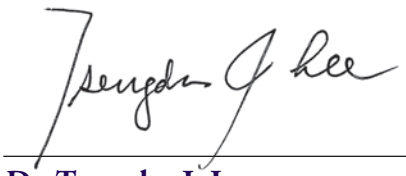
NASA High-End Computing Community and Stakeholders:

We are delighted to present this report on NASA's High-End Computing (HEC) Program, covering the years 2007 and 2008. This publication captures significant science, engineering, and technical achievements from across NASA, enabled by the Agency's world-class high-end computing resources and services.

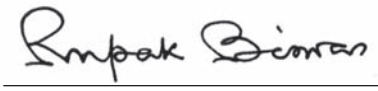
High-end computing serves an increasingly important role in NASA's missions—helping to safeguard our Space Shuttle fleet and astronauts, design next-generation space exploration vehicles, advance understanding of human impact on Earth's climate, expand knowledge of the origin and evolution of our universe, improve aerospace modeling, and much more. The HEC Program continues its commitment to maintaining a productive, service-oriented computing environment for all of these endeavors.

In this report, you will read about the technologies and services that make the HEC Program increasingly essential to NASA missions, and about the pioneering work that these resources have enabled. The report describes the Program's integrated environments that include premier computing systems, high-speed networks, a vast data management and archive capability, and an array of support services. During the past two years, the Program's facilities have achieved a 10-fold increase in computational capacity. Such technological accomplishments support principal investigators (PIs) from all four NASA mission directorates. In this report, 43 of these PIs relate their HEC-enabled science and engineering successes.

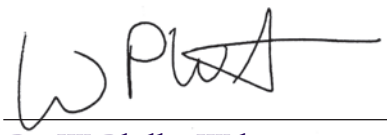
As NASA expands the frontiers of space exploration, scientific discovery, and aeronautics research, the HEC Program continues its dedication to providing computing resources and services to advance the Agency's missions. By delivering a reliable, service-driven computing environment that maximizes scientific discovery and engineering productivity, we will help ensure the success of America's space program for generations to come.



Dr. Tsengdar J. Lee
Scientific Computing Portfolio Manager
NASA Science Mission Directorate
NASA Headquarters



Dr. Rupak Biswas
High-End Computing Capability (HECC) Project Manager
NASA Advanced Supercomputing Division
NASA Ames Research Center



Dr. W. Phillip Webster
NASA Center for Computational Sciences (NCCS) Project Manager
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center

TABLE OF CONTENTS

EXECUTIVE SUMMARY 1

PROGRAM OVERVIEW 3

Introduction 3

Facilities 4

High-End Computing Support Services 5

HEC Operations 5

Data Management and Archive 7

High-Speed Networking 9

User Support 11

Data Analysis and Visualization 13

Future Mission Challenges 15

SCIENCE AND ENGINEERING HIGHLIGHTS 19

Aeronautics Research Mission Directorate 21

High-Resolution Navier-Stokes Code Development for Rotorcraft Aeromechanics 22

Integrated Inlet/Fan Simulation (IISIM) 24

Large Eddy Simulation for Highly Loaded Turbomachinery 26

Receptivity and Stability of Hypersonic Boundary Layers 28

Toward Improved Radiative Transport in Hypersonic Reentry 30

Turbomachinery Aeroacoustics: Turbine Noise Generation in Turbofan Engines 32

USM3D Analysis of the HRRLS Configurations 34

X-51 Aerodynamics 36

Exploration Systems Mission Directorate 39

Air Rig Testing of the Heritage J-2X Fuel Turbine 40

Ares I Roll Control System Jet Effects on Control Rolling Moment in Flight 42

Ares I-X Aerodynamics Database Development 44

Crew Exploration Vehicle Aerosciences Program 46

Computational Aeroelastic Simulation of Ground Wind-Induced Oscillation of the Ares I-X and the Checkout Model 48

Computational Support for Orion CEV Heatshield TPS Design and Analysis 50
J-2X Fuel and Oxidizer Turbine Simulations Including Disc and Tip Cavities 52
Proximity Aerodynamics of the Ares I Launch Vehicle During Stage Separation Maneuvers 54
Thrust Oscillation Focus Team Fluid Dynamics Analysis Support 56

Science Mission Directorate 59
Aerothermal Analysis in Support of Mars Science Laboratory Heatshield Qualification 60
CFD Support for Mars Science Laboratory Entry, Descent, and Landing 62
Computational Study of Relativistic Jets 64
Cosmology and Galaxy Formation 66
Coupled Ocean and Atmosphere Data Assimilation Systems for Climate Studies 68
Detailed Signatures of Cosmological Reionization 70
GEOS-5/Modern Era Retrospective-Analysis for Research and Applications (MERRA) 72
GEOS-5 Support of NASA Field Campaigns: TC4 • ARCTAS • TIGERZ 74
Global Modeling of Aerosols and Their Impacts on Climate and Air Quality 76
High-Resolution Modeling of Aerosol Impacts on the Asian Monsoon Water Cycle 78
High-Resolution Simulations of Coronal Mass Ejections 80
High-Resolution Wind Fields for Constraining North American Fluxes of Carbon Dioxide 82
Non-Boussinesq Ocean General-Circulation Model and GRACE Applications 84
Numerical Simulation of the Historical Martian Dynamo 86
Observing System Experiments: Evaluating and Enhancing the Impact of Satellite Observations 88
Plasma Redistribution During Geospace Storms: Processes and Consequences 90
Simulation of Coalescing Black Hole Binaries 92
Solar Surface Magneto-Convection 94

Space Operations Mission Directorate 97
Automated Aero-Database Creation for Launch Vehicles 98
CFD Analysis of Shuttle Main Engine Turbopump Seal Cracks 100
Numerical Analysis of Boundary Layer Transition Flight Experiments on the Space Shuttle 102
Space Shuttle Ascent Aerodynamics and Debris Transport Analyses 104
SSME High Pressure Fuel Pump Impeller Crack Investigation 106
Time-Accurate Computational Analyses of the Launch Pad Flame Trench 108

National Leadership Computing System 111
Modeling the Rheological Properties of Concrete 112
Transition in High-Speed Boundary Layers: Numerical Investigations Using DNS 114

INDEX 116



EXECUTIVE SUMMARY

High-fidelity modeling and simulation, enabled by supercomputing, are increasingly important to NASA's mission "to pioneer the future in space exploration, scientific discovery, and aeronautics research." While scientific and engineering advancements used to rely primarily on theoretical studies and physical experiments, today computational modeling and simulation are equal partners in such achievements. As a result, the use of high-end computing is now integral to the Agency's work in all four mission directorates.

NASA's High-End Computing (HEC) Program provides high-level oversight and coordination of the Agency's two HEC projects: the High-End Computing Capability (HECC) Project implemented by the NASA Advanced Supercomputing (NAS) Division at Ames Research Center, supporting HEC users in the four Mission Directorates; and the NASA Center for Computational Sciences (NCCS), implemented by the Computational and Information Sciences and Technology Office (CISTO) at Goddard Space Flight Center, supporting HEC users in the Science Mission Directorate. Funded by the Strategic Capabilities Assets Program and the Science Mission Directorate, these projects facilitate hundreds of computational projects from across the Agency. The HEC Program's Board of Advisors represents the strategic interests of each mission directorate.

This report presents the Program's successes in dramatically enhancing its facilities, resources, and services to meet the mission directorates' escalating demands, and describes advancements that Agency scientists and engineers have made using HEC Program resources in 2007 and 2008.

The *Program Overview* section (page 3) describes the range of resources and support services provided at the two HEC Program facilities and how they benefit the Agency's high-end computing users. Integrated HEC services are essential to making the Agency's supercomputing resources operate at peak efficiency and to helping users be as productive as possible. Following are some of the major technical achievements over the last two years.

- Facilities: The NAS facility has undergone significant expansions and upgrades—almost doubling its computer floor space and power and cooling infrastructure to accommodate new supercomputers. NCCS completed the necessary support infrastructure upgrades for a major increase in computing power and a six-fold increase in storage.
- HEC Operations: NAS installed the Pleiades supercomputer, currently the third fastest in the world. NCCS increased their supercomputing capacity five-fold with installation of the Discover system.
- Data Management and Archive: NCCS is expanding scientific collaboration and advancement through development of a portal for intelligent data access and use.
- High-Speed Networking: A joint NAS-NCCS effort resulted in a 54-fold speedup in data transfer for a major weather modeling project.
- User Support: NAS supported engineers doing time-critical launch and reentry computations for Space Shuttle missions, and saved 2 million processor-hours by optimizing three important aerodynamics codes. NCCS supported U.S. and international scientific field campaigns, nationally important climate research, and spacecraft environment simulations for mission engineering efforts.
- Data Analysis and Visualization: In spring 2008, NAS debuted the hyperwall-2 visualization system, enabling unprecedented large-scale data analysis and concurrent visualization. NCCS installed an interactive, large-memory data analysis system that gives users direct access to NCCS' 1.2-petabyte global filesystem and data archive.

NASA's HEC users currently number over 1,500 and come from every major NASA center, as well as universities, industry, and other agencies. The *Science and Engineering Highlights* section of this report (page 19) features the accomplishments of over 40 computational projects, selected based on their impact to the Agency over the past two years. For example:

- Aeronautics Research: Engineers are modeling innovative space exploration vehicles, aircraft components, and combustion and propulsion technologies. A Boeing team assessed the aerodynamic characteristics of the X-51 hypersonic vehicle, scheduled to fly in fall 2009. This project extends previous work completed by NASA on the X-43 Program in pursuit of practical hypersonic flight.
- Exploration Systems: NASA is designing America's next-generation spacecraft to take humans back to the Moon and ultimately to Mars. Engineers at Marshall Space Flight Center used over 7 million processor-hours to simulate the J-2X engine, which will power the upper stages of the next-generation Ares I and Ares V vehicles.
- Science: Researchers are advancing models and analyzing observations to better understand Earth's system and its impact on humankind, our planet's relationship with the Sun, and the evolving solar system and universe. Astrophysicists from NASA Goddard made fundamental discoveries about black hole mergers that are essential to the success of the Laser Interferometer Space Antenna, the first instrument expected to directly measure gravitational radiation from space.
- Space Operations: Research and engineering activities support the Space Shuttle and International Space Station programs, as well as launch, space transportation, and space communications work in both human and robotic exploration programs. NASA Johnson engineers led a project to model the ascent aerodynamics and debris transport of the Space Shuttle to protect Agency assets and, most importantly, its people.

The report also looks at future directions in Agency use of HEC technologies over the next several years. The Program will support strategic directions set for NASA by the new administration to further climate change research and monitoring; mount a strong program of human and robotic space exploration; support safe flight of the Space Shuttle to complete assembly of the Internation Space Station; and continue NASA's commitment to aeronautics research. The HEC Program's continued dedication to providing the best performance, usability, productivity, and security of its resources and services will propel computational science and engineering advances and help assure success for NASA missions.

PROGRAM OVERVIEW

INTRODUCTION

Over the past several years, NASA's mission "to pioneer the future in space exploration, scientific discovery, and aeronautics research" has been greatly enhanced through the contributions of high-fidelity modeling and simulation, powered by the best available supercomputing resources. As computational analysis is now an equal partner with theoretical study and physical experiments, high-end computing (HEC) has become an integral part of NASA's efforts in every key mission area. Major advancements are increasingly emerging from computational modeling first, and later validated by experimental studies and observations, or explained by theoretical work.

NASA's HEC Program, now in its fourth year, provides high-level oversight and coordination of the Agency's two HEC projects: the High-End Computing Capability (HECC) Project, formed under the Strategic Capabilities Assets Program (SCAP), implemented by the NASA Advanced Supercomputing (NAS) Division at Ames Research Center, and supporting HEC users in the four Mission Directorates; and the NASA Center for Computational Sciences (NCCS), implemented by the Computational and Information Sciences and Technology Office (CISTO) at Goddard Space Flight Center, and supporting HEC users in the Science Mission Directorate (SMD).

This program-level coordination ensures that the HEC facilities managed by these two projects provide a comprehensive set of supercomputing resources and services addressing the requirements of NASA, its external collaborators, and the nation. Overall HEC resource allocations are made annually to the mission directorates by the Program, and each mission directorate sub-allocates and prioritizes its own computational projects. A HEC Board of Advisors represents the strategic interests of the mission directorates, SCAP, and the office of the chief information officer. The Program is managed by SMD and funded through SMD and SCAP, which receives funding from all mission directorates to serve as an Agency-wide resource.

Since the 2006 *High-End Computing at NASA* report, usage of the Agency's HEC resources has grown rapidly—to more than 119 million processor-hours per year. During that time, the Program has also increased its overall computing capacity more than 10-fold. This expanded capacity has allowed the HEC Program's more than 1,500 scientific and engineering users at NASA centers, government laboratories, universities, and corporations to tackle nearly 500 computational projects in diverse disciplines.

The impact of users' computational work on mission success has never been greater. For example, Space Shuttle Orbiter support engineers have been simulating the transition from laminar to turbulent flow caused by protrusions such as gap fillers or thermal blankets. Using NASA's Columbia supercomputer, these engineers can quickly compute and predict the heating effects of such protuberances on the shuttle's thermal tiles during reentry into Earth's atmosphere. This capability is vital for real-time shuttle safety assessments during a mission. Moreover, such simulations are being used by the Exploration Systems Mission Directorate to optimize the design of future spacecraft. Increases in available supercomputing resources have also enabled NASA scientists to address their key Agency role in the national climate research initiative.

Achievements such as these rely on the tools provided by a comprehensive computing infrastructure including: HEC operations; data management and archive; high-speed networking; user support and application optimization; and data analysis and visualization. The pages that follow describe these services—provided by the NAS and NCCS facilities—along with examples of their positive impact in solving NASA's scientific and engineering challenges.



FACILITIES

The NAS and NCCS facilities are home to the HEC Program's most valuable assets—its computing resources and the people who support those resources and their users. Together, these facilities provide NASA users with more than 800 teraflops of computational power from six systems, distributed across 40,000 square feet of computer room floor space. In addition to high-end computing systems, each facility houses auxiliary equipment required to run the systems including front-end systems, disk arrays, massive data archive systems, and high-speed network routers and switches. Each site also provides specialized systems for data analysis and visualization. While serving their respective user communities, NAS and NCCS are integrating processes such as requests for computing time and user account management: and building capabilities for large file transfer and cross-center data backup and recovery.

NASA Advanced Supercomputing Facility

The NAS facility at NASA Ames Research Center, known worldwide for its innovation and expertise in HEC, was built in 1986, following a long history of computing leadership at Ames extending back to the early 1950s. The NAS facility provides high-end computing resources and services for all mission directorates and the NASA Engineering and Safety Center, and serves time-critical Agency needs such as in-orbit Space Shuttle analysis.

Over the past two years, the NAS facility has undergone a major transformation; by taking advantage of vacant computer rooms at Ames, it has expanded from one computer floor to four, and nearly doubled supercomputing floor space, electrical power, and cooling. We have also transitioned to a multi-vendor environment, and expanded peak computational capacity from 62 to over 700 teraflops. In the process, both the primary and new computer rooms were retrofitted to handle the dramatically increased power consumption and cooling requirements. New 90- and 450-ton chillers were installed with associated pumps and plumbing, and the power complex was upgraded to include vast new wiring arrays and the largest power switch west of the Mississippi River.

Following these facility upgrades, NAS now operates four supercomputers: Pleiades, Columbia, RTJones, and Schirra, described in the *HEC Operations* section (page 5). In early 2008, after an extensive upgrade of the its visualization laboratory, NAS also installed the hyperwall-2 visualization and data analysis system, highlighted in the *Data Analysis and Visualization* section (page 13).

NAS will continue to enhance its high-end computing and visualization resources, supported by comprehensive user-focused services, to ensure high productivity and further mission success for HEC Program users and the Agency.

NASA Center for Computational Sciences Facility

The NCCS facility at NASA Goddard Space Flight Center was formed in 1990 with the arrival of its first Cray supercomputers. NCCS now supports modeling and analysis activities for Science Mission Directorate (SMD) users in Earth Science, Astrophysics, Heliophysics, and Planetary Science. Carrying on a role dating from the 1960s to provide computing and data services to the Agency's science community, the facility enables users to run the sophisticated models needed to make the best use of NASA satellite observations, prepare for upcoming flight missions, and support national and international scientific field campaigns. The NCCS facility operates a high-performance computing environment comprising hardware, software, networks, storage capabilities, and tools—and supports user consulting and training activities.

The primary NCCS supercomputer resource is the Discover cluster. Installed in 2006 and augmented several times since, this system has increased NCCS computing power five-fold and storage capacity nearly six-fold. Discover includes a multi-tiered storage system that supports data-intensive scientific computation through a large, online General Parallel File System, and a data migration facility for long-term storage and preservation of valuable project and user data. In addition, NCCS has implemented a Data Portal to distribute data to a broader set of scientists and engineers, and to enhance data sharing; and upgraded its analysis and visualization system to provide tools for model development and validation, and for performing science using model results.

NCCS continually refines and updates its data-centric HEC system architecture and service offerings to support the missions of advancing scientific research and understanding Earth, the solar system, and the universe.



HIGH-END COMPUTING
SUPPORT SERVICES

NASA's HEC Program provides a full complement of services to support the Agency's scientists and engineers through the entire life cycle of their projects. Users require both capacity and capability computing, as well as batch and interactive service, including robust job scheduling and monitoring. Program support includes HEC operations, data management and archival storage, high-speed networking, 24x7 user support, application services, and data analysis and visualization—all targeted to help users make effective use of the HEC resources. Each facility optimizes its delivery of these services to best meet the specific needs of its user base.

HEC OPERATIONS

NASA's HEC operations services create the foundation for conducting large-scale modeling and simulation projects. These services provide not only massive computing power but also a wide array of tools ranging from compilers and debuggers, to job and workflow management. These services must address the challenges presented by increasingly complex models running on massively parallel architectures, and must help users make the most of both multi-core and specialty processors, such as graphics processing units. We meet these challenges through careful architecture planning and by implementing system hardware and software upgrades only after meticulous testing, to ensure maximum value with minimal disruption to users.

NAS HEC Operations

To keep pace with the growing needs of Agency users, NAS regularly evaluates new and emerging supercomputing architectures, and acquires and integrates those systems deemed to provide the best value to NASA. In early 2006, with the 10,240-processor Columbia supercomputer (62 teraflops, peak performance) at maximum utilization after about 18 months of operation, the NAS Technology Refresh (NTR) team began a four-phase process of enhancing NASA's HEC capacity. As part of the formal NTR evaluation process, we acquired a next-generation IBM POWER5+ system, named Schirra (320 dual-core processors), in spring 2007 to gain insight into the performance of this architecture on the NASA workload and to determine its feasibility to meet upcoming requirements. Later in 2007, an SGI Altix ICE system, named RTJones (1,024 quad-core Intel processors), was added to the NAS environment to support the Aeronautics Research Mission Directorate. In addition, Columbia was augmented to 14,336 processors and 89 teraflops (Tflops) peak performance—a 40% increase in capacity.

By spring 2008, NAS experts had gained extensive experience with all candidate HEC architectures, and determined that the SGI Altix ICE would provide the best overall value to the Agency. That fall, NAS and industry partners finished building the resulting Pleiades supercomputer, and integrating it with RTJones, establishing a system with 51,200 processor cores and 609 Tflops peak performance. At over eight times the capacity of Columbia's initial configuration, Pleiades exceeds the NTR goal to increase total sustained computing capacity at least four-fold every three years. Pleiades also achieved 487 Tflops on the HEC industry's LINPACK benchmark, making it the third most powerful supercomputer on the November 2008 TOP500 list. This ranking, combined with the November 2008 Green500 list ranking, made Pleiades the second-most powerful and energy-efficient general-purpose supercomputer in the world.



The Pleiades supercomputer, with a peak performance of 609 teraflops, is the third-fastest general-purpose supercomputer in the world (November 2008 TOP500 List).

Pleiades is already having a tangible impact on NASA's time-critical science and engineering challenges. For example, the Agency is saving thousands of hours and millions of dollars in experimental tests through computational analyses of Ares I Crew Launch Vehicle stage separation events and possible designs of the Orion Crew Exploration Vehicle thermal protection system.

NAS also refines its HEC environment for maximum reliability. For example, a unique weighting approach has been implemented to optimize data traffic movement to and from Pleiades through the nearly 20 miles of InfiniBand (IB) fabric that includes new fiber optic IB technology. Engineers customized the IB routing algorithm to reduce message contention and improve system transfer rates within a modified 10-dimensional hypercube. Due to the scale of Pleiades'

IB fabrics, a failover mechanism was implemented in MPT (SGI's version of Message Passing Interface), which can ride through temporary or permanent hardware failures between computer node end points. This also facilitates robust interconnections between Pleiades, Columbia, the mass storage systems, and the hyperwall-2 visualization system.

To ensure that users have access to the most effective resources, our experts also develop custom software tools and work with security experts on advanced methods to thwart regular attempts to break into NASA computers. NAS has developed and deployed a password system for the secure HEC enclave that enforces strict, uniform password rules to conform to NASA and federal regulations. The security team conducted a trade study and purchased a customizable, enterprise security management system for processing and correlating complex, interrelated, security-relevant information (such as logs and monitoring). This system feeds an extensible dashboard and sends alerts to flag potential intrusions.

NAS' near-term goal is to make Pleiades a stable production system by early 2009. In the longer term, NAS is preparing for a 10-petaflops HEC environment in 2012—a more than 10-fold increase over today's capability. Emerging innovative architectures and systems will be strategically leveraged to offer users the most effective supercomputing platforms and environments for NASA's computational challenges.

NCCS HEC Operations

NCCS provides computing resources for the Science Mission Directorate's (SMD) large-scale scientific and engineering models and simulations. NCCS computing services are uniquely configured to meet the data-intensive nature of these computations. SMD users require services to support development and execution of modeling and simulation codes.

During the last two years, NCCS has greatly increased its computing power, evolving to a more homogeneous, Linux-based cluster environment with vast, shared online disk storage. Compared with previous NCCS systems, this commodity approach offers greater computing and storage resources to the SMD user community, and reduces systems management complexity by implementing a single operating system. The current production platform, a Linux cluster named Discover, has 6,656 processor cores delivering 65 Tflops, with 10.8 terabytes of main memory and a 20-gigabit-per-second InfiniBand interconnect. This approach also allows greater flexibility in implementing incremental upgrades annually within the available budget. The resulting system enables large-scale ensemble runs, and storage and distribution of the data output. We also provide a test environment to evaluate system configuration changes and new user applications. Systems experts continually use benchmark information to determine required processor and I/O performance for model runs, and establish file organization approaches to optimize computational throughput.



The Discover cluster at NCCS. Discover's 6,656 processor cores yield a combined peak performance of 65 teraflops. During 2009, NCCS anticipates more than doubling Discover's capacity with an additional 8,192 processor cores, all based on Intel's new "Nehalem" processor.

Code development, the first stage in the life cycle of large-scale scientific applications (models), places great demands on scientists' time. NCCS provides a development environment where evolving models can be run. This environment includes a full complement of tools for developing, managing, debugging, and testing the code base of these applications. It also provides performance analysis tools and libraries, consulting services, and next-generation knowledge-based collaboration tools that enable developers to share their experiences. NCCS also maintains a code repository for SMD models, along with code versioning software to support modifications and upgrades. Access to the repository is granted to both NASA code development personnel and external collaborators. Our effective code development services support complex community-based model development and validation activities, and move codes more quickly from development to the production environment.

With the growing national interest in climate change research, contributions of NASA Earth Science programs to this field will continue to increase. NCCS has evaluated SMD's scientific and engineering initiatives and determined the HEC resource capability and capacity needed to meet those requirements. NCCS will continue to expand its Linux-based cluster(s) and enhance job-scheduling capabilities to address the increased processor count and diversity of job submissions. We are preparing to deliver unique and improved capabilities with increased computational and I/O performance to support complex model execution and data production, and meet the real-time execution needs of mission support operations. NCCS is also exploring enhanced workflow tools to allow seamless linkage between models, as well as use of "model recipes" and lessons learned from other scientists and engineers. We will support these workflows with global HEC resource management and automatic data migration on and off compute platforms. As time-to-solution for modeling and simulation remains a fundamental limitation, NCCS and their research partners will continue investigating opportunities to exploit multi-core parallelism on commodity processors and integration of powerful numerical accelerators.

DATA MANAGEMENT AND ARCHIVE

HEC data management and archive services address users' needs to store, retrieve, move, share, and preserve the vast amounts of invaluable scientific and engineering data produced by NASA and partner agencies. Data policies, including those for data aging and disposition, together with user needs for increased storage capacity to keep pace with computing power, guide planning for future data management and archive capabilities and capacity.

NAS Data Management and Archive

The NAS facility provides users with 16 petabytes (PB) of tape archive storage and 300 terabytes (TB) of online disk cache spread across three main archive systems. This is in addition to approximately 3 PB of disk storage attached to NAS supercomputing and data analysis systems. NAS' goal is to create the most robust and flexible high-performance storage environment possible.

For special projects, NAS storage experts create custom file systems to hold large amounts of temporary data, typically exceeding many terabytes. As an example, Science Mission Directorate (SMD) users working on ECCO2 global high-resolution ocean data syntheses required dozens of terabytes of disk storage to integrate their models and analyze results. We greatly expanded primary and secondary storage allocations and cross-mounted these to the archive storage system, substantially reducing bottlenecks during ECCO2 data generation and analysis. Custom training is also available to all NAS users requiring help with data management.



The Lou archive system housed at the NAS facility stores data generated on HEC Program supercomputers. In 2009, new archive systems will increase storage capacity seven-fold.

NAS has seen a steep growth in demands for computing time over the last several years—approximately 85% each year since 2004. This growth, in turn, places a heavy strain on the existing archive file systems. This trend is largely due to the increasing resolution of simulation runs and larger-capacity, higher-throughput supercomputers such as Pleiades and Columbia. In 2004, users transferred data into the archive at an average rate of 1.8 TB per day—today that number is more than 12 TB per day.

Over the last two years, our data management experts have made several key changes to improve long-term storage and transfer of data to archive systems. One improvement was to increase the size of the archive disk cache 10-fold by reusing older disk arrays from HEC platforms. Another improvement was to split the archive system, Lou, into three separate systems—one for SMD, one for the other three NASA mission directorates, and the third dedicated to testing and hyperwall-2 data analysis and visualization. This change helped streamline user access to data and improved transfer reliability. NAS also recently installed an InfiniBand (IB) network connection between Pleiades and the archive systems. As a result, users can now use batch job scripts to automate a file transfer from any computing node to the archive (via a file transfer protocol such as Secure Copy or bbftp). Previously, Pleiades users could only do a network file transfer from the front-end or bridge nodes.

In addition to hardware upgrades, both NAS and NCCS have worked to establish better redundancy and disaster recovery capabilities. NAS sends backups of system data to NCCS to ensure it is available in a second physical location in the event of a natural disaster or other unforeseen event. In turn, NCCS sends a large portion of their data to NAS. Each facility periodically stores encrypted files for one another to provide redundancy of key system information.

Several upgrades to the archive systems are in progress, including replacement of tape drives and storage silos no longer supported by the vendor. These upgrades will make more efficient use of space and take advantage of the latest tape densities and technologies. Following extensive testing, NAS has selected new hardware manufactured and supported by Spectra Logic, which will significantly reduce operating costs and floor space required—from approximately 800 square feet to only 70 on the main computer room floor, while increasing storage capacity seven-fold.

NCCS Data Management and Archive

NCCS addresses the full range of data management and stewardship requirements of the SMD user community. Data created and used by large-scale scientific codes and engineering simulations is accessible by all of the facility's HEC platforms and users, regardless of their physical location. This allows output from one computational run to be input to another without the user having to explicitly copy data from one system to another. Likewise, core observational and engineering datasets are available to SMD users, external researchers and collaborators, and the general public.

NCCS allows for sharing of data by users other than the data owners. This is especially beneficial when analyzing scientific data from emerging production codes, and when drawing from simulation and satellite observation data. To provide concurrent high-speed access to user files from every platform in the NCCS environment, a data store of 1.2 PB is visible to all systems via the Global Parallel File System (GPFS).

To facilitate large collaborative projects, we make it easy for users to set up project data spaces accessible by designated team members from various institutions, including public access for those outside NASA. Access to these data spaces is provided from within the Discover cluster, and distribution of project datasets are supported via the Data Portal. Currently, a copy of public data is stored to the local Data Portal within a local GPFS environment. With an increasing demand for sharing scientific results pertaining to climate change research, NCCS has implemented a method for the Data Portal to distribute data without needing to make local copies.

Increases in computational model complexity and number of job runs generate tremendous data management and archive requirements. NCCS combines heuristic evaluations of future requirements with actual storage usage statistics to develop capacity plans, and uses benchmark information to determine required I/O performance and establish file organization/striping approaches.

NCCS archives datasets using its Data Migration Facility (DMF) mass storage system, which currently holds up to 16.5 PB of data. Users define data requiring long-term storage, and DMF automatically migrates data from online disk cache to tape. Since 2006, NCCS has more than doubled the size of the DMF cache to better support increased archive demand commensurate with the increased computational capabilities. The addition of some 200 TB of online disk cache



The NCCS data archive infrastructure includes two linked StorageTek SL8500 silos. With an additional nine silos in other locations, NCCS has a total archive capacity of 16.5 petabytes.

to the DMF system, along with an additional 850 TB of online disk to the GPFS, allows users to maintain critical datasets online and tailor I/O performance to significantly reduce data retrievals from tape. Users manage and monitor archived data for longer-term need. The DMF allows deleting user datasets when no longer needed; transferring ownership when users leave projects; and proper disposal of data at the completion of projects.

The NCCS facility also has a unique stewardship requirement to provide long-term storage capacity to preserve invaluable scientific and engineering data produced by NASA and partner agencies. Project teams manage these datasets to provide continuity of responsibility, with formal agreements to ensure appropriate data retention.

Although limited search and query capabilities exist today, NCCS will incorporate additional data sharing and publishing capabilities in the future. These services will enhance data discovery, making it easier for scientists and engineers to locate relevant data to support projects. Implementation of these services will include enhanced metadata and supporting metadata search capabilities. Sharing and publishing rules and the related metadata will be managed and modified though a database, and will not require scientists and engineers to modify individual files.

Other upcoming plans include establishing capabilities and procedures for failover of critical applications from one facility to another.

HIGH-SPEED NETWORKING

The HEC Program's high-speed networking services provide fast, large-capacity connections to support exponentially increasing data transfers between computing resources and NASA's distributed user base. High-performance connectivity is supplied to and utilized by NASA centers and our university and research partners. Through in-depth analysis of traffic flows, the Program provides end-to-end (supercomputer-to-desktop) networking analysis and support for NASA and its partners. Networking experts develop solutions to meet the exponential growth in HEC network traffic, making multi-terabyte data transfers seamless for users.

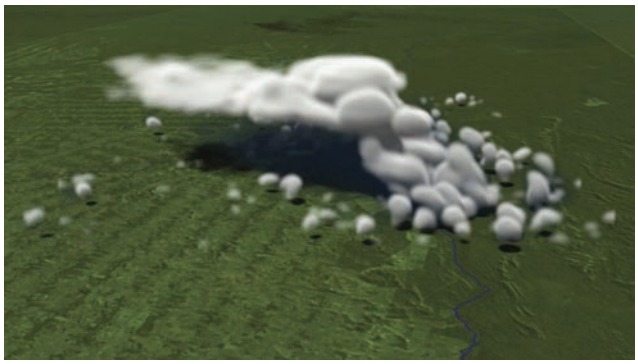
NAS High-Speed Networking

Network engineers at the NAS facility, working together with NAS systems engineers, have for the past two years focused on developing innovative approaches for boosting network performance for the Agency's challenging applications and larger data transfers. This work involves the exploration of new technologies capable of maximizing network performance to and from NAS resources.

NAS network engineers have also been working closely with users, both remotely and through site visits, to help them use existing bandwidth more efficiently by: optimizing multiple aspects of end-to-end flows; tuning user systems; working with wide-area network (WAN) service providers such as National LambdaRail, the NASA Integrated Services Network, and Internet2; and working with user site infrastructure support teams to identify and remove bottlenecks along the network path. NAS-developed tools also help users take better advantage of existing HEC network capabilities. In one case, NAS engineers, in collaboration with engineers from NCCS, helped scientists working on the 3D Cloud-Resolving Model project (Goddard Cumulus Ensemble) increase data transfer rates by deploying NAS' version of the open-source file transfer application bbftp and making some end-system adjustments. This assistance resulted in a 54-fold improvement in network throughput performance and dramatic time-savings for the project.

Recently, our network engineers have tested 10-gigabit-per-second (Gbps) and faster firewall devices, which will help improve the security of HEC resources without creating bottlenecks in HEC traffic flows. Engineers are proactively looking for technologies to accelerate the supercomputer traffic flows to and from users. The latest development on this front is with the Cisco Wide Area Application Services, an appliance system capable of helping users achieve high performance with minimal adjustment to their desktop or local server system.

Network monitoring is another focus area for NAS. By using both custom-designed and off-the-shelf network analysis tools, engineers monitor network traffic conditions in real-time, and to look for historical trends. Recent innovations in the monitoring arena from cPacket and NetOptics include 10 Gbps "taps" that passively monitor network activity for trending, performance monitoring, and troubleshooting needs. The team is also investigating more effective ways to measure and evaluate increasingly large flows, using approaches that can readily scale above 10 Gbps while retaining high fidelity for smaller flows.



This visualization shows a three-dimensional, high-resolution simulation of a convective cloud system over South America. NASA Ames Research Center and NASA Goddard Space Flight Center network engineers helped scientists working on NASA's 3D Cloud-Resolving Model project dramatically increase their data transfer rate.

NAS is also focusing on enhanced network management. Using custom-designed and commercial off-the-shelf (COTS) network management tools, engineers manage a highly complex network environment. Examples of custom-designed tools include: the web-based Network Inventory Management System that allows more effective management of network resources; a web-based Access Control List (ACL) Management tool to more efficiently manage and troubleshoot ACLs; a Dynamic Host Configuration Protocol script that automates notifications to networking staff regarding users that exceed the allowed time on the system hardening network; and a packet size mismatch checking script to identify Maximum Transmission Unit errors. COTS tool examples include the ManageEngine network analysis tool suite, which has advanced network monitoring and management features, and Netflow Tracker to enable real-time monitoring of NetFlow data.

The NAS networking team recently completed three-year roadmaps for both WAN and local area network (LAN) environments. Required technology enhancements include lambda switching, disaster recovery, user-driven performance assessment tools, secure InfiniBand (IB) over WAN, and IB management. These advances in emergent network technologies are specifically targeted to supply the high-performance networking necessary to keep pace with the increasing computing capabilities. Further, tools such as these support effective management of complex networking environments.

By continuing collaborations with other research and development (R&D) partner networks, advanced networking vendors, and companies in Silicon Valley, our network engineers will remain on the forefront of emerging technologies to ensure that NASA's unique computational requirements are met. Important contributions include involvement in and representation within the Joint Engineering Team (a subcomponent of the interagency Networking and Information Technology Research and Development program), which coordinates networking activities, operations, and plans among multiple federal operational and research networks. NAS also reports on issues that affect the networking research community and makes recommendations to help mitigate those issues, especially related to security across government networks.

NCCS High-Speed Networking

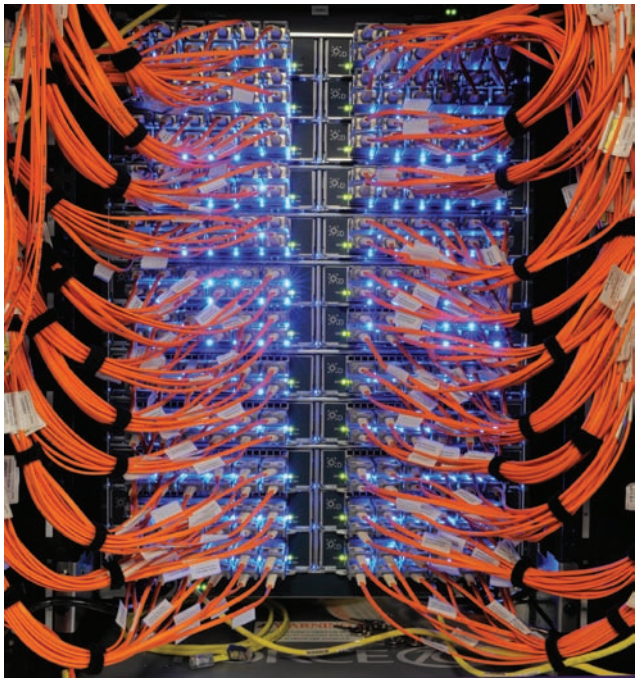
NCCS networking efforts include both direct support of existing HEC user activities, and R&D in advanced communications technologies and protocols for future HEC production data flows. Over the past two years, several significant networking R&D achievements have come to fruition. A 10-Gbps coast-to-coast network was established between the University of California, San Diego and NCCS in Greenbelt, Maryland. This initial 10-Gbps capability is an important milestone toward establishing connectivity from NASA Goddard to the external high-performance network community. This effort earned a NASA group achievement award.

Also in 2006 at the annual Supercomputing Conference (SC06) in Tampa, Florida, the network team demonstrated use of the National Science Foundation's DRAGON Project's Xnet capabilities to dynamically stream uncompressed high-definition video from NASA Goddard to the exhibit floor. This was critical to showing the flexibility of new dynamic circuit switch technology to provide off-campus connections with immediate and temporary access to high-speed WANs. At SC07 in Reno, Nevada, this team supported real-time 3D imagery, further illustrating the ability of dynamic networks to properly synchronize multiple data streams required for 3D presentations.

Beyond these R&D activities, since 2006, NCCS has been providing direct support for existing user activities. To improve high-speed data transfer for NASA's HEC users, the NCCS Science and Engineering Network (SEN) infrastructure was upgraded to 10 Gbps. Networking teams also offer specialized support for individual Science Mission Directorate projects. The joint effort of the NAS and NCCS teams to help network users from Goddard's 3D Cloud Resolving Model project produced the 54-fold improvement in data transfer performance through enhanced file transfer tools and end-system tuning.

Near-term network support activities are focused on further improving network performance and data transfer rates for HEC users. We are working closely with the user community and the NASA Integrated Services Network to support the transition of intra-NASA HEC production flows. NCCS plans to further enhance the SEN infrastructure, and is researching opportunities to provide further performance gains for HEC projects.

Building on the existing 10-Gbps Ethernet WAN and 20-Gbps InfiniBand switching fabric, the NCCS network service team will continue evaluating new technologies to support NASA's HEC user community. Of particular interest are the assessments of emerging 40- to 100-Gbps communication technologies. Industry leaders are collaborating on a 40-Gbps Live Network trial. NASA Goddard will also be testing the use of a 100-Gbps Ethernet testbed that complements and interconnects with Internet2, ESnet, Infinera, Juniper, and Level3 networks. Additionally, NASA Goddard is sponsoring a Small Business Innovation Research/Small Business Technology Transfer opportunity for an n x 10-Gbps Offload network interface card for NASA, National LambdaRail, and grid computing to implement new technologies and protocols that support HEC science and engineering applications.



This tiered interconnect switch on the NCCS Discover Linux cluster manages data traffic over the 20 gigabit-per-second (Gbps) InfiniBand internal network. NCCS and the NAS facility maintain local-area networks and connect to each other via wide-area networks with 10-Gbps performance.

USER SUPPORT

User services staff provide direct assistance to the scientific and engineering user communities in every facet of their interaction with HEC resources—from setting up user accounts, to disseminating system information, to one-on-one problem solving and group training. Support ranges from resolving simple system usage issues to consulting on complex code optimization and advanced software techniques. The goal is to make effective use of the HEC resources and remove any obstacles in the way of user satisfaction.

NAS User Support

NAS provides tiered levels of support to all NASA mission directorate users: A 24x7 control room staff resolves basic system usage questions and coordinates end-to-end support (tier 1). Scientific consultants provide assistance in troubleshooting execution issues with user applications (tier 2), and applications experts perform significant code modifications using advanced software tools and techniques (tier 3).

In addition to providing tier 1 support, the control room team monitors the physical facility and HEC resources and manages computer jobs and queues—all to ensure a stable, productive computing environment for users. An important aspect of this team's work is the operational support provided for the aerothermal and debris analysis teams during Space Shuttle missions. Before and during each shuttle launch, NAS prepares and tests all HEC components to ensure engineers can provide computational analyses to mission managers to clear shuttles for landing. The ability to quickly reconfigure and reprioritize computing resources to assess potential shuttle launch damage is a key service to the shuttle program. The team also provides support for the hyperwall-2 visualization system.

NAS' scientific consulting services include application performance optimization; evaluating, installing, and customizing performance analysis software and tools; and specialized benchmarking of current and future HEC architectures to identify and leverage those best suited for the NAS computing environment. Application specialists examine performance characteristics of scientific and analysis codes and optimize them to enhance utilization of HEC resources and technologies. Small but high-impact code adjustments are provided routinely, and on request, NAS provides detailed analysis and advanced optimization services.

Over the last year, this work has greatly benefitted research applications for over 30 projects in all mission directorates. For example, a comprehensive optimization of USM3D—an important computational fluid dynamics (CFD) code used by the Aeronautics Research and Exploration Systems mission directorates for intensive aerodynamic analyses—improved the code's runtime two-fold, and reduced memory requirements by factor of 2 to support very large computational grids (over 100 million cells) as required for high-fidelity solutions.



HEC Program user services staff members are available 24x7 to answer basic system usage questions and coordinate end-to-end support.

Also within the last year, the applications team optimized two other important CFD codes, Phantom and OVERFLOW, which are heavily used by ESMD and SOMD for aerodynamic analyses of vehicle designs. In response to time-critical ESMD analysis needs, NAS applications experts reduced the run-times of these codes—saving nearly 2 million processor-hours and effectively freeing up more than an entire 512-processor node of the Columbia supercomputer. This work enabled completion of calculations that otherwise could not have been run.

Recently, NAS tool developers launched a web-based application for monitoring computing activity. Based on the more complex Heads Up Display (HUD) developed for NAS control room operators, this new "miniHUD" provides users an at-a-glance overview of the state of NAS supercomputers, enabling them to make informed decisions to manage their jobs. Users can "drill down" to get details on various nodes and subsystems, such as processor usage and job queue status.

With a new generation of supercomputers in place, NAS is now planning for a 10-petaflops HEC environment. To continue giving users world-class support as resources scale up, the user services staff is developing new and improved processes, tools, documentation methods, and training programs. We continue to work with other NASA centers on an integrated approach to handling user account creation and requests for computing time, and incorporating Agency initiatives (such as the NASA Account Management System and NASA Consolidated Active Directory) to ensure a more uniform and secure environment for users.

NCCS User Support

NCCS user services provides tiered levels of support to meet the full range of Science Mission Directorate (SMD) needs. Level 1 services, provided by NCCS and the Software Integration and Visualization Office (SIVO), include account support, help desk service, and access to system documentation along with system information services such as email notices, teleconferences, and forums. Level 2 and 3 services include consulting and training to help users make the most effective use of HEC resources. This includes supporting the use of modeling frameworks for greater portability and/or extensibility and improving code performance through optimization. NCCS User Services also coordinates Agency and center-provided services that are outside of NCCS but critical to providing a complete and responsive computing environment for advancing SMD projects.

NCCS and SIVO consultants provide expertise in computational science and scientific and engineering code development, including support in numerical techniques, software design and implementation, code parallelization using MPI or OpenMP, code porting, and performance optimization. Consulting services include both level 2 response to complex user queries and level 3 assistance to scientists and engineers requiring more extensive support.

NCCS training activity addresses all HEC services, platforms, software, and tools as well as a broad range of topics related to developing accurate and efficient scientific codes. Support is also provided for understanding and using data resources both within NCCS and at external data centers.

Over the last two years, significant user support efforts have been applied to evolving the system configuration to meet growing SMD computational needs. These needs include greater job scaling, increased memory, longer execution periods, and meeting product delivery timelines in support of operational projects and spaceflight opportunities; NCCS has determined the system requirements to support these needs. The user services team has also modified job submission and queuing policies to address the increased processor requirements of the models, the processor range for concurrent job execution, and job completion expectations.

Our user services experts also took the lead in providing a timely and successful migration of scientific codes from a shared memory system to a new distributed memory cluster environment. They contacted each user being impacted, provided training, assisted in code porting and/or redevelopment, and tracked the migration of applications in incremental steps. Additionally, this team proactively addressed user data needs as they migrated from SGI's clustered filesystem, CXFS, to the Global Parallel File System; and all user codes and data were successfully migrated.

In response to the Agency's evolving policies, NCCS has proactively worked with NASA Goddard personnel to clarify procedures for granting foreign nationals access to HEC Program supercomputers. These efforts have demystified and streamlined the procedural requirements and enabled the NCCS help desk to more easily establish user accounts. NCCS participates in NASA's security policy team and will continue to update its practices accordingly, and coordinate with the NAS facility to ensure consistency across the Program.

Another user services role is ensuring that users can adequately exploit future NCCS capabilities. Multi-core processors present a challenge for the entire high-end computing community, and NCCS is working to identify and evaluate performance enhancements these technologies could provide. Also, migration to newer versions of the Earth System Modeling Framework greatly improved memory management for complex codes, allowing greater scalability and better performance. As time-to-solution remains a fundamental research limitation, NCCS and SIVO are investigating using numerical accelerators for computationally intensive portions of NASA's scientific models. NCCS also supports activities to extend the use of NASA models and data to the broader scientific community.

Looking ahead, our user services staff will continue to reach out to the SMD user community to address their needs for high-end computing, data storage, visualization and analysis, and data sharing. NCCS will provide users with access to its knowledge-based trouble tracking system to give them more insight into problem resolution activities. Moving beyond static reporting will facilitate better user participation and expedited problem resolution.

Both the NCCS and NAS facilities will also be involved in forming a user board to represent the interests of HEC users. Following the example of an SMD Computational Modeling Capabilities Workshop in July 2008, workshops for other mission directorates are being planned to broadly understand and prepare for NASA's upcoming HEC needs.

DATA ANALYSIS AND VISUALIZATION

The HEC Program's data analysis and visualization services enable NASA scientists and engineers to find meaning in the vast amounts of data in their computational models and observational datasets. Analysis tools harness the power of computer processing to filter, sort, search, and compare datasets, as well as to apply advanced statistical and other types of algorithms—all with the goal of discovering useful information in datasets. In turn, visualization tools leverage the human brain's ability to identify interesting features and patterns in images and animations, allowing scientists and engineers to more deeply explore data and more clearly convey results to colleagues and the public.

NAS Data Analysis and Visualization

To help users understand and interpret their results, NAS visualization experts capture, process, and render enormous amounts of data to produce high-resolution images and videos. These experts also develop and adapt specialized visualization solutions for the Agency's unique science and engineering problems. Working closely with users, the team customizes and creates new tools and techniques to expose the intricate temporal and spatial details of computational models.

The NAS visualization team has developed special technologies for moving large datasets directly to graphics hardware as they are generated so that they can be displayed and analyzed in real time. A cornerstone of this capability is the hyperwall-2 visualization system installed in spring 2008. This powerful system provides a high-speed, fully interactive environment that enables users to visualize, analyze, and explore high-resolution results and pinpoint critical details in large, complex datasets. The hyperwall-2 is a matrix of functionally interconnected graphics workstations and displays coupled directly to the NAS facility's high-end computers via InfiniBand. The 128-screen, quarter-billion-pixel flat panel display system measures 23 feet wide by 10 feet high, giving users a supercomputer-scale environment to handle the very large datasets produced by high-end computers and observing instruments. Powered by 128 Nvidia 8800GTX graphics processing units (GPUs) and 1,024 Opteron processor cores, the hyperwall-2 has 74 teraflops of peak processing power and a storage capacity of 475 terabytes.

The hyperwall-2 is also integrated with a NAS-developed, state-of-the-art concurrent visualization framework, which enables real-time graphical processing and display of data while applications are running. This capability is critical to supporting huge datasets that are difficult to store, transfer, and view as a whole, and delivers results that are immediately available for analysis. Most importantly, concurrent visualization makes it feasible to render and store animations showing every simulation timestep, which allows users to see rapid processes in their models—often for the first time. With concurrent visualization users can also do on-the-fly identification of computational problems or parameter adjustments needed for applications. Live production data can be pulled from the supercomputer to hyperwall-2 without hindering code performance, allowing users to reap the benefits of visualization services without slowing turnaround time of their analyses.



The 128-screen hyperwall-2 visualization system enables users to view, analyze, and explore their high-resolution modeling and simulation results.

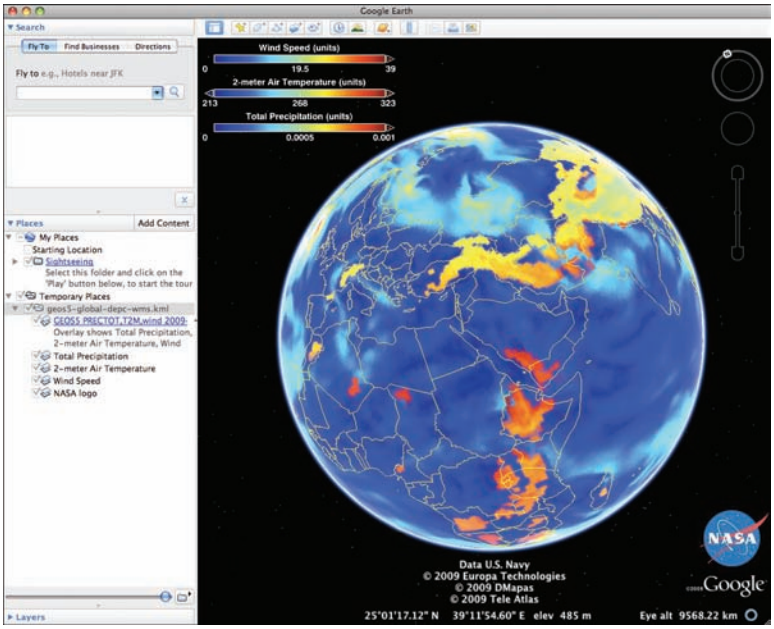
Other inventive visualization capabilities being explored at NAS include efficient GPU techniques, multivariate data visualization, and out-of-core techniques. Efficient GPU computation and rendering techniques enable both concurrent and rapid-iteration post-processing visualizations. Multivariate techniques include "linked derived spaces," which allows linking and selecting subsets of 2D scatterplots to facilitate visual tracking of key data points across all variables. Out-of-core data management techniques enable exploration of datasets that are too large to fit in memory.

Over the past couple of years, NAS experts have applied these capabilities to many applications, benefitting projects in all mission directorates. For example, detailed animations created from time-accurate, 3D simulations of ignition conditions within the flame trench of the Space Shuttle launch pad have provided NASA scientists with valuable insights into both repair criteria for damage caused during a shuttle liftoff, and potential modifications for the single-booster Ares I vehicle. Concurrent visualization methods have played a key role in identifying inherently unsteady flow structures in data-intensive simulations of V-22 Osprey rotors; and advanced GPU techniques have accelerated rendering of particle traces for massive simulations of convection on the Sun's surface.

Future advancements, such as streaming of visualizations to remote users and development of large GPU clusters, will expand the capacity and breadth of these valuable services. Visualization will continue to become more tightly integrated into the traditional environment of supercomputing, storage, and networks to offer an even more powerful tool to scientists and engineers.

NCCS Data Analysis and Visualization

NCCS provides a spectrum of data analysis and visualization services to help scientists access and manipulate large amounts of data produced by simulations, experiments, and satellite observations. For visualization, NCCS provides dedicated computing resources and access to large display systems. In addition, analysis services go beyond visualization to support commercial analysis software and tailored analytic tools developed by scientists. NCCS tools enable seamless movement of data to the analysis, visualization, and Data Portal environments, and data discovery tools help users assemble the required datasets for analysis. Each of these services is also supported by consulting and training for users.



The NCCS Data Portal hosts the Web Map Service (WMS), which is provided by NASA Goddard's Software Integration and Visualization Office (SIVO). In this view, WMS visualizes multiple datasets from a run of the Goddard Earth Observing System Model, Version 5 (GEOS-5) and displays the results using the Google Earth interface.

Together with our scientific visualization partners, NCCS also provides specialized visualization support for Science Mission Directorate (SMD) projects. Several multi-panel hyperwall systems display data continuously at multiple locations around the NASA Goddard campus. These displays show current meteorological phenomena and highlight special features from SMD research. Multi-panel imagery allows scientists to visually assess before-and-after scenarios, and images can be transferred directly to user workstations for exploring the intricacies of their models and simulations. SMD users also employ NASA Goddard's Scientific Visualization Studio facilities to produce high-quality movies in support of education and public outreach. Complete with professional narration, these movies have appeared on many scientific television programs.

NCCS supports data analysis directly on the high-end computing platforms, and allows scientists to copy results of interest to their own servers or desktops. Using the NCCS Data Portal, collaborators can access data and perform limited analysis. While analysis on the supercomputers has proven extremely successful for model developers, and the Data Portal is a productive resource for many smaller applications, NCCS has also planned for capability enhancements to meet users' evolving scientific analysis challenges. Greater model complexity and resolution have increased data output volumes, making it more difficult for scientists to transfer data to their local systems. Analyzing these large datasets now requires more sophisticated, parallel analysis tools. Scientists often need to analyze data generated elsewhere in conjunction with external datasets, such as the NASA Goddard repositories of observing system data from many Agency satellites.

These growing requirements, combined with greater emphasis on collaborative research, led NCCS to install a new interactive, large-memory analysis platform with direct access to the entire NCCS global filesystem and data archive. This gives users fast access to large datasets needed for scientific analysis. Based on user interviews and knowledge of analytic techniques used at other supercomputing centers, NCCS has established a phased data analysis and visualization service development approach to incrementally address scientific user needs. This approach allows the flexibility for prioritizing capability enhancements with feedback from the user community. It allows scientists to continue with their current statistical model verifications using established tools such as Fortran, IDL, Matlab, or GrADS, while also forming the framework from which to explore new analytic paradigms for comprehensive scientific research.



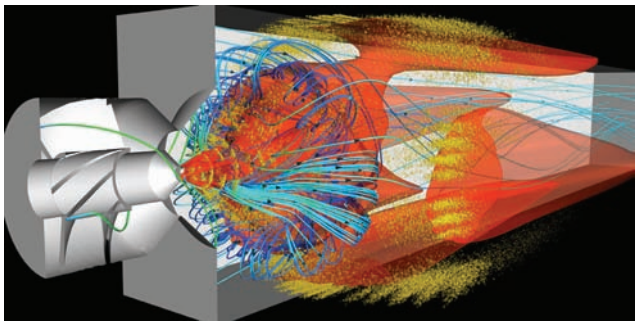
FUTURE MISSION CHALLENGES

As an Agency-wide resource, NASA's HEC Program will support strategic directions set by the new presidential administration to: advance global climate change research and monitoring; mount a robust program of space exploration involving humans and robots; support the safe flight of the Space Shuttle to complete assembly of the International Space Station; and renew NASA's commitment to aeronautics research. The NAS facility serves all four of the Agency's mission directorates, while NCCS focuses on the Science Mission Directorate. Below are brief summaries of anticipated mission directorate plans for the near-term use of HEC Program resources and services.

Aeronautics Research Mission Directorate (ARMD)

As ARMD's largest consumer of HEC resources, the Fundamental Aeronautics Program (FAP) will continue to fully utilize its shares of the Pleiades, RTJones, and Columbia supercomputers at NAS. FAP uses these systems to enhance development of physics-based multidisciplinary design, analysis, and optimization tools for evaluating radically new vehicle designs and assessing the potential impact of innovative technologies on overall vehicle performance. High-end computing is enabling FAP's long-term, cutting-edge research addressing the concerns of modern air transportation. This research includes: improving aircraft performance while reducing noise and emissions; eliminating environmental and performance barriers to practical supersonic vehicles; improving mobility to meet greater demand for air transportation; and providing technologies to enable enhanced future space exploration capabilities.

Under ARMD's Airspace Systems Program, NASA's NextGen Airportal Project is harnessing the Pleiades supercomputer to develop technologies that will maximize single-airport capacity and improve the efficiency of multi-airport operations while maintaining or enhancing safety. Among the factors limiting runway capacity are the large spacing distances currently required on final approach to ensure that all aircraft avoid wake vortices from other aircraft. Researchers are developing a "fast-time" model that will accurately predict the precise location, movement, and decay rate of these wake vortices and provide datasets as simulation input for wake detection sensors. Research groups in the Aviation Safety Program are using the Columbia supercomputer to develop Advanced Satellite Aviation-Weather products for predicting in-flight icing; and to model polymer-based construction materials at the atomistic scale as part of the Aircraft Aging and Durability Project.



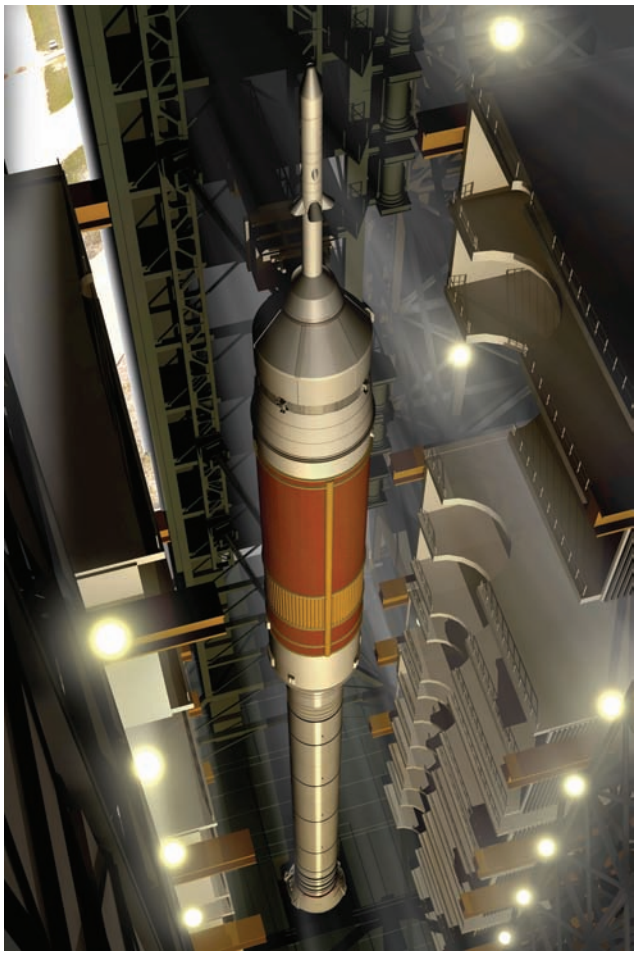
To help reduce harmful combustion emissions produced by aircraft, the Aeronautics Research Mission Directorate is assessing the effectiveness of software tools in predicting the presence of nitrogen oxide (NOx) and other emissions. The National Combustion Code was used to analyze the air flow through a lean-direct-injection combustor. Results show that the flow produces a recirculation zone that is critical to combustion stability but also produces NOx.

Exploration Systems Mission Directorate (ESMD)

In the coming years, HEC resources will continue assisting ESMD vehicle design, engineering, and mission planning. NASA supercomputers are supporting the Constellation Program's development of three next-generation space exploration vehicles: the Orion Crew Exploration Vehicle (CEV), the Ares I Crew Launch Vehicle (CLV), and the heavy-lift Ares V Cargo Launch Vehicle. As development ramps up, Constellation will need larger-scale simulations with finer grid resolutions and complex, time-accurate flow interactions to assess more intricate geometries and aerodynamic conditions.

Future support for CEV development will include aerothermal analyses for heatshield design, and assessments of design factors and abort conditions for launch abort system control. As the CLV design continues to be refined during later development stages, computational challenges will include aerodynamic analyses of: the vehicle's functional details such as fuel feed-lines, brackets, and umbilicals; guidance and attitude control system performance; and stage separation maneuvers with plume effects. Ares V aerodynamic analysis support will grow as the vehicle enters its major development phases. Computations will involve extensive aerodynamic database generation for each design cycle, as well as multi-species analyses of plume interaction and base heating effects for various engine configurations.

Additionally, the HEC Program has begun supporting ESMD's Lunar Precursor Robotic Program, which manages path-finding robotic missions to the Moon, leading the way



This artist's rendition shows the next-generation Ares I rocket being stacked in the Vehicle Assembly Building (VAB) at NASA's Kennedy Space Center in Florida. Simulations are helping to determine whether the VAB can safely manage combustion scenarios for the greater amounts of solid rocket booster fuel that Ares I and V will require.

to sustained human exploration of our solar system. The Lunar Crator Observatory and Sensing Satellite (LCROSS) mission—scheduled to launch with the Lunar Reconnaissance Orbiter in summer 2009—will determine whether water ice is present in a permanently shadowed crater at the Moon's south pole.

Science Mission Directorate (SMD)

Using Columbia and Pleiades at NAS and Discover at NCCS, SMD will analyze growing streams of data from new Earth-observing satellites and space missions. Across SMD's four divisions, models will need ever-higher spacial, temporal, and spectral resolution to match improvements in data resolution and to help NASA plan future observations.

Earth Science uses high-end computing to address problems in climate change and prediction. With greater resolution, NASA expects to better predict conditions such as hurricane intensity. Global climate simulations for the next Intergovernmental Panel on Climate Change assessment will include fully interacting atmospheric chemistry. Four-dimensional variational data assimilation is important for accurately incorporating precipitation observations into models and running Observing System Simulation Experiments for satellite design. Forecasting earthquakes and other solid Earth hazards requires daily simulations processing at least near-real-time, terabyte-sized data streams.

A high priority for Heliophysics is moving from today's 2D models towards 3D kinetic models of magnetic reconnection, a phenomenon that, among other effects, provides the energy release in solar eruptions. Forecasting space weather from initiation at the Sun to interaction with Earth and other solar system bodies means resolving each sub-domain and coupling domains across boundaries, with a minimum computational requirement equaling the sum of all domain calculations.

Planetary Science mission engineering employs high-fidelity space vehicle Entry, Descent, and Landing (EDL) simulations and landing site safety assessments. Expanding on techniques that have proven successful with the Phoenix Mars Lander, the Mars Science Laboratory (MSL) EDL simulations run four times as many particle trajectories at a higher frequency. Landing site safety analysis for MSL and a future Mars Sample Return mission entails hazard maps with 100 times the resolution used for Phoenix.

Astrophysics challenges include modeling the first stars and simulating how solar systems evolve from proto-planetary disks. Greater computing power is necessary for predicting gravitational waveforms from mergers of black holes with size differences of 20:1 to 100:1. Likewise, capturing star formation and other smaller-scale phenomena within cosmological simulations requires many more particles than is possible today.



New Earth-observing satellites and space missions being launched by NASA's Science Mission Directorate will generate more data at higher resolutions than ever before, presenting challenges for Earth and space science modelers. The Solar Dynamics Observatory shown above will help scientists better understand the Sun's influence on Earth and near-Earth space by studying the solar atmosphere on small scales of space and time and in many wavelengths simultaneously.

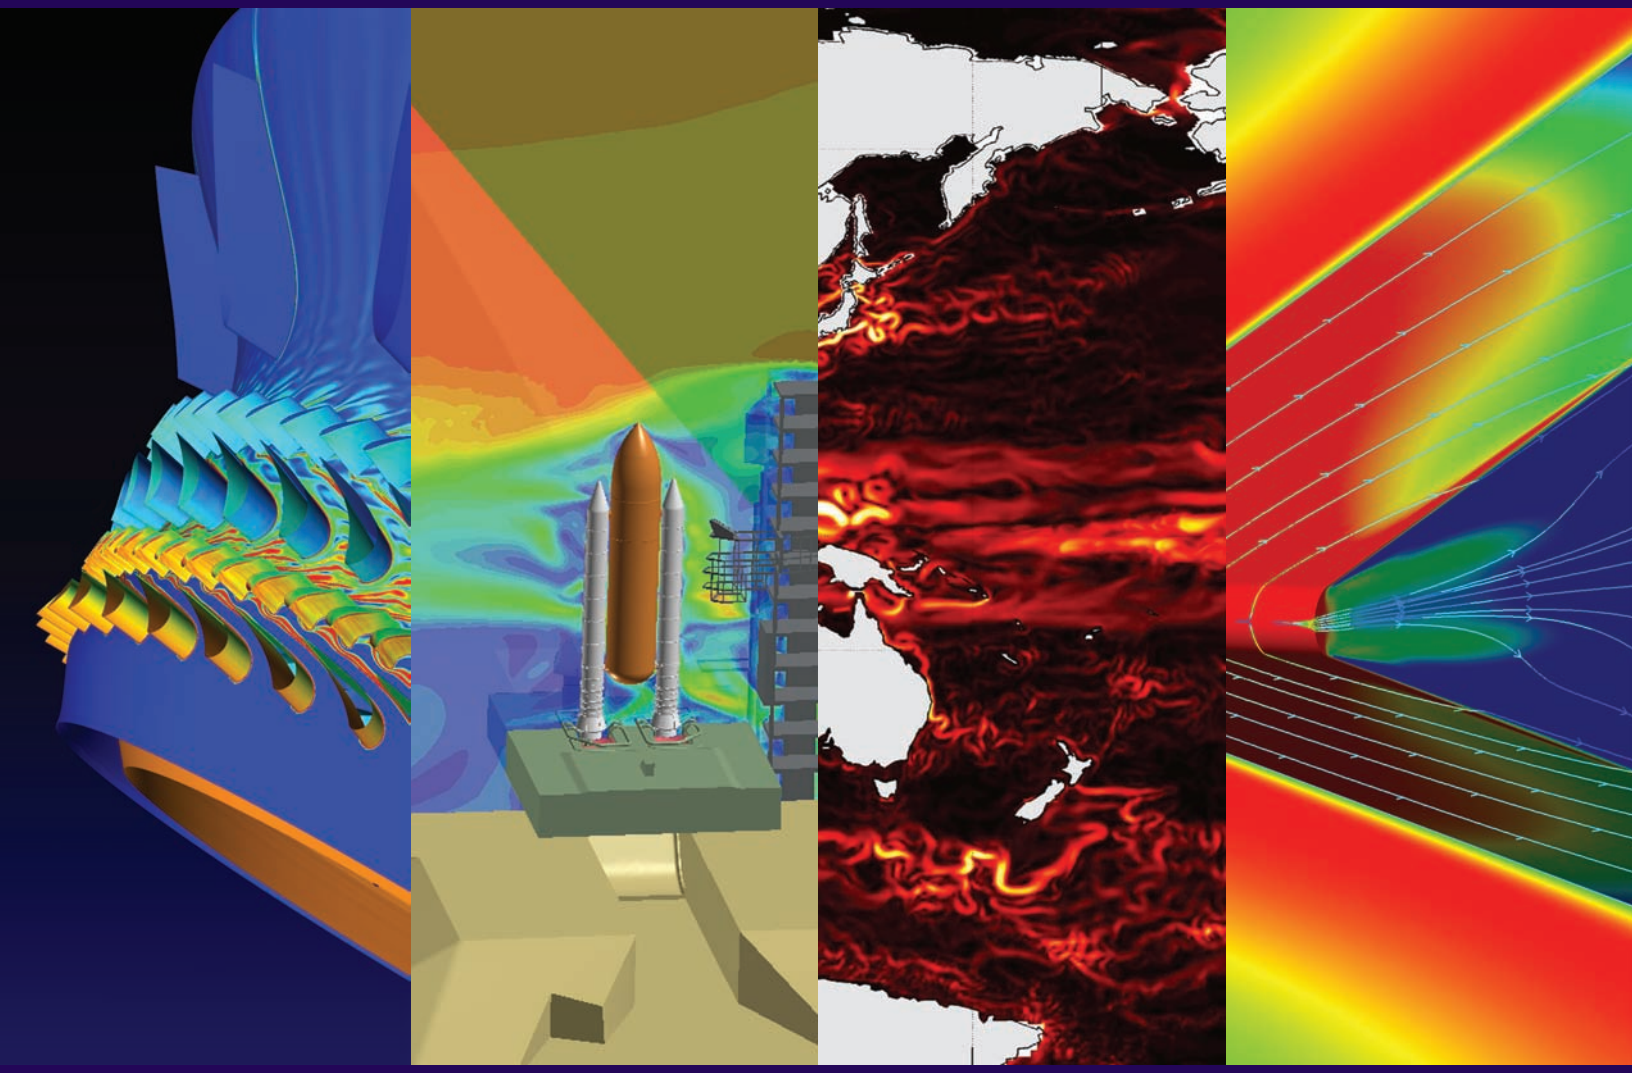
Space Operations Mission Directorate (SOMD)

With the approaching retirement of the Space Shuttle and a new generation of space vehicles on the way, the next few years are a pivotal time for SOMD. Throughout the remaining shuttle missions—currently planned through 2010—dedicated portions of NAS supercomputers and support staff must be on call from launch to landing to evaluate potential reentry risks posed by specific ice formation or foam debris damage sites, and to calculate heating on the thermal protection system to ensure a safe reentry. A mirrored Return to Flight data warehouse, co-located at NASA's Ames and Langley Research Centers, supports these on-the-fly calculations by rapidly transferring and disseminating hundreds of gigabytes of computational fluid dynamics data. SOMD will also rely on high-end computing resources for any future component redesigns or fuel tank processing improvements needed to keep the shuttle running safely and efficiently through its last missions. These efforts will become more important if the new administration opts to extend shuttle service beyond the planned retirement date.

Furthermore, high-end computing will contribute to the conversion of shuttle ground operations infrastructures for next-generation launch vehicles. For example, simulations are helping to determine whether the existing Vehicle Assembly Building at NASA Kennedy Space Center can safely manage potential combustion scenarios for the significantly greater amounts of solid rocket booster fuel that Ares I and V will require. Engineers are also performing analyses of shuttle and Ares ignition environments to determine whether modifications should be made to the current launch platform and flame trench configuration. Combining intensive, time-accurate plume and combustion simulations with large-scale models of ground operations infrastructures, such analyses present some of the most resource-demanding computations supported by NASA's High-End Computing Program.



This artist's rendition shows an Ares I rocket at Launch Pad 39B at NASA's Kennedy Space Center in Florida. Computational analyses of Space Shuttle and Ares ignition environments are helping NASA engineers to determine whether modifications to the current launch platform and flame trench configuration will be required to accommodate Ares.

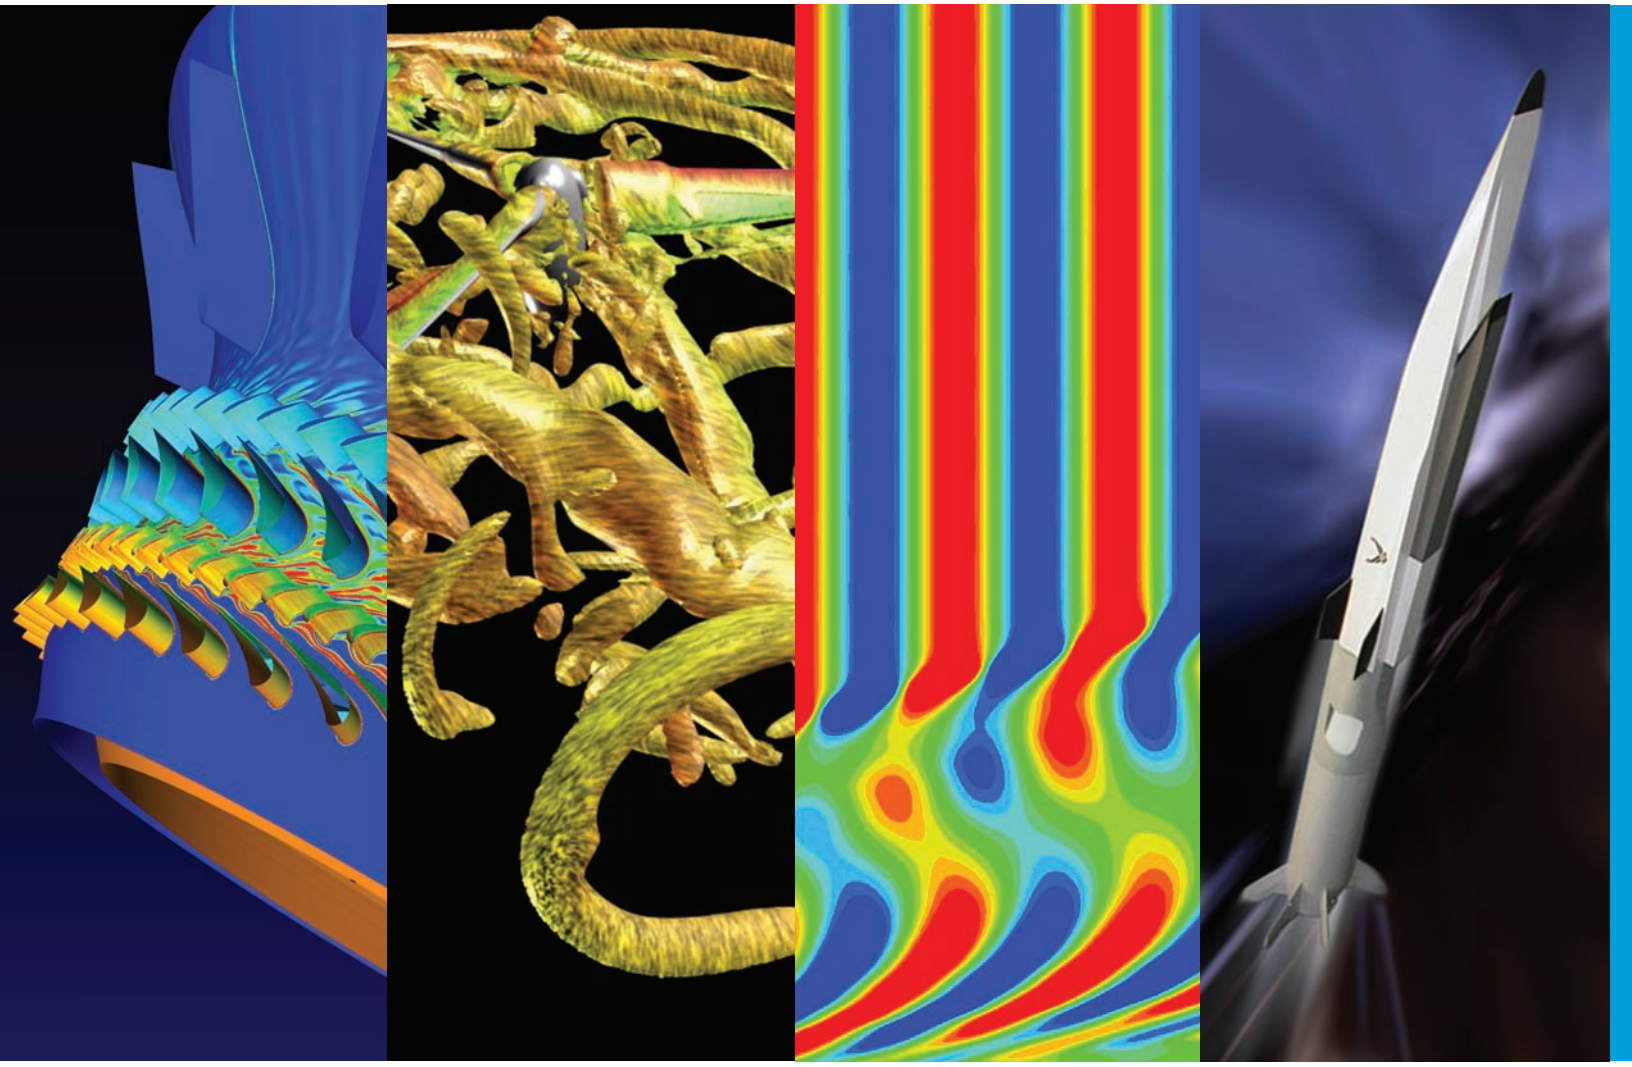


HIGHLIGHTS

SCIENCE AND ENGINEERING

This section presents 43 user projects from NASA's Aeronautics Research, Exploration Systems, Science, and Space Operations Mission Directorates, and the National Leadership Computing System Initiative, chosen because of their importance to the Agency, their impact during 2007 and 2008, and their technical maturity.

- **Aeronautics Research Mission Directorate** 21
- **Exploration Systems Mission Directorate** 39
- **Science Mission Directorate** 59
- **Space Operations Mission Directorate** 97
- **National Leadership Computing System** 111



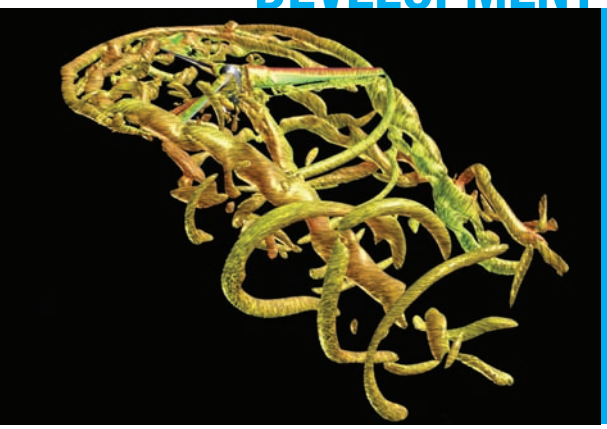
MISSION DIRECTORATE

AERONAUTICS RESEARCH

The Aeronautics Research Mission Directorate conducts cutting-edge, fundamental research in traditional and emerging disciplines to help transform the nation's air transportation system, and to support future air and space vehicles. Our goals are to improve airspace capacity and mobility, improve aviation safety, and improve aircraft performance while reducing noise, emissions, and fuel burn. Our world-class capability is built on a tradition of expertise in aeronautical engineering and its core research areas, including aerodynamics, aeroacoustics, materials and structures, propulsion, dynamics and control, sensor and actuator technologies, advanced computational and mathematical techniques, and experimental measurement techniques.

DR. JAIWON SHIN
Associate Administrator
http://www.aeronautics.nasa.gov

HIGH-RESOLUTION NAVIER-STOKES CODE DEVELOPMENT FOR ROTORCRAFT AEROMECHANICS



AERONAUTICS RESEARCH MISSION DIRECTORATE

NEAL M. CHADERJIAN
NASA Ames Research Center
(650) 604-4472
Neal.Chaderjian@nasa.gov

◀ **Figure 1:** Improved rotor tip vortex resolution using high-order spatial accuracy and grid adaption.

Project Description: Helicopters and tiltrotor aircraft provide many crucial services, including emergency medical and rescue evacuation, security patrols, offshore oil platform access, heavy-lift capability, and military operations. Some of the phenomena associated with rotorcraft flight include aerodynamic performance and noise, vortex wakes generated from the rotating blades, and rotor blade flexibility and vibration. Blade-Vortex Interaction (BVI) also occurs when a rotor blade interacts with, and in some cases slices through, the vortices generated by other rotor blades. This not only affects the aerodynamic performance of the vehicle, but it is responsible for much of the noise generated by the rotor blades.

Many of these phenomena are poorly understood and difficult to accurately predict. One of the goals of the Subsonic Rotary Wing (SRW) Project, part of NASA's Fundamental Aeronautics Program, is to develop improved physics-based computational tools to address these issues. The long-term objective of this effort is to develop a more accurate aeromechanics computational tool that couples computational fluid dynamics (CFD) and computational structural dynamics (CSD) flow simulation tools with a rotor blade trim code. A trim code prescribes blade motions such that the resultant forces and moments are in balance for a desired flight condition.

This project combines the efforts of several NASA Ames researchers into two broad categories: code development and application support. The latter uses current CFD/CSD/trim capability to support wind tunnel tests. This report focuses on the code development portion of this work, where the primary objective is to improve accuracy of the OVERFLOW-2 Reynolds-averaged Navier-Stokes (RANS) flow solver and to explore new methods of unsteady flow visualization.

Relevance of Work to NASA: This code development effort directly supports the SRW Project's goal to conduct long-term, cutting-edge research in the core competencies of the subsonic rotary wing regime. More specifically, the focus is on improving our prediction capability in rotorcraft aeromechanics through research and development of physics-based, high-fidelity computational tools. Specific project deliverables and milestone metrics are met by validating these new tools with wind tunnel measurements.

Computational Approach: Flow simulations have been carried out for an isolated V-22 Osprey rotor. This simple, rigid, three-blade rotor geometry, along with the spinner hub, is an ideal case for exploring new methods for improving simulation accuracy of the generated vortex wake. A Rotor Grid Assistant (RGA) script is used to automate the generation of overset structured computational grids and OVERFLOW-2 input files so that different grid resolutions can be readily explored. Current state-of-the-art CFD methods use algorithms that are second-order accurate in time, third-order accurate in space, and grid resolution in the vortex wake region with grid spacing on the order of one vortex core diameter. Total grid size typically consists of tens of millions of grid points. Straightforward refinement of the grid in the wake region (to achieve ten grid cells across a vortex core diameter) would result in a grid system consisting of several billion grid points. This is not practical with current supercomputers, so the approach adopted here is to improve OVERFLOW-2's spatial accuracy up to sixth order, and use grid adaption to locally improve the vortex wake resolution by four to six times. This will result in a grid system that consists of a few hundred million grid points rather than billions of grid points. Since rotor flows are inherently unsteady, new concurrent visualization methods have been used to identify flow structures and determine cause and effect for predicted quantitative values.

Results: Preliminary results have been obtained for a V-22 isolated rotor in hover with a 14° collective and a rotor tip Mach number of M_{tip} = 0.625. The OVERFLOW-2 CFD code was modified to include up to sixth-order spatial accuracy and grid adaption using higher-resolution Cartesian grids embedded into the uniform coarser Cartesian background grid.

A sensor function based on vorticity magnitude is used to identify vortex cores and position the embedded Cartesian grids to better resolve these vortices. Figure 2 shows the predicted vorticity magnitude contours using second- and fifth-order spatial accuracy. It is also apparent that the vortex strength is much stronger using the fifth-order method. Second-order methods cause too much dissipation and dispersion errors. Figure 2 also shows an even greater improvement in the vortex strength using a third-order accurate method with Cartesian grid adaption.

It is anticipated that combining grid adaption with high-order spatial accuracy will improve overall accuracy of the CFD code while controlling the computer time required for a solution. Figure 1 shows a concurrent visualization of the V-22 rotor in forward-descending flight, where the advance ratio is 0.1 and the angle of descent is 6°. The vortex wake is visualized using iso-surfaces based on the Q-criterion (the second invariant of the velocity gradient tensor). Strong BVI is evident as the rotor blades slice through the rotor tip vortices of other blades. The vortices also roll up, forming two super-vortices similar to wing-tip vortices found on fixed-wing aircraft. A time-dependent animation produced during this project reveals the complex nature of the flow. The NASA Advanced Supercomputing (NAS) Division's Visualization group added functionality to the OVERFLOW-2 flow solver so that visualization extracts were written out at every time step, while the solution was evolving in time. The old paradigm of writing out the grid and solution files every 20 or so time steps for later post-processing would result in tens of terabytes of disk usage.

A texture mapping technique is also employed to indicate the instantaneous velocity field on the vortex iso-surfaces.

Role of High-End Computing: The NAS facility provides state-of-the-art computational resources needed to address this compute-intensive problem. The forward flight case was run on the Columbia supercomputer using 20 million grid points, and required 11 hours of wall-clock time per revolution using 64 processors. A hover case was run on the Pleiades supercomputer using 150 million grid points, and required 17 hours of wall-clock time per revolution using 512 processor-nodes. A rotor simulation typically takes 10 revolutions from impulsive start conditions to achieve dynamic equilibrium.

Future: During the next 18 months, further improvements to the grid-adaption method will be implemented. Furthermore, a Subsonic Rotary Wing milestone will be met by coupling this improved CFD tool with a flexible rotor (CSD) code and a rotor-blade trim algorithm.

Co-Investigators

- Thomas Pulliam, Terry Holst, Jasim Ahmad, David Kao, Guru Guruswamy, Ethan Romander, and I-Chung Chang, all of NASA Ames Research Center

Publications

[1] Holst, T. L. and Pulliam, T. H., "Overset Solution Adaptive Grid Approach Applied to Hovering Rotorcraft Flows," To be presented at the 27th AIAA Applied Aerodynamics Conference, *AIAA Paper No. 2009-3519*, San Antonio, TX, June 22–25, 2009.

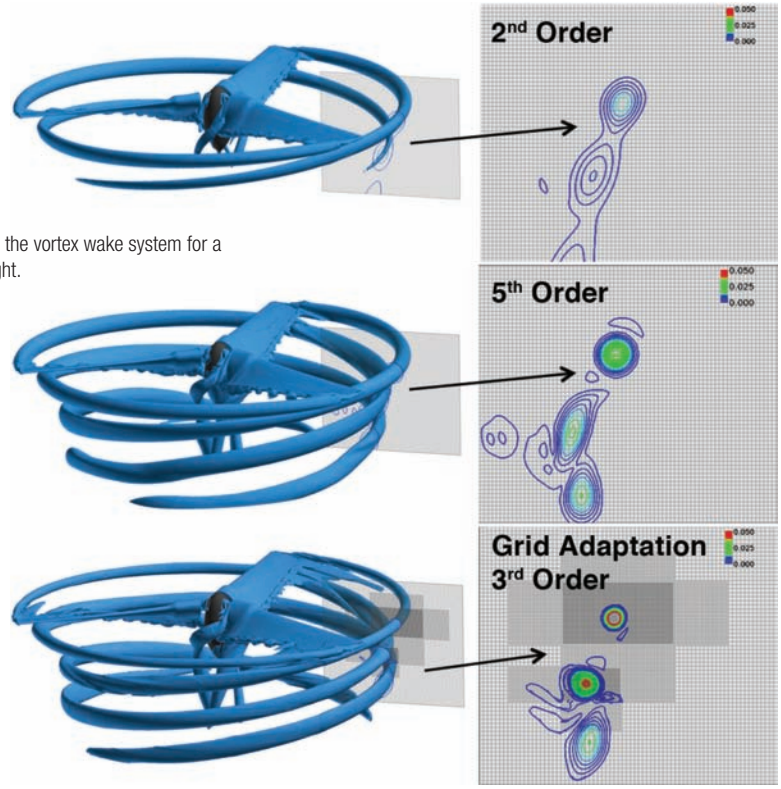
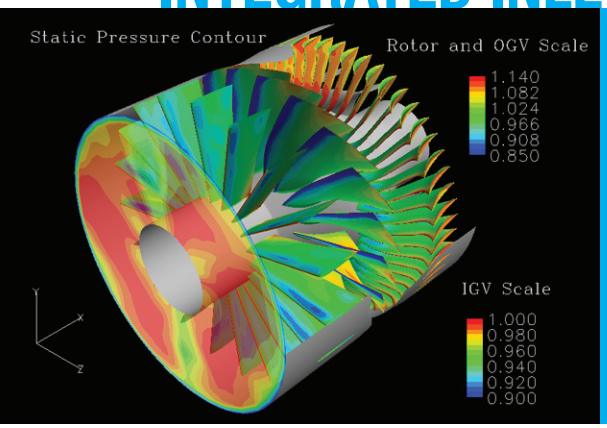


Figure 2: Concurrent visualization of the vortex wake system for a V-22 rotor in forward-descending flight.

INTEGRATED INLET/FAN SIMULATION (IISIM)



AERONAUTICS RESEARCH MISSION DIRECTORATE

MICHAEL HATHAWAY
Army Vehicle Technology Directorate
NASA Glenn Research Center
(216) 433-6250
Michael.D.Hathaway@nasa.gov

◀ **Figure 1:** Instantaneous "snapshot" in time from an unsteady, 3D, full-annulus TURBO simulation of the 1½ stage fan. The case shown here captures results of a simulation which includes a two-per rotor revolution total pressure distortion induced by two diametrically opposed rods placed in the flow upstream of the fan inlet. The simulation results were used to predict the unsteady blade pressure loading from which the blade response characteristics where calculated and compared to measurements.

Project Description: This project was a collaborative effort among NASA, the Army Vehicle Technology Directorate, Honeywell Aerospace, and AVETEC, Inc. (Springfield, OH), to support development of Advanced Virtual Test Cell computational tools for simulating operation of a complete gas turbine engine with multi-component and multidisciplinary interactions. Traditionally, gas turbine engine subsystems (e.g., inlet and fan) have been designed, analyzed, and tested as isolated components. With trends toward more compact, higher-power, higher-density engines and future vehicle concepts with embedded engines and more compact, low-observable inlets, significant flow distortions are generated that the fan/engine must accommodate. With engine components becoming more closely coupled, conventional means of accounting for component interactions may be inadequate. Consequently, decreased performance, stall margin, and even life of the fan (and engine) may result if interacting are not properly addressed.

Development and validation of a consistent computational methodology for integrated inlet-fan/engine simulations enables analysis of the interaction and component matching effects between the inlet and fan/engine. For example, the effect of the fan/engine flow distribution can be included in the inlet analysis, as it can directly affect the inlet performance if the components are not properly matched, especially during flight operations wherein the bypass ratio can change significantly. Flow distortions generated by the inlet or ingested by the inlet are passed to the fan/engine, affecting performance, stability, stage, component matching, and fan aeromechanical response. Aerodynamic response of the fan/engine to the inlet distortion conversely impacts the inlet. Inlet distortions can comprise not only total pressure distortions, but also thermal and swirl distortions, as well as constituent-based distortions such as those resulting from steam ingestion. All of these distortions can potentially be analyzed with a validated integrated inlet-fan/engine simulation capability to better understand the impact on the inlet-fan/engine system, thereby leading to design improvements.

The objectives of this work are to:

- Develop a capability to simulate an integrated inlet/fan geometry, both with and without flow control, by coupling state-of-the-art computational fluid dynamics (CFD) codes for the inlet (Wind-US) and fan (TURBO).
- Validate capability of the fan simulation to support fan aeromechanical response analysis using predicted aero forcing function and fan damping based on the results of a modal analysis of blade motion from an ANSYS analysis to prescribe blade deflections for the TURBO aeromechanical simulations.
- Demonstrate a capability to generate coupled inlet (Wind-US) and fan (TURBO) simulations of an integrated inlet/fan geometry, both with and without flow control.

Relevance of Work to NASA: Inlet/fan interactions have been identified as a challenging problem for all vehicle classes within the NASA Fundamental Aeronautics Program (FAP). In addition, FAP's research philosophy is to develop technology and capabilities to enable validated, multi-component, multidisciplinary analysis-leading simulations of complete engine and vehicle systems. This work was performed under the Subsonic Fixed-Wing Project and supports level 2, integrated methods and technologies to develop multidisciplinary solutions.

Computational Approach: The primary computational approach of this work was to couple the Wind-US and TURBO CFD codes. Both are capable of simulating unsteady, 3D, turbulent flows by solving the compressible Reynolds-Averaged Navier-Stokes equations. Wind-US was developed to simulate inlet flows, while TURBO was developed to simulate turbomachinery flows. The reason for this choice of codes is that they were specifically developed and validated for their respective components, and both solvers use similar numerical approximations and algorithms, which simplify the integration effort. Both solvers are launched simultaneously in simulating the integrated inlet/fan flow. Parallel execution of the two solvers requires an interface that allows transference of information between Wind-US and TURBO. This interface, called the Aerodynamic Interface Plane (AIP), lies at a convenient distance upstream from the fan. It divides the integrated

inlet/fan computational domain into an inlet domain and a fan domain. Wind-US computes flow through the inlet while TURBO computes flow through the fan. At the end of a time step, flow information at the interface is exchanged between the two solvers via Message Passing Interface (MPI) libraries. A simulation of a compressor with a long upstream duct was used to test the Wind-US and TURBO coupling capability.

A suitable inlet/fan geometry for which test data were available was selected for independent validation of Wind-US and TURBO predictions of the inlet and fan geometries. The geometry selected for simulation is a Lockheed inlet coupled to a Honeywell fan tested under the Air Force-sponsored Versatile Active Highly Integrated Inlet/Fan for Affordability Performance and Durability (VAIIPR) Program. The Lockheed inlet incorporated flow control to produce a uniform total pressure at the AIP between the inlet and fan. A known disturbance was then imposed at the AIP by inclusion of two diametrically opposed rods inserted in the flow path to generate a two-per revolution disturbance to the fan for which the blade response was measured (Figure 1). Fan performance and blade response were measured for two different orientations of the two-rod disturbance generator. Unsteady full-annulus TURBO simulations of the Honeywell fan stage were then conducted with the measured total pressure distortion resulting from both orientations of the two-rod disturbance prescribed at the fan inlet boundary, the AIP plane. Comparisons were made of the measured and predicted performance and blade response.

Results:

- Completed simulations of the baseline fan geometry at high- and low-speed conditions without the two-rod distortion.
- Completed simulations of the baseline fan geometry at low-speed conditions with the two-rod distortion at two orientations, with the second rotated 180 degrees from the first (Figure 1), and completed forced response analyses from the results.
- Completed coupling of the Wind-US and TURBO codes and subsequently demonstrated the code coupling capability on a compressor stage with an extended inlet (Figure 2).

Role of High-End Computing: These simulations would not have been possible without the computational resources provided by the NASA Advanced Supercomputing (NAS) Division. Due to contractual obligations with industry special resources (Toucan), high-priority queues (armd_spl) were provided to accomplish the large simulations (335 gigabytes of memory, 291 processors, 65 hours per rotor revolution) in a timely manner and to accommodate the significant storage requirements (1 terabyte of temporary storage each for three researchers) for multiple cases. Responsiveness of the NAS team in supporting our computational requirements was especially important in enabling adequate progress to meet our contractual obligations.

Future: This project was terminated due to reprioritization of research funding.

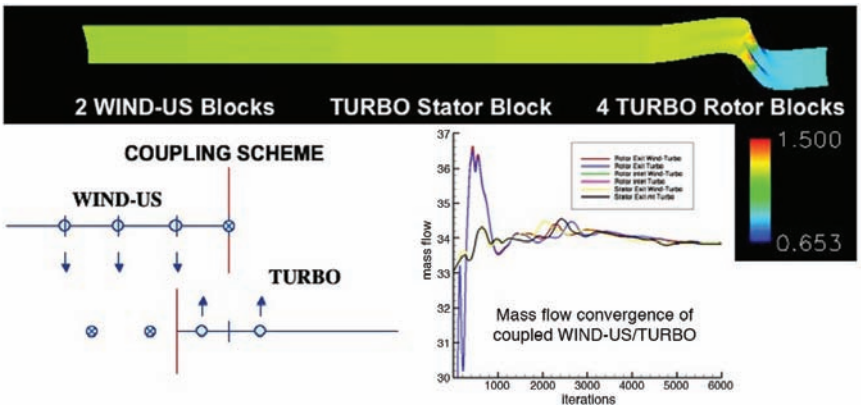
Co-Investigators

- Wai Ming To, University of Toledo
- Ambady Suresh, General Electric Company Global Research Center
- Rakesh Srivistava, Honeywell Engine Co.
- Milind Bakhle and John Lytle, NASA Glenn Research Center
- T. S. Reddy, University of Toledo
- Jeff Dalton, AVETEC, Inc.

Publications

[1] To, WM., Hathaway, M. D., "TURBO Simulation of a Honeywell Fan," to be published as a NASA TM with limited distribution.

[2] Reddy, T. S., Bakhle, M. A., "Forced Response Analysis of a Low-Aspect Ratio Fan," NASA internal publication.



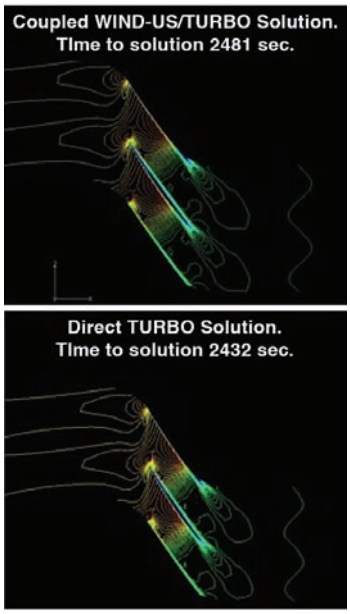
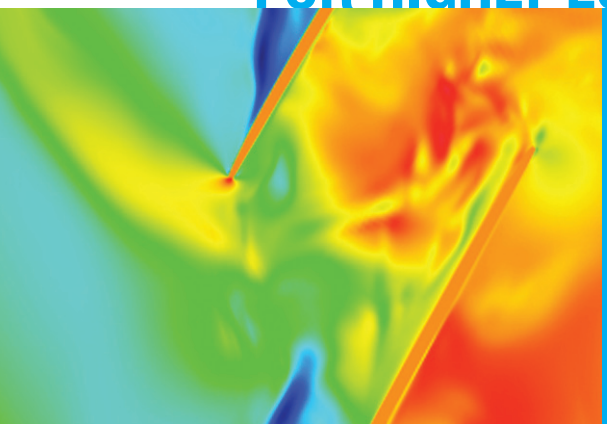


Figure 2: Demonstrated integrated inlet/fan simulation capability coupling of the Wind-US code, for simulation of an engine inlet duct, and the TURBO code for simulating a fan stage. Massflow convergence and comparison to direct TURBO simulation is shown.

LARGE EDDY SIMULATION FOR HIGHLY LOADED TURBOMACHINERY



AERONAUTICS RESEARCH MISSION DIRECTORATE

CHUNILL HAH
NASA Glenn Research Center
(216) 433-6377
Chunill.Hah-1@nasa.gov

◀ Close-up from Figure 1.

Project Description: Modern high-speed compressors operate with increasing aerodynamic blade loading. With this increased aerodynamic loading, it is critical to maintain a suitable stall margin to avoid engine stall during operation. Detailed unsteady flow structures in compressors operating near stall are not well understood, and it is generally believed that the conventional Reynolds-Averaged Navier-Stokes (RANS) approach does not predict the flowfield adequately. This fundamental aeronautics research project focuses on development and validation of a Large Eddy Simulation (LES) for this type of flowfield. The developed simulation tools can be applied to calculate detailed unsteady flow features so that advanced compressor designs can be studied to maintain a wide stall margin.

To date, the LES has been successfully applied to study self-induced flow instability and resulting non-synchronous blade vibration in axial compressors. The flow instability originates from interactions among tip vortex oscillation, vortex shedding, and passage shock (Figures 1 and 2). The simulation created with the LES tool clearly explains physics of flow instability in axial compressors for the first time. Further developments of the simulation are aimed at performing simulations of unsteady flow interactions between inlet and compressor stages with various flow control devices. A detailed understanding of the unsteady flowfield can contribute to better design of the compressor and any possible flow control devices.

Relevance of Work to NASA: This work is funded by NASA's Aeronautics Research Mission Directorate and supports NASA Fundamental Aeronautics Program goals to develop methods of subsonic fixed-wing simulations, specifically pertaining to inlet/fan interaction under supersonic conditions. The developed tool will be used to optimize both advanced fan design and flow control in the supersonic inlet/fan.

Computational Approach: A LES module has been integrated into the H3D code, a turbomachinery flow analysis code developed at NASA Glenn Research Center. The code employs a third-order accurate interpolation scheme for the convection terms and a central-differencing method for diffusion terms. A standard two-equation turbulence closure is used for the steady and unsteady RANS analysis. H3D is widely used throughout the U.S. aeronautics research community for the analysis of compressors and turbines. A Smagorinsky-type eddy-viscosity model is used for the sub-grid stress tensor for large eddy simulations. Dynamic models by Germano and Vreman are also implemented in the code. H3D has been parallelized for large-scale computations—up to 250 million grid-nodes have been used to simulate unsteady flowfields in an isolated compressor rotor.

Results: The H3D code was successfully applied to calculate combined flowfield in a fan stage with an ultra-compact inlet. The numerical results fairly accurately represent major flow characteristics including influence on stall margin.

The code was also applied to study flow instability in a transonic compressor at near-stall conditions. H3D calculated measured frequency of flow instability very well (Figure 3), and the underlying flow physics were explained with the simulated flowfield. This was the world's first-ever successful calculation of flow instability phenomena in a transonic compressor.

The code has been applied to investigate flowfield in a transonic rotor known as NASA Rotor 37. Numerical results from the LES show many improvements from the RANS simulations. It is thought that the LES module calculates flow interaction between the passage shock and tip leakage vortex much more realistically. Results from LES match the measured data very well, especially near the casing where flow interaction is strong.

Role of High-End Computing: The current LES of unsteady flowfields in highly loaded turbomachinery requires large-scale computations. Experts at the NASA Advanced Supercomputing (NAS) facility performed code optimization and parallelization on the H3D code to maximize computational efficiency on NAS supercomputers. Timely execution of the numerical analyses require parallel processing with many processors available.

Future: The H3D code will be extended to conduct unsteady flow simulations in multi-stage compressors and turbines. The code will be applied to simulate and develop optimum flow control strategy in a compact inlet/fan stage. Further improvement in computational efficiency is necessary to complete this work.

Co-Investigators

- Haoqiang H. Jin, NASA Ames Research Center

Publications

[1] Hah, C., Bergner, J., and P. Schiffer, "Tip Clearance Vortex Oscillation, Vortex Shedding and Rotating Instabilities in an Axial Transonic compressor Rotor," ASME paper, *GT2008-50105*, 2008.

[2] Hah, C., "Aerodynamic Study of Circumferential Grooves in a Transonic Axial compressor," ASME paper, *FED-55232*, 2008.

[3] Hah, C., "Self-Induced Flow Unsteadiness and Non-Synchronous Vibrations in Axial Compressors," ISOROMAC paper, *ISROMAC12-2008-20032*, 2008.

[4] Mueller, M., Schiffer, H., and Hah, C., "Interaction of Rotor and Casing Treatment Flow in an Axial Single-Stage Transonic Compressor," ASME paper, *GT2008-50135*, 2008.

[5] Hah, C., Bergner, J., and Schiffer, P., "Short Length-Scale Rotating Stall Inception in a Transonic Axial Compressor," ASME paper, *GT2006-90045*, 2006.

[6] Hah, C. and Lee, Y., "Unsteady Tip Leakage Vortex Phenomena in a Ducted Propeller," *International Journal of Transport Phenomena*, Vol. 9, No. 3, 169–176, 2007.

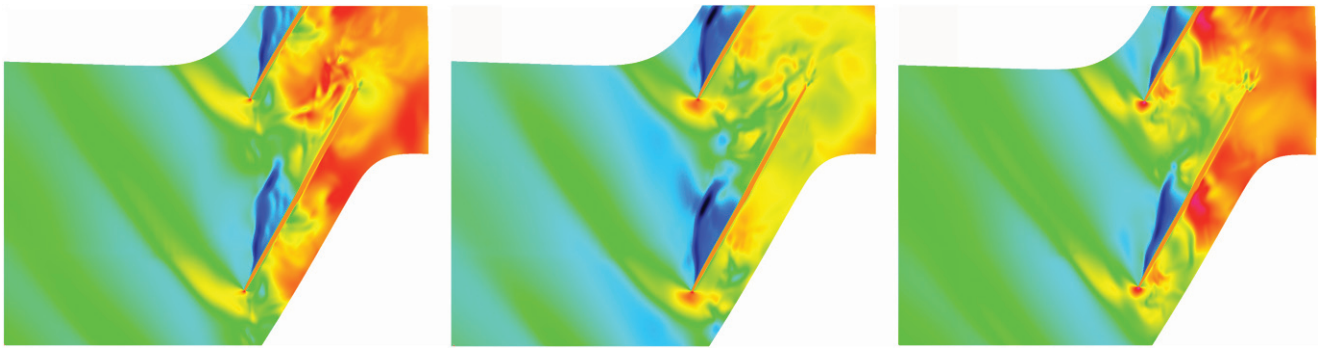


Figure 1: A Large Eddy Simulation (LES) calculates the oscillation of tip clearance vortex, as the measurements show. The synchronized tip vortex oscillation across several blade passages creates flow instability that can cause structural failure of the blade at certain operating conditions.

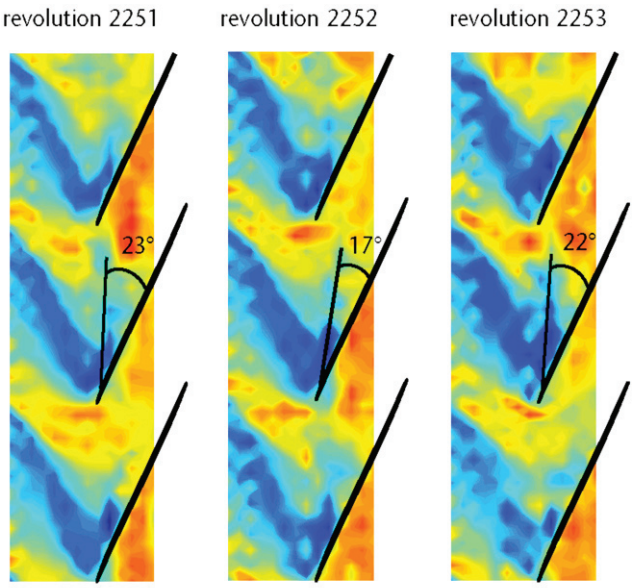


Figure 2: This measured instantaneous pressure field on the casing reveals oscillation of tip clearance vortex, due to interaction between this vortex and the passage shock.

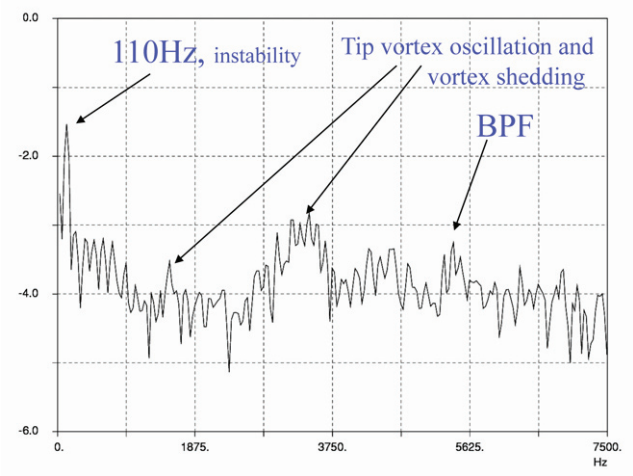
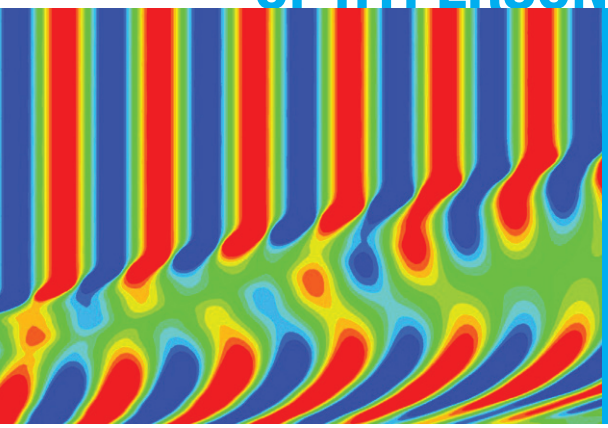


Figure 3: Wall pressure spectrum from LES shows dominant frequencies. The 110 Hz represents flow instability due to flow interaction between the passage shock and tip clearance vortex. The calculated frequency agrees well with the value from measurement.

RECEPTIVITY AND STABILITY OF HYPERSONIC BOUNDARY LAYERS



AERONAUTICS RESEARCH MISSION DIRECTORATE

PONNAMPALAM BALAKUMAR
NASA Langley Research Center
(757) 864-8453
Ponnampalam.Balakumar-1@nasa.gov

◀ Close-up of Figure 2.

Project Description: Accurately predicting transition onset and transition end-points, modeling this transitional region, and modeling the turbulence region are major challenges in accurately computing the aerodynamic quantities using computational fluid dynamics codes. The transition process depends primarily on the boundary layer characteristics and on the frequency and wave number distributions of the disturbances that enter the boundary layer. The difficulty is computing, predicting, or prescribing the initial spectral, amplitude, and phase distribution of the disturbances inside the boundary layer. In any new transition prediction strategy, one should quantify these two quantities and determine the minimum amount of information necessary to predict the transition onset accurately. The objectives of this research are to overcome some of these difficulties, and to eventually come up with an improved transition prediction method. Accurate transition onset prediction will help compute the heating and skin friction loads on the vehicle accurately and will improve design of thermal protection systems and structural components.

To understand and quantify receptivity coefficients and stability characteristics of hypersonic boundary layers, interaction of acoustic waves with hypersonic boundary layers over sharp and blunt flat plates, wedges, and cones were numerically simulated at different freestream and wall conditions. The importance of slow and fast acoustic waves, unit Reynolds number effects, the bluntness effects, and wall cooling effects on the receptivity and stability were systematically investigated. The receptivity coefficients, stability properties, and the transition Reynolds numbers were obtained for different cases.

Relevance of Work to NASA: One of the NASA Fundamental Aeronautics Program's objectives is to develop physics-based models to predict transition in hypersonic flows. Understanding the transition process from the first principle will lead to improved predictive capabilities in flows over hypersonic vehicles. NASA's interest in space exploration requires development of vehicles that fly through the hypersonic regime. Efficient, reliable, and reusable hypersonic vehicles will benefit NASA's space exploration mission.

Computational Approach: Three-dimensional compressible Navier-Stokes equations are solved using the fifth-order-accurate weighted essentially non-oscillatory (WENO) scheme for space discretization, and using the third-order total-variation-diminishing (TVD) Runge-Kutta scheme for time integration. These methods are suitable in flows with discontinuities or high-gradient regions. After the steady mean flow is computed, acoustic disturbances are superimposed at the outer boundary of the computational domain and time-accurate simulations are performed.

Results: A number of fundamental studies have been numerically performed to evaluate the effects of different parameters such as slow and fast acoustic waves; bluntness; wall cooling on the receptivity; and stability of hypersonic boundary layers over plates, wedges, and cones. Findings include:

- *Receptivity of Hypersonic Boundary Layers Over Cones and Wedges to Acoustic Disturbances:* The receptivity coefficient of the instability waves generated by the slow acoustic wave is about four times the amplitude of the freestream acoustic waves. The amplitude of the instability waves generated by the slow acoustic waves is about 60 times larger than that for the case of fast acoustic waves (Figure 1).
- *Effects of Nose Bluntness on Receptivity and Stability of Hypersonic Boundary Layers Over Cones:* The bluntness has a strong stabilizing effect on the boundary layers. This is due to the entropy layers that persist for longer distances with increasing bluntness. The receptivity coefficients for large bluntness are much smaller—on the order of 10^{-3} (Figure 2).
- *Effects of Wall Cooling on Receptivity and Stability of Hypersonic Boundary Layers Over Cones:* Wall cooling stabilizes the first mode and destabilizes the second mode and shifts the

transition onset further upstream. The fast mode is not affected by the wall cooling, and the receptivity coefficient of the fast wave is about 50 times larger than for the slow wave (Figure 3).

Role of High-End Computing: Performing a time-accurate simulation of receptivity and stability processes in hypersonic boundary layers is computationally demanding. One computation for a two-dimensional case takes about one week of computer time on 24 processors. Performing parametric studies for several cases requires Columbia supercomputer resources to obtain the results in a reasonable amount of time.

Future: Continuation of this work includes simulating the transition process in three-dimensional hypersonic boundary layers such as flow over cones at angles of attack and ellipsoids. Another goal is to simulate the roughness-induced transition in hypersonic boundary layers. These are three-dimensional simulations and require months of computational time for one case on 64 processors.

Co-Investigators

- Kursat Kara and Osama A. Kandil, Dept. of Aerospace Engineering, Old Dominion University

Publications

[1] Kara, K., Balakumar, P., and Kandil, O. A., "Effects of Wall Cooling on Hypersonic Boundary Layers Receptivity over a Cone," 38th AIAA Fluid Dynamics Conference and Exhibit, Seattle, WA, *AIAA 2008-3734*, 2008.

[2] Kara, K., Balakumar, P., and Kandil, O. A., "Effects of Nose Bluntness on Stability of Hypersonic Boundary Layers Receptivity over a Blunt Cone," 37th AIAA Fluid Dynamics Conference and Exhibit, Miami, FL, *AIAA 2007-4492*, 2007.

[3] Kara, K., Balakumar, P., and Kandil, O. A., "Receptivity of Hypersonic Boundary Layers Due To Acoustic Disturbances over Blunt Cone," 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, *AIAA 2007-0945*, 2007.

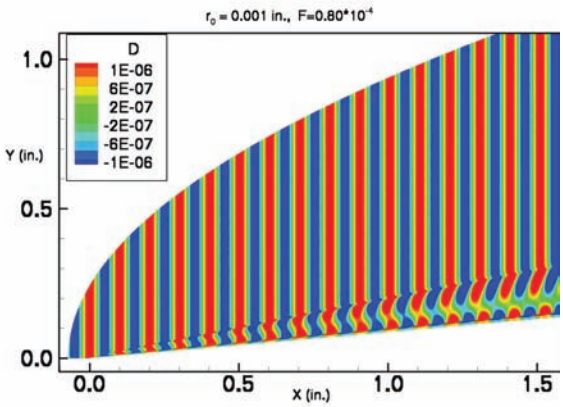


Figure 1: Contours of the unsteady density fluctuations due to the interaction of a slow acoustic wave over a 5-degree half-angle sharp cone at a Mach number of M = 6.

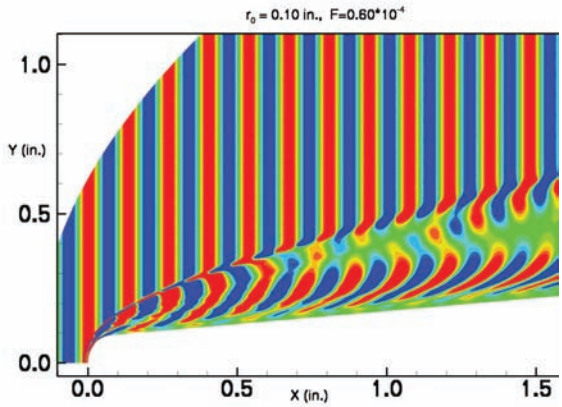


Figure 2: Contours of the unsteady density fluctuations due to the interaction of a slow acoustic wave over a 5-degree half-angle blunt cone.

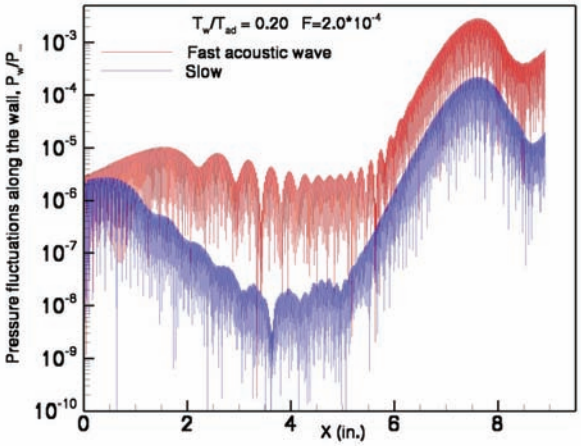
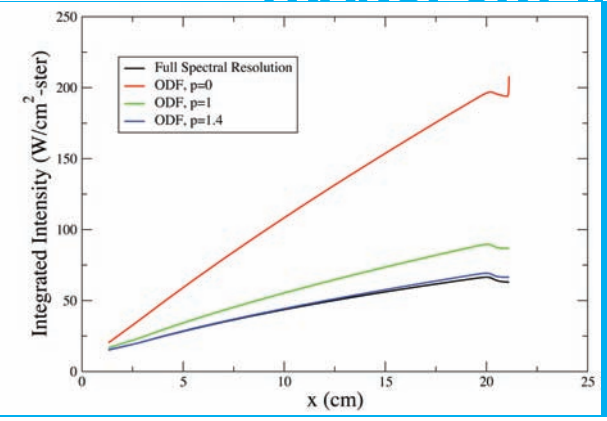


Figure 3: Unsteady pressure fluctuations along the wall in a log scale generated by the slow and fast acoustic waves in the cold wall case.

TOWARD IMPROVED RADIATIVE TRANSPORT IN HYPERSONIC REENTRY



AERONAUTICS RESEARCH MISSION DIRECTORATE

ALAN WRAY
NASA Ames Research Center
(650) 604-6066
Alan.A.Wray@nasa.gov

◀ **Figure 1:** Comparison of simulations of radiation along the peak radiative heating ray in the reentry of a crew exploration vehicle type of body into Earth's atmosphere. The black curve required 20 million points in wavelength space, and the blue curve only 10 million points, using an opacity distribution function method that will be incorporated into the HyperRad program.

Project Description: Large bodies reentering planetary atmospheres undergo substantial heating from radiation produced by high-temperature gas that results as the body decelerates. This heating must be accounted for in designing spacecraft to be both safe and efficient in carrying out their missions. Current software to compute such effects is hampered by large uncertainties resulting from various approximations in both physical modeling and numerical methods. These approximations were necessary in an earlier era of computational power, but now can be largely removed.

Our new software program, dubbed HyperRad, will bring the computation of radiative effects in hypersonic flow to a new level of accuracy and quantifiable uncertainty consistent with current computer hardware. Design robustness and confidence levels will thereby be increased and costs associated with wind tunnel and flight testing will be reduced.

Relevance of Work to NASA: Spacecraft design engineers must balance the requirements of safe mission completion with payload size and mass. This means that the weight and design of the thermal protection system (TPS) must be optimized to ensure integrity of the spacecraft on reentry without an excessive allocation of mass. Such designs require both experimental and computational studies to ensure accurate engineering specifications. Traditional fluid mechanics simulations must be augmented by radiative transport effects in order to gain complete knowledge of the thermal environment of the spacecraft.

Computational Approach: Several ongoing computational efforts are required to accomplish the project goals. First, accurate chemical databases must be constructed to obtain the radiative emission and absorption properties of the fluid, as functions of its composition, history, and thermodynamic state. These databases are obtained both from experiment and from quantum and classical computational algorithms. Second, these databases are being put in a form that can be efficiently utilized at runtime of the HyperRad radiative transport code (Figures 1 and 2). Finally, HyperRad will be optimized to compute the radiation throughout the flowfield with enough speed that the total runtime of the simulation is not greatly increased from that of non-radiative cases.

Results: A large database of molecular and atomic states, energy levels, lifetimes, collision cross sections, and coefficients required for radiative transport has been assembled from the literature and from massive new quantum chemistry calculations done on the Columbia supercomputer. The database includes information about the ground and excited states of molecular, atomic, and ionic species found in the atmospheres of Earth and Mars. This data forms the foundation for the tables required for efficient simulation of the radiation. These run-time data structures have been designed and are being programmed into Fortran 90. Other results include the following:

- Techniques for reducing the cost of carrying large numbers of points in wavelength-space have been designed and tested by comparison to existing experimental and simulation data.
- Techniques to allow efficient transport of radiation over many angular directions have been tested and compared to theoretical results and to existing approximations.
- Effective parallelization of these methods has been designed and tested.

Role of High-End Computing: The Columbia supercomputer has been essential for the quantum chemical calculations used to update and extend chemical databases required to complete the development of HyperRad. It is also on the high-end computing (HEC) systems that large-scale radiation-hydrodynamic simulations of spacecraft reentering the atmosphere will ultimately be done; thus, all program design considerations are based on such hardware. Testing of

the components in HyperRad has been and will continue to be done on the Columbia and Pleiades systems.

Future: We are moving into the testing phases of the most computationally intensive portion of the radiative transport algorithm. Columbia, Pleiades, and future HEC systems will be essential to completing this project in a timely manner. Large computational meshes will be required for the computational fluid dynamics component, and equally fine, but different, meshes will be used for the radiation computation. Conversion of data between these two mesh classes will also be an important programming and execution challenge in the coming year.

Co-Investigators

- David Schwenke, Richard Jaffe, Yen Liu, Duane Carbon, Galina Chaban, all of NASA Ames Research Center
- Winifred Huo, Huo Consulting LLC
- Dinesh Prabhu, NASA Ames, ELORET
- Thierry Magin, Stanford University

Publications

[1] Jaffe, R., Schwenke, D., and Chaban, G., "Theoretical analysis of N2 collisional dissociation and rotation-vibration energy transfer," AIAA Aerospace Sciences meeting, January 2009.

[2] Jaffe, R., Schwenke, D., Chaban, G., "Theoretical analysis of N2 collisional dissociation and rotation-vibration energy transfer," AIAA Aerospace Sciences meeting, January 2009.

[3] Magin, T., Caillault, L., Bourdon, A., and Laux, C., "Nonequilibrium radiative heat flux modeling for the Huygens entry probe", *Journal of Geophysical Research*, 111, 2006.

[4] Huo, W., "Electron Recombination and Collisional Excitation in Air," Lecture Series, von Karmen Institute for Fluid Dynamics, 2008.

[5] Bourdon, A., Panesi, M., Brandis, A., Magin, T., Chaban, G.,Huo, W., Jaffe, R., and Schwenke, D., "Simulation of flows in shock-tube facilities by means of a detailed chemical mechanism for nitrogen excitation and dissociation," Proceedings, Stanford-NASA Center for Turbulence Research, 2008.

[6] Graille, B., Magin, T., and Massot, M., "Modeling of reactive plasmas for atmospheric entry flows based on kinetic theory," Proceedings, Stanford-NASA Center for Turbulence Research, 2008.

[7] Wray, A., Prabhu, D., and Ripoll, J-F., "Opacity Distribution Functions Applied to the CEV Reentry," AIAA Thermosciences Meeting, 2007.

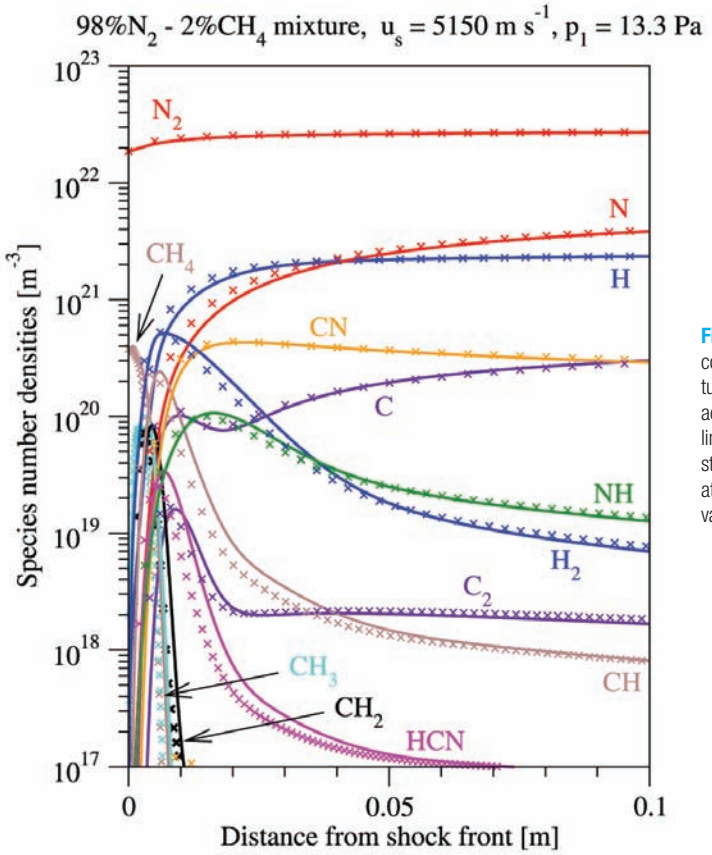
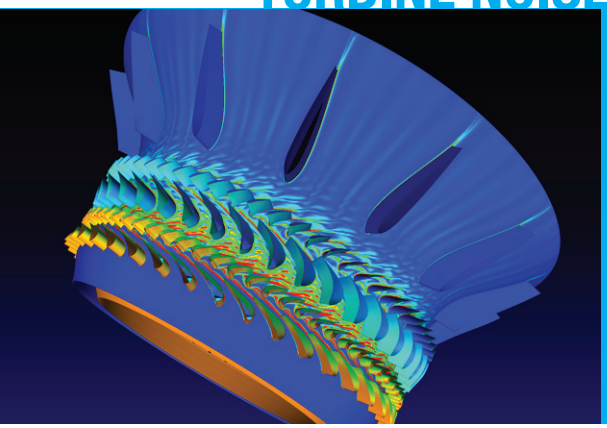


Figure 2: Results (Magin et al., 2006) comparing an Ames simulation of a shock-tube experiment (solid lines) and a time-accurate collisional-radiative model (xxxxxx lines) for species number densities downstream of a reentry shock in a Titan-like atmosphere. Such simulations are used to validate the models used in HyperRad.

TURBOMACHINERY AEROACOUSTICS: TURBINE NOISE GENERATION IN TURBOFAN ENGINES



AERONAUTICS RESEARCH MISSION DIRECTORATE

DALE VAN ZANTE
NASA Glenn Research Center
(216) 433-3640
Dale.E.VanZante@nasa.gov

◀ **Figure 1:** Instantaneous view of the five-blade-row, high-pressure turbine coarse mesh simulation with flow colored by vorticity to show velocity non-uniformities. The blade surfaces are colored by static pressure. The interaction of velocity non-uniformities with the blade surfaces is one noise generation mechanism in turbines.

Project Description: As aircraft engine bypass ratios continue to increase, the relative contribution of the turbine to the overall aircraft noise signature increases. Due to the dominance of fan and jet sources in moderate- and high-bypass ratio engines, turbine noise has typically been ignored, resulting in minimal effort focused on noise reduction technology development for turbines. Additionally, the desire for lower-cost, lower-weight, and higher-performance engines has resulted in turbine design changes that have typically increased turbine noise.

Understanding the relative importance of various turbine noise generation mechanisms and the characteristics of turbine acoustic transmission loss are essential ingredients in developing robust models for predicting the turbine noise signature. Typically, turbine noise models in use today are semi-empirical in nature and not suitable for detailed design and analysis studies. A computationally based investigation has been undertaken to help guide development of a more robust turbine noise prediction capability that does not rely on empiricism and is capable of addressing general design changes.

The Acoustics Discipline of NASA's Fundamental Aeronautics Subsonic Fixed Wing Project is pursuing technologies to reduce aircraft noise, with the ultimate goal of containing objectionable noise from aircraft to within the airport boundary. Aircraft noise is an amalgam of propulsion and airframe sources whose relative contributions depend on the aircraft type and operating condition. Generally speaking, propulsion noise (that is, engine noise) is a significant contributor to the total aircraft noise signature. Of the various sources of engine noise, fan and jet sources have received much attention in the past, but with the advent of ultra-high-bypass ratio engines (like the Pratt & Whitney geared turbofan), turbine noise is emerging as an important source of noise that must be mitigated in future low-noise propulsion systems. The Turbine Noise Project's objective is to develop an understanding of how and where noise is produced inside the turbine. Based on this knowledge, noise models can be formulated that will aid in evaluation of new engine designs and in the development of turbine noise mitigation technologies.

Relevance of Work to NASA: NASA is responsible for maintaining and improving an aircraft system noise prediction code called Aircraft NOise Prediction Program (ANOPP). This code is used by the Federal Aviation Administration, and engine and aircraft manufacturers to assess the impact of noise from contemporary aircraft on communities, as well as to evaluate how changes in design of the aircraft system will alter noise impact. Results from the Acoustics Discipline research activities are used to continuously improve models used in the ANOPP tool.

Computational Approach: The NASA turbomachinery aerodynamics solver TURBO is used to calculate the time-varying pressure field inside a turbine (Figures 2 and 3). A portion of this unsteady pressure field will propagate through the turbine blade rows and will emerge as noise from the engine exhaust. Frequency, modal content, and other characteristic information may be extracted from the pressure data by postprocessing and then used to construct reduced-order models for noise generation and propagation. To properly capture important pressure wave features inside the turbine, the numerical mesh must be about ten times denser than what is typically used for aerodynamic performance calculations. Such dense meshes, required by the wave propagation physics, result in very large computational resource requirements.

Results: The turbine noise work supports the Subsonic Fixed Wing Project goal of developing and/or improving and validating the next-generation multi-fidelity component and aircraft noise prediction capability. Accomplishments to date include:

- A converged time-accurate solution for a single turbine stage (two blade rows) with a mesh density appropriate for capturing twice the blade-passing-frequency (2xBPF) tone has been completed. The computational domain contains 80

million nodes and the solution dataset is 965 gigabytes in size. The simulation required 11 days of run-time on 200 processors of the RTJones supercomputer at NASA Ames Research Center to converge.

- Spectral and modal analysis of unsteady pressure data from the single-stage simulation shows that the simulation captures the anticipated acoustic modes properly. These results have been reported to joint working groups involving industry, academia, and government.

Role of High-End Computing: The availability of computing resources including large numbers of processors, high-speed networks, and parallel visualization capabilities have all been enabling technologies for the turbine noise effort. Prior to recent additions to NASA high-end computing resources (namely RTJones), a simulation of this scope would not have been practical.

Future: With the proof-of-concept phase nearly complete, the next phase is to complete a simulation of an entire aircraft engine high-pressure turbine (HPT) consisting of five blade rows. The additional solution domain will permit a more comprehensive assessment of the upstream and downstream traveling acoustic waves within the turbine. Preliminary coarse mesh simulations of the HPT are already complete and will help guide setup for the fine mesh simulation (Figure 1). The fine mesh case will require 200 million nodes and 600 processors to converge the solution in a reasonable period of time.

Co-Investigators

- Edmane Envia, NASA Glenn Research Center

Publications

[1] Van Zante, D. and Envia, E, "A Numerical Investigation of Turbine Noise Source Hierarchy and Its Acoustic Transmission Characteristics," Invited presentation at AeroAcoustics Research Consortium Turbine Noise Workshop, Vancouver, BC, May 2008.

[2] Van Zante, D. and Envia, E., "A Numerical Investigation of Turbine Noise Source Hierarchy and Its Acoustic Transmission Characteristics: Proof-of-concept progress," Acoustics Technical Working Group, Williamsburg, VA, Sept. 2008.

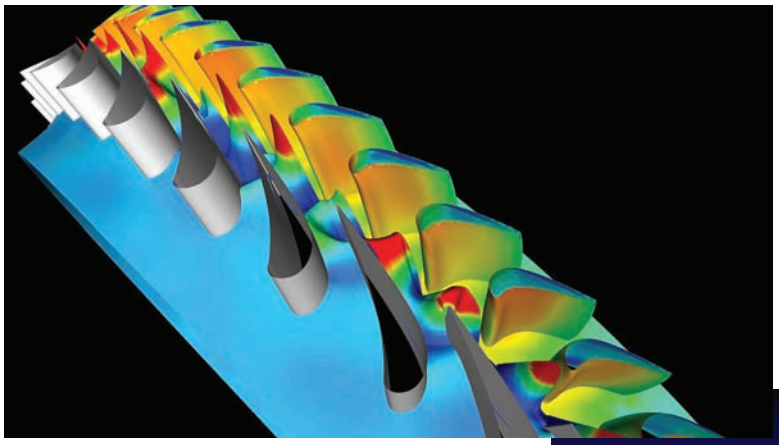


Figure 2: Pressure wave formation: In this view of the first stage of the turbine, the rotor (colored by static pressure) cuts through the wake of the vane (gray blade row) and forms a series of pressure waves shown by the color shaded "wavy" surface. These pressure waves propagate upstream and downstream from the rotor. Only the upstream propagating wave is shown here. A portion of this fluctuating pressure will emerge from the engine nozzle as noise.

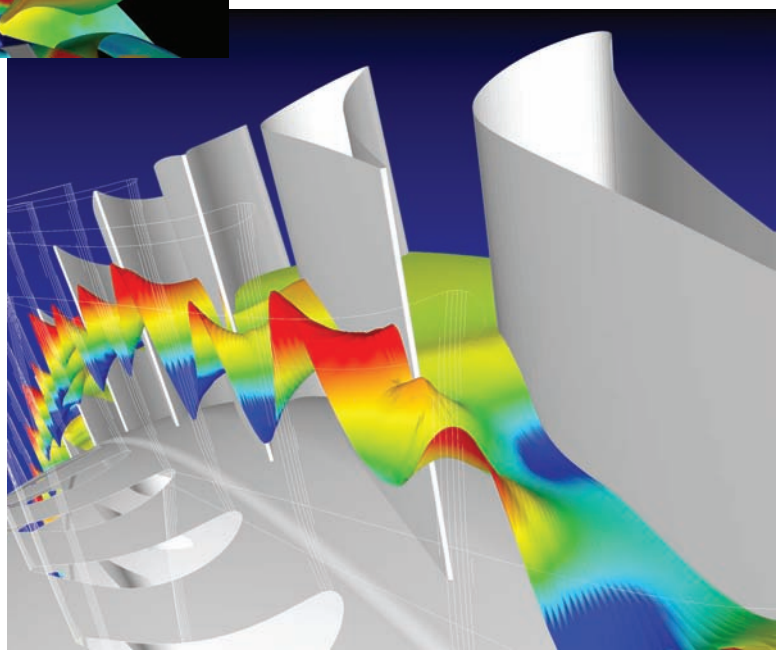
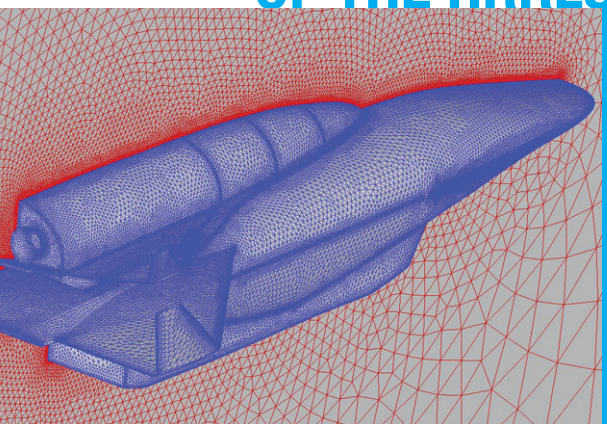


Figure 3: View of the pressure waves propagating upstream through the vane. The wave is attenuated as it moves upstream against the high subsonic Mach number flow coming through the vane passage.

USM3D ANALYSIS OF THE HRRLS CONFIGURATIONS



AERONAUTICS RESEARCH MISSION DIRECTORATE

JENN LOUH PAO
NASA Langley Research Center
(757) 864-3765
Jenn.L.Pao@nasa.gov

◀ **Figure 1:** Surface grid of the Highly Reliable Reusable Launch Systems configuration.

Project Description: The leading roles of the Multi-Disciplinary Analysis Optimization (MDAO) team within the Hypersonics Project of NASA's Fundamental Aeronautics Program are to develop and analyze reference vehicle concepts to determine potential system capabilities, and to establish research and technology goals and requirements. One of the primary reference missions for MDAO is the Highly Reliable Reusable Launch Systems (HRRLS). The objectives of the MDAO system studies for the HRRLS mission are to provide: reference concepts for project disciplines to analyze/exercise tools and apply technologies; a means to exercise and evaluate MDAO tool development progress; and reference concepts for technology assessment and investment guidance. The specific objective of this task is to compute a longitudinal aerodynamic database for the HRRLS reference vehicle, from subsonic to hypersonic Mach numbers. This database was used in the trajectory analysis and optimization for the HRRLS mission.

To develop and analyze the HRRLS reference vehicle concepts, extensive computational fluid dynamics computations were performed for both mated and first stage configurations from subsonic to hypersonic speeds. These computations include: preliminary coarse grid solutions; grid convergence studies to determine final grids with adequate grid density; final grid solutions; and analysis of grid density impact on the longitudinal force and moment coefficients.

Relevance of Work to NASA: The work presented here is closely aligned with two of the Aeronautics Research Mission Directorate's primary goals: to gain advanced knowledge in the fundamental discipline of aeronautics, and to develop access to space technologies for safer long-range, high-speed aerospace transportation systems.

Computational Approach: In the HRRLS reference vehicle study, a longitudinal unpowered aerodynamic matrix was created for both mated and first stage configurations. The Tetrahedral Unstructured Software System (TetrUSS) and Navier-Stokes flow solver USM3D, developed at NASA Langley Research Center, were used to generate unstructured grids and compute the aerodynamic matrix. The surface grid for the HRRLS mated configuration is depicted in Figure 1. The final aerodynamic matrix contains more than 100 points over a Mach number range of 0.5–8.0, with an angle of attack range of -2–12 degrees to cover the three sigma trajectory dispersions. The flow simulation was performed using full viscous calculations with the Spalart-Allmaras turbulence model.

Results: In the early stages of the HRRLS configuration analysis, a set of preliminary subsonic and supersonic USM3D solutions were computed using a coarse grid. Soon after, extensive USM3D grid convergence studies were performed for the HRRLS mated and first stage configurations at subsonic, supersonic, and hypersonic Mach numbers. At the conclusion of the grid convergence study, the final subsonic, supersonic, and hypersonic grids were determined from a group of 34 grids. By applying these final grids, the USM3D aerodynamic database was generated for the HRRLS mated and first stage configurations. The flowfield solution for a representative case of the HRRLS mated configuration at Mach 0.8 and 8 degrees angle of attack is shown in Figure 2. Mach contours on the plane of symmetry are depicted in Figure 2a, the two-dimensional streamline plot on the plane of symmetry at the aft end is shown in Figure 2b, the three-dimensional isosurface of zero axial velocity in the external nozzle region and behind the base is shown in Figure 2c, the surface pressure contours on the leeward and windward sides are depicted in Figures 2d and 2e, and the surface pressure distribution along the centerline is shown in Figure 2f.

Role of High-End Computing: All computational results were generated on the Columbia supercomputer. Each job used 64 processors and required 60–70 gigabytes of system memory to execute all computations. Each subsonic case took 30–48

processor-hours to complete, while each supersonic case took 96–144 processor-hours to complete, and each hypersonic case took 160–200 processor-hours to complete, due to oscillatory behavior. NASA high-end computing experts helped to meet this challenge by making the most efficient use of the computer time allocations and supporting the largest volume of data transfer in the NASA Advanced Supercomputing (NAS) Division's history across the wide-area network between Columbia and NASA Langley. The NAS Division's high-end computing capability was a key component to the success of the HRRLS reference vehicle development task.

Future: The HRRLS reference vehicle was developed with a wing of fixed size to generate a longitudinal aerodynamic matrix. To achieve longitudinal stability for the HRRLS configuration, a parametric study of wings with common planform shape but with different areas to trim the HRRLS vehicle is planned. In addition, an assessment of turbine engine flow impact in the external nozzle region of the HRRLS configuration is also planned.

Publications

[1] Pao, J. L., "Presentation: Preliminary USM3D Results for the TSTO-mated and TSTO-1st-Stage Configurations at Subsonic and Supersonic Mach Numbers," TN 08-517 (NNL07AA00B), NASA Langley Research Center, August 2008.

[2] Pao, J. L., "USM3D Grid Convergence Study for the TSTO-mated and TSTO-1st-Stage Configurations at Subsonic, Supersonic and Hypersonic Mach Numbers," TN 08-518 (NNL07AA00B), NASA Langley Research Center, August 2008.

[3] Pao, J. L., "USM3D Computed Aerodynamic Database for the TSTO-mated and TSTO-1st-Stage Configurations at Subsonic, Supersonic and Hypersonic Mach Numbers," TN 08-520 (NNL07AA00B), NASA Langley Research Center, September 2008.

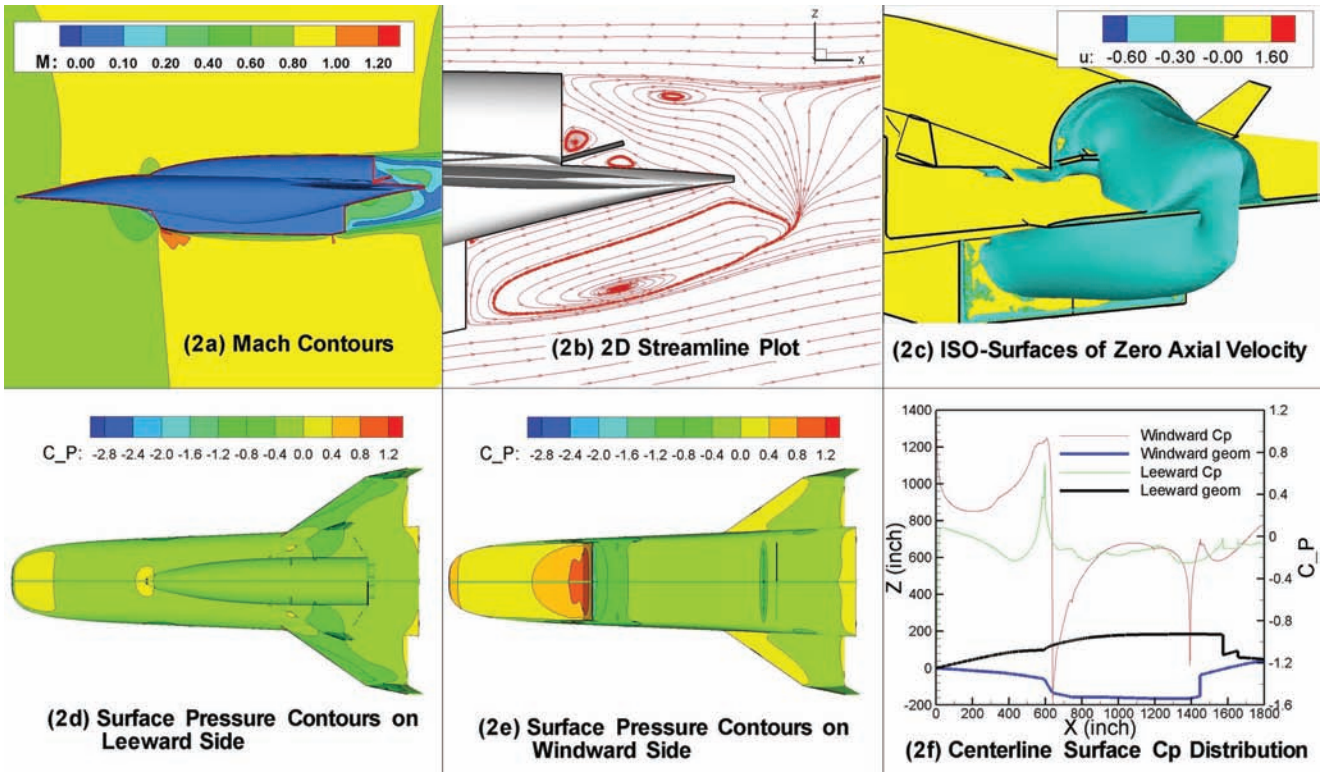
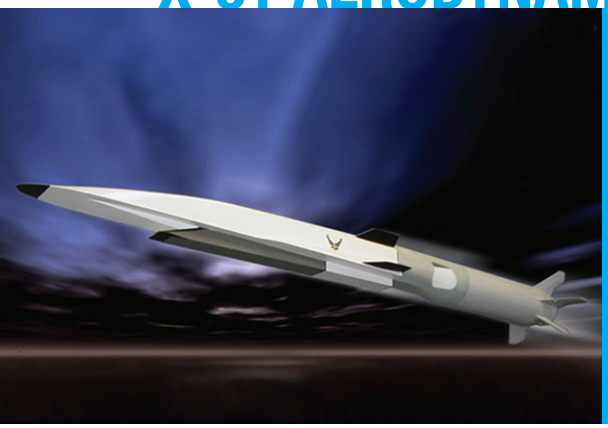


Figure 2: Flowfield solution for the Highly Reliable Reusable Launch Systems configuration at Mach = 0.8 and angle of attack = 8 degrees.

X-51 AERODYNAMICS



AERONAUTICS RESEARCH MISSION DIRECTORATE

TODD MAGEE
Boeing Research & Technology
(714) 896-1134
Todd.Magee@boeing.com

◀ **Figure 1:** The goal of the X-51 Program is to flight-demonstrate an endothermic hydrocarbon-fueled scramjet engine. The work is sponsored by the Air Force Research Laboratory (AFRL) and Defense Advanced Research Projects Agency. Many team members are involved in the project, including individuals from Boeing, Pratt & Whitney, and NASA.

Project Description: The goal of this project is to utilize high-end computing to assess aerodynamic characteristics of the X-51 vehicle in preparation for its first flight, scheduled for fall 2009. The X-51 vehicle is an Air Force Research Laboratory- and Defense Advanced Research Projects Agency (DARPA)-sponsored flight demonstrator of the first hydrocarbon-fueled scramjet engine (see Figure 1). The flight profile for the X-51 covers several phases and flight regimes. The vehicle will be dropped from a B-52 carrier aircraft at Mach = 0.8, boosted to the scramjet starting condition (Mach ~4.5) using an Army ATACMS booster, then separate from the booster and ignite the scramjet engine, and accelerate up to a Mach 6+ cruise. After the vehicle runs out of fuel, it will decelerate to the mission ending point. Aerodynamic data is required over all flight regimes to support performance, stability and control, and loads analyses, as well as fin actuator and structure sizing. Utilizing wind tunnels to generate all of this data would be very impractical and costly—which ties in with a secondary project goal: to minimize the use of wind tunnel data by leveraging the use of computational fluid dynamics (CFD) and high-end computing.

This work entailed the generation of CFD results (Euler and Navier-Stokes) for both the X-51 hypersonic cruiser and stack vehicles over a flight condition range of Mach = 0.6 to 7, angles-of-attack = -10 to 25 degrees, sideslip angles = -4 to 4 degrees, and fin deflections = -25 to 25 degrees. The cruiser is the portion of the vehicle that contains the scramjet engine while the stack consists of the cruiser, inter-stage, and booster combined. The solutions were then processed to generate force and moment data for the cruiser and stack to populate various databases. The databases included skin friction versus altitude, general vehicle force and moment, fin forces and moments, beta vane calibration, B-52 influence, and stage separation. Many of the cases were utilized to fill in where wind tunnel data was not available. In addition, CFD pre-test analysis was conducted on several wind tunnel models to support wind tunnel tests at the Arnold Engineering Development Center (AEDC) Von Karman Tunnel B and the NASA Langley Research Center Unitary Plan Wind Tunnel (UPWT).

Relevance of Work to NASA: This project is supported by the Hypersonics Project under NASA's Fundamental Aeronautics Program. It leverages and extends the work completed by NASA on the X-43 Program in pursuit of practical hypersonic flight, and is fundamental to the X-51 Program meeting Critical Design Review delivery dates, and the subsequent hardware manufacturing/procurement. Therefore, it supports technology development for future access to space, hypersonic atmospheric vehicles, and high-performance weapons advancement.

Computational Approach: Three CFD tools were used in this project: the NASA Cart3D (Euler) code, the NASA OVERFLOW (Navier-Stokes) code, and the Boeing BCFD (Navier-Stokes) code. Integrated force and moment, surface pressure, viscous shear stress, and various flowfield data were utilized from the CFD simulation results.

Results:

The project generated a number of databases:

- *Skin friction database*: The skin friction database captures the effect of altitude on skin friction drag. Navier-Stokes CFD was conducted for both the cruiser and stack at various altitudes and Mach numbers (see Figure 2).
- *Flight polars for X-51 cruiser thrust/drag bookkeeping*: This database provides a determination of what forces are attributed to aerodynamics or propulsion. Navier-Stokes CFD was conducted at several supersonic and hypersonic Mach numbers and angles of attack (see Figure 2).
- *B-52 captive/carry and launch database*: This Euler CFD-based influence database was generated for conducting 6-degrees of freedom (6-DOF) separation analyses. The X-51 stack vehicle was analyzed at a matrix of locations below the B-52, and at various flight conditions (see Figure 3).
- *B-52 viscous captive/carry and launch cases*: Navier-Stokes CFD was utilized at several conditions that were also analyzed with Euler CFD to assess the effects of viscosity on the results (see Figure 3).

- *Pre-test predictions for AEDC VKF wind tunnel test*: Euler CFD was used to generate pre-test predictions for an AEDC Von Karman Facility (VKF) Tunnel B wind tunnel test. These pre-test predictions were used to ensure quality of the test data (see Figure 2).
- *Additional fin deflection data*: Euler CFD was utilized to provide fin effectiveness data for flight conditions that were not acquired through wind tunnel testing (see Figure 2).
- *Beta Vane calibration matrix*: CFD was utilized to generate a database to calibrate the beta vane sensor on the X-51 vehicle.
- *Stage separation CFD*: Navier-Stokes CFD was utilized to construct an aerodynamic database for use in 6-DOF separation analyses. Several flight conditions and various separation locations between the cruiser and interstage/booster sections were analyzed (see Figure 2).
- *Boat-tail CFD analyses*: Utilized Navier-Stokes CFD to determine the effect on aerodynamics by boat-tailing the backend of the X-51 cruiser vehicle.
- *Sting Effects CFD*: Navier-Stokes CFD was utilized to determine effects of the sting, used in several wind tunnel tests, on the wind tunnel force and moment results (see Figure 2).
- *Pre-test predictions for NASA UPWT wind tunnel test*: Pre-test predictions using Euler and Navier-Stokes CFD were generated to ensure quality of the test results.

A total of 5,000 CFD runs were completed over the course of the project—approximately 4,500 Euler runs and 500 Navier-Stokes runs.

Role of High-End Computing: The X-51 Program would not have been able to complete the aerodynamic database work (within its time and budget constraints) without NASA's Columbia supercomputer. Many of the grids utilized in the work were too large to run on Boeing's systems. In addition, Boeing's systems did not have enough processors to provide the throughput for completing the work before the X-51 Critical Design Review. A total of ~1,400,000 processor-hours were consumed on Columbia to complete the project tasks.

Future: Future project goals include generation of CFD data to support the manufacturing and flight test phases of the project.

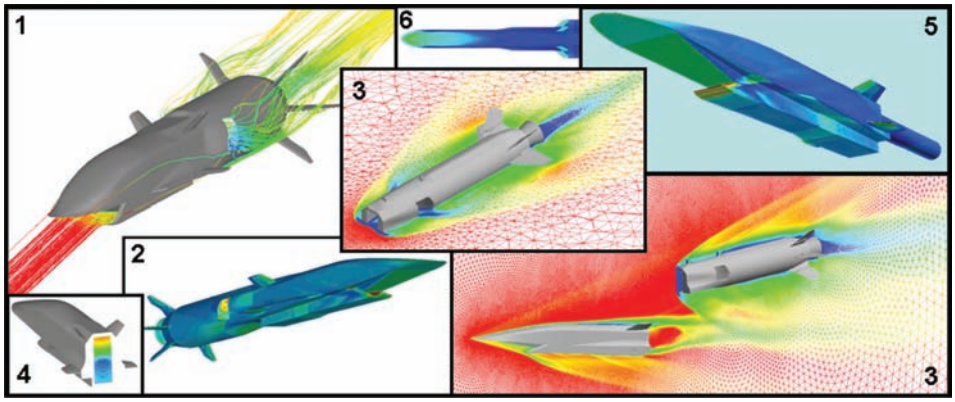


Figure 2: The figure shows a montage of CFD results that have been computed for the X-51 aerodynamics project using Columbia. Starting at top left: 1) OVERFLOW Navier-Stokes result for the X-51 stack at Mach 5; 2) Pressure coefficient (Cp) contours for the X-51 stack computed using the Cart3D code; 3) Stage separation results using the Boeing BCFD code; 4) Exit pressure contours using the OVERFLOW code, which were used in determining the proper thrust/drag bookkeeping; 5) BCFD results to determine influence of the sting on wind tunnel force and moments; 6) OVERFLOW N-S results for the cruiser at Mach 5.

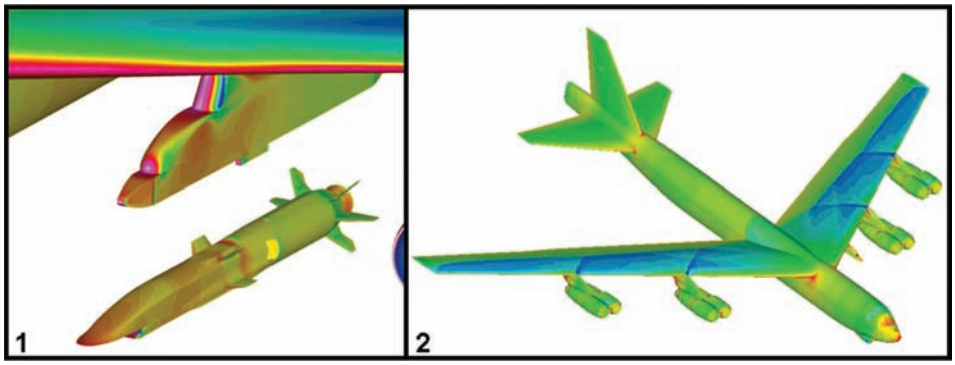
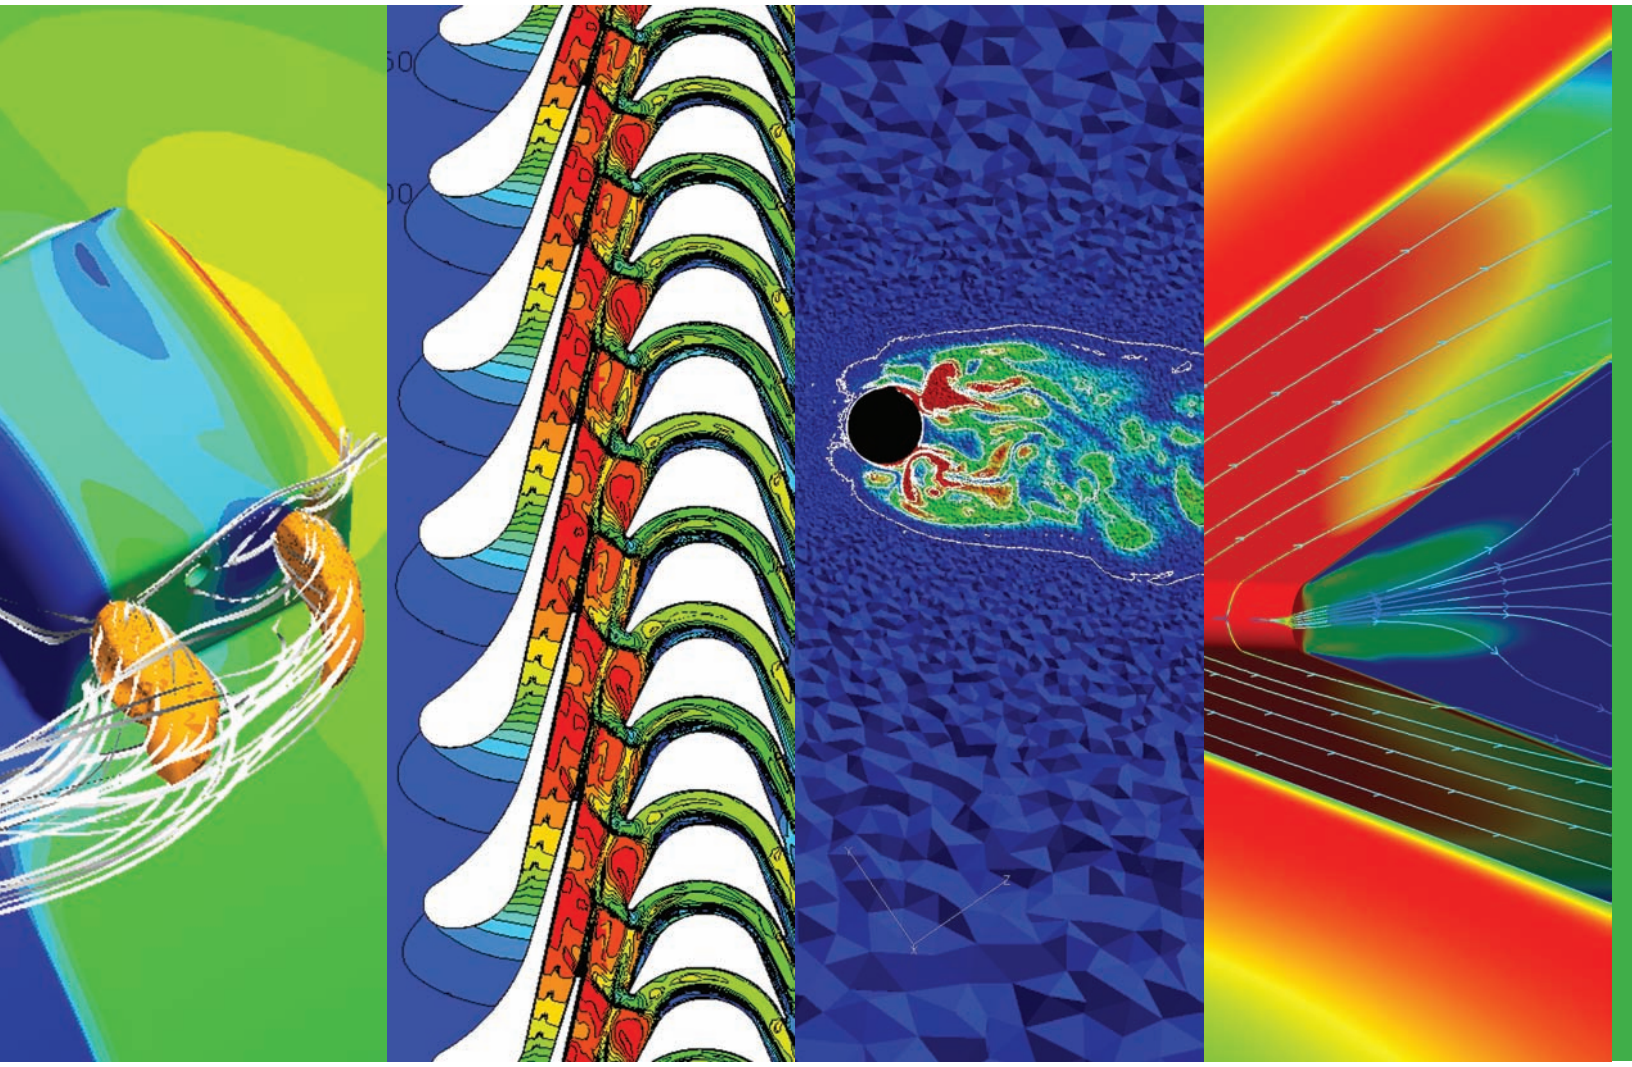


Figure 3: Two B-52 captive/carry and launch CFD cases are shown for the work conducted for the X-51 aerodynamics project: 1) Shows an OVERFLOW Navier-Stokes solution for the X-51 stack, 36 inches below the pylon at the Mach = 0.8 drop condition; 2) Shows the Cart3D-generated Cp contours for the B-52 with the X-51 stack attached at the Mach = 0.8 flight condition.



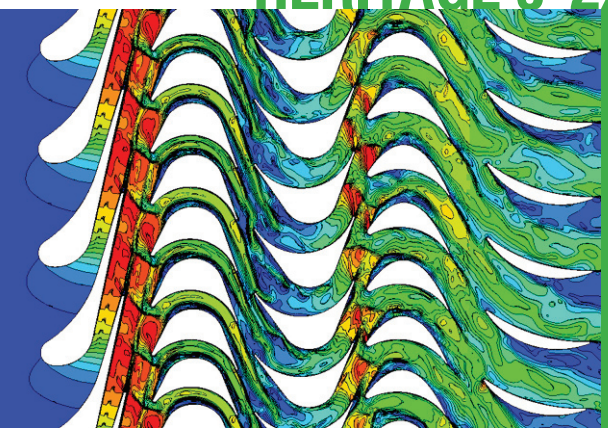
MISSION DIRECTORATE

EXPLORATION SYSTEMS

The Exploration Systems Mission Directorate is developing new vehicles, capabilities, supporting technologies, and foundational research that will enable sustained human and robotic exploration of the Moon and other destinations beyond low-Earth Orbit. At the heart of NASA's exploration efforts are the Constellation Program's next-generation space vehicles: the Orion Crew Exploration Vehicle that will carry astronauts through space, the Ares I Crew Launch Vehicle that will launch Orion and the crew into orbit, and the heavy-lift Ares V Cargo Launch Vehicle that will carry additional vehicle components and equipment into orbit for rendezvous with Orion. Other key areas of research and technology development include robotic missions to the Moon, health and safety of crews on long-duration space missions, and risk mitigation for exploration projects.

DOUG COOKE
Associate Administrator
http://www.nasa.gov/exploration/home/index.html

AIR RIG TESTING OF THE HERITAGE J-2X FUEL TURBINE



◀ Close-up of Figure 1.

EXPLORATION SYSTEMS MISSION DIRECTORATE

PRESTON SCHMAUCH
NASA Marshall Space Flight Center
(256) 544-1218
Preston.B.Schmauch@nasa.gov

Project Description: Currently, all turbine blade stresses on the Ares I J-2X upper stage engine are determined using blade loading obtained from unsteady computational fluid dynamics (CFD) analyses. The J-2X fuel turbine is a 2½-stage supersonic engine—and there is currently no on-blade data for the second stage of a supersonic engine. To fill this data gap, the Heritage Fuel Air Turbine Test (HFATT) will provide one of the most instrumented air rig turbine tests ever performed. The air rig test runs the turbine at scaled conditions in air as opposed to actual engine conditions. It will utilize substantial on-rotor, unsteady pressure instrumentation and will include various instrumentation on the second stage. This test will help reduce risk for the J-2X turbopumps.

To assess and modify the run box that envelopes all run conditions that the turbine could experience during the upcoming HFATT testing, various test points were simulated. Unsteady CFD simulations were performed to help determine the outer corners of the run box and the best transient locations to investigate. Torque, power, and speeds from the simulations were examined to ensure that the test box was enveloped by the limitations of the test facility at NASA Marshall Space Flight Center (MSFC). In addition, the engine design point pressure ratio, flow coefficient, and various other points were run and interrogated to help produce design curve estimations for different pressure ratios and to formulate detailed pre-test predictions. Temperature and pressure-range data obtained from the unsteady CFD simulations were used to calibrate the various time-accurate instruments that will be used in testing.

One area of particular interest in the HFATT and CFD simulations is unsteady pressure frequencies. Currently, unsteady blade loading is determined by doing a Fourier decomposition of the pressures at every computational node solved in the unsteady CFD simulation. A stress team then uses those coefficient results in a forced-response analysis. Several modes show low factors of safety for the engine running conditions and are consequently an area of interest in the HFATT. Comparisons will be made to determine any conservatism that might exist in the CFD simulations and will subsequently be used for code validation. NASA is currently having tip dampers employed for some of the J-2X turbine blades to increase factors of safety.

Relevance of Work to NASA: The Ares I J-2X upper stage engine is a critical component in reaching NASA's current goal of returning to the Moon. Due to the higher thrust requirements of the Ares I upper stage, J-2X turbopump running conditions were increased to accommodate the additional thrust. The increased rotational speeds of the J-2X turbopumps have led to increased stresses and lower factors of safety. The HFATT project will increase model fidelity and reduce risk associated with the J-2X engine design and analysis. The information obtained in the testing will also help improve safety and design quality for future programs.

Computational Approach: Phantom, a NASA-developed code that has been anchored and validated for supersonic turbines, is being used for the HFATT unsteady CFD simulations. Phantom uses three-dimensional, unsteady Navier-Stokes equations as the governing equations. The Baldwin-Lomax turbulence model is used for turbulence closure. An overset O and H grid topology with moving grids to model blade motion is employed for the simulations. To simulate several run points in and on the run box, a periodic 1/7 sector of the turbine was modeled for all off-design running points, and a full annulus simulation was run for the design point.

Results: All simulations needed for pre-test support of HFATT have been completed. The run box was sculpted using results obtained from the unsteady CFD analyses, such as those shown in Figures 1 and 2. One example of this is the run conditions at one corner of the run box corresponding

to a low flow rate and low rotational speed. The simulation results showed that this preliminary point would "unchoke," and the run box was altered to avoid this phenomenon. We are currently awaiting the instrumented hardware before testing can begin. Testing is projected to begin in May 2009.

Role of High-End Computing: In total, 23 different run points were simulated and analyzed in support of the HFATT. To run the high number of test points, significant resources were required. Without NASA's HEC resources and the Columbia supercomputer, we would not have been able to provide such a comprehensive level of support within a reasonable time frame. Approximately 2 million processor-hours were needed to complete all simulations for the HFATT project.

Future: Pre-test run conditions do not always match actual running conditions, and the HFATT is no exception. After HFATT testing, test points with actual run conditions will be simulated using NASA HEC resources. These simulations will provide comparisons to assess the quality of the CFD methods currently used for stress analysis. Alternate test configurations could also be simulated using NASA HEC resources.

Co-Investigators

- Daniel Dorney, NASA Marshall Space Flight Center

Publications

[1] Dorney, D., Griffin, L., and Schmauch, P., "Unsteady Flow Simulations for the J-2X Turbopumps," 54th JANNAF Propulsion Meeting, May 2007.

[2] Marcu, B., Tran, K., Dorney, D., and Schmauch, P., "Turbine Design and Analysis for the J-2X Engine Turbopumps," 44th AIAA Joint Propulsion Conference, *AIAA 2008-4660*, July 2008.

[3] Marcu, B., Zabo, R., Dorney, D., and Zoladz, T., "The Effect of Acoustic Disturbances on the Operations of the Space Shuttle Main Engine Fuel Flowmeter," 43rd AIAA Joint Propulsion Conference, *AIAA 2007-5534*, July 2007.

[4] Dorney, D., Griffin, L., Marcu, B., and Williams, M., "Unsteady Flow Interactions between the LH2 Feedline and SSME LPFP Inducer," 42nd AIAA Joint Propulsion Conference, *AIAA 2006-5073*, July 2006.

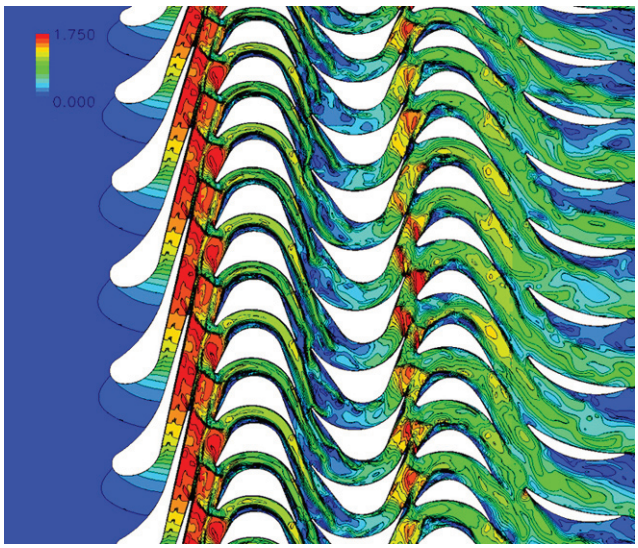


Figure 1: Mach number contours from simulation of the J-2X fuel turbine air rig.

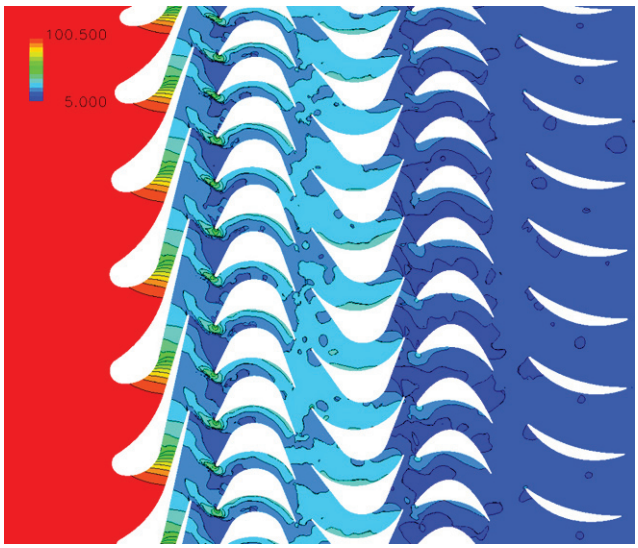
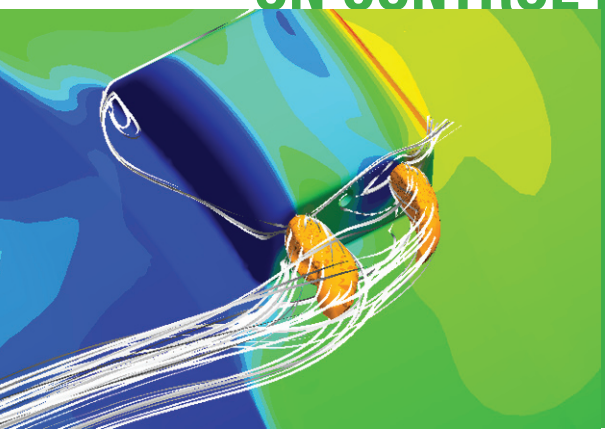


Figure 2: Static pressure contours from simulation of the J-2X fuel turbine air rig (psi).

ARES I ROLL CONTROL SYSTEM JET EFFECTS ON CONTROL ROLLING MOMENT IN FLIGHT



EXPLORATION SYSTEMS MISSION DIRECTORATE

KHALED S. ABDOL-HAMID
NASA Langley Research Center
(757) 864-8224
K.S.Abdol-Hamid@nasa.gov

◀ Close-up of Figure 1.

Project Description: It is important to determine effects of the roll control system (RoCS) jets on the rolling moment of the Ares I Crew Launch Vehicle. This challenging task includes determining how best to model the chemically reactive jet plumes from the RoCS and how to capture the plume interactions with the freestream airflow around a complex vehicle with many detailed protuberances. Computational fluid dynamics (CFD) simulations help accomplish this task by providing efficient, detailed aerodynamic data to supplement the more limited and costly experimental wind tunnel test data. Several computational analyses have previously revealed the impacts of installing the RoCS on different configurations.

The objective of this project is to quantify jet effects on the net efficiency of the Ares I vehicle's RoCS in flight. These jet effects vary with flow conditions such as flight Mach number, local atmospheric pressure, angle of attack, and the roll angle of the vehicle. Therefore, a practical database construction requires many hundreds of flow computations for a given vehicle outer mold line (OML) and RoCS thruster configuration. The present task is divided into three phases. The first phase included a preliminary study of computational best practices along with examinations of jet flow properties in a cross-flow and roll control thruster plume interaction with the vehicle. An existing computational mesh for an Ares I-X flight test vehicle configuration with a simplified OML was used in the preliminary study. The second phase was a limited parametric study of jet effects for the Ares I configuration at selected Mach numbers, angles of attack, and roll angles. The third phase will be a parametric CFD study of approximately 125 flow conditions along a nominal Ares I ascent trajectory. Results of these computations will provide an aerodynamic database for guidance and control simulation applications.

Relevance of Work to NASA: Complementary to ground-based wind tunnel testing, CFD methods are being used extensively to support design analysis and aerodynamic database development for NASA's next-generation of space exploration vehicles. The Ares Project has designated the CFD flow solver USM3D as the primary code for developing the computational aerodynamic database for the vehicle. Results of this study have provided critical insight into the process by which jet plumes affect RoCS performance in flight.

Computational Approach: Flow solution development uses the USM3D Navier-Stokes solver on a tetrahedral, unstructured mesh in which the flow variables are all calculated at cell centers. Pre-processing of mesh cell connectivity, partitioning for parallel computing, and flow input conditions are conducted using NASA's HEC resources, as are extensive post-processing steps required to extract information from the solutions for Ares I design applications. The overall computational process is known as the Tetrahedral Unstructured Solver System (TetrUSS). Although the thruster and freestream flow conditions were established at full-scale flight vehicle values, the jet effect CFD computations were performed at wind tunnel Reynolds numbers to conserve computer resources by reducing the mesh size and the number of iterations needed to converge each solution.

Results: The solutions were developed over a highly refined mesh system to capture the interaction of jet plumes with cross-flows, as well as the subsequent flow interactions with the vehicle and its protuberances downstream of the RoCS location. The total mesh contained 70 million cells to resolve relevant flow interaction details. The assessment of jet effects at each flow condition required two solutions: one with thrusters idle and one with thrusters firing for positive or negative roll. The differences between the force and moment coefficients of the two solutions provided a quantitative measure of jet effect on RoCS efficiency in flight. An example of the complex interaction between the freestream flow and the RoCS plumes is shown in Figure 1.

Role of High-End Computing: All of the solutions were computed on the NASA Advanced Supercomputing (NAS) facility's Columbia supercomputer at NASA Ames Research Center. The scale of parallel computing required for this project involved 256 processors for each solution and a significant amount of total computer hours. The USM3D flow solver uses approximately 10 microseconds per cell per iteration on a single processor core, and each solution used 256 processors in parallel to solve the flow equations. Even with an immense demand for resources from this and other high-priority projects within the Agency, NASA's HEC Program provided an exceptional level of support for this project with timely allocation of the supercomputer nodes needed to run several jobs simultaneously, and the runtime hours to enable this project to meet its deadline.

Future: This RoCS jet effects study will be part of the aerodynamics database supporting future Ares I design cycles. To achieve precise and smooth vehicle roll control, a subsequent task will provide jet effect data for a large parametric matrix to support guidance, navigation, and control studies. Continuing HEC resource support will be critical to the successful completion of Ares Project mission objectives. We are planning to complete 400 CFD solutions using the Columbia supercomputer. The resources needed to complete this task are estimated to be 5 million processor-hours.

Co-Investigators

- S. Paul Pao and Karen A. Deere, both of NASA Langley Research Center

Publications

[1] Deere, K.A., Pao, S.P., and Abdol-Hamid, K.S. "A Computational Investigation of the Roll Control System Jet Effects on Rolling Moment of the Ares I-X Clean Configuration," NASA TP 2008 (review and publication pending).

[2] Deere, K.A., Pao, S.P., and Abdol-Hamid, K.S. "A Computational Investigation of the Roll Control System Jet Effects on Rolling Moment of the Ares I A103 Full Protuberance Configuration," NASA TP 2008 (review and publication pending).

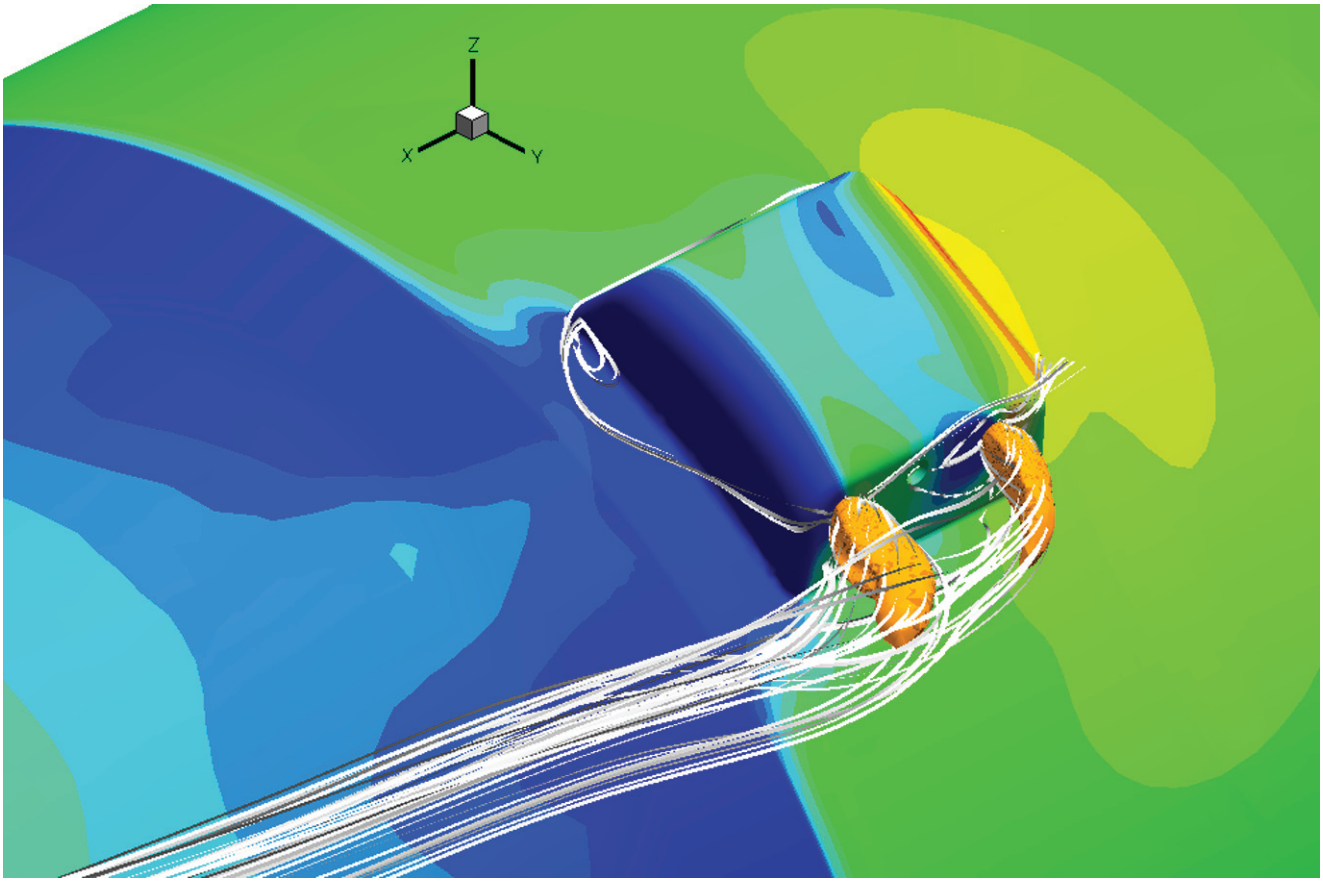
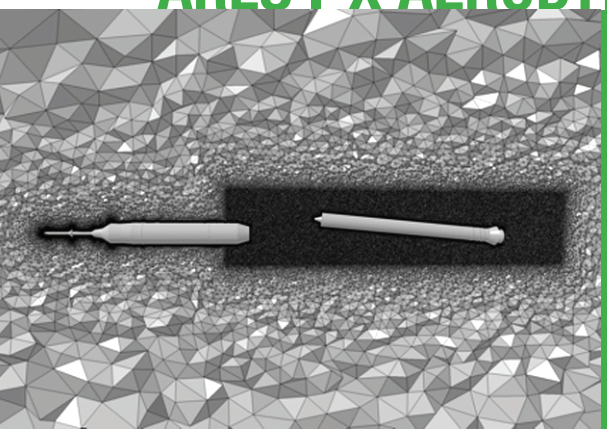


Figure 1: Interaction between the freestream flow and the roll control system plumes on the Ares I Crew Launch Vehicle.

ARES I-X AERODYNAMICS DATABASE DEVELOPMENT



EXPLORATION SYSTEMS MISSION DIRECTORATE

STEVEN BAUER
NASA Langley Research Center
(757) 864-5946
Steven.X.Bauer@nasa.gov

◀ **Figure 1:** Topology of a typical computational grid in the vicinity of the Ares I-X upper stage simulator and the first stage.

Project Description: As a crucial step in the development of the new Ares I Crew Launch Vehicle (CLV) design, NASA will be launching the Ares I-X Flight Test Vehicle (AIX FTV) in 2009. The Ares I-X FTV will be similar in mass and size to the actual Ares I CLV, and will imitate the first two minutes of the Ares I launch trajectory and flight conditions. The test vehicle will launch through a speed of Mach 4.7 and separate from its first stage (FS) solid rocket booster (SRB) at 130,000 feet. The core objectives of this first test flight are to demonstrate performance of the flight control systems, and to characterize the flight environment during stage separation. In support of these efforts, extensive computational fluid dynamics (CFD) simulations have been performed to generate and/or populate databases of aerodynamic conditions for the Ares I-X test flight configurations and parameters. These aerodynamic databases are being used by the Ares I-X Guidance, Navigation & Control (GN&C) community to help assess aerodynamic performance and flight control functionality of the AIX FTV.

Relevance of Work to NASA: This project is an important part of the Agency's development of the Ares I CLV. The extensive simulation data will help supplement and interpret the experimental data obtained from the actual AIX test flight to better evaluate and understand aerodynamic factors of the Ares I CLV design.

Computational Approach: A majority of the CFD simulations for this project were conducted using USM3D (Figure 2), a Reynolds-averaged Navier-Stokes code for unstructured grids developed at NASA Langley Research Center. A second CFD code, OVERFLOW, was also used to perform simulations of first stage descent cases. Grid sizes ranged from 40 million grid cells for the simplest "clean" configurations at wind tunnel Reynolds number conditions, to 90 million grid cells for configurations including the full set of protuberances at flight Reynolds numbers.

Results: Hundreds of CFD solutions (such as that shown in Figure 1) have been generated to build the aerodynamic databases needed to assess key aspects of Ares I-X flight performance. For the main part of the launch trajectory before first-stage separation, simulation cases covered 11 Mach numbers from 0.9–4.5, angles of attack from 0°–90°, and roll angles from 0°–360°. Various combinations of these parameters were simulated for a simplified, clean version of the vehicle geometry and for a more detailed version of the geometry including protuberances. A number of specific protuberance cases were run on a simplified configuration to determine pressure coefficients and heating rates for various instruments. Selected cases were also run at both wind tunnel condition Reynolds numbers, and full flight condition Reynolds numbers.

CFD simulations were also performed to analyze separation of the AIX vehicle's first-stage (FS) rocket booster from the upper stage simulator (USS), and descent of the FS and USS components as they fall back to Earth. Forty-two stage separation cases were conducted at Mach 4.5 at various separation distances and orientation angles of the FS and USS components. Both USS and FS descent cases were run at Mach numbers from 0.5–4.5 with angles of attack from 0°–180°.

Additional sets of cases were also generated to assess plume effects and control authority for the vehicle's roll control system (RoCS), booster deceleration motors (BDMs), and booster tumble motors (BTMs). Thirty-three RoCS simulations were conducted for three Mach numbers at various roll angles. These cases were run with RoCS not firing, with clockwise RoCS jet pairs firing, and with counterclockwise RoCS jet pairs firing to determine the interference effects of RoCS plumes on the overall generated forces and moments for the vehicle. Twenty-four BDM and two BTM firing cases were conducted for Mach 4.5 at various pitch and roll angles to determine plume effects on the vehicle and the effects of one or more motor "out" cases.

Results of the CFD simulations for all these cases have provided extensive data on the forces, moments, surface pressure coefficients, line loads, and in some cases, heating rates for various conditions of the AIX FTV.

Role of High-End Computing: HEC resources were used to predict forces and moments, surface pressures, line loads, and some heating rates on the AIX FTV. These resources were also used to visualize the flowfield around the vehicle at various flight conditions and around various control system jets as they are firing.

While the majority of computations were run on the Columbia supercomputer, recent work has been moved to the Pleiades system. Extensive efforts to improve the performance of USM3D on Columbia were undertaken by USM3D code developers and NASA Advanced Supercomputing facility staff, and a similar effort is currently underway to improve USM3D efficiency on Pleiades. Typical runs on either system used from 128 processors for the smallest grids to 256 processors for the largest grids. Computations required as few as 12,000 iterations to converge for steady-state solutions on clean configurations, and as many as 180,000 iterations for time-accurate solutions on configurations with extensive flow separation, translating to runtimes from 11 hours to as many as 160 hours.

Future: Aerodynamic database generation will continue in order to complete all of the RoCS, BDM, BTM, and main engine firing cases for the AIX FTV. Additional cases will also be run to fill in gaps where no data exists or where required by GN&C to evaluate and refine control of the vehicle. After launch of the FTV in 2009, additional cases will be needed to match exact flight conditions of interest, which will be used to assess the veracity of the computational results.

Co-Investigators

- Steven Krist, William Compton, Craig Hunter, and Karen Deere, all of NASA Langley Research Center

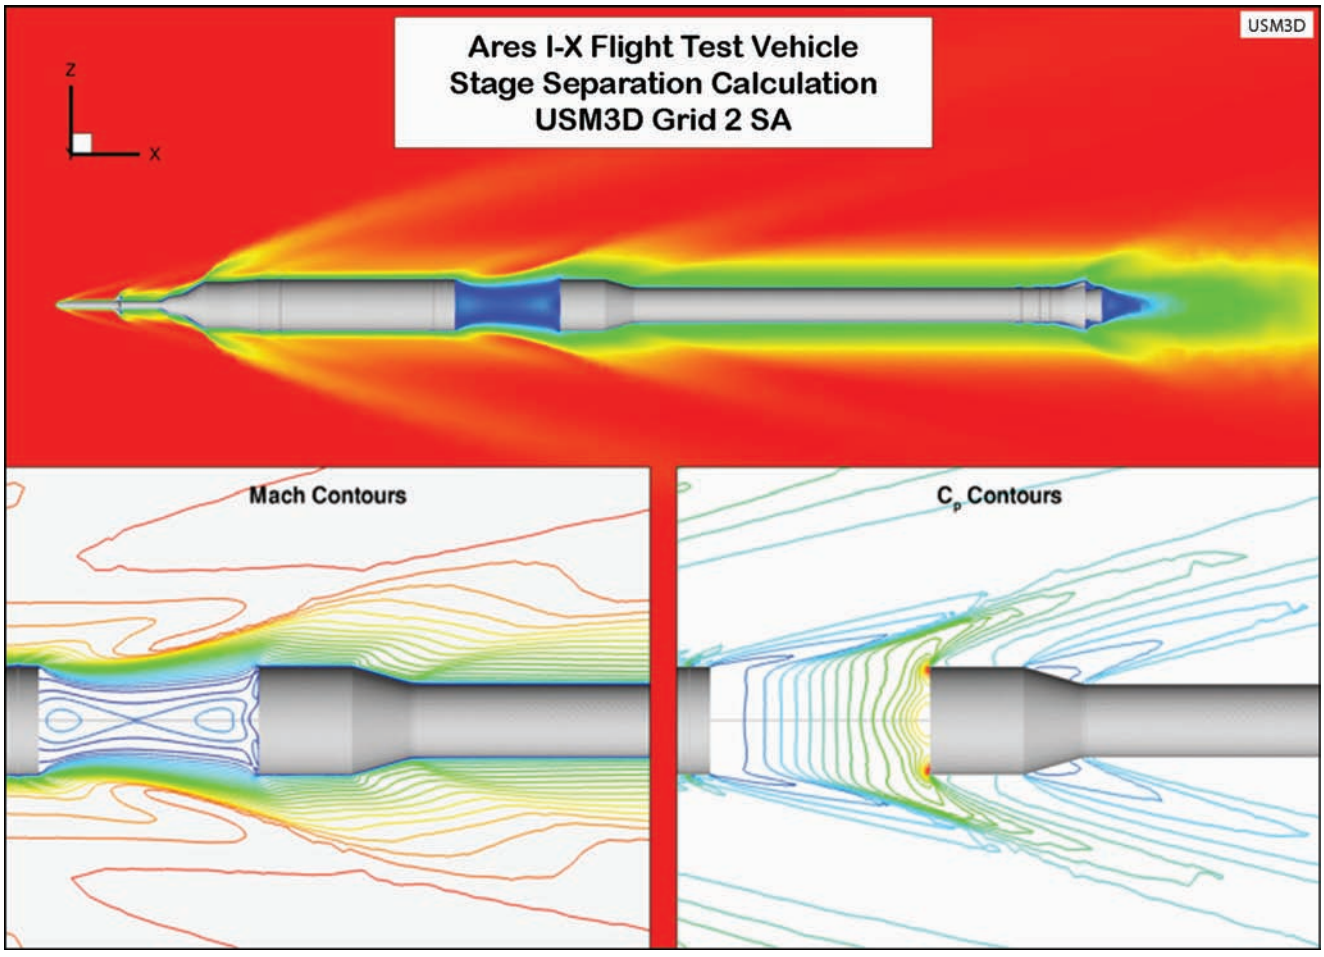
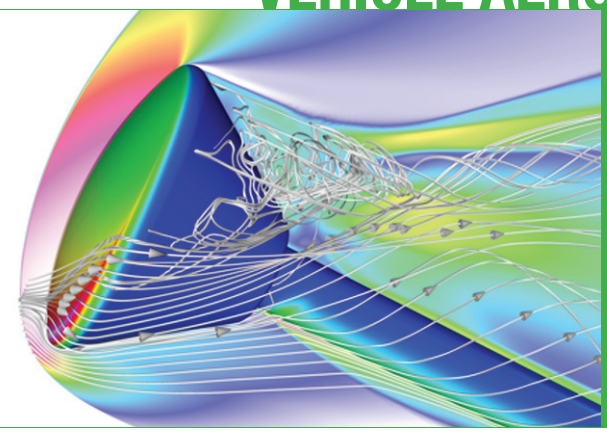


Figure 2: Typical USM3D flowfield for a stage separation case.

CREW EXPLORATION VEHICLE AEROSCIENCES PROGRAM



EXPLORATION SYSTEMS MISSION DIRECTORATE

JOSEPH OLEJNICZAK
NASA Ames Research Center
(650) 604-3252
Joseph.Olejniczak@nasa.gov

◀ **Figure 1:** Simulation of the Orion crew module and wind tunnel sting.

Project Description: The Crew Exploration Vehicle (CEV) Aerosciences Program (CAP) is developing complete aerodynamic and aerothermodynamic databases for the Orion crew module (CM) and launch abort system (LAS) covering the range of all possible operating conditions.

Accurate aerodynamic data such as lift, drag, pitching moment, and dynamic stability derivatives are required to design the flight control system and ensure that the pinpoint landing requirement can be met. The aerodynamic database covers the entire CEV operational envelope including nominal ascent, ascent abort scenarios, on-orbit plume environments, reentry flight from the hypersonic through subsonic regimes, and the terminal landing approach including parachute deployment.

The aerothermodynamic database covers the portion of atmospheric flight that produces significant aeroheating for the vehicle. The ascent heating environment must be quantified to ensure vehicle integrity during nominal and off-nominal ascent conditions. Thermal protection system (TPS) design requires convective and radiative heating environments for the entire vehicle surface during reentry, including localized heating rates on penetrations and protuberances.

The CAP databases are built using a combination of computational fluid dynamics (CFD) results and wind tunnel data, which together provide higher-fidelity databases at lower cost than either could alone. The databases will require thousands of high-fidelity numerical solutions modeling the flowfield around the CM and LAS for all flight regimes. The CFD solutions are also critical to understanding how to extrapolate wind tunnel data to extreme flight conditions that cannot be replicated by ground tests. Conversely, ground test data are also used to help quantify the uncertainty of the CFD solutions. CFD is also used to assess local geometric features, such as the compression pads that attach the crew module to the launch vehicle. For previous programs, heating rates on such features would be determined from an extensive set of ground test data that would be extrapolated to flight conditions using engineering models. For Orion, however, a small set of ground test data was obtained to quantify the uncertainty in the CFD; and CFD solutions were used to develop the engineering heating model that the designers could use.

Relevance of Work to NASA: The CAP aerodynamic and aerothermodynamic databases are critical to the design and operation of Orion—a key component of NASA's Vision for Space Exploration and Constellation Program objectives. The databases will be provided to the Orion prime contractor as Government Furnished Material (GFM) and will be used to both design and operate the vehicle. These databases are the largest GFM components of the Orion Project and represent a significant investment of NASA resources.

Computational Approach: We use a number of high-fidelity codes to compute the flowfield around the CM and LAS. Using multiple, independent codes for the same flight conditions increases our confidence in the computed results. The DPLR and LAURA codes are reacting Navier-Stokes solvers that include thermochemical nonequilibrium. These codes are used to compute aerothermodynamic heating rates and aerodynamic coefficients in the hypersonic regime. The NEQAIR radiation solver is a first-principles physics code that computes production of radiation by gas in the hot shock layer, transport of photons through the shock layer, and radiative heating to the CEV surface. Aerodynamic coefficients in the subsonic, transonic, and supersonic regimes are computed using four different CFD tools. The OVERFLOW and USM3D codes solve the Reynolds-averaged Navier-Stokes equations using multiple overset structured grids. Cart3D is an inviscid, compressible flow analysis package that uses Cartesian grids to solve flow problems over complex geometries such as the Orion LAS with abort motor (AM) and attitude control motor (ACM) plumes. The unstructured Euler CFD code FELISA is also being used.

Results: CAP has generated thousands of three-dimensional CFD solutions for the CM and LAS geometries. These solutions, ranging from Mach 0.3 ascent abort conditions to Mach 40 reentry conditions, have been used to populate multiple aerodynamic and aerothermodynamic databases supporting Orion design analysis cycles and the Orion Preliminary Design Review (PDR). Two specific examples of our simulations are given below.

Figure 1 shows a DPLR CFD calculation of the flow around the CM with the wind tunnel sting. The surface is colored with pressure contours, and streamlines are shown to help visualize the flow. In this example, aeroheating data on the front of the CM helps validate the CFD results for flight conditions, and the CFD quantifies the effect of the sting on the wake flow and afterbody heating—allowing more accurate extrapolation of the wind tunnel data to flight conditions.

Figure 2 shows snapshots of a time-accurate OVERFLOW simulation of the LAS with four of the ACMs firing and the plumes flowing around the LAV. The surfaces are colored by pressure, and the image at lower-left shows the pressure change due to the ACM plume interference effects. Predicting the aerodynamics of the LAS requires accurate modeling of the ACM plumes over a range of freestream Mach numbers from 0–6, a range of angles-of-attack from 0–180 degrees, and various combinations of ACM thrust levels from each of the eight nozzles.

Role of High-End Computing: While each individual solution may only take a few hundred to a few thousand processor-hours depending on the analysis code and geometric complexity modeled, the thousands of high-fidelity CFD solutions needed to populate the CEV databases could not be completed without access to a supercomputer such as Columbia. The availability of such high-end computing resources, coupled with advancements in CFD fidelity, have allowed the CEV Project to create databases using computational results and provide data complimentary to wind tunnel testing. In addition, high-fidelity simulations of geometrically complex configurations provide insights into flow physics and vehicle response that cannot be obtained with any other method.

Future: Through Orion development and flight operations, CAP will compute thousands of high-fidelity numerical solutions to populate the aerodynamic and aerothermodynamic databases over the next few years. As detailed design work is completed, the geometric models of the Orion vehicles will become increasingly more complex and the computations will become more memory-intensive and time-consuming. Terabytes of complete flowfield data will have to be analyzed and stored for future use.

Co-Investigators

- Stuart Rogers, NASA Ames Research Center
- Benjamin Kirk, NASA Johnson Space Center
- Richard Thompson, NASA Langley Research Center

Collaborating Organizations

- NASA Ames Research Center, NASA Johnson Space Center, and NASA Langley Research Center

Publications

[1] Hollis, B., et al., "Aeroheating Testing and Prediction for Project Orion CEV at Turbulent Conditions," AIAA Paper No. 2008-1226, January 2008.

[2] Amar, A., et al., "Protuberance Boundary Layer Transition for Project Orion Crew Entry Vehicle," AIAA Paper No. 2008-1227, January 2008.

[3] Bibb, K., et al., "Aerodynamic Analysis of Simulated Heat Shield Recession for the Orion Command Module," AIAA Paper No. 2008-0356, January 2008.

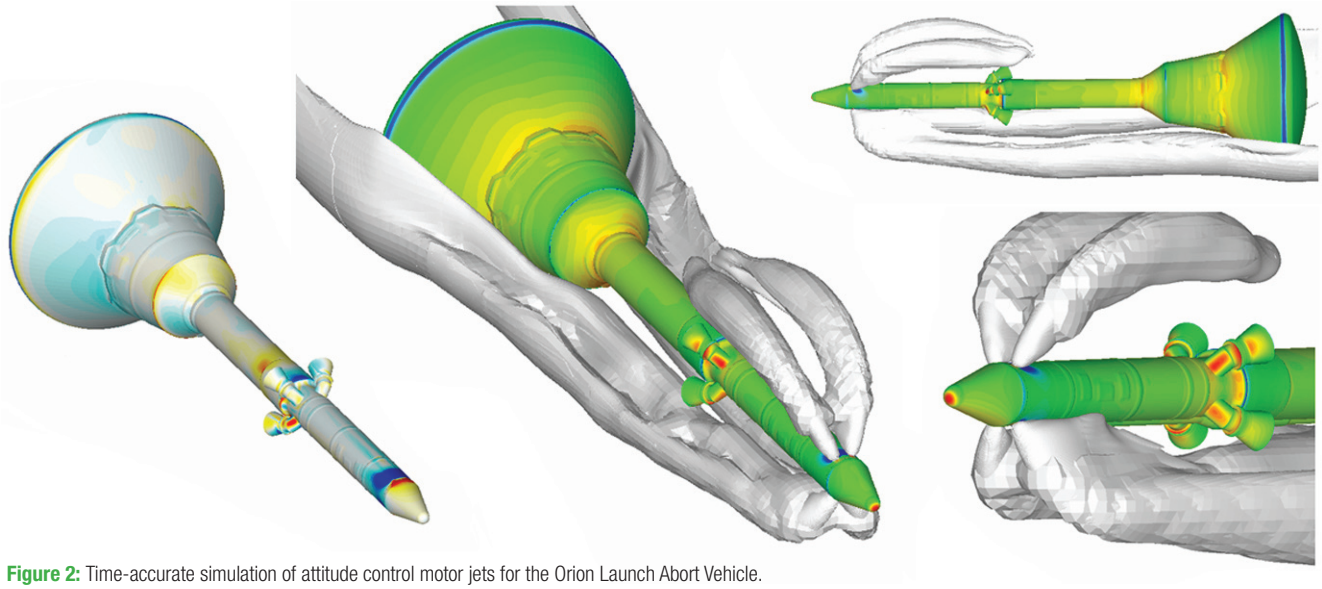
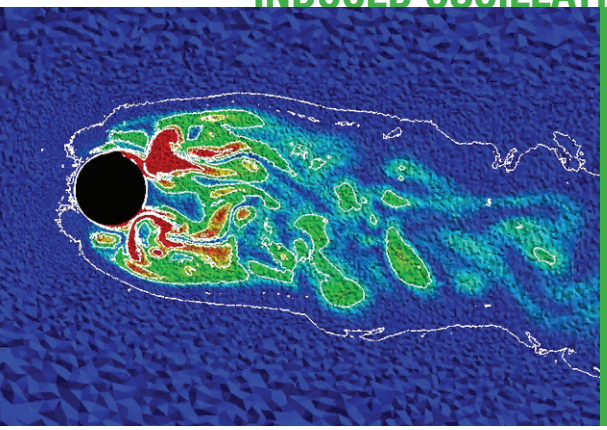


Figure 2: Time-accurate simulation of attitude control motor jets for the Orion Launch Abort Vehicle.

COMPUTATIONAL AEROELASTIC SIMULATION OF GROUND WIND-INDUCED OSCILLATION OF THE ARES I-X AND THE CHECKOUT MODEL



◀ Detail of Figure 1.

EXPLORATION SYSTEMS MISSION DIRECTORATE

ROBERT BARTELS
NASA Langley Research Center
(757) 864-2813
Robert.E.Bartels@nasa.gov

Project Description: For any vehicle on the launch pad, a poorly quantified dynamic ground wind loads environment can result in excessive and potentially catastrophic motion of the launch vehicle. For this reason, the NASA Langley Research Center Aeroelasticity Branch has been tasked with providing the Ares I Project with computational data quantifying the dynamic ground wind environment around a flexible vehicle on the pad. The Ares computational aeroelastic (CAE) analysis team is performing simulations of the Ares I Crew Launch Vehicle (CLV) both on the ground and during ascent. Analyses include static aeroelastic increments for Guidance, Navigation & Control (GN&C), flutter, buffeting, and ground wind loads. In particular, the CAE team performed simulations of ground wind-induced oscillation for both a flexible Ground Wind Loads (GWL) checkout model and the Ares I-X flight test vehicle on a launch pad. The goal of this computational analysis was to provide data to confirm that the proper wind tunnel model environment was achieved.

For a flexible launch vehicle the size of the Ares CLV, accurate wind tunnel-to-flight scaling in the presence of ground winds is difficult because the majority of the dynamic forcing is due to ground wind-induced vortex shedding—an unsteady flow phenomenon where vortices are created at the back of a body and detach periodically from either side. Vortex shedding about a smooth cylindrical body is sensitive to Reynolds number, especially within the critical Reynolds number range. Outside of the critical range the shedding pattern has a well-defined Strouhal frequency, but within this range the shedding is chaotic. The wind tunnel and the flight scale Reynolds numbers for this vehicle fall within the low and high ends of the critical Reynolds number range, making wind tunnel data difficult to scale to full-size flight conditions. Additionally, there are inherent uncertainties in the structural vibration properties of any wind tunnel model, especially for a ground turntable mount system like that used in the wind tunnel test for this case. To help account for these uncertainties, we performed computational aeroelastic simulations of the ground wind environment surrounding the flexible launch vehicle. The checkout wind tunnel model was used to develop best practices and calibration data for subsequent wind tunnel testing of Ares configurations. Nine wind tunnel conditions at which peak bending moments were observed for the checkout model, and the baseline test condition for the Ares I-X vehicle were simulated. The computed unsteady pressures, tie-down bending moments, and shedding frequencies were compared with wind tunnel data.

Relevance of Work to NASA: This work supports the Ares Program and is important to achieving a successful launch of NASA's next-generation space exploration vehicles. The computational data from this project have provided a benchmark against which wind tunnel data can be compared, and have reduced the uncertainty in wind tunnel and empirical design data, resulting in a better vehicle design.

Computational Approach: The Navier-Stokes computational fluid dynamics (CFD) code FUN3D was used to simulate the static and dynamic flexible vehicle responses to ground wind-induced oscillation. FUN3D is an unstructured, finite-volume node-based CFD code developed by the Computational Aerosciences Branch at NASA Langley. The simulations were performed with a one-equation turbulence model and with a Detached Eddy Simulation (DES) turbulence model. FUN3D is capable of performing distributed Message Passing Interface (MPI) parallel computation and was compiled using the libraries Metis, ParMetis, mvapich, and the Intel compiler. The launch pad was modeled as a symmetry plane with computational grids containing 10–40 million grid points. Turbulent and laminar flowfields with dynamic fluid/structure interaction (FSI) were simulated using an algorithm that includes dynamic, time-accurate deformation of the volume grid. The structural model used was an MSC.Nastran finite-element model.

Results: The GWL checkout model simulation confirmed the levels of vehicle motion observed in the wind tunnel tests. It also provided insight into the shedding mechanism that was

producing high tie-down bending moments. The checkout model simulations showed that strong, coherent shedding patterns occurred mainly over the crew module and first stage rocket. However, the simulations also showed that the largest bending moment dynamics occurred in conditions for which there were the largest levels of shedding dynamics over the crew module and launch abort tower regions. Because of this result, the Ares I-X wind tunnel model was populated with additional unsteady pressure transducers in the crew model region. Figures 1 and 2 show sample vorticity contours from the simulations, and Figure 3 shows an example of resulting surface pressures on the vehicle.

Role of High-End Computing: Computations for this work were performed on the NASA Advanced Supercomputing (NAS) Division's Columbia supercomputer, with each simulation distributed over hundreds of processors. Because the vortex shedding of a full-scale vehicle occurs in a critical Reynolds number range, these computations were exceedingly challenging, requiring hundreds of processors and typical runtimes of one to four weeks. Data post-processing was performed using Tecplot on the NAS platform.

Future: Ares computational aeroelastic analyses will continue to quantify the flexible response of the Ares I and Ares V vehicles. Ascent static analysis will provide flex-to-rigid increments for Guidance, Navigation and Control development. Flutter and buffet analyses will be performed to ensure that the vehicle retains structural integrity throughout ascent. Ground wind loads analyses will continue to further reduce wind tunnel data uncertainties.

Co-Investigators

- Pawel Chwalowski, Ray Mineck, Robert Biedron, and Steve Massey, all of NASA Langley Research Center

Publications

[1] Bartels, R. E., "Development of Advanced Computational Aeroelasticity Tools at NASA Langley Research Center," NATO RTO Specialists Meeting on Advanced Aeroelasticity AVT-154, paper 003, May 3-6, 2008.

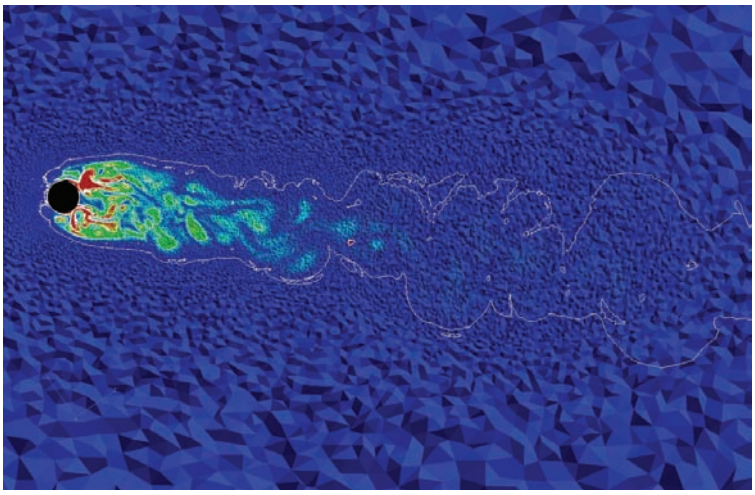


Figure 1: Crinkle cut of grid cross-section and contours of constant vorticity.

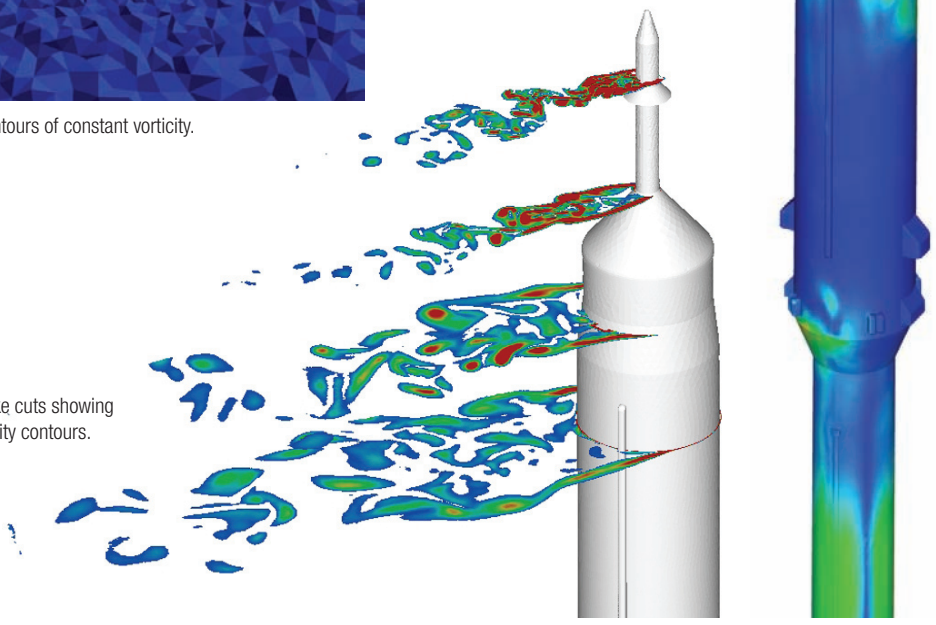
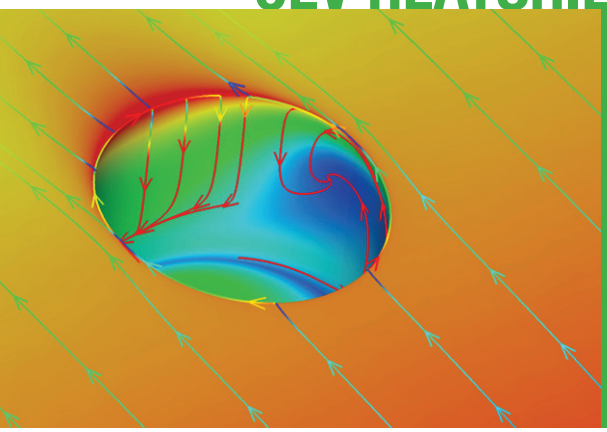


Figure 2: Wake cuts showing constant vorticity contours.



Figure 3: Contours of pressure coefficient peak dynamics for the ground wind loads checkout model.

COMPUTATIONAL SUPPORT FOR ORION CEV HEATSHIELD TPS DESIGN AND ANALYSIS



EXPLORATION SYSTEMS MISSION DIRECTORATE

MICHAEL J. WRIGHT
NASA Ames Research Center
(650) 604-4210
Michael.J.Wright@nasa.gov

◀ Close-up of Figure 2.

Project Description: A thermal protection system (TPS) will protect NASA's Orion Crew Exploration Vehicle (CEV) as it returns from space. In order to understand the extreme environments that the TPS must withstand, aerothermal computational fluid dynamics (CFD) is used to simulate both ground tests and the actual flight environment. NASA's CEV Aerosciences Program (CAP) is responsible for generating aerothermal databases to be used for sizing the TPS system for Orion, while the TPS Advanced Development Project (TPS-ADP) is responsible for simulating high-energy TPS tests and damaged TPS materials. The primary goals of the TPS-ADP are to develop two heat shield concepts to be presented at the Orion TPS downselect review in spring 2009, and to recommend a single design to the CEV Project Office. Aerothermal analyses of high-energy arc jet tests and the effects of in-orbit micrometeorite damage are key components of the overall risk and reliability assessment of the TPS system.

Arc jets are powerful, high-energy facilities used for testing TPS materials in environments similar to those encountered during atmospheric entry. The largest NASA arc jet can operate at a rated power of 60 megawatts per kilogram of gas for up to an hour at a time to simulate extreme entry conditions. The TPS-ADP conducts testing in arc jet facilities at NASA Ames Research Center and NASA Johnson Space Center, as well as in a Department of Defense facility in Tullahoma, Tennessee. Because of the high energy levels and great degree of thermo-chemical non-equilibrium in these tests, analysis is required to understand these facilities and the experimental results. Therefore, CFD simulations are combined with measured data to determine important flow quantities that cannot be ascertained from experiment alone. These simulations are a vital step in both designing the experimental program and understanding how materials perform in intense, chemically reacting flows. Results from the testing and analysis help to drive material selection and overall heatshield design, and will form a cornerstone of the TPS selection.

The potential effects of TPS damage due to micrometeoroid orbital debris (MMOD) strikes are simulated using CFD and material response codes. The required analysis includes: high-velocity impact testing of the TPS materials to determine what types of damage can occur after an MMOD impact; aerothermal CFD analysis to determine the heating augmentation that would result from this type of damage; and ablative material response analysis to determine whether the damaged TPS would still maintain the required internal temperatures.

Relevance of Work to NASA: This work enables both experimental support and assessment of localized damage effects that are pertinent to the design and materials selection for the CEV heat shield. The CEV is NASA's next-generation vehicle for human space operations, both for missions to low-Earth orbit and to eventually carry astronauts to the Moon and return them safely to Earth.

Computational Approach: Aerothermal environments are simulated using the Data Parallel Line Relaxation (DPLR) hypersonic real-gas flow solver developed at NASA Ames. Researchers from Ames' Space Technology Division have developed best practices for use of both the parallel DPLR code and supporting utility codes to accurately characterize chemically reacting arc jets and test specimens. Such simulations rely heavily on experimental measurements for boundary conditions. These simulations are geometrically simple but physically complex, requiring the simultaneous solution of up to 17 coupled partial differential equations, finite-rate chemistry, and thermo-chemical non-equilibrium. A representative arc jet simulation, such as the one shown in Figure 1, may use upwards of 180 processors for up to 5 hours, and several of these solutions are often necessary to match the experimental measurements.

The DPLR code is also used for orbital debris damage simulations, which are split between larger capsule-only runs, and

local damage site runs. This process leverages existing techniques that have been used regularly on the Columbia supercomputer during Space Shuttle missions. Each typical orbital debris damage simulation may require over 15 to 20 hours of computing using up to 200 processors. Figure 2 shows a simulation of the full CEV with small-scale damage to the shoulder region.

Results: Analysis results from this project have contributed to the last three design cycles of the CEV spacecraft, as well as to the documentation delivered to Lockheed Martin as part of the transition of primary responsibility for heat shield development. The TPS-ADP is currently on track to deliver two heat shield design concepts for Phenolic Impregnated Carbon Ablator and Avcoat TPS materials, including an MMOD damage risk assessment for each, for the Orion Preliminary Design Review (PDR).

Role of High-End Computing: The NASA Advanced Supercomputing (NAS) facility supercomputers are used for the CFD simulations of damaged TPS materials—the most numerically intensive part of these analyses. The availability of the Columbia supercomputer was an important contributor to the TPS design team's success.

Future: The TPS Advanced Development Project will continue to use HEC resources to support the final TPS down-selection in the spring of 2009. The TPS Insight/Oversight team will use HEC resources to independently verify Lockheed Martin analysis results between the PDR and the Critical Design Review.

Note: On April 7, 2009, NASA selected the Avcoat ablator system for the Orion crew module.

Co-Investigators

- Todd White, Mike Barnhardt, Tahir Gökçen, Dinesh Prabhu, all of ELORET Corporation

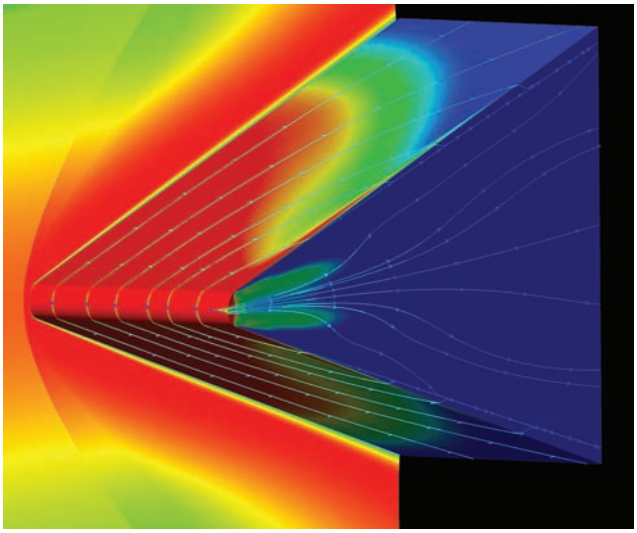


Figure 1: Data Parallel Line Relaxation (DPLR) simulation of a thermal protection system sample wedge in an arc jet flow.

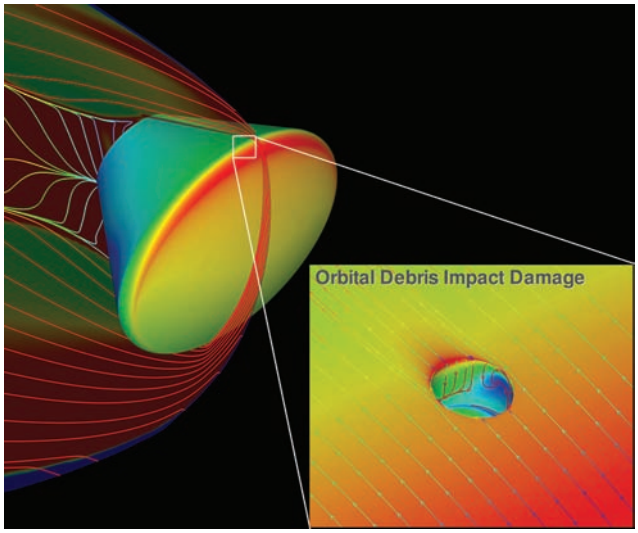
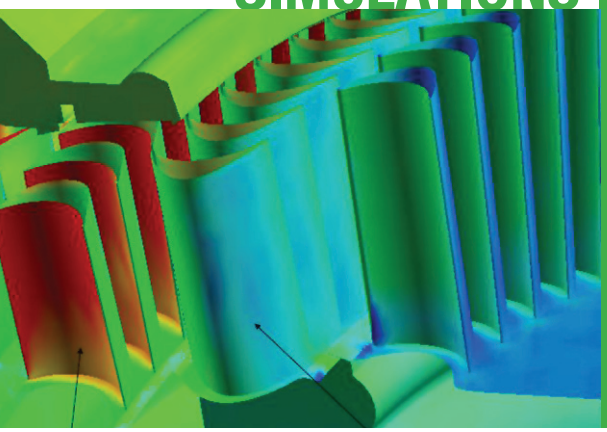


Figure 2: DPLR simulation of the full Crew Exploration Vehicle, with small-scale damage to the shoulder region.

J-2X FUEL AND OXIDIZER TURBINE SIMULATIONS INCLUDING DISC AND TIP CAVITIES



EXPLORATION SYSTEMS MISSION DIRECTORATE

PRESTON SCHMAUCH
NASA Marshall Space Flight Center
(256) 544-1218
Preston.B.Schmauch@nasa.gov

◀ Close-up of Figure 1.

Project Description: The Ares I upper stage turbopumps operate in a highly dynamic environment that can be very sensitive to small changes in the flow path. The heritage J-2 turbines were designed with tip seals that minimized flow between the main flow path and the disc and tip cavities. The current J-2X configuration, however, has the tip seals removed and the main flow path is free to interact with any flows in the cavities. Our objective is to simulate the current J-2X turbine configurations and determine what performance and fluid interaction effects will result from the removal of the tip seals.

To better understand the flow effects that will result from removing the tip seals in the current configurations, we are performing full-annulus, unsteady computational fluid dynamics (CFD) simulations of the oxidizer and fuel turbines, both with and without cavities. The simulations must be run full-annulus because flow in the cavities is not a periodic phenomenon. By comparing fluid property histories between the simulations with and without cavities, we will better understand the magnitude of the impact the design change will have on the performance and life of the parts. The results will also be used to increase our current understanding of the conservatisms that are built into modeling and analysis methods.

Relevance of Work to NASA: This project directly affects the Ares I vehicle by reducing risk associated with the J-2X engine and, thereby, the upper stage. Its results will help NASA and contractors make more educated decisions on possible design modifications for current turbopump configurations. These simulations could also reveal benefits or areas of concern that might not be considered otherwise. Knowledge from these simulations will also be applied to other engines and models used in NASA programs.

Computational Approach: Phantom, a NASA code that has been anchored and validated for supersonic turbines, is being used for the unsteady CFD simulations. Phantom uses three-dimensional, unsteady Navier-Stokes equations as the governing equations. The Baldwin-Lomax turbulence model is used for turbulence closure. An overset O and H grid topology with moving grids to model blade motion is employed for the simulations. With all blades modeled in these simulations, more than 65 million computational nodes were used on the turbine alone. During this project, the Phantom code was also adapted to allow for the addition of the multiple, attached tip and disc cavity grids necessary for these simulations.

Results: Simulations of the J-2X oxidizer turbine with and without cavities were recently completed. Preliminary post-processing has shown significant blade-loading and dynamic pressure environment differences between the cavity and non-cavity cases. Figure 1 shows surface pressure contours from simulation of the J-2X oxidizer turbine with cavities. Preliminary stress analysis shows an increase in the factor of safety on stress when running with the cavities attached. It is believed that the cavities reduce the highly unsteady blade-loading due to the row-to-row blade passing, and convert that energy into a more broadband unsteadiness. More analysis is underway to investigate other implications this effect could have on stress margins. The J-2X fuel turbine model currently has the first disc cavity added and the simulation is still being worked.

Role of High-End Computing: The size of the simulations for this project required a minimum of about 400 gigabytes of memory. The J-2X oxidizer turbine simulation with cavities alone required 256 processors on the Columbia supercomputer running for approximately 200 days, or 1.2 million processor-hours. More than 2 million additional processor-hours were used to simulate various power balances and configurations of the turbines without the cavities. Without the high-end computing resources at the NASA Advanced Supercomputing (NAS) facility, we would not have had the capability to run these simulations.

Future: The computational grids for the J-2X fuel turbine with cavity are currently being generated and will be simulated using NAS HEC resources. It is expected that the results of these simulations will call for further analyses with refined focus on specific topics. Currently, old power balances are being used for the cavity simulations so that a one-to-one comparison can be done for the case without cavities. Future analyses may be done with updated power balances.

Co-Investigators

- Daniel Dorney, NASA Marshall Space Flight Center

Publications

[1] Dorney, D., Griffin, L., and Schmauch, P., "Unsteady Flow Simulations for the J-2X Turbopumps," 54th JANNAF Propulsion Meeting, May 2007.

[2] Marcu, B., Tran, K., Dorney, D., and Schmauch, P., "Turbine Design and Analysis for the J-2X Engine Turbopumps," 44th AIAA Joint Propulsion Conference, *AIAA 2008-4660*, July 2008.

[3] Marcu, B., Zabo, R., Dorney, D., and Zoladz, T., "The Effect of Acoustic Disturbances on the Operations of the Space Shuttle Main Engine Fuel Flowmeter," 43rd AIAA Joint Propulsion Conference, *AIAA 2007-5534*, July 2007.

[4] Dorney, D., Griffin, L., Marcu, B., and Williams, M., "Unsteady Flow Interactions between the LH2 Feedline and SSME LPFP Inducer," 42nd AIAA Joint Propulsion Conference, *AIAA 2006-5073*, July 2006.

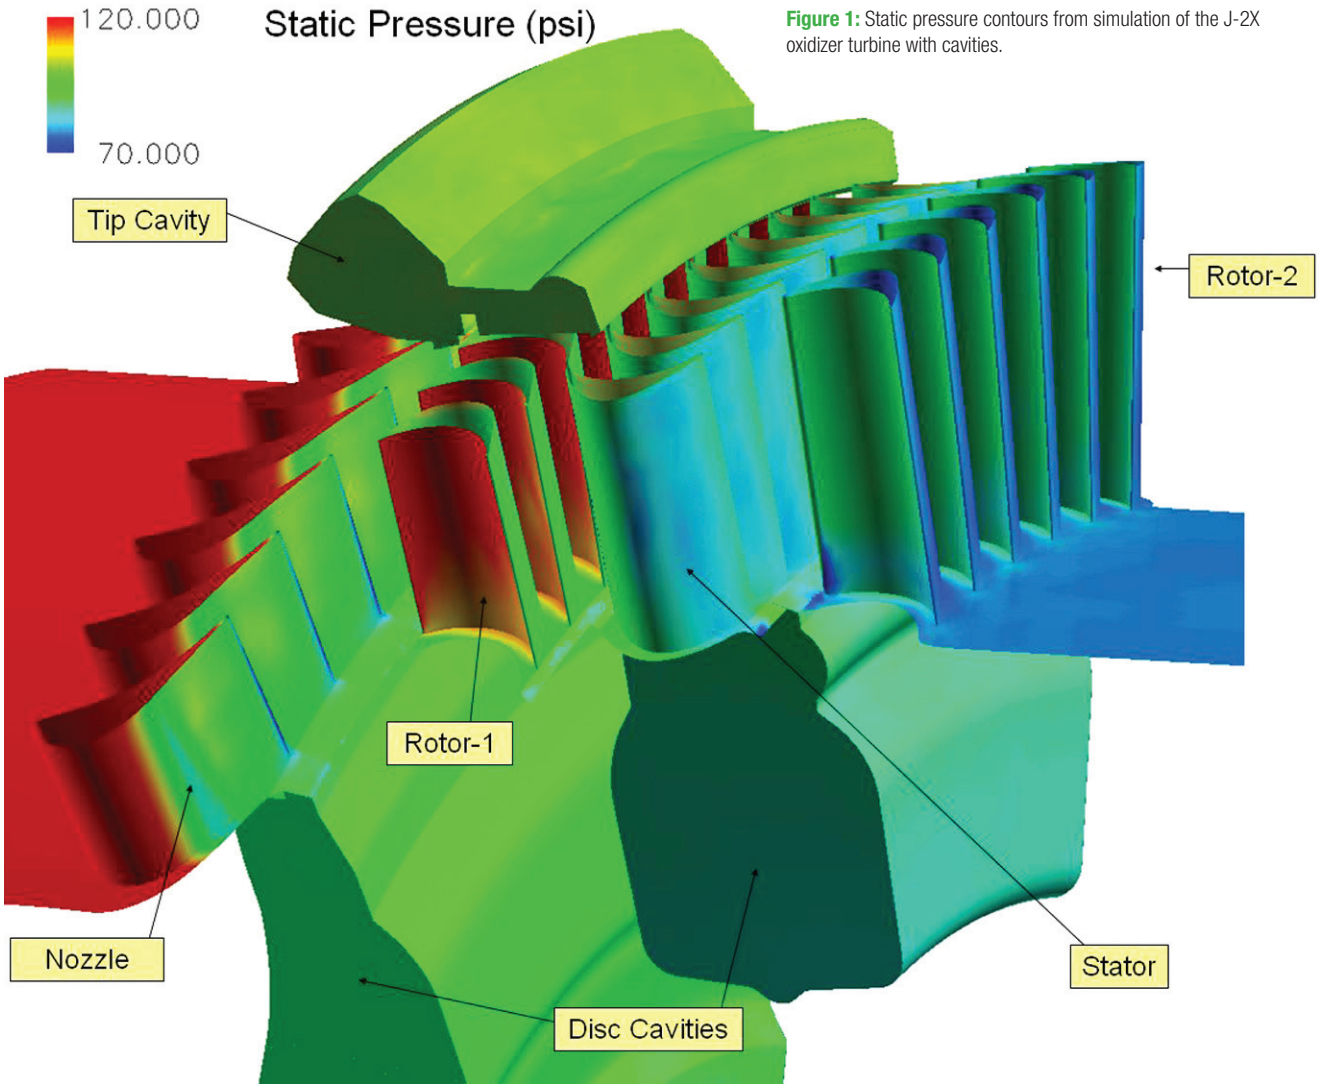
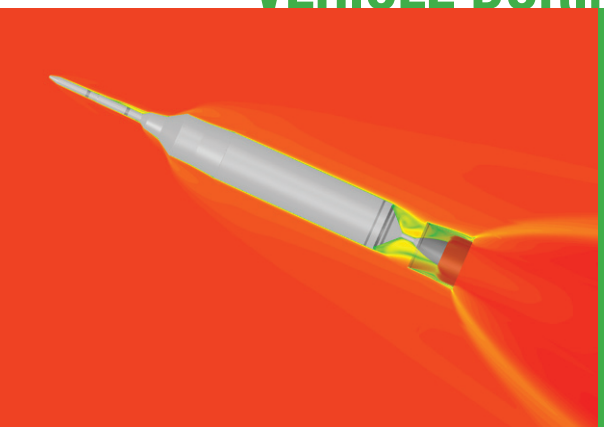


Figure 1: Static pressure contours from simulation of the J-2X oxidizer turbine with cavities.

PROXIMITY AERODYNAMICS OF THE ARES I LAUNCH VEHICLE DURING STAGE SEPARATION MANEUVERS



EXPLORATION SYSTEMS MISSION DIRECTORATE

GOETZ H. KLOPFER
NASA Ames Research Center
(650) 604-3993
Goetz.H.Klopfer@nasa.gov

◀ Close-up of Figure 1.

Project Description: The objective of this project is to quantify proximity effects on the aerodynamics of the Ares I Crew Launch Vehicle (CLV) during inflight separation maneuvers such as rocket booster stage separation and crew capsule separation in the course of an emergency abort. This can be challenging because the proximity effects vary with flow conditions such as flight Mach number, local atmospheric pressure, angle of attack, and Reynolds numbers. In addition, proximity effects are strongly dependent on the relative positions of the separating components of the CLV, such as separation distance, relative angle of attack, and lateral offsets. Computational fluid dynamics (CFD) analyses are used along with wind tunnel data to develop aerodynamic databases for use in designing the Ares I launch vehicle. A practical database construction requires many hundreds of computations for a given launch vehicle configuration and set of flow conditions.

This project was divided into several phases. First, a preliminary study of crew capsule separation during aborts using Apollo-based configurations established computational best practices and examined the flow properties of the proximity effects. In the second phase, extensive numerical simulations of several types of stage separation scenarios were generated to provide data for a staging trade study conducted at NASA Marshall Space Flight Center. The third and ongoing phase is a parametric study of proximity effects on the Ares I configuration, with the staging process selected by the Ares Staging Trade Study group. The next phase will involve a similar parametric study with the inclusion of plume effects for the upper stage settling motors, booster separation motors, and upper stage main engine.

Relevance of Work to NASA: Results of this study have provided critical insight into the proximity effects on the aerodynamic performance and risk assessment of the Ares I launch vehicle. The proximity aerodynamic data with plume effects that will be generated for a large parametric matrix will support Guidance, Navigation, and Control (GN&C) and risk assessment studies to achieve a precise and safe stage separation process.

Computational Approach: Most of the flow simulations were conducted on the Columbia supercomputer using the OVERFLOW Navier-Stokes solver with overset structured meshes. Both steady-state and time-accurate simulations were required, depending on the separation distances and other parameters. For small separation distances between the stages, the flowfield was typically steady. At larger distances, however, flows were usually unsteady and time-accurate simulations were required. In addition, inviscid simulations were also conducted using the Cart3D Cartesian mesh generation code to provide guidance in establishing grid resolution requirements for the viscous OVERFLOW simulations. Prescribed motion and fully coupled six-degrees-of-freedom (6-DOF) simulations were also used when appropriate.

Results: The numerical simulations provided important data for the Staging Trade Study group and enabled the Constellation Engineering Management Council to select and finalize the stage separation process to be used for the CLV. The numerically derived data provided both guidance for setting up wind tunnel tests and data impossible to obtain from wind tunnel tests, such as proximity effects at very small or very large separation distances. Figure 1 shows a typical result from the staging trade study effort, capturing the interaction of the J-2X upper stage main engine plume with the interstage (IS) as it falls through the main plume. The concern for this type of stage separation was whether the IS would collide with the J-2X nozzle bell.

Role of High-End Computing: The scale of parallel computing required for this work—involving upwards of 256 processors for each solution, grid sizes of 100–200 million grid points, and runtimes of up to 20,000 processor-hours per case for the

steady-state problems—was quite large by production CFD standards. The NASA Advanced Supercomputing (NAS) Division provided an exceptional level of support for this project, with timely allocations of Columbia supercomputer nodes needed for running several jobs simultaneously, and computing hours needed to meet project deadlines.

Future: The next phase of this project will be to study plume effects on proximity aerodynamics. This data will be part of aerodynamic database generation in support of future Ares I design cycles. The computing requirements will ramp up substantially due to increased grid resolution requirements and the increased number of parameters involved. Continuing use of NAS high-end computing resources will be critical for achieving the Ares I Project mission objectives.

Co-Investigators

- Jeffrey Onufer, Shishir Pandya, and William Chan, all of NASA Ames Research Center
- Veronica Hawke-Wong, ELORET Corp.

Publications*

[1] Chan, W.M., Klopfer, G.H., Onufer, J.T., and Pandya, S.A."Proximity Aerodynamics Analyses for Launch Abort Systems," *AIAA Paper 2008-7326*, 26th AIAA Applied Aerodynamics Conference, Honolulu, Hawaii, August 2008.

* *Most of this work is International Traffic in Arms Regulations (ITAR)-controlled and is not published.*

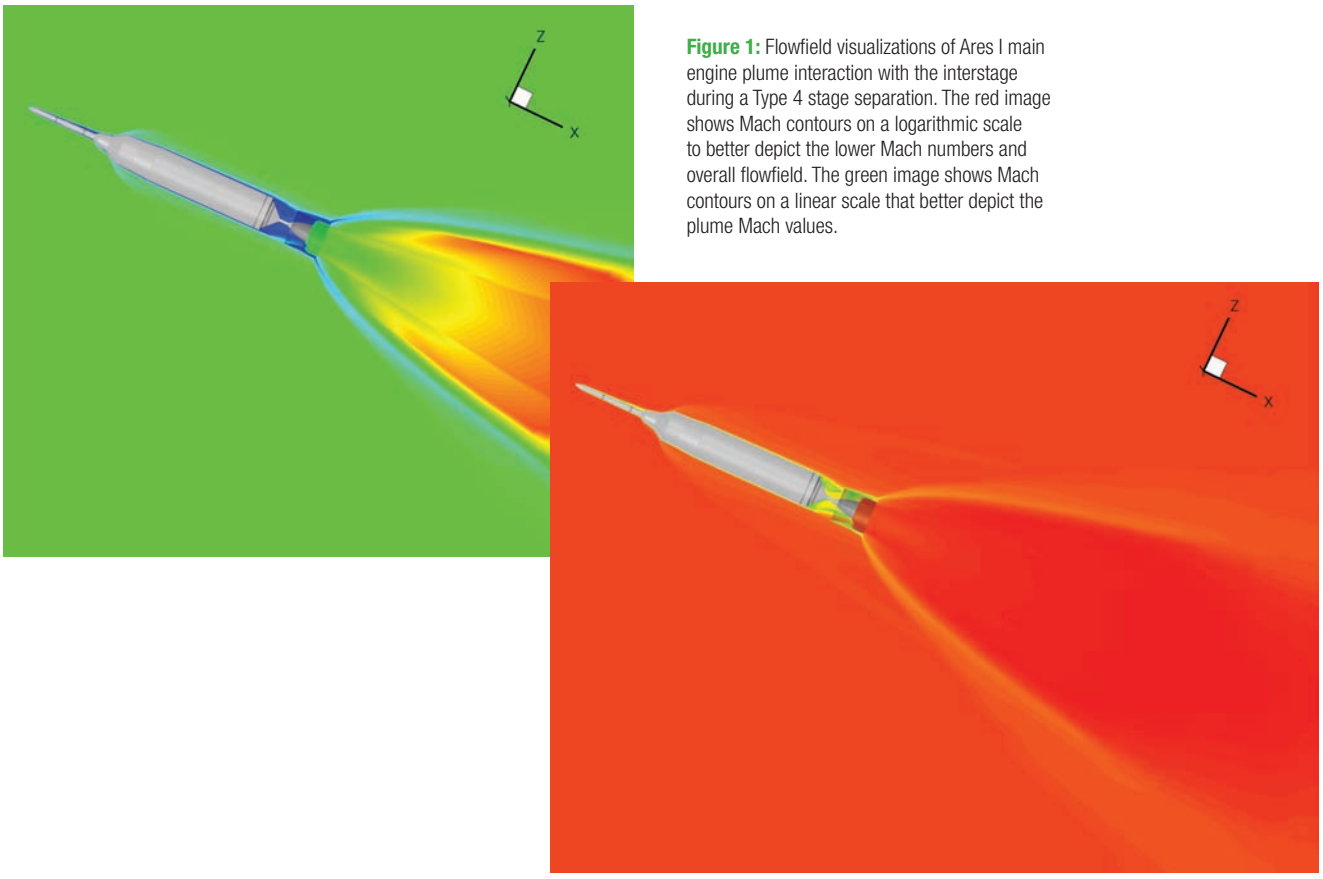
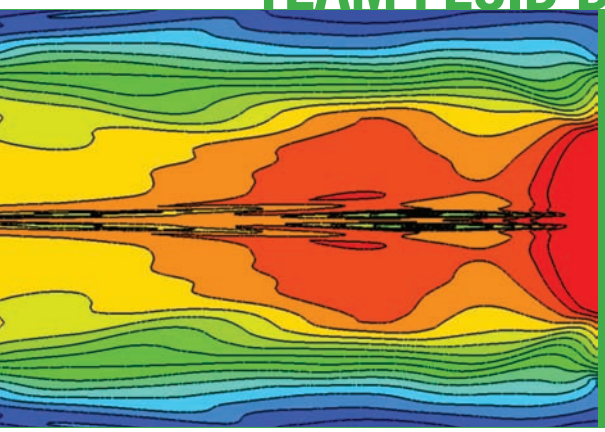


Figure 1: Flowfield visualizations of Ares I main engine plume interaction with the interstage during a Type 4 stage separation. The red image shows Mach contours on a logarithmic scale to better depict the lower Mach numbers and overall flowfield. The green image shows Mach contours on a linear scale that better depict the plume Mach values.

THRUST OSCILLATION FOCUS TEAM FLUID DYNAMICS ANALYSIS SUPPORT



EXPLORATION SYSTEMS MISSION DIRECTORATE

PHILIP DAVIS
NASA Marshall Space Flight Center
(256) 544-7728
Philip.A.Davis@nasa.gov

◀ High-resolution close-up of Figure 2.

Project Description: A major issue facing the Ares I Crew Launch Vehicle (CLV) is the presence of thrust oscillations that are predicted to occur at the vehicle's acoustic resonance frequency, potentially causing problems for the vehicle and the astronauts aboard. NASA engineers are currently investigating strategies to reduce or eliminate the thrust oscillations. This project aims to assist those engineers by examining the internal flow dynamics of the Space Shuttle Reusable Solid Rocket Motor (RSRM) to learn the cause of the thrust oscillations in the current four-segment RSRM and determine how these problems can be corrected for the future five-segment booster.

Computational fluid dynamics (CFD) techniques are being applied to RSRM geometries at burn times of 80 and 110 seconds using Phantom, a NASA-developed, three-dimensional flow solver. These times were chosen because they exhibit the greatest thrust oscillations in data acquired from the firing of test motors and past shuttle motors. This study centers on the vortex shedding that occurs downstream of the three inhibitors that protrude into the flow within the motor. The number of vortices at each location, the frequency of creation, and the relative strength of each vortex is under investigation, as well as how the vortices affect the acoustics of the motor. Figures 1 and 2 show a shadowgraph and flow velocity contours for the 80-second RSRM burning profile. Additionally, fluid dynamics analysis is being performed at NASA Marshall Space Flight Center's Cold Flow Testing Facility. This facility has been used in the past to test the internal acoustics of the RSRM using a 10% scale model with air as the fluid. The current CFD analysis focuses on past testing performed at the facility, and will be used to help determine what testing should be done in the future. Figure 3 shows a scaled entropy plot of the Cold Flow Testing Facility main chamber.

Relevance of Work to NASA: This analysis will provide NASA with insight into their current technology, including the performance of the Space Shuttle RSRM and how to perform future testing inside the NASA Marshall Cold Flow Testing Facility. It is expected that this information will help to resolve the thrust oscillation problems currently facing the Ares I vehicle.

Computational Approach: Phantom is a fully three-dimensional, unsteady, finite-difference code that is third-order accurate in space and second-order accurate in time. It utilizes the general equations set for liquids and gases, and is capable of handling two-phase flows. The turbulence model is a modified Baldwin-Lomax model. The code also has the ability to use A^*p^n boundary conditions for burning surfaces.

Results: The first simulations for this project were begun in December 2007. Inhibitor heights were adjusted to correct levels in March 2008, and corrected head-end burning was implemented in June 2008. In July 2008, the project began collecting data for analysis with the final iteration of geometries. Currently, simulations have generated one second of data for each of the cases that are still running. Four simulations for the Cold Flow Facility geometry are running with varying flow conditions to attempt to match heritage data.

Role of High-End Computing: Each of the RSRM geometries consists of 31,130,190 computational grid cells, which are divided into 120 passages run in parallel on 120 processors using the Message Passing Interface (MPI) library on the Columbia supercomputer. This system has been vital and necessary to achieving the grid density needed to provide the best analysis possible, as even using 120 processors per simulation only accumulates approximately 0.025 seconds of data in 24 hours. At least 2 seconds of CFD data must be generated for the RSRM simulations, which would be impossible without the computing resources that have been provided by the NASA Advanced Supercomputing Division.

Future: The current CFD simulations for the RSRM geometries are not yet complete and will continue to accumulate data. The Cold Flow Facility CFD simulations are still at the beginning stages of data collection and will continue to run to accumulate more data. Once these CFD simulations are completed, the next phase of analysis could include assessing mitigation strategies for reducing thrust oscillations. One example of this would be to alter the inhibitor heights at one or several of the joint locations. Other analyses that could be performed include using the Ares I solid booster geometry as well as future Cold Flow Facility setups.

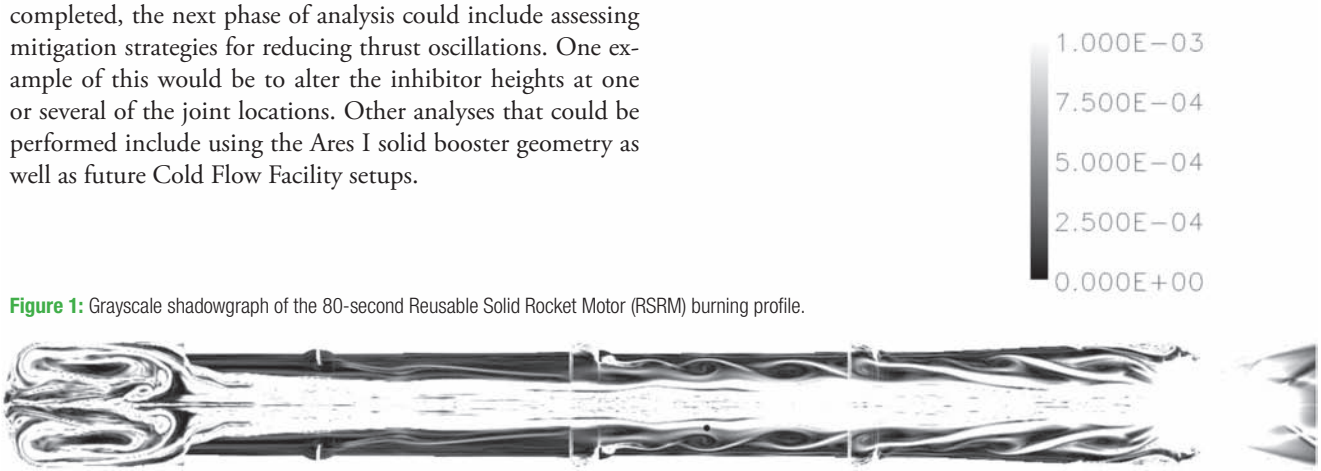


Figure 1: Grayscale shadowgraph of the 80-second Reusable Solid Rocket Motor (RSRM) burning profile.

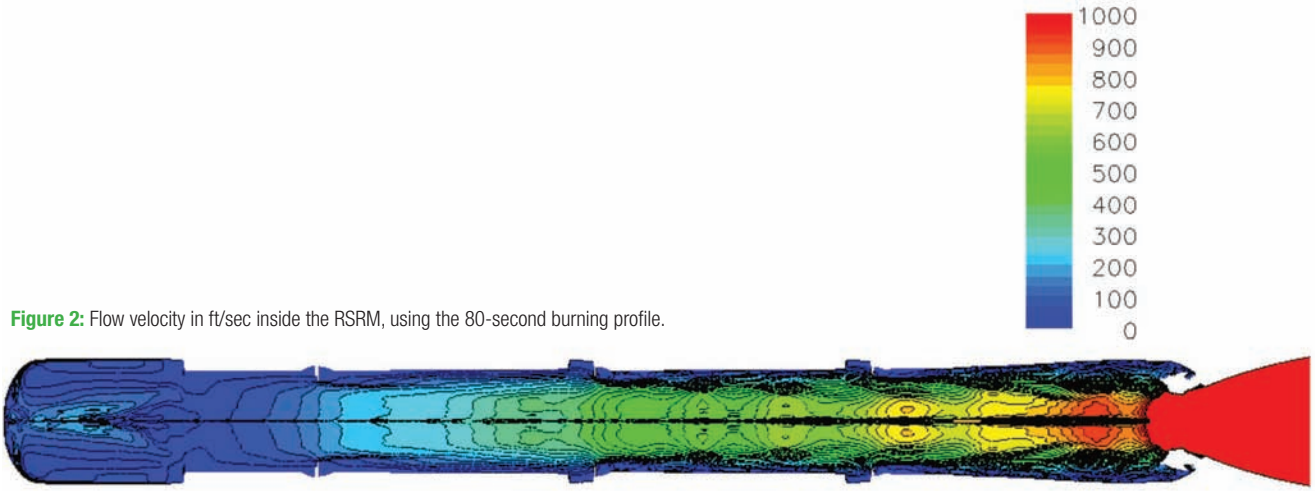


Figure 2: Flow velocity in ft/sec inside the RSRM, using the 80-second burning profile.

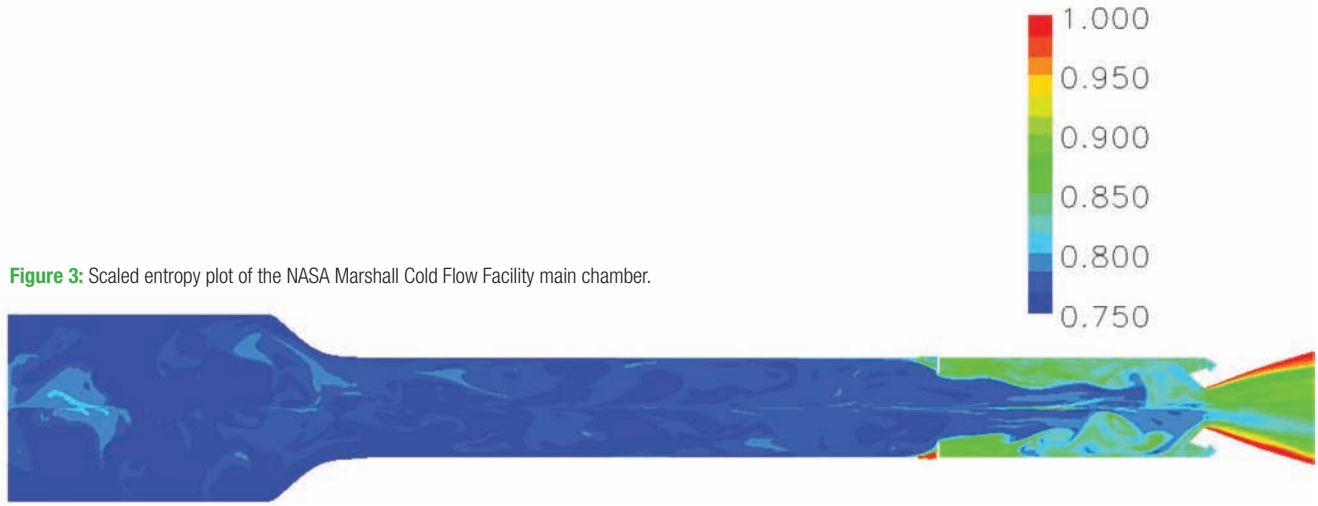
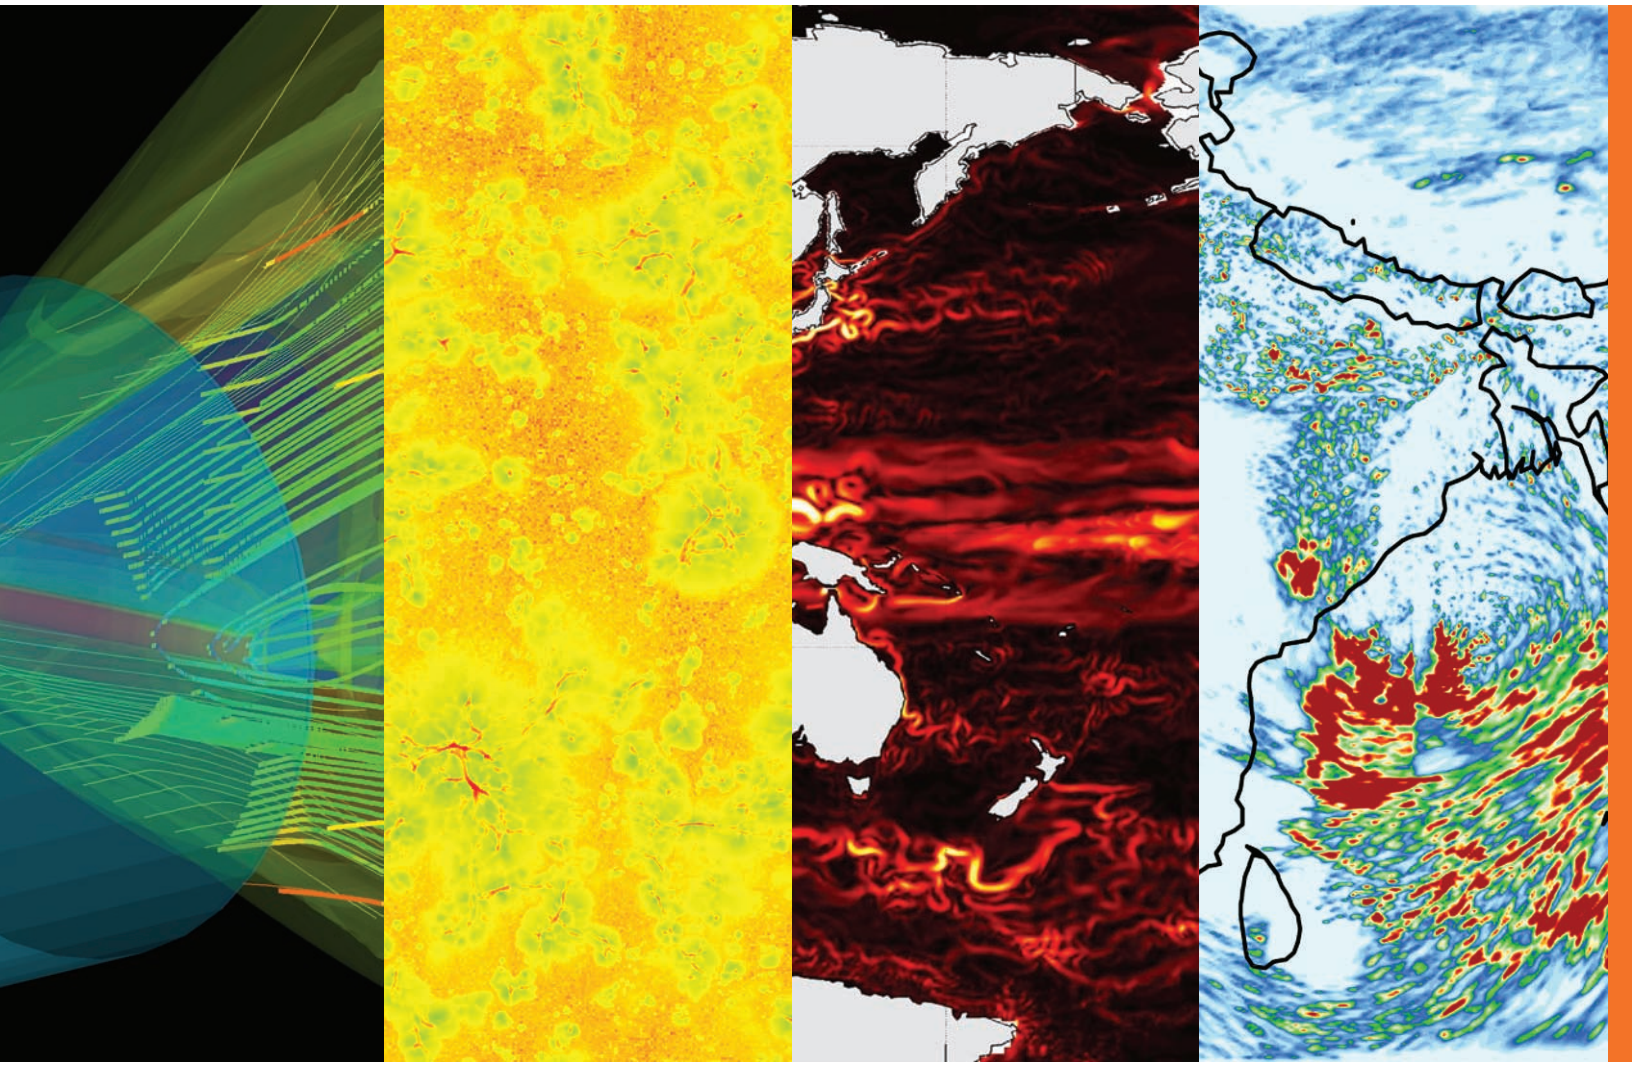


Figure 3: Scaled entropy plot of the NASA Marshall Cold Flow Facility main chamber.



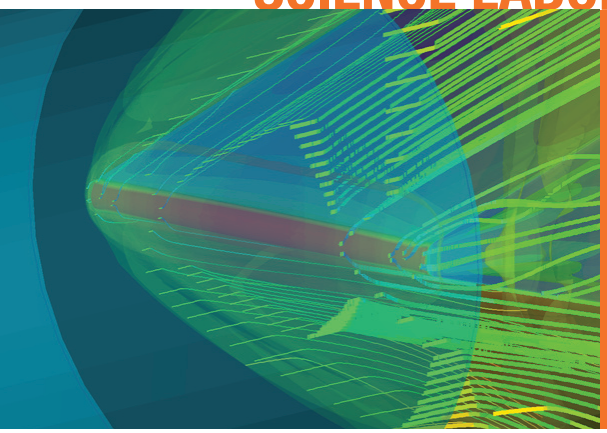
MISSION DIRECTORATE

SCIENCE

NASA's Science Mission Directorate conducts scientific exploration that is enabled by access to space. We project humankind's vantage point into space with observatories in Earth orbit and deep space; spacecraft visiting the Moon, Mars, and other planetary bodies; and robotic landers, rovers, and sample return missions. From space, in space, and about space, NASA's science vision encompasses questions as practical as hurricane formation, as enticing as the prospect of lunar resources, and as profound as the origin of the universe. The Science Mission Directorate organizes its work into four broad scientific pursuits: Earth Science, Planetary Science, Heliophysics, and Astrophysics.

DR. EDWARD J. WEILER
Associate Administrator
http://nasascience.nasa.gov/

AEROTHERMAL ANALYSIS IN SUPPORT OF MARS SCIENCE LABORATORY HEATSHIELD QUALIFICATION



SCIENCE MISSION DIRECTORATE

MICHAEL J. WRIGHT
NASA Ames Research Center
(650) 604-4210
Michael.J.Wright@nasa.gov
http://mars.jpl.nasa.gov/msl/

◀ Close-up of Figure 1.

Project Description: The Mars Science Laboratory (MSL) is a flagship-class rover mission scheduled for launch in winter 2011. The purpose of this project is to provide computational fluid dynamics (CFD) analysis to support the design, testing, and qualification of the Phenolic Impregnated Carbon Ablator (PICA) heatshield, Super Lightweight Ablator (SLA) backshell, and Acusil-II Parachute Closeout Cone (PCC) thermal protection system (TPS) that will protect the spacecraft from the harsh entry environment at Mars.

This is an applied engineering—as opposed to a research—project. The work consists of aerothermal CFD simulations of the chemically reacting, non-equilibrium hypersonic flowfield around the TPS material samples tested in an arc jet environment. We also analyze the hypersonic flowfield about the MSL spacecraft during Martian entry. Key areas of interest are the turbulent heating and shear stress on the heatshield and the local transient heating rates and flow topologies produced by firing reaction control system (RCS) thrusters on the backshell. The simulations use the Data-Parallel Line Relaxation (DPLR) and the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) codes. DPLR, developed at NASA Ames Research Center by this project's principal investigator and collaborators, was a co-winner of NASA's 2007 Software of the Year Award.

Relevance of Work to NASA: Material testing after the Critical Design Review uncovered a catastrophic failure mode in the baseline heatshield material for MSL. The uncovered failure mode forced the project to change the heatshield TPS material to one with a lower technology readiness level (TRL) because, at the time, there was less than one year to develop, design, and qualify the new concept before launch (the original launch date for MSL was set for October 2009). When MSL management switched the heatshield TPS material, they identified the TPS system as the largest risk item to the entire MSL program. A critical part of the engineering analysis required to reduce this risk is CFD simulations of the arc jet testing performed on the material over the range of expected entry environments, and corresponding analysis of the expected flight environment.

Computational Approach: We test TPS materials using a 60-megawatt test bay at the NASA Ames Arc Jet Complex, as well as smaller facilities at Ames and the Arnold Engineering Development Center. At these energy levels, the supersonic gas that flows over the test articles is highly dissociated and in a state of extreme non-equilibrium. Simulations of this complex flow require state-of-the-art computational tools that can model both the fluid dynamics and the chemical processes involved. We use the DPLR aerothermal CFD code to perform these simulations. Full 3D, non-equilibrium reacting flow Navier-Stokes calculations of the test article in the arc jet flow predict the incident heating, pressure, and shear stress on the model. The simulations are very complex, involving millions of grid points and up to 16 chemically reacting species. We then use these predicted quantities as boundary conditions for another code, which predicts the response of the ablative material. We perform flight simulations using similar physical models in DPLR and LAURA. Similar tools are necessary because the energy content of the flow impacting the flight environment is similar to that encountered in the arc jet.

Results: This project began in October 2007, when MSL management changed the TPS material. Between October 2007 and July 2008, we performed dozens of simulations of both the flight and arc jet environments (Figures 1 and 2). These results successfully demonstrated the adequacy of the PICA heatshield for MSL at the final TPS review in July 2008. Since July, work has continued in support of final developmental tests and in the planning and execution of the qualification test program. Although the mission has been delayed, all elements of the aeroshell TPS were designed, tested, and built in time to support the original launch date.

Role of High-End Computing: The computational power of the Columbia supercomputer at the NASA Advanced Supercomputing (NAS) facility was an enabling capability for this work. The project was extremely schedule-driven given the original, fixed launch date of the spacecraft. It was only with priority access to Columbia that we were able to generate the engineering data to demonstrate that the heatshield would perform.

Future: This project will terminate with the launch of MSL in winter 2011.

Co-Investigators

- Chun Tang, Todd White, Dinesh Prabhu, all of ELORET Corp.
- Karl Edquist, Artem Dyakonov, NASA Langley Research Center
- James Brown, NASA Ames Research Center

Publications

[1] Wright, M., "Sizing and Margins Assessment of the Mars Science Laboratory Aeroshell Thermal Protection System," submitted to the 41st AIAA Thermophysics Conference, June 2009.

[2] Tang, C., "Numerical Simulations of Protruding Gap-Fillers on the Mars Science Laboratory Heatshield," submitted to the 41st AIAA Thermophysics Conference, June 2009.

[3] Prabhu, D., "CFD Analysis Framework for Arc-Heated Flowfields I: Stagnation Testing in Arc Jets at NASA ARC," submitted to the 41st AIAA Thermophysics Conference, June 2009.

[4] Prabhu, D.,"CFD Analysis Framework for Arc-Heated Flowfields II: Shear Testing in Arc Jets at NASA ARC," submitted to the 41st AIAA Thermophysics Conference, June 2009.

[5] White, T., "CFD and Material Response Framework for Wedge Testing in AEDC H2," submitted to the 41st AIAA Thermophysics Conference, June 2009.

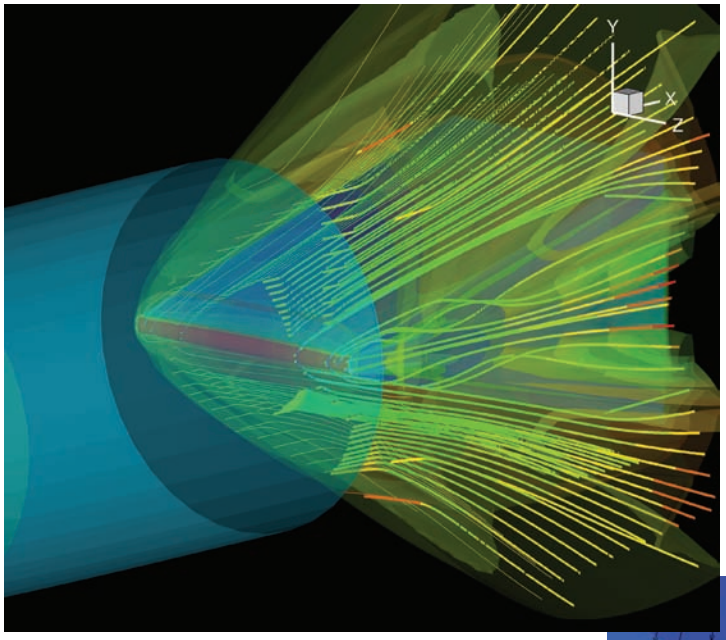


Figure 1: Simulation of a Phenolic Impregnated Carbon Ablator (PICA) wedge sample in shear at NASA Ames Research Center's Interaction Heating Facility. PICA is the heatshield material selected for the Mars Science Laboratory (MSL).

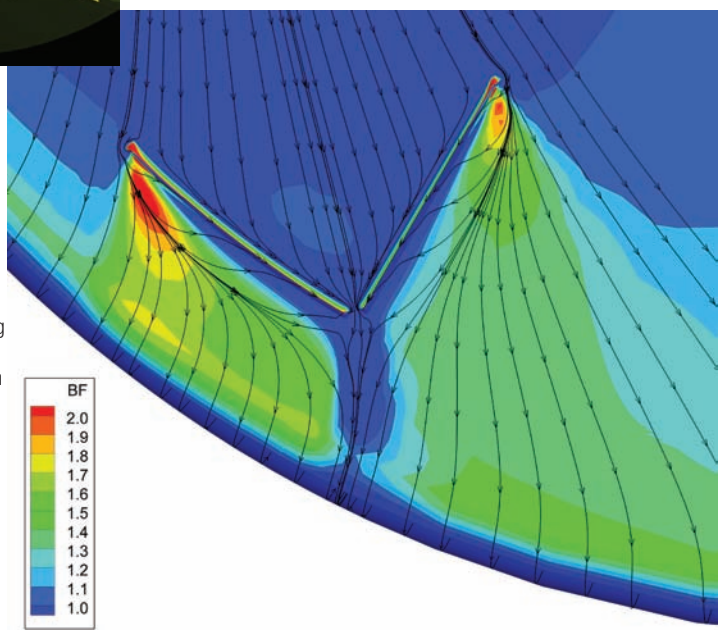
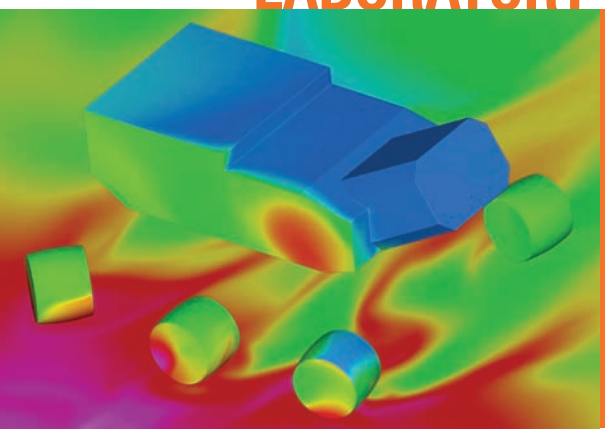


Figure 2: Simulation of an MSL flight heatshield showing augmented heating due to gap-filler protrusion. BF is the bump-factor, defined as the ratio of heating to that which would be encountered with zero protrusion.

CFD SUPPORT FOR MARS SCIENCE LABORATORY ENTRY, DESCENT, AND LANDING



SCIENCE MISSION DIRECTORATE

JOHN W. VAN NORMAN
Analytical Mechanics Associates, Inc.
(757) 864-7997
John.W.Vannorman@nasa.gov
http://mars.jpl.nasa.gov/msl/

◀ Close-up of Figure 2.

Project Description: The Mars Science Laboratory (MSL) is a large rover being designed to perform various planetary science tasks on Mars towards the objective of determining whether life existed or could have existed there. The MSL payload and entry system are each at least three times as massive as any previously flown to Mars, but the mission will need to perform a precision landing within an ellipse several times smaller than that of prior missions.

To achieve precision landing capability, the MSL capsule will fly a lifting trajectory with coordinated banked turns to reduce velocity and align with the landing target. A reaction control system (RCS) consisting of eight small rocket engines will steer the vehicle and damp out oscillations. Depending upon vehicle attitude and flow conditions, RCS plumes may interact with external flow to reduce control authority and induce unintended motions. Such phenomena are termed "aero/RCS interactions." Once near landing, the MSL lander-stage will separate from the backshell and parachute to begin descent under power of eight lander-stage rocket engines. Then, an umbilical will lower the rover to the ground in a "skycrane" maneuver during which the Mars lander-stage engine (MLE) plumes impinge on the ground.

The goal of this project is to provide analysis of fluid dynamics phenomena during MSL entry, descent, and landing (EDL). These phenomena include static aerodynamics, aero/RCS interactions, and MLE plume/ground interactions.

We perform computational fluid dynamics (CFD) simulations of MSL EDL elements to attain a better understanding of EDL element behaviors. Of particular interest have been (i) the characterization of MSL and Phoenix capsule (a 2007–08 Mars lander mission) static aerodynamics and aero/RCS interactions (Figure 1); and (ii) MSL MLE plume-induced rover environments (Figure 2). To this end, we run simulations at flight or wind tunnel conditions using the FUN3D, LAURA, and/or OVERFLOW-2 flow solver codes. We compare numerical results against experimental data where available.

Relevance of Work to NASA: This work is part of the MSL EDL analysis being performed at the Atmospheric Flight and Entry Systems Branch of NASA Langley Research Center. Results and lessons learned will help to advance the state of the art in EDL analysis.

Computational Approach: The research uses both unstructured and structured grid methods. Either approach begins with the surface definition of a flight article or wind tunnel model in the form of Initial Graphics Exchange Specification (IGES) files either acquired through the NASA Jet Propulsion Laboratory or constructed from engineering drawings using the commercial code Gridgen. For the unstructured grid approach, we generate grids using the NASA Langley-developed GridTool and VGrid packages and solve flows with the FUN3D solver (also developed at NASA Langley). For the structured approach, we generate grids using Gridgen and the NASA Ames Research Center-developed Chimera Grid Tool package and solve flows with the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) or OVERFLOW-2.

Results: These CFD simulations have made significant contributions to MSL mission design:

- Results from steady-state Phoenix aero/RCS analyses revealed a general lack of robust control authority and possible control reversal about one axis. As a result, the mission team increased the controller deadbands and did not use the RCS during the successful EDL [1].
- Computations showed that early, tangentially firing MSL RCS designs caused large aero/RCS interactions and enhanced aeroheating. These findings led to a configuration change.

- Static aerodynamic coefficients from simulations agreed closely with those measured in the NASA Langley Unitary Plan Wind Tunnel at supersonic Mach numbers.
- MLE plume simulations produced engineering estimates of particle impingement and thermal environments for the MSL rover [2].

Role of High-End Computing: Simulating viscous multispecies flowfields about complex geometries would be daunting without high-end computing. Some simulations were demanding in terms of sheer size, with grids exceeding the 50-million-cell mark, while others required running several combinations of freestream conditions, vehicle orientation, and RCS configuration to yield clear trends. Use of the Columbia supercomputer at the NASA Advanced Supercomputing (NAS) facility made quick turnaround possible, reducing delivery time from weeks or months to hours or days.

Future: Work is continuing on various MSL EDL elements, including subsonic and transonic aero/RCS interaction analysis and MLE plume impingement studies. The overall goal is to provide high-fidelity CFD analyses to help ensure a successful mission.

Co-Investigators

- Pawel Chwalowski, Analytical Mechanics Associates, Inc.
- Pieter Buning, NASA Langley Research Center

Publications

[1] Dyakonov, A., Glass, C., Desai, P., and Van Norman, J., "Analysis of Effectiveness of Phoenix Entry Reaction Control System," AIAA/AAS Astrodynamics Specialist Conference and Exhibit, 2008.

[2] Sengupta, A., Kulleck, J., Van Norman, J., Mehta, M., and Pokora, M., "Main Landing Engine Plume Impingement Environment of the Mars Science Laboratory," IEEE, in press, 2008.

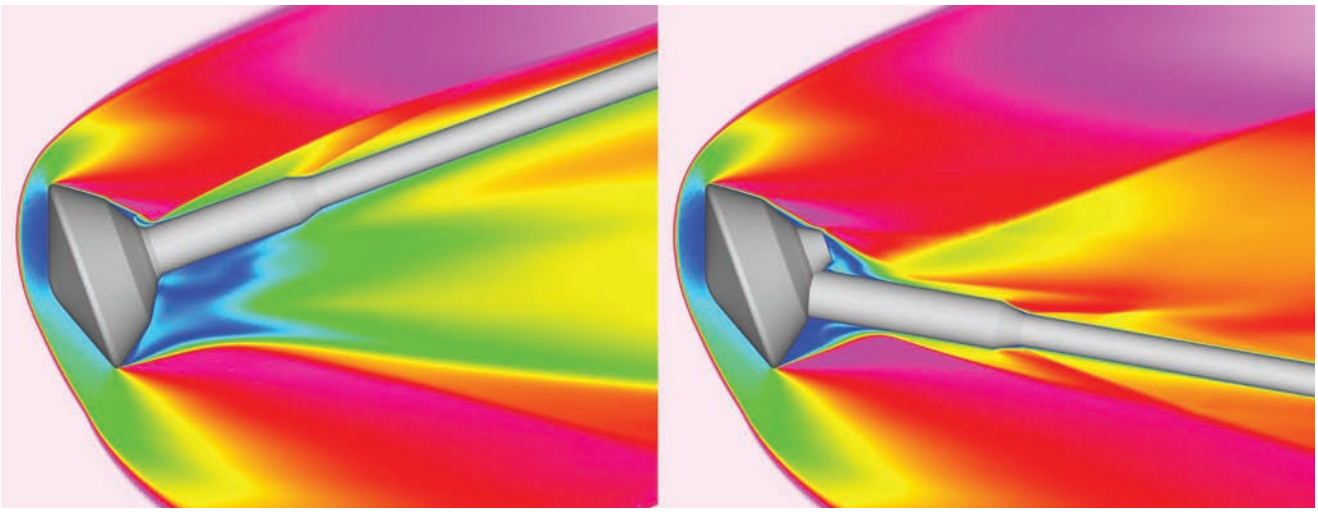


Figure 1: Mach contours surrounding Mars Science Laboratory (MSL) wind tunnel models with axisymmetric (left) and 30° (right) sting configurations at Mach 4.5 and -20° angle of attack.

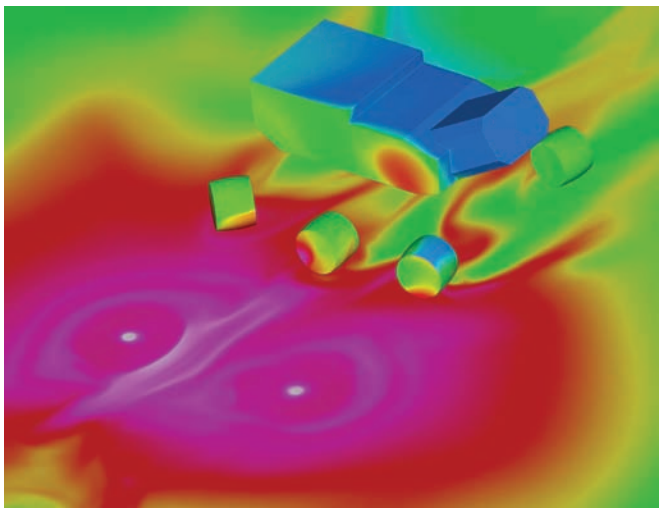
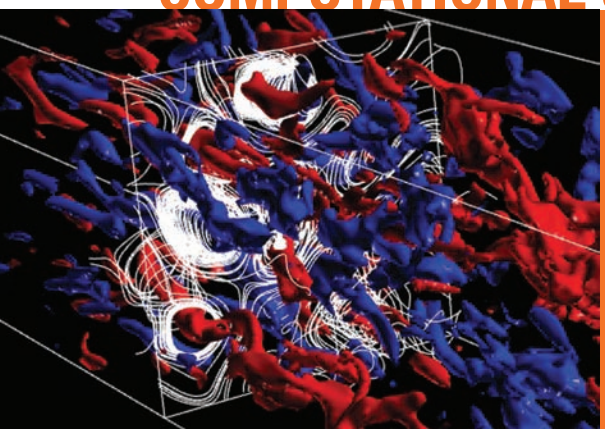


Figure 2: Instantaneous temperature contours surrounding the MSL rover exposed to plumes from the Mars lander-stage engines.

COMPUTATIONAL STUDY OF RELATIVISTIC JETS



SCIENCE MISSION DIRECTORATE

KEN-ICHI NISHIKAWA
National Space Science and Technology Center
(256) 961-7614
Ken-Ichi.Nishikawa-1@nasa.gov

◀ Detail of Figure 1.

Project Description: This project seeks a better understanding of the formation and dynamics of relativistic jets and their emission by supernova remnants, active galactic nuclei, and gamma-ray bursts (GRBs). We are studying how jets form from black holes, how relativistic jets propagate, and how various free-energy sources and their associated heating, acceleration, and plasma transport lead to the excitation of instabilities. We are also investigating ultra-relativistic jets associated with GRBs and their prompt and afterglow emission.

We are using relativistic and general relativistic particle-in-cell (RPIC/GRPIC) and magnetohydrodynamics (MHD) codes with large 3D systems. This research is designed to provide a fundamental understanding of macroscopic relativistic MHD processes, microscopic plasma processes, and observed emissions. Our codes allow us to study the effect of inherent nonlinear processes on plasma dynamics and to compare simulation results with observations and analytical predictions.

Relevance of Work to NASA: Our studies are driven by anticipated new science from present and future NASA missions. The Swift mission has excelled at rapid burst afterglow follow-ups, generating a catalog of GRB redshifts and, in some cases, the properties of their associated supernovae. The Fermi Gamma-Ray Space Telescope investigates the spectra of GRBs over an unprecedented energy span. Our work will help to calculate the emission efficiency and self-consistent spectra expected from GRB relativistic shocks observed by Fermi. Future advanced X-ray telescopes and the Beyond Einstein Program's Black Hole Finder Probe will study other sources of relativistic jets; our work will be crucial to understanding these emissions. Funding from this research comes from NASA's Astrophysics Theory and Fundamental Physics Program.

Computational Approach: Thanks to the NASA Columbia supercomputer's unique structure and support for large 3D systems, we are systematically investigating the dynamics of relativistic jets with three codes:

3D RPIC code: We rewrote an earlier Fortran 77 code using the Message Passing Interface (MPI) and Open Multi-Processing (OpenMP). We have used it for several large simulations, which require terabytes of memory to achieve the necessary resolution and allow the full development of nonlinearities.

3D GRPIC code: The particle motion follows the contravariant form of the Newton-Lorentz equation. The acceleration is a function of the spacetime curvature defined by the metric and the Lorentz force due to the electromagnetic field. The local field is described by the Maxwell field tensor, whose components follow the contravariant general relativistic form of Maxwell's equations. The simulation moves the particles using an adaptive 5th-order Runge-Kutta scheme and calculates the fields and currents self-consistently. We have parallelized this code using MPI.

3D General Relativistic MHD code: "RAISHIN" is a conservative, high-resolution, shock-capturing scheme based on a 3+1 formalism of the General Relativistic MHD equations in a curved spacetime. RAISHIN computes numerical fluxes using the Harten-Lax-van Leer (HLL) approximate Riemann solver scheme. A flux-interpolated, constrained transport (flux-CT) scheme maintains a divergence-free magnetic field. We have vectorized this code and parallelized it using OpenMP.

Results: Several projects are using our simulation models, with the following outcomes:

- *Generation of shock with plowing ambient plasma:* New studies with a larger simulation domain show that continuously injected jets excite the Weibel instability (Figure 1), pile up ambient plasma, and generate a shock.
- *Parallelizing the RAISHIN code with OpenMP:* In test simulations, this code runs 24 times faster using 64 processors.
- *Relativistic MHD simulations of kink instability in force-free helical field:* Using 128 processors, we have run simulations with different magnetic pitch profiles. Preliminary results

show linear growth of kink instability from initial small perturbations and saturation in the non-linear stage (Figure 2). The growth and structure of kink instability is quite different with each pitch profile. Further simulations will study the effect of relativistic jets on kink instability.

- *Magnetic turbulence production by isotropic cosmic-ray ions streaming from supernova remnant shocks:* We have confirmed that the drift of cosmic-ray ions in the upstream plasma generates a turbulent magnetic field. However, field perturbations grow much more slowly than estimated using an MHD approach.
- *Jet formation from a black hole with kinetic processes:* We have developed a GRPIC code to perform simulations in curved space-time near black holes.

Role of High-End Computing: The significant processing power, memory, and data storage available on the NASA Advanced Supercomputing (NAS) facility's Columbia supercomputer meet the demands of our 3D codes. For example, we completed a recent RPIC simulation in less than 10 days using 320 processors. NASA's High-End Computing personnel are helping us to optimize the MPI code, maintaining temporary disk storage for diagnostics, and supporting 3D visualizations.

Future: We will continue to perform studies with larger simulation domains, which are essential to understanding the physics involved.

Co-Investigators

- Yosuke Mizuno, University of Alabama in Huntsville
- Jacek Niemiec, Polish Academy of Sciences
- Michael Watson, Fisk University

Publications

[1] Hardee, P., Mizuno, Y., and Nishikawa, K.-I., "GRMHD/RMHD Simulations and Stability of Magnetized Spine-Sheath Relativistic Jets," *Astrophysics & Space Science*, vol. 311, pp. 281–286, 2007.

[2] Ramirez-Ruiz, E., Nishikawa, K.-I., and Hededal, C.B., "e± Loading and the Origin of the Upstream Magnetic Field in GRB Shocks," 2007. *Astrophysical Journal*, vol. 671, pp. 1877–1885, 2007.

[3] Mizuno, Y., Hardee, P. and Nishikawa, K.-I., "3D Relativistic Magnetohydrodynamic Simulations of Magnetized Spine-Sheath Relativistic Jets," *Astrophysical Journal*, vol. 662, pp. 835–850, 2007.

[4] Wu, K., Fuerst, S.V., Mizuno, Y., Nishikawa, K.-I., Brandurdi-Raymont, G., and Lee, K.G., "General Relativistic Radiative Transfer: Applications to Black-Hole Systems," *Chinese Journal of Astronomy and Astrophysics, Supplement*, vol. 8, pp. 226–236, 2008.

[5] Mizuno, Y., Hardee, P. Hartmann, D., Nishikawa, K.-I., and Zhang, B., "Magnetohydrodynamic Boost for Relativistic Jets," *Astrophysical Journal*, vol. 672, pp. 72–82, 2008.

[6] Niemiec, J., Pohl, M., Stroman, T., and Nishikawa, K.-I., "Production of Magnetic Turbulence by Cosmic Rays Drifting Upstream of Supernova Remnant Shocks," *Astrophysical Journal*, vol. 684, pp. 1174–1189, 2008.

[7] Mizuno, Y., Zhang, B., Giacomazzo, B., Nishikawa, K.-I. , Hardee, P., Nagataki, S., and Hartmann, D., "Magnetohydrodynamic Effects in Propagating Relativistic Jets: Reverse Shock and Magnetic Acceleration," *Astrophysical Journal Letters*, vol. 690, pp. L47–L51, 2009.

[8] Nishikawa, K.-I., Niemiec, J., Medvedev, M., Sol, H., Hardee, P., Mizuno, Y., Zhang, B., Pohl, M., Oka, M., Hartmann, D.H., "Long Lasting Weibel Instability and Strong Magnetic Fields Associated with a Relativistic Shock System," *Astrophysical Journal Letters*, to be submitted, 2009.

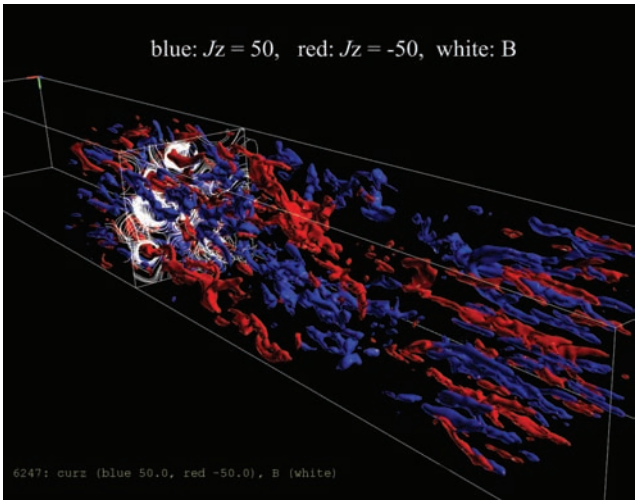


Figure 1: A snapshot viewed from the front of a relativistic jet at t = 59.8/ω_{pe} showing an isosurface of the z-component of the current density (±Jz) and the magnetic field lines (white) at the linear stage.

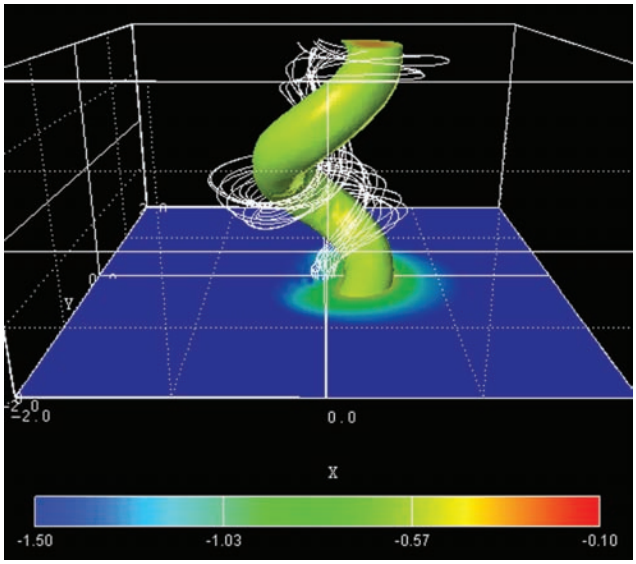
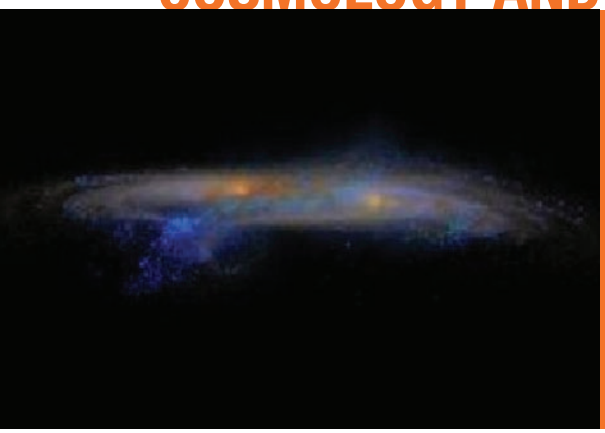


Figure 2: A snapshot of 3D isovolume density with magnetic field lines (white) at the non-linear stage of current-driven kink instability.

COSMOLOGY AND GALAXY FORMATION



SCIENCE MISSION DIRECTORATE

JOEL PRIMACK
University of California, Santa Cruz
(831) 459-2580
joel@scipp.ucsc.edu
http://scipp.ucsc.edu/personnel/profiles/primack.html

◀ **Figure 1:** A side-view snapshot from a simulated merger of two Sbc spiral galaxies, showing the first pass of the galaxies 0.59 Gigayears (590 million years) into the simulation.

Project Description: According to the now-standard Lambda Cold Dark Matter (ΛCDM) "double dark" theory, almost all of the universe is invisible dark matter and dark energy. This theory successfully predicted the distribution of temperature anisotropies in the cosmic background radiation measured by NASA's Wilkinson Microwave Anisotropy Probe (WMAP), and the distribution of nearby and high-redshift (distant) galaxies. Our project attempts to understand the structure and distribution of dark matter halos in ΛCDM and the formation and evolution of galaxies within the cosmic web. This requires modeling the complex hydrodynamics at cosmological and smaller scales (including the formation of stars and supermassive black holes) and their effects in heating the surrounding gas and providing the heavy elements from which planets form.

We are conducting high-resolution hydrodynamic simulations of mergers of gas-rich disk galaxies, which may form many of the fast-rotating elliptical galaxies. We have also created an analytical model that correctly predicts the properties of such elliptical galaxies, allowing us to interpolate between simulated cases and extrapolate beyond them. This work lets us calculate the entire evolving population of early-type galaxies. We are also doing hydrodynamic simulations of cold gas inflows and multiple mergers at high redshifts $z \sim 2$, which may form the ~25% of elliptical galaxies classified as slow rotators. Finally, we are running large dissipationless simulations, including a constrained realization of our local region of the universe and our "Bolshoi" simulation of a volume approximately 1 billion light years (1 Gigalightyear) on a side. The latter uses the new 5th-year cosmological parameters from WMAP and has mass and force resolution an order of magnitude better than the European Virgo Consortium's Millennium Run.

Relevance of Work to NASA: A key challenge in astronomy is to explain how the structures in today's universe formed within the ΛCDM framework and to test these new theories against observational evidence, e.g., from NASA's Hubble, Chandra, Spitzer, and Fermi space telescopes. The theories we are developing and simulating help to predict and interpret observations, and to design future missions such as the James Webb Space Telescope and Joint Dark Energy Mission. We are providing the main theoretical support for the Deep Extragalactic Evolutionary Probe (DEEP) and All-wavelength Extended Groth strip International Survey (AEGIS), which incorporate extensive data from NASA's space observatories. Primary funding for this research comes from NASA's Astrophysics Theory and Fundamental Physics Program, with additional funding from Spitzer and Hubble theory grants.

Computational Approach: Our large dark matter cosmological simulations use the dissipationless Adaptive Refinement Tree (ART) code. For hydrodynamic simulations, we use both ART-Hydro and the smooth-particle hydrodynamics code known as GADGET. Our binary galaxy merger simulations use GADGET. Our cosmological merger simulations start from a large ART-Hydro simulation; from this we map regions into GADGET, splitting particles and using other techniques to achieve an order of magnitude higher resolution. In these simulations, we use our Sunrise code to predict the effects of cosmological dust, which absorbs about 9/10 of the light of the bright new stars produced in galaxy mergers and re-radiates it at longer (infrared) wavelengths.

Results: We have characterized star formation in galaxies formed from binary mergers involving a variety of mass ratios [1]; characterized the morphologies of the resulting early-type galaxies including the effects of dust using our Sunrise code [2] (Figures 1, 2); and compared merger predictions with AEGIS data [3]. We developed an analytic model that correctly predicts the properties of elliptical galaxies formed from binary mergers of disk galaxies [4], and found that mergers of gas-rich disk galaxies with properties given by recent semi-analytic models lead to formation of elliptical galaxies with the observed size-mass relations from high redshift $z \sim 3$ to low redshift $z < 0.5$ [5]. We have simulated multiple mergers of galaxies at redshifts $z \sim 2$, including higher-resolution resimulations of regions identified in a large hydrodynamic

simulation. In these simulations, the resulting galaxies resemble the observed slowly rotating elliptical galaxies that are not produced by binary mergers [6].

Role of High-End Computing: The Columbia and Pleiades supercomputers at the NASA Advanced Supercomputing (NAS) facility have been extremely helpful in running dissipationless and hydrodynamic simulations of large-scale structure evolution and hydrodynamic simulations of galaxy mergers, including development and use of our Sunrise code. In particular, our Bolshoi Gigalightyear simulation harnesses the power of Pleiades. Collaboration with the NAS Visualization Group has been crucial in visualizing and interpreting the results of these simulations. These visualizations are helping astronomers and the wider public (e.g., planetarium visitors) to understand the evolving cosmos.

Future: The Bolshoi simulation will become the basis for a higher-resolution merger tree to enable models of unprecedented detail, which will capture earlier stages in galaxy formation and follow the dark matter substructure to small scales where galaxies merge. We intend to make public not only the analyzed results but also the entire merger tree, so that groups around the world can base models on it.

Co-Investigators

- Anatoly Klypin, New Mexico State University

Publications

[1] Cox, T., Jonsson, P., Somerville, R., Primack, J., and Dekel, A., "The Effect of Galaxy Mass Ratio on Merger-Driven Starbursts," *Monthly Notices of the Royal Astronomical Society*, Vol. 384, pp. 386–409, 2008.

[2] Lotz, J., Jonsson, P., Cox, T., and Primack, J., "Galaxy Merger Morphologies and Time-Scales from Simulations of Equal-Mass Gas-Rich Disc Mergers," *Monthly Notices of the Royal Astronomical Society*, Vol. 391, pp. 1137–1162, 2008.

[3] Lotz, J. et al. "The Evolution of Galaxy Mergers and Morphology at z < 1.2 in the Extended Groth Strip," *Astrophysical Journal*, Vol. 672, pp. 177–197, 2008.

[4] Covington, M., Dekel, A., Cox, T., Jonsson, P., and Primack, J., "Predicting the Properties of the Remnants of Dissipative Galaxy Mergers," *Monthly Notices of the Royal Astronomical Society*, Vol. 384, pp. 94–106, 2008.

[5] Covington, M., *The Production and Evolution of Scaling Laws Via Galaxy Merging* (Ph.D. dissertation, University of California, Santa Cruz. Supervisor: J.R. Primack), 2008.

[6] Novak, G., *Simulated Galaxy Remnants Produced by Binary and Multiple Mergers* (Ph.D. dissertation, University of California, Santa Cruz. Supervisor, J.R. Primack), 2008.

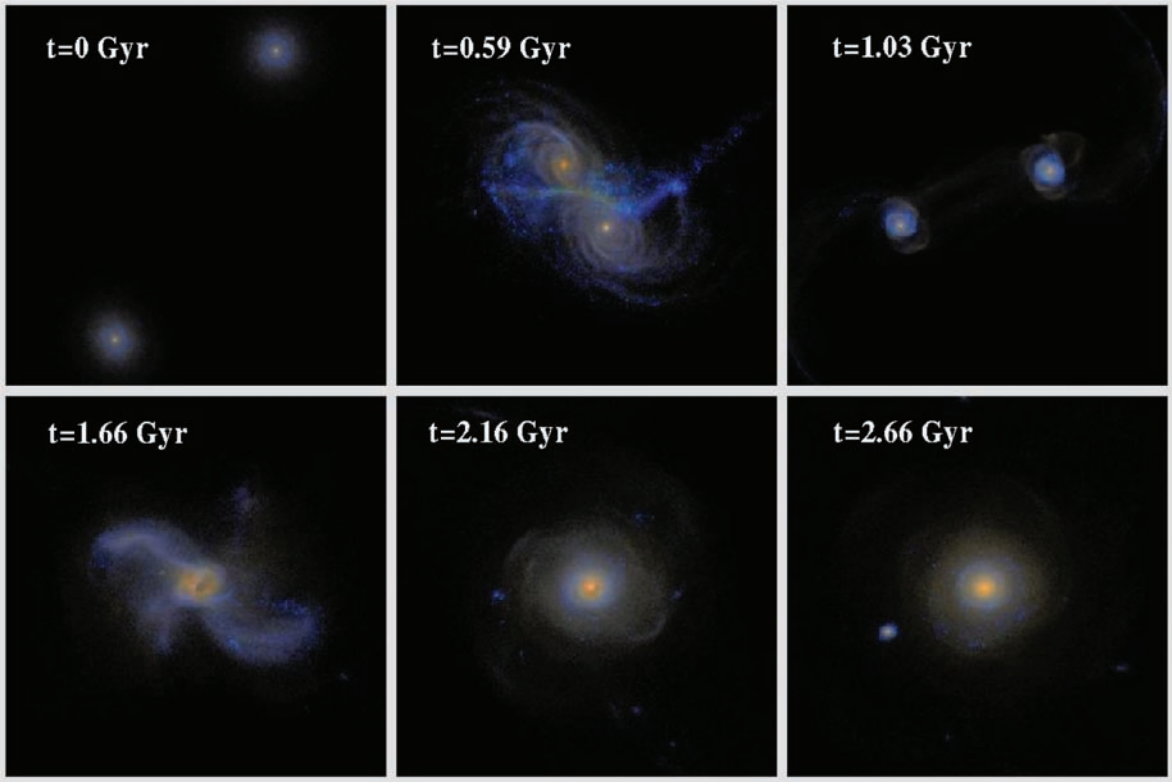
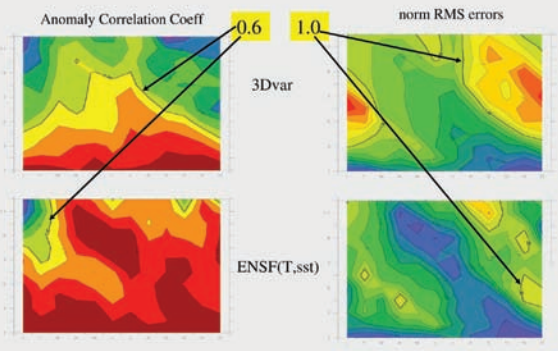


Figure 2: These front-view, u-r-z composite color images with dust extinction come from a simulated merger of two Sbc spiral galaxies over 2.66 Gigayears (Gyr). Top row: initial pre-merger galaxies, the first pass, and subsequent maximal separation. The first and third images are 200 kiloparsecs (kpc) across; the second is 100 kpc. Bottom row: 100-kpc views of the merger and post-merger 0.5 Gyr and 1 Gyr later. Star-forming regions in the initial discs, tidal tails, and outer regions of the remnant appear blue; dust-enshrouded star-forming nuclei appear red [2].

COUPLED OCEAN AND ATMOSPHERE DATA ASSIMILATION SYSTEMS FOR CLIMATE STUDIES



SCIENCE MISSION DIRECTORATE

VENKATRAMANI BALAJI
NOAA Geophysical Fluid Dynamics Laboratory
(609) 452-6516
v.balaji@noaa.gov

◀ **Figure 1:** The El Niño-Southern Oscillation forecast skills, including sea-surface temperature anomaly correlation coefficients (left) and normalized Root-Mean-Square errors over the East Pacific area (right). The coupled model ensemble is initialized from the coupled data assimilation products for the last quarter of the 20th century using both atmospheric and oceanic observations.

Project Description: Climate changes occur in a coupled Earth system that includes the atmosphere, ocean, land, and sea-ice components. Due to incomplete understanding of the dynamical and physical processes, modeling is always uncertain, and generated simulations drift away from real-world scenarios. Climate modeling includes predicting future changes as well as assessing historical variations. The necessary estimates of climate states and initial conditions come from data assimilation—blending observational data with coupled models. Assimilation requires massive computational resources.

With computational support from the NASA Advanced Supercomputing (NAS) Facility, the National Oceanic and Atmospheric Administration (NOAA) Geophysical Fluid Dynamics Laboratory (GFDL) has developed a coupled data assimilation (CDA) system consisting of an ensemble filter applied to a fully coupled global climate model (CGCM). Within the coupled framework, the assimilation provides a self-consistent, temporally continuous estimate of the model state and its uncertainty. This estimate takes the form of discrete ensemble members that can be used to directly initialize probabilistic climate forecasts with minimal initial coupling shocks. GFDL's CDA system serves as an estimator of historical climate variations and a predictor of future climate changes. Compared to traditional methods, the CDA system has several advantages:

- Directly solves a temporally evolving, joint-distribution function of climate states under observational data constraints.
- Uses a multi-variate analysis scheme maintaining physical balances among state variables and coupled components.
- Has minimal initial shocks for numerical climate forecasts.

Relevance of Work to NASA: Climate studies and predicting future climate changes are one of the common missions between NOAA and NASA. Efforts such as GFDL's CDA system support the research objectives of NASA's Earth Science Division, particularly the Climate Variability and Change Focus Area objective to "Understand the role of oceans, atmosphere, and ice in the climate system and improve predictive capability for its future evolution."

Computational Approach: To meet the need for accurately assessing historical climate variations and predicting future climate changes, we have developed an ensemble CDA system. Our implementation views the evolution of climate states as a continuous, stochastic, and dynamic process. The filtering assimilation combines an observational probability density function (PDF) with a prior PDF derived from the CGCM to produce an analyzed PDF. Using a super-parallelization configuration, the coupled assimilation is a continuous data-incorporation process that includes atmospheric and oceanic data assimilation components (Figure 3).

Results: Climate Detection Experiments using GFDL's CDA system on NASA's High-End Computing (HEC) systems indicate that the assimilation with the greenhouse gas and natural aerosol radiative forcing at fixed pre-industrial (1860) levels produces a consistent multidecadal warming trend in almost all oceans with its own interannual variability (Figure 2). For oceans that have reasonable observation coverage (e.g., the Pacific and North Atlantic Oceans), the ocean data assimilation process retrieves the trend and the variability quite well, with faster spin-up times and reduced uncertainty.

We have initialized climate estimates and forecasts from observed atmospheric and oceanic data (Figure 1). Hindcast statistics show that this ensemble climate state estimate and prediction system improved ENSO forecast skills dramatically. This improvement happens mainly because the self-consistent ensemble initial conditions from this coupled assimilation system keep all components of the coupled model in a physically balanced state, which helps model dynamics project initial signals onto a seasonal-interannual time-scale.

Role of High-End Computing: Our flagship climate models, CM2.0 and CM2.1, are based on the GFDL Flexible Modeling System, a software environment for developing new physics and new algorithms concurrently, and for expressing them on a variety of HEC architectures, spanning distributed and shared memory, as well as vector architectures. Results from these models served as inputs to the International Panel on Climate Change Fourth Assessment Report. NASA HEC resources at the NAS facility have enabled us to run a significant number of trials and tests of these models. Our efforts have included testing and evaluating the CDA system's performance and scaling. Access to NAS' Columbia supercomputer has been essential for meeting the demands of huge computational resources for climate estimation and prediction using our CDA system.

Future: In the future, we will reorient the CDA system to focus on multi-decadal-scale climate predictions that require assimilating a greater number of observations coherently into a more advanced coupled model—including higher resolutions and more comprehensive physical processes. This undertaking will require a much more powerful computational resource.

Co-Investigators

- Anthony Rosati, Shaoqing Zhang, NOAA Geophysical Dynamics Fluid Laboratory

Publications

[1] Zhang, S., Rosati, A., Harrison, M.J., Gudge, R., and Stern, W., "GFDL's Coupled Ensemble Data Assimilation System, 1980–2006 Oceanic Reanalysis and Its Impact on ENSO Forecasts," Third World Climate Research Programme International Conference on Reanalysis, Tokyo, Japan, Jan. 27–Feb. 2, 2008, 2008.

[2] Zhang, S., Harrison, M.J., Rosati, A., and Wittenberg, A.,"System Design and Evaluation of Coupled Ensemble Data Assimilation for Global Oceanic Climate Studies," *Monthly Weather Review,* Vol. 135, pp. 3541–3564, 2007.

[3] Zhang, S., Harrison, M.J., Wittenberg, A.T., Rosati, A., Anderson, J.L., and Balaji, V., "Initialization of an ENSO Forecast System Using a Parallelized Ensemble Filter," *Monthly Weather Review*, Vol. 133, pp. 3176–3201, 2005.

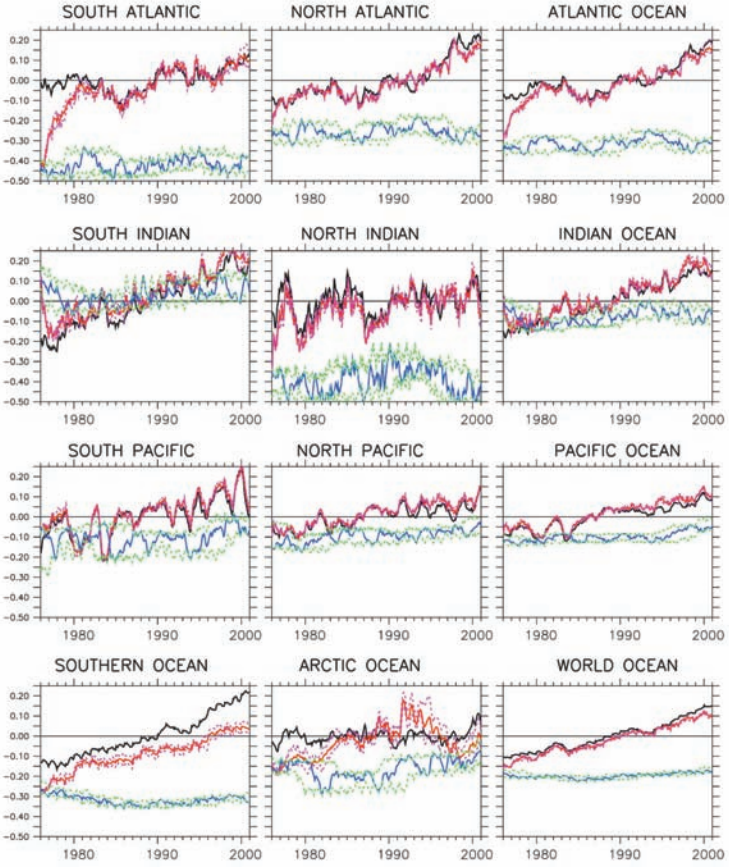


Figure 2: Time series showing the anomalies of the top-500-meter ocean heat content (averaged temperature) in different oceans for the observed "truth" (in black; based on time-varying radiative forcings of greenhouse gases and natural aerosols); the coupled data assimilation (CDA) (red); and the control (blue). The green/pink dashed lines plot the upper/lower bounds of the control vs. CDA spread, which are estimated by the model/assimilation ensemble. All anomalies are computed using observed climatology.

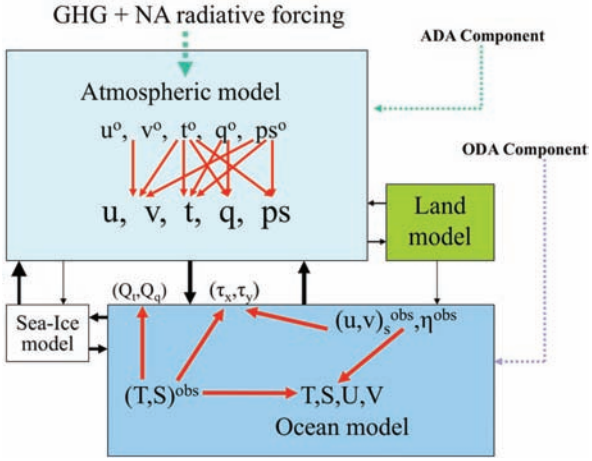
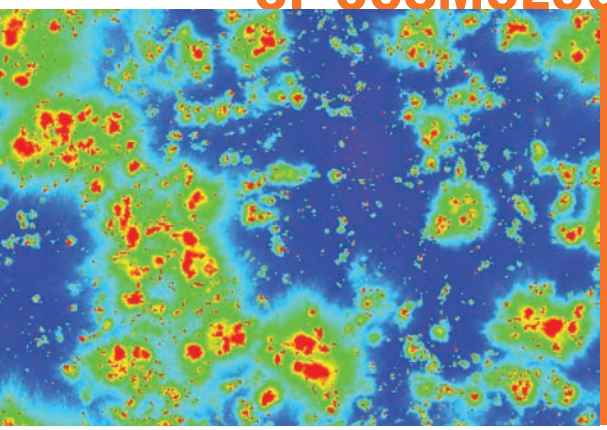


Figure 3: A Geophysical Fluid Dynamics Laboratory coupled climate model exchanges fluxes between model components (atmosphere, land, ocean, and sea-ice models); constraints of atmospheric temperature and wind; and oceanic temperature (T), salinity (S), and currents (U,V) from atmospheric and oceanic data assimilations (ADA/ODA). To isolate 20th century anthropogenic effects, the atmospheric component models the effects of radiative forcing due to both contemporaneous (time-varying) and pre-industrial (fixed 1860) levels of greenhouse gases (GHG) and natural aerosols (NA).

DETAILED SIGNATURES OF COSMOLOGICAL REIONIZATION



◀ Close-up of Figure 1.

SCIENCE MISSION DIRECTORATE

RENYUE CEN
Princeton University
(609) 258-3806
cen@astro.princeton.edu
http://www.astro.princeton.edu/~cen/

Project Description: The primary goals of this project are to understand the complex cosmological reionization process using detailed radiative transfer hydrodynamics simulations and to provide concrete, detailed observables to confront with current and upcoming observations of the high-redshift universe ($z > 6$). Sometimes called "cosmic dawn," the reionization period began roughly 300 million years after the Big Bang. As the earliest stars appeared, they generated enough ultraviolet light to turn hydrogen atoms back into protons and electrons. These regions of reionization continually expand until they overlap, marking the end of cosmological reionization.

Our cosmological reionization simulations are among the most advanced of their kind, featuring a simulation box on the order of 100 megaparsecs (Mpc) in size, a very high mass resolution (around 1 million solar masses), accurate 3D radiative transfer (using a ray-tracing method), and 3D hydrodynamics. They resolve ionizing, photon-producing galaxies using at least 26 billion particles. The computations couple a total variation diminishing (TVD) hydrocode, to follow hydrodynamics of the cosmic gas, and a 3D ray-tracing code, to follow the propagation of cosmological reionization fronts.

Such simulations serve two primary purposes: First, they provide the most accurate characterization of the reionization process, which allows direct comparison to a wide range of observations spanning the entire electromagnetic spectrum. These include 21-centimeter observations of the neutral hydrogen evolution, as well as forthcoming observations of the first galaxies by the James Webb Space Telescope (JWST), and of intervening free-electron polarization of the cosmic microwave background (CMB) by the European Planck mission. Second, we can use our simulations to calibrate faster, semi-numerical methods, which are necessary for exploring the vast parameter space that reflects our limited knowledge of the high-redshift universe.

Relevance of Work to NASA: These simulations will provide a quantitative framework for interpreting observations by NASA and others, including the Wilkinson Microwave Anisotropy Probe (WMAP), JWST, and Planck. Our simulations will help explore the last, high-redshift ($z > 6$) frontier of the universe, where several cosmic landmarks occur, including: the formation of the first stars, the appearance of the first galaxies; and the first enrichment of the pristine cosmic gas by the metals synthesized in the stars, which will shape the subsequent evolution of galaxies and the intergalactic medium. Our research program supports NASA's mission to "pioneer the future in space exploration, scientific discovery, and aeronautics research," and directly advances the research objectives of the Astrophysics theme "Origin and Evolution of Cosmic Structure" in the *Science Plan for NASA's Science Mission Directorate 2007–2016*. NASA funding for our project comes from the Astrophysics Theory and Fundamental Physics Program.

Computational Approach: RADHYDRO is a hybrid code that combines a very-high-resolution N-body code, a shock-capturing TVD hydrodynamics code, and a ray-tracing radiative transfer code. It allows us to simultaneously compute the formation of low-mass, high-redshift galaxies and the evolution of the intergalactic medium and the fluctuating ionizing radiation background. RADHYDRO is fully parallelized using OpenMP.

Results: We carried out the world's largest 3D radiative transfer hydrodynamics simulations of cosmological reionization. One simulation, run on the NASA Advanced Supercomputing (NAS) facility's Columbia supercomputer, tracked nearly 29 billion dark matter particles on a computational mesh with more than 1.5 trillion cells. We also showed that an inhomogeneous reionization process imprints important signatures on the intergalactic medium, even in the lower-redshift regions already accessible to observations [1] (Figures 1 and 2).

From our results, the NAS visualization team has produced visualizations that have received broad exposure in venues such as the Museo di Storia Naturale (Natural History Museum), Trieste, Italy, and SC08: The International Conference for High-Performance Computing, Networking, Storage, and Analysis, Austin, Texas.

Role of High-End Computing: The RADHYDRO code requires a large symmetric multiprocessing (SMP) machine; the Columbia supercomputer at NAS is the best available platform for this application. The NAS visualization team led by Chris Henze provided an extremely valuable service in helping us visualize the complex data produced by our simulations to yield a form with high public outreach value, in addition to the intrinsic scientific value.

Future: We will expand our usage of the HEC resources continuously and expect to make still larger and better simulations in the next several years, in anticipation of and preparation for the launch of major missions, including Planck and JWST.

Publications

[1] Trac, H., Cen, R., Loeb, A., "Imprint of Inhomogeneous Hydrogen Reionization on the Temperature Distribution of the Intergalactic Medium," *The Astrophysical Journal Letters,* Volume 689, Issue 2, pp. L81–L84, 2008.

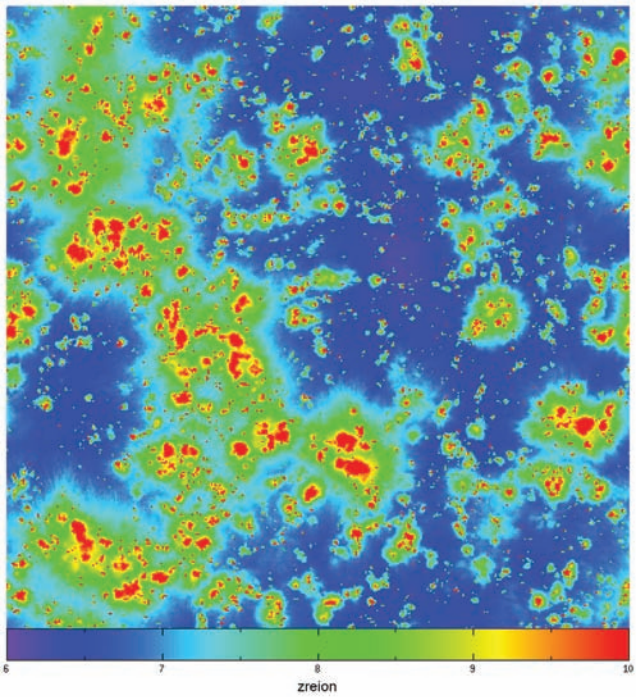


Figure 1: This visualization shows a 100-megaparsec-squared (100 Mpc)^2 slice with a thickness of two hydrodynamics cells (130 kiloparsecs) from the late reionization model of cosmological formation. Coloring traces the redshifts of reionization in the individual cells.

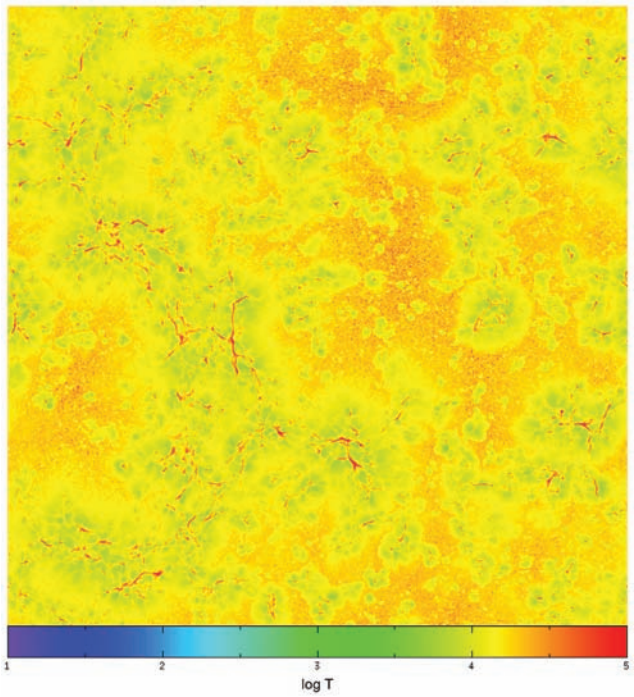
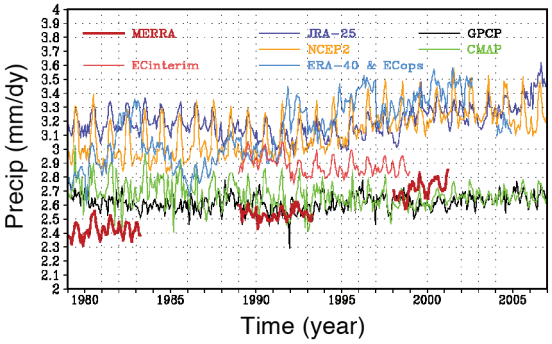


Figure 2: This image is a visualization of the same simulation domain; colors indicate the temperature at the end of reionization.

GEOS-5/MODERN ERA RETROSPECTIVE-ANALYSIS FOR RESEARCH AND APPLICATIONS (MERRA)



SCIENCE MISSION DIRECTORATE

MICHELE RIENECKER
NASA Goddard Space Flight Center
(301) 614-6142
Michele.Rienecker@nasa.gov
http://gmao.gsfc.nasa.gov/

◀ **Figure 1:** Global monthly mean precipitation (millimeters/day) from the three MERRA streams (as of November 13, 2008) compared with observations from the Global Precipitation Climatology Project (GPCP) and the NOAA Climate Prediction Center (CPC) Merged Analysis of Precipitation (CMAP). Also shown are other centers' previous (NCEP and ERA-40) and ongoing (ECInterim and JRA-25) reanalyses. MERRA, focused on the hydrological cycle, is closer to the observations. All reanalyses are sensitive to changes in the observing system.

Project Description: Retrospective analyses (or reanalyses) synthesize temporally and spatially irregular observations from the historical observational databases to provide a gridded record of essential climate variables. The model-data synthesis (analysis) uses a "frozen" data assimilation system (DAS) to provide a consistent view, in space and time, of observations from different sources and of different types. The synthesis also provides a view of unobserved variables consistent with those observed. The Modern Era Retrospective-analysis for Research and Applications (MERRA) is producing an atmospheric retrospective-analysis of the satellite era (1979 to present) in order to improve upon previous reanalyses of the hydrological cycle and minimize the influence of observing system changes on the representation of climate variability and trends.

The DAS used for MERRA consists of the Goddard Earth Observing System, Version 5 (GEOS-5) atmospheric model coupled to the Grid-point Statistical Interpolation (GSI) analysis scheme being developed by the National Centers for Environmental Prediction's Environmental Modeling Center (NCEP/EMC) and NASA's Global Modeling and Assimilation Office. Unlike earlier reanalyses, MERRA uses the GSI's online satellite radiance bias correction and includes updated cross-satellite calibrations for the Special Sensor Microwave/Imager (SSM/I) and Microwave Sounding Unit (MSU) to reduce artificial variability associated with the change of satellite platforms. In addition to satellite radiances, MERRA assimilates several remotely sensed retrieved datasets (e.g., SSM/I surface winds and cloud track winds) and conventional observations from radiosonde, dropsonde, aircraft, and surface pressure instruments. Figure 2 depicts the observing system at various points over recent decades.

Relevance of Work to NASA: MERRA supports NASA's climate science by placing current research satellite observations in a climate context. It provides a gridded historical record of meteorology for performing climate diagnostics, initializing and validating climate predictions, and undertaking studies with atmospheric constituent transport and ocean and land surface models. By focusing on an improved representation of the water cycle, MERRA also supports NASA's programs to characterize and predict Earth's energy and water cycles. GEOS-5 and MERRA are supported by the NASA Modeling, Analysis, and Prediction (MAP) Program. Additional MERRA support comes from NASA's Research, Education, Applications Solutions Network (REASoN).

Computational Approach: GEOS-5 uses finite-volume dynamics on a spherical grid. The MERRA configuration has a 2/3-degree longitude by 1/2-degree latitude grid with 72 vertical levels to 0.01 hectopascals (hPa), with assimilation analyses every 6 hours. The GSI uses a 3D variational approach to solve the least-squares fit of GEOS-5 analysis states to the model (background) states and to the available satellite and *in situ* data. Recursive filters are the basic building blocks used to create background error covariance structures.

Results: For MERRA, the GEOS-5 DAS has undergone tuning, with a focus on the hydrological cycle rather than weather prediction skill as for other systems used for reanalysis (Figure 1). We validated the MERRA system by comparing analyses for each season in 2004 and for January 2006 against independent data and other reanalyses. The validation activity was comprehensive and included investigation of climate phenomena such as monsoons, low-level jets across the U.S., the diurnal cycle of precipitation, radiative forcing diagnostics, precipitation distributions, and many other aspects [3]. We also conducted limited sensitivity experiments to gauge the impact of satellite bias estimates and new observing systems such as SSM/I. These experiments showed that we should expect a 10% increase in tropical precipitation in MERRA solely from the introduction of SSM/I [4]. In May 2008, after validation and external review, MERRA began delivering product collections to the Goddard Earth Science Data

Information and Services Center (GES-DISC - http://disc.sci.gsfc.nasa.gov/MDISC/), where the data are available on-line.

Role of High-End Computing: To process the 30-year observational record (in three 10-year streams), MERRA requires continuous access to 432 processors of the Discover system at the NASA Center for Computational Sciences (NCCS) over 18 months. MERRA will generate about 70 terabytes of data products and a similar volume of intermediate archive data.

Future: MERRA will complete the 30-year reanalysis to 2008 in about August 2009. We are conducting observing system sensitivity experiments by withholding the Earth Observing System (EOS) data streams to evaluate the impact of these research data on the inferred climate. Other sensitivity experiments will help to estimate uncertainty due to model and assimilation configurations (e.g., resolution, covariance models, improved estimates of emissivity for radiative transfer calculations) and datasets (e.g., using only data types available prior to the introduction of SSM/I in August 1987).

Co-Investigators

- Max Suarez, Ron Gelaro, Julio Bacmeister, Emily Hui-Chun Liu, Ricardo Todling, Michael Bosilovich, Siegfried Schubert, Gi-Kong Kim, Junye Chen, all of NASA Goddard Space Flight Center

Publications

[1] Rienecker, M.M., et al., "The GEOS-5 Data Assimilation System – Documentation of Versions 5.0.1, 5.1.0 and 5.2.0," *NASA GSFC Technical Report Series on Global Modeling and Data Assimilation*, NASA/TM-2008-104606, Vol. 27, 101 pp. (http://gmao.gsfc.nasa.gov/pubs/docs/GEOS5_104606-Vol27.pdf), 2008.

[2] Bosilovich, M.G., Schubert, S.D., Rienecker, M., Todling, R., Suarez, M., Bacmeister, J., Gelaro, R., Kim, G.-K., Stajner, I., and Chen, J.,"NASA's Modern Era Retrospective-analysis for Research and Applications," *U.S. CLIVAR Variations*, Vol. 4, No. 2, pp. 5–8 (http://www.usclivar.org/Newsletter/VariationsV4N2.pdf), 2006.

[3] Schubert, S.D., et al., "Assimilating Earth System Observations at NASA: MERRA and Beyond," 3rd WCRP International Conference on Reanalyses, January 2008, Tokyo, Japan (http://wcrp.ipsl.jussieu.fr/Workshops/Reanalysis2008/abstract.html), 2008.

[4] Bosilovich, M.G., et al.,"Evaluation of Variations in the Late-80s Observing System and the Impacts on the GEOS-5 Data Assimilation System," 3rd WCRP International Conference on Reanalyses, January 2008, Tokyo, Japan (http://wcrp.ipsl.jussieu.fr/Workshops/Reanalysis2008/abstract.html), 2008.

[5] Bosilovich, M.G., "NASA's Modern Era Retrospective-analysis for Research and Applications: Integrating Earth Observations," *Earthzine*, September 26, 2008 (http://www.earthzine.org/2008/09/26/nasas-modern-era-retrospective-analysis/), 2008.

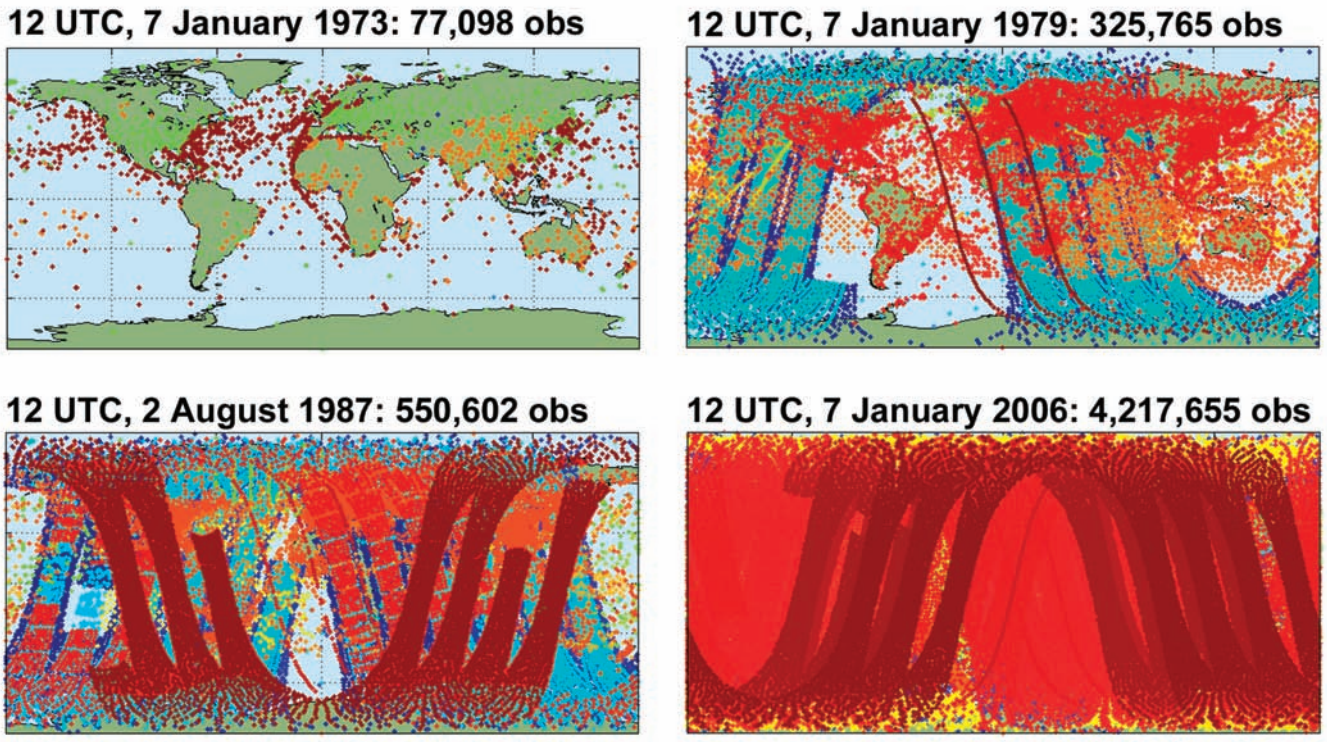
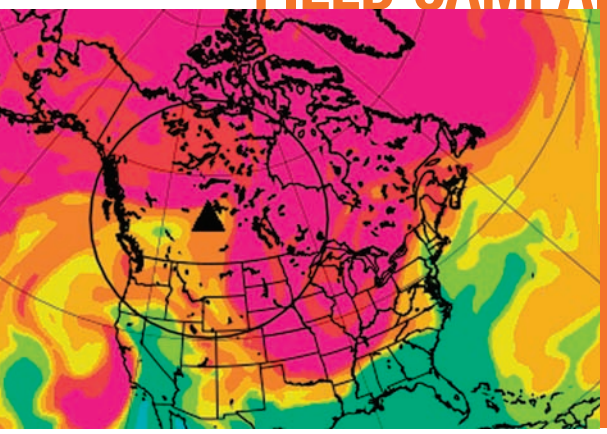


Figure 2: The number of atmospheric observations to be assimilated has increased dramatically over the last few decades. The panels show the evolution of observing systems from 1973 (pre-satellite) to 1979 (TIROS Operational Vertical Sounder [TOVS]) to 1987 (add Special Sensor Microwave Imager [SSMI]) to 2006 (add Atmospheric Infrared Sounder [AIRS]). Each color represents a different observing system, and the titles list the number of observation points for a single 6-hour period.

GEOS-5 SUPPORT OF NASA FIELD CAMPAIGNS: TC4 • ARCTAS • TIGERZ



SCIENCE MISSION DIRECTORATE

MICHELE RIENECKER
NASA Goddard Space Flight Center
(301) 614-6142
Michele.Rienecker@nasa.gov
http://gmao.gsfc.nasa.gov/

◀ Detail of Figure 2.

Project Description: Weather forecasts are an important part of planning aircraft flights for NASA's field campaigns, as is ancillary information such as aerosol distributions. A team at NASA Goddard Space Flight Center supports NASA field campaigns with real-time products and forecasts from the Goddard Earth Observing System Model, Version 5 (GEOS-5) to aid in flight planning and post-mission analysis. The objective is to provide a wide variety of mission-specific information on weather and atmospheric composition, summarized for ease of access through the Data Portal hosted by the NASA Center for Computational Sciences (NCCS). This support requires the close collaboration of mission planners and several organizations at NASA Goddard: the Global Modeling and Assimilation Office, the Atmospheric Chemistry and Dynamics Branch, the Software Integration and Visualization Office, and NCCS.

Three recent examples of GEOS-5 support for field campaigns were the TC4 (Tropical Composition, Cloud and Climate Coupling), ARCTAS (Arctic Research of the Composition of the Troposphere from Aircraft and Satellites), and TIGERZ missions. The TC4 mission, from July 12 to August 12, 2007, investigated atmospheric structure, properties, and processes in the tropical Eastern Pacific. The ARCTAS field campaign was undertaken during Spring and Summer 2008 to investigate the atmospheric transport pathways from mid-latitudes to the Arctic, and the relative contributions of different source regions to Arctic air pollution. Thus, for ARCTAS, the GEOS-5 system was instrumented with a set of tagged tracers to track transport: hydrophobic/hydrophilic organic carbon tag tracers driven by boreal and non-boreal biomass burning; carbon monoxide (CO) tag tracers driven by boreal and non-boreal biomass burning and by non-biomass emissions over Northern and Southern Asia, Europe, and North America; and chlorofluorocarbon (CFC) tag tracers with tropospheric and stratospheric origins. We determined biomass-burning sources of carbonaceous aerosols, CO, and sulfur dioxide (SO_2) in near real-time from MODIS imagery and land mapping. These were persisted forward in time for forecasts. We identified dust and sea-salt aerosol sources from the local meteorology, and specified other emissions from climatological databases. The ARCTAS system was continued during May and June 2008 to support the TIGERZ campaign, for which the AERONET (Aerosol Robotic Network) project deployed ground-based instruments to look at clouds and aerosols in the Indo-Gangetic basin.

Products for the field campaigns were based on meteorological analyses from the GEOS-5 data assimilation and forecast system. This system was complemented by the GEOS-5 Aerosol/Chemistry (AeroChem) components, including global CO and carbon dioxide (CO_2) tracers and aerosols (dust, sea-salt, organic carbon, black carbon, and sulfates) from the Goddard Chemistry Aerosol Radiation and Transport (GOCART) model. We provided the analyses and 5-day forecasts to the flight-planning teams in near real-time. Examples of meteorological and aerosol/chemical products available through a multi-faceted data delivery system can be found at http://gmao.gsfc.nasa.gov/projects/arctas/.

Relevance of Work to NASA: NASA undertakes aircraft field campaigns as part of its science mission strategy. These campaigns integrate satellite and aircraft observations for scientific analysis and satellite algorithm validation. GEOS-5 analyses and forecasts are one of the information sources used by satellite and model science teams in pre-mission flight planning and post-mission data interpretation. This work is funded by NASA's Modeling, Analysis, and Prediction (MAP) Program.

Computational Approach: GEOS-5 uses finite-volume dynamics on a spherical grid. The configuration used for these field campaigns has a 2/3-degree longitude by 1/2-degree latitude grid with 72 vertical levels to 0.01 hectopascals (hPa), with assimilation analyses conducted every 6 hours. The analysis uses a three-dimensional variational approach to solve the least-squares fit of GEOS-5 analysis states to the model

(background) states and to the available satellite and *in situ* data. Advection, diffusion, and convective transport of the CO, CO_2, and GOCART tracers are performed on-line within GEOS-5. For TC4, we conducted an additional ¼-degree 2-day meteorological forecast to aid flight planning.

Results: Figures 1 and 2 show examples of GEOS-5 products for ARCTAS. The mission support was successful, with GEOS-5 products delivered on time for most of the mission duration thanks to the NCCS ensuring timely execution of job streams and supporting the data portal. For example, a DC-8 flight on June 29, 2008 sampled the Siberian fire plume transported to the region in the mid-troposphere as predicted by GEOS-5.

Role of High-End Computing: The GEOS-5 systems were run on 128 processors of the NCCS Explore supercomputer, with a continuous job stream allowing timely delivery of products to inform flight planning.

Future: GEOS-5 products with improved systems (e.g., the ARCTAS system was updated from the TC4 system) will be used for post-mission analysis. The system will also be used to support HIPPO (High-performance Instrumented Airborne Platform for Environmental Research [HIAPER] Pole-to-Pole Observations), a mission that will measure cross-sections of atmospheric concentrations approximately pole-to-pole from the surface to the tropopause. The program will provide the first comprehensive, global survey of atmospheric trace gases, covering the full troposphere in all seasons and multiple years. We will also support other field campaigns as they arise.

Co-Investigators

- Peter Colarco, Arlindo da Silva, Max Suarez, Ricardo Todling, Larry Takacs, Gi-Kong Kim, Eric Nielsen, all of NASA Goddard Space Flight Center

Publications

[1] Rienecker, M.M., et al., "The GEOS-5 Data Assimilation System – Documentation of Versions 5.0.1, 5.1.0 and 5.2.0," *NASA GSFC Technical Report Series on Global Modeling and Data Assimilation*, NASA/TM-2008-104606, Vol. 27, 101 pp, 2008.

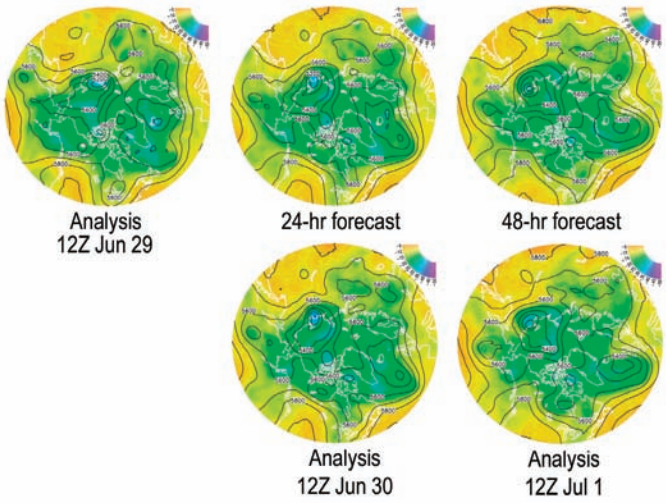


Figure 1: This image shows 500-hectopascal (hPa) temperatures (shading) and heights (contours) during NASA's ARCTAS (Arctic Research of the Composition of the Troposphere from Aircraft and Satellites) mission. An analysis from the Goddard Earth Observing System Model, Version 5 (GEOS-5) is shown with 24- and 48-hour forecasts and validating analyses. These fields, with the accompanying atmospheric chemistry fields, were used to help plan a DC-8 flight on June 29, 2008.

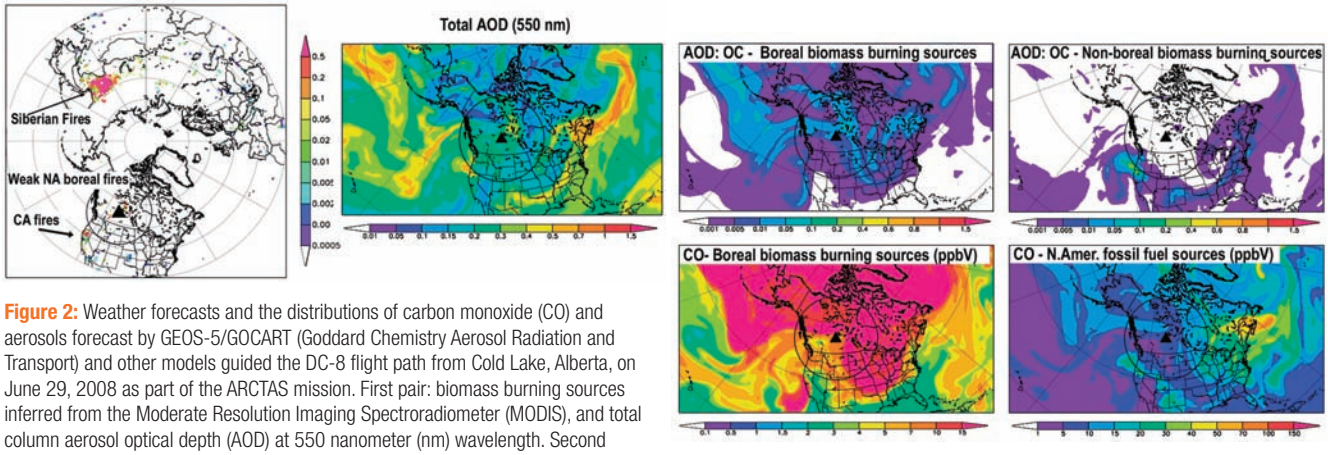
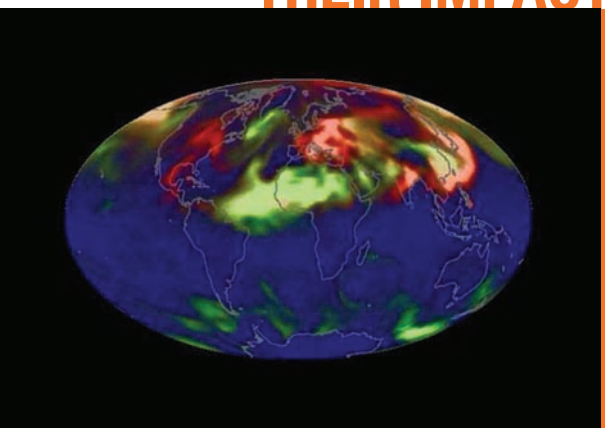


Figure 2: Weather forecasts and the distributions of carbon monoxide (CO) and aerosols forecast by GEOS-5/GOCART (Goddard Chemistry Aerosol Radiation and Transport) and other models guided the DC-8 flight path from Cold Lake, Alberta, on June 29, 2008 as part of the ARCTAS mission. First pair: biomass burning sources inferred from the Moderate Resolution Imaging Spectroradiometer (MODIS), and total column aerosol optical depth (AOD) at 550 nanometer (nm) wavelength. Second pair: AOD contributions from boreal and non-boreal biomass emissions. Third pair: CO distributions (parts per billion by volume, ppbV) at 550 hPa from boreal biomass burning and North American fossil fuel emissions.

GLOBAL MODELING OF AEROSOLS AND THEIR IMPACTS ON CLIMATE AND AIR QUALITY



SCIENCE MISSION DIRECTORATE

MIAN CHIN
NASA Goddard Space Flight Center
(301) 614-6007
mian.chin@nasa.gov
http://croc.gsfc.nasa.gov/gocart

◀ Snapshot from Figure 1.

Project Description: We have several projects supported by NASA's Science Mission Directorate, with an overall goal to develop a global aerosol/chemistry/transport model to study aerosols and their impacts on climate and air quality, through the use and analysis of satellite and other observational data. We seek to understand the role of aerosols in radiative forcing and climate change from pre-industrial time to the present, and to investigate regional and global change of aerosols and related gases over multi-decadal time-scales. We are also assessing the effect of long-range aerosol transport and anthropogenic emissions on surface air quality, as well as supporting NASA field experiments and satellite retrievals.

Collectively, our research projects focus on interactive use of the Goddard Chemistry Aerosol Radiation and Transport (GOCART) model alongside satellite and *in situ* observations. We use satellite-based fire data to improve simulations of biomass burning emissions and land-cover/vegetation data to account for dust source variations. We use the model to conduct multi-year simulations of aerosols and trace gases, which we compare with data from NASA's satellites and sensors, including the Moderate Resolution Imaging Spectroradiometer (MODIS), Multiangle Imaging SpectroRadiometer (MISR), Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation (CALIPSO), Ozone Monitoring Instrument (OMI), Measurements of Pollution in the Troposphere (MOPITT), and Atmospheric Infrared Sounder (AIRS).

We investigate the relationships between aerosols and carbon monoxide (CO) in terms of sources, chemistry, and long-range transport. We calculate the aerosol radiative forcing and estimate its climate effects in different regions.

In addition, we evaluate the application of satellite data to air quality studies by examining the quantitative link between the remotely sensed data and surface pollutant concentrations and testing the applicability of combining aerosols with CO. We also provide model information on aerosol composition and vertical distributions to satellite retrieval teams.

Relevance of Work to NASA: Our projects are supported by several NASA programs, including the Modeling, Analysis, and Prediction (MAP) Program; the Atmospheric Composition Modeling and Analysis Program (ACMAP); the Earth Observing System (EOS); CALIPSO; and the Tropospheric Chemistry Program (TCP). Our activities are highly relevant to NASA's science objectives, particularly to understanding the mechanisms that drive Earth system changes and the Earth system's response to natural and human-induced changes.

Computational Approach: GOCART is a global model of atmospheric processes, including emission, chemistry, dry deposition, wet removal, advection, convection, and radiative forcing. Our projects use supercomputing for many model runs of different scenarios.

Results:

- *Intercontinental transport of aerosols and implications for regional air quality:* This study assesses the impact of long-range transport of aerosols (Figure 1).
- *Possibilities and challenges in using satellite aerosol optical depth (AOD) data for air quality studies:* This study investigates the relationship of column AOD to surface concentrations of fine-grained ($PM_{2.5}$) airborne particulate matter.
- *Aerosol absorption:* Aerosol absorption is a key parameter for estimating the climate forcing of aerosols. This project examines the global distribution of absorbing aerosols and compares the quantities from NASA's ground-based Aerosol Robotic Network (AERONET) (Figure 2).
- *Long-term trend of atmospheric aerosols:* This project simulates global aerosol changes from 1980 to the present and examines the relationship between emission, atmospheric loading, and AOD.

Role of High-End Computing: Our research critically depends on the NASA Center for Computational Sciences (NCCS), whose supercomputers have performed all the simulations, and whose personnel have helped our group to optimize model executions and to transition from one platform to another.

Future: We aim to conduct global model simulations at much higher spatial and temporal resolutions and spanning multiple decades; our computing needs will grow significantly.

Co-Investigators

- Thomas Diehl, Huisheng Bian, Hongbin Yu, Qian Tan, Tom Kucsera, all of NASA Goddard Space Flight Center/University of Maryland, Baltimore County

Publications

[1] Chin, M., Diehl, T., Ginoux, P., and Malm, W., "Intercontinental Transport of Pollution and Dust Aerosols: Implications for Regional Air Quality," *Atmospheric Chemistry and Physics*, Vol. 7, pp. 5501–5517, 2007.

[2] Bian, H., Chin, M., Kawa, S.R., Duncan, B., Arellano, A., and Kasibhatla, P., "Sensitivity of Global CO Simulations to Uncertainties in Biomass Burning Sources," *Journal of Geophysical Research*, Vol. 112, D23308, doi:10.1029/2006JD008376, 2007.

[3] Yu, H., Remer, L.A., Chin, M., Bian, H., Kleidman, R.G., and Diehl, T., "A Satellite-Based Assessment of Trans-Pacific Transport of Pollution Aerosol," *Journal of Geophysical Research*, Vol. 113, D14S12, doi:10.1029/2007JD009349, 2008.

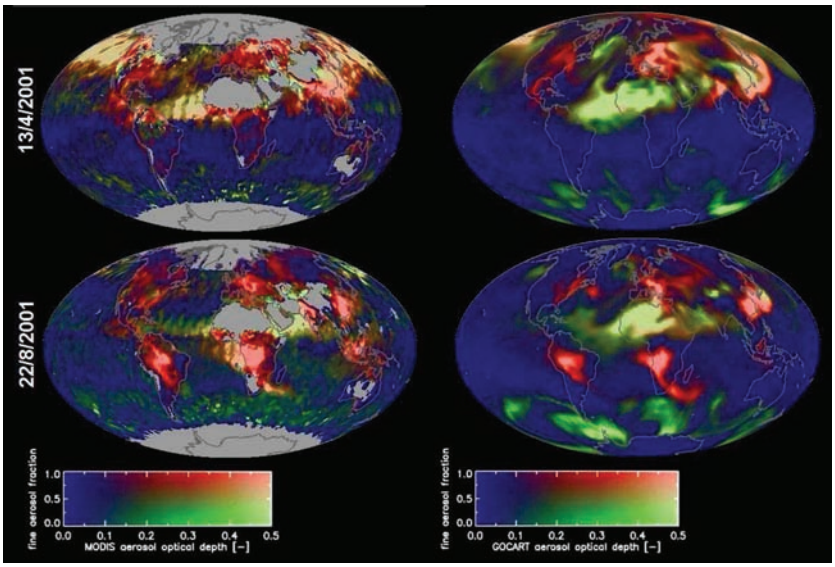


Figure 1: Aerosol optical depth observed by the Moderate Resolution Imaging Spectroradiometer (MODIS) (left) and simulated by the Goddard Chemistry Aerosol Radiation and Transport (GOCART) model (right) for 13 April 2001 (top) and 22 August 2001 (bottom). Red indicates fine mode aerosols (e.g., pollution and smoke); green indicates coarse mode aerosols (e.g., dust and sea-salt); color brightness is proportional to the aerosol optical depth. 13 April sees transport of heavy dust and pollution from Asia to the Pacific, and dust transport from Africa to the Atlantic. 22 August sees large smoke plumes over South America and Southern Africa.

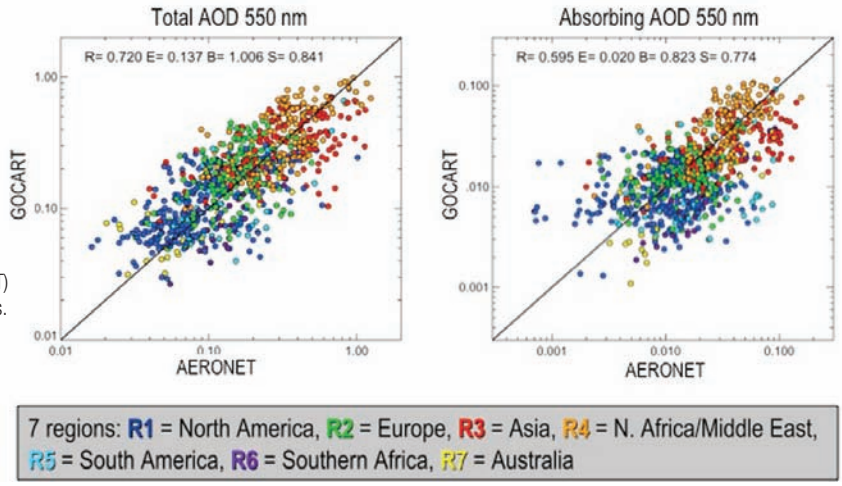
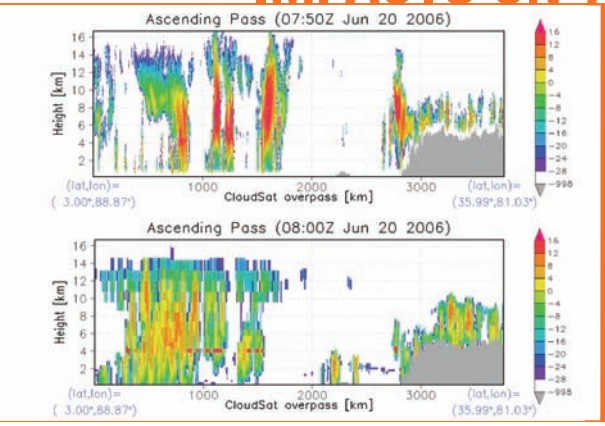


Figure 2: Comparisons of total (left) and absorbing (right) aerosol optical depth from Aerosol Robotic Network (AERONET) observations and GOCART simulations in seven global regions. Data are monthly averaged values in 2004.

HIGH-RESOLUTION MODELING OF AEROSOL IMPACTS ON THE ASIAN MONSOON WATER CYCLE



SCIENCE MISSION DIRECTORATE

WILLIAM K. LAU
NASA Goddard Space Flight Center
(301) 614-6332
William.K.Lau@nasa.gov

◀ **Figure 1:** Cross-sections of radar echo observed by the CloudSat Cloud Profiling Radar (top) and simulated by the Weather Research and Forecasting (WRF) Model (bottom) for June 20, 2006. The model echoes are based on the Satellite Data Simulation Unit (SDSU) radar simulator. Units are in dBZ, a measure of reflectivity.

Project Description: Severe floods and droughts caused by monsoon fluctuations have always impacted Asian societies. As monsoons now interact with aerosols from industrial and urban pollution, they threaten the water supply, human health, and biodiversity of the Asian monsoon region. This interaction may also exacerbate the global effects of climate change, given the region's dense population and its role in the global water cycle.

This project aims to clarify the interactions between aerosols and the monsoon water cycle, and how they may modulate the regional climatic impacts of global warming. The project is testing various feedback hypotheses using high-resolution climate models and satellite and *in situ* observations.

We are using the regional-scale Weather Research and Forecasting (WRF) Model to simulate the impacts of absorbing aerosols (dust and black carbon) on the Indian monsoon water cycle and to test several hypotheses [5, 6]. Thanks to its flexible horizontal and vertical resolution, WRF can realistically represent a wide range of physical processes, topography, radiative transfer, precipitation and cloud processes, and land/atmosphere hydrology and energy coupling. We have implemented a radiation module that computes short- and longwave radiative flux and atmospheric heating. This scheme links to the Goddard Chemistry Aerosol Radiation and Transport (GOCART) aerosol module to estimate aerosol concentrations, optical depth, optical properties, and effects on cloud-precipitation microphysics, structure, and dynamics.

Relevance of Work to NASA: With funding from the Earth Science Division's Interdisciplinary Investigation Program, this work serves NASA's goal to understand climate change impacts on the water cycle. It uses satellite products from NASA's Tropical Rainfall Measuring Mission (TRMM), Moderate Resolution Imaging Spectroradiometer (MODIS), CloudSat, and Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation (CALIPSO) and outputs from a NASA chemistry transport model to construct aerosol-forcing functions and validate model results.

Computational Approach: In order to resolve the microphysics of cloud and rain formation, we run WRF at very high resolution—less than 10-kilometer (km) horizontal grid spacing with 31 vertical layers. To mitigate the large computational demand, we use a triple-nest grid with horizontal resolutions of 27, 9, and 3 km. Using the Discover supercomputer at the NASA Center for Computational Sciences, we conducted a model integration for the period May 1 to July 1 (covering the pre-monsoon and the monsoon onset) in both 2005 and 2006. This provides a good case study, as observations show heavier loading of absorbing aerosols (dust and black carbon) in 2006.

Results: We recently documented the "elevated-heat-pump" (EHP) hypothesis linking aerosols to the monsoon cycle; this highlights the role of the Himalayas and Tibetan Plateau in trapping aerosols over the Indo-Gangetic Plain. Studies using satellite and reanalysis data show preliminary evidence of aerosol impacts on the variability of the Indian monsoon. We have also implemented the radiative codes associated with different aerosol species, so that we can use the model to study the impacts of aerosol forcing. And through control and anomalies experiments, we have tested the model's sensitivity to the domain design and the cumulus parameterization, and ensured implementation of the proper aerosol radiation codes.

In the 2005/6 model integration, simulated cloud vertical profiles compare fairly well with CloudSat observations (Figure 1), and simulated rainfall agrees reasonably well with daily rainfall distribution (Figure 2), producing heavy rain (red) over the Bay of Bengal and the western coast. Preliminary results show significant promise for the WRF-GOCART module in studying aerosol effects on the monsoon cycle.

Role of High-End Computing: A cloud-resolving mesoscale model is computationally demanding; our model domain contains over 200,000 grid cells. By using 256 Intel Xeon processors on Discover, we can finish a 1-day integration in less than 3 hours. The challenge now becomes data migration or transfer, as network bandwidths have not kept pace with computing speeds.

Future: We expect to complete the first integration for the two simulated months over WRF simulation Domains 1 (27-km grid) and 2 (9-km grid), using several schemes for aerosol forcing and cumulus parameterization. We will then carry out a 2-week model integration using Domain 3 (3-km grid) to examine diurnal variability, orographic effects, and possible microphysics effects on land convection immediately before and after the monsoon onset. We expect model outputs on the order of 50 terabytes.

Co-Investigators

- Wei-Kuo Tao, Mian Chin, NASA Goddard Space Flight Center
- Kyu-Myong Kim, Jainn J. (Roger) Shi, Toshihisa Matsui, all of University of Maryland, Baltimore County

Publications

[1] Gautam, R., Hsu, N.C., Kafatos, M., Tsay, S.-C., "Influences of Winter Haze on Fog/Low Cloud over the Indo-Gangetic Plains," *Journal of Geophysical Research*, Vol. 112, D05207, doi:10.1029/2005JD00703, 2007.

[2] Lau, K.M., Kim, M.K., and Kim, K.M., "Asian Summer Monsoon Anomalies Induced by Aerosol Direct Forcing: The Role of the Tibetan Plateau," *Climate Dynamics*, Vol. 26 (7-8), pp. 855–864, doi:10.1007/s00382-006-0114-z, 2006.

[3] Lau, K.M., and Kim, K.M.,"Observational Relationships Between Aerosol and Asian Monsoon Rainfall, and Circulation," *Geophysical Research Letters*, Vol. 33, L21810, doi:10.1029/2006GL027546, 2006.

[4] Lau, K.M, Ramanathan, V., Wu, G-X., Li, Z., Tsay, S.C., Hsu, C., Siika, R., Holben, B., Lu, D., Tartari, G., Chin, M., Koudelova, P., Chen, H., Ma, Y., Huang, J., Taniguchi, K., and Zhang., R., "The Joint Aerosol-Monsoon Experiment: A New Challenge in Monsoon Climate Research," *Bulletin of the American Meteorological Society*, Vol. 89, pp. 369–383, DOI:10.1175/BAMS-89-3-369, 2008.

[5] Lau, K.-M., and Kim, K.-M., "Absorbing Aerosols Enhance Indian Summer Monsoon Rainfall," *iLEAPS Newsletter*, No. 5, pp. 22–24, 2008.

[6] Lau, K.-M., Kim, K.-M., Hsu, C.N.-Y., and Singh, R.P., "Seasonal Co-Variability of Aerosol and Precipitation over the Indian Monsoon and Adjacent Deserts," *GEWEX News*, 18(1), pp. 4–6, 2008.

[7] Shi, J.J., Matsui, T., Tao, W.-K., Chin, M., and Peters-Lidard, C.,"Implementation of the Updated Goddard Longwave and Shortwave Radiation Packages into WRF," 2007 WRF Users' Workshop, Boulder, CO, 2007.

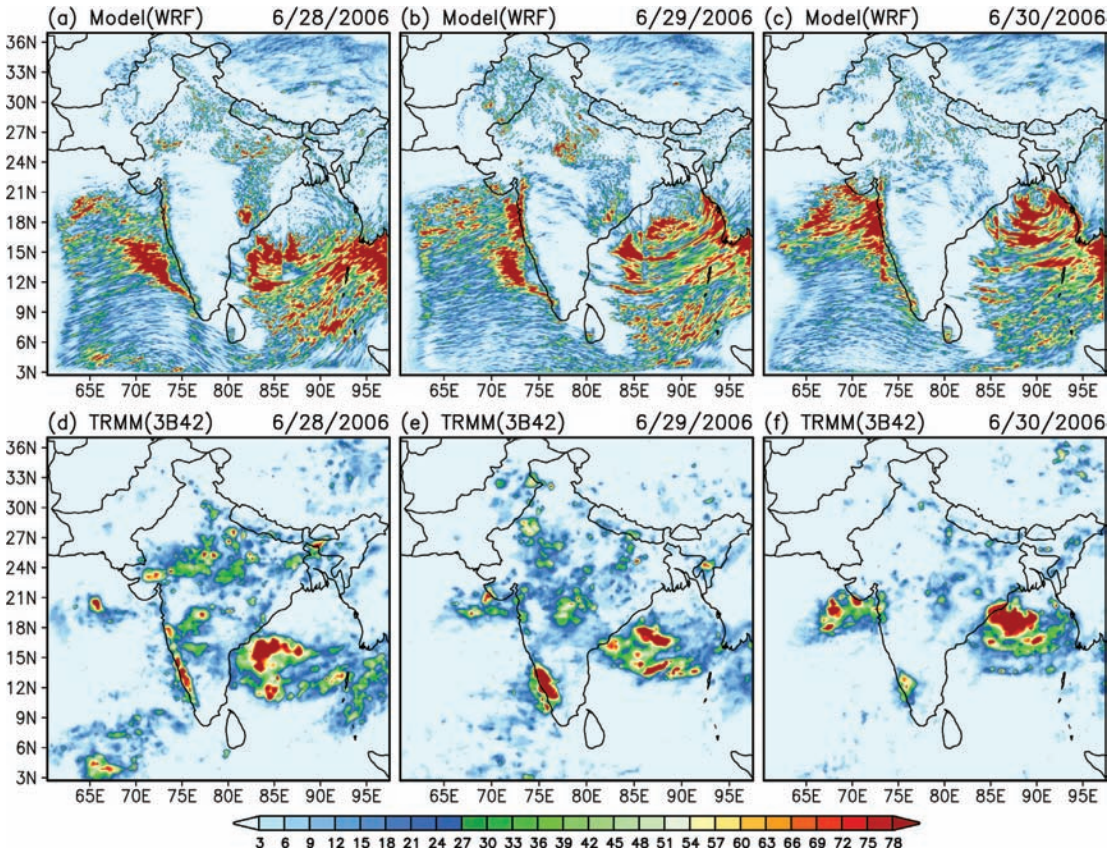
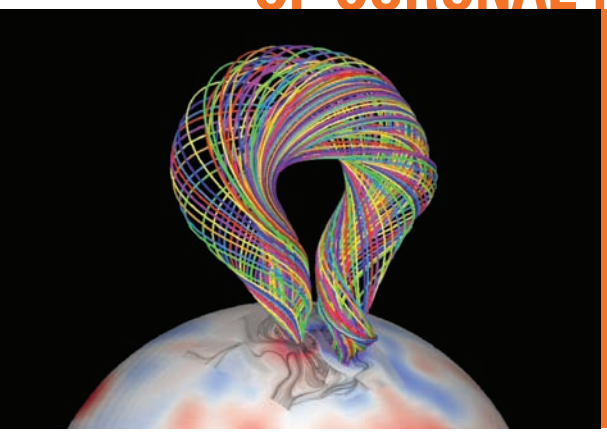


Figure 2: Rainfall distributions from Weather Research and Forecasting (WRF) Model simulations at 9-kilometer resolution (top row) and from Tropical Rainfall Measurement Mission (TRMM) satellite estimates (bottom row). Units are in millimeters per day.

HIGH-RESOLUTION SIMULATIONS OF CORONAL MASS EJECTIONS



SCIENCE MISSION DIRECTORATE

ZORAN MIKIC
Predictive Science, Inc.
(858) 450-6493
mikicz@predsci.com
http://www.predsci.com

◀ Detail from Figure 1.

Project Description: Coronal mass ejections (CMEs) are violent eruptions that send giant clouds of solar plasma into space. The radiation and energetic particles from the solar flares that frequently accompany CMEs can cripple satellites, disrupt communications and power systems, and endanger humans in space. We believe that we can best understand the mysterious properties of CMEs using numerical simulations of solar magnetic fields. Models have matured enough that we can use them to interpret CME observations; understand the complex behavior of the Sun, especially the structure of the high-temperature corona and the dynamic events that occur within it; and determine how that structure and those events flow outward and manifest themselves in the surrounding heliosphere.

We are using numerical simulations of the solar corona to investigate what powers CMEs, how they are initiated, and how they propagate in the inner heliosphere. This task requires the development of detailed models of the corona and solar wind. We use "event studies" of CMEs in which we apply models that are as realistic as possible to make direct comparisons between models and observations. A long-term goal is to improve our ability to predict when flares and CMEs will occur. Such predictions are necessary for undertaking NASA's Vision for Space Exploration.

Relevance of Work to NASA: A central goal of NASA's Heliophysics program is to understand the influence of the Sun and its activity on the inner heliosphere. Present and upcoming NASA missions—including the Solar and Heliospheric Observatory (SOHO), Hinode, the Solar Terrestrial Relations Observatory (STEREO), and the Solar Dynamics Observatory (SDO)—produce massive quantities of high-resolution observations of the Sun. We need sophisticated 3D numerical models to interpret these data and, in turn, understand solar physics. Understanding CMEs, including their initiation, propagation, and interaction with the Earth's magnetosphere, is a central goal of the National Space Weather Program, which aims to protect the nation's space assets. Funding for this research comes from NASA's Heliophysics Theory and Living With a Star (LWS) Programs.

Computational Approach: We use 3D magnetohydrodynamic (MHD) numerical models to simulate in detail how the magnetic field of the Sun behaves. Our simulations use tens of millions of mesh points and tens of thousands of time-steps to evolve the solar magnetic field. Our 3D code uses implicit time-differencing, requiring iterative solvers to invert the resulting very large sparse matrices.

Results: We have been able to simulate the initiation and propagation of a CME from an active region observed on May 12, 1997, including its signature in extreme ultraviolet (EUV) and X-ray emissions (Figure 1). The simulated and observed CMEs have strikingly similar characteristics. We studied the topology of the magnetic field in detail to understand the role of magnetic reconnection in the eruption process. We have embarked on a new event study of the May 13, 2005 CME as part of a Focused Science Team for NASA's LWS Program. We have also predicted the structure of the solar corona for the August 1, 2008 total solar eclipse (Figure 2). These studies are crucial tests of the model's predictive capability and give us important clues on how to improve the model. A version of our 3D code, Magnetohydrodynamics outside A Sphere (MAS), is available to the heliophysics community through NASA's Community Coordinated Modeling Center (CCMC).

Role of High-End Computing: We need massively parallel supercomputers such as Columbia at the NASA Advanced Supercomputing (NAS) facility to solve the stiff equations that describe the evolution of magnetic fields on the Sun. Our code uses domain decomposition and runs particularly well on the fast-communication interconnect on Columbia using the standard Message Passing Interface (MPI). The code works on

a wide variety of architectures and scales well on thousands of processors. It is suitable for transition to the next generation of massively parallel petascale machines.

Future: The high-resolution images that will come from NASA's SDO mission will require simulations with even more mesh points. In addition, SDO's multi-spectral observations will provide opportunities to improve our understanding of coronal heating mechanisms by comparing models with EUV and X-ray emission observations. We plan to continue to improve our model to meet these challenges.

Co-Investigators

- Jon A. Linker, Pete Riley, Roberto Lionello, Viacheslav Titov, all of Predictive Science, Inc.
- Yung Mok, University of California, Irvine

Publications

[1] Mikic, Z., and Lee, M.A., "An Introduction to Theory and Models of CMEs, Shocks, and Solar Energetic Particles," *Space Science Review*, 123, 57, 2006.

[2] Riley, P., Linker, J.A., Mikic, Z., and Odstrcil, D., "Modeling Interplanetary Coronal Mass Ejections," *Advances in Space Research*, 38, 535, 2006.

[3] Mikic, Z., Linker, J.A., Lionello, R., Riley, P., and Titov, V., "Predicting the Structure of the Solar Corona for the Total Solar Eclipse of March 29, 2006," in *Solar and Stellar Physics Through Eclipses* (O. Demircan, S. O. Selam, and B. Albayrak, eds.), Astronomical Society of the Pacific Conference Series, 370, 299, 2007.

[4] Amari, T., Aly, J.J., Mikic, Z., and Linker, J.A., "Coronal Mass Ejection Initiation and Complex Topology Configurations in the Flux Cancellation and Breakout Models," *Astrophysical Journal Letters,* 671, L189, 2007.

[5] Riley, P., Lionello, R., Mikic, Z., and Linker, J.A., "Using Global Simulations to Relate the Three-Part Structure of Coronal Mass Ejections to In Situ Signatures," *Astrophysical Journal*, 672, 1221, 2008.

[6] Titov, V.S., Mikic, Z., Linker, J.A., and Lionello, R., "1997 May 12 Coronal Mass Ejection Event. I. A Simplified Model of the Preeruptive Magnetic Structure," *Astrophysical Journal*, 675, 1614, 2008.

[7] Mok, Y., Mikic, Z., Lionello, R., and Linker, J.A., "The Formation of Coronal Loops by Thermal Instability in Three Dimensions," *Astrophysical Journal Letters*, 679, L161, 2008.

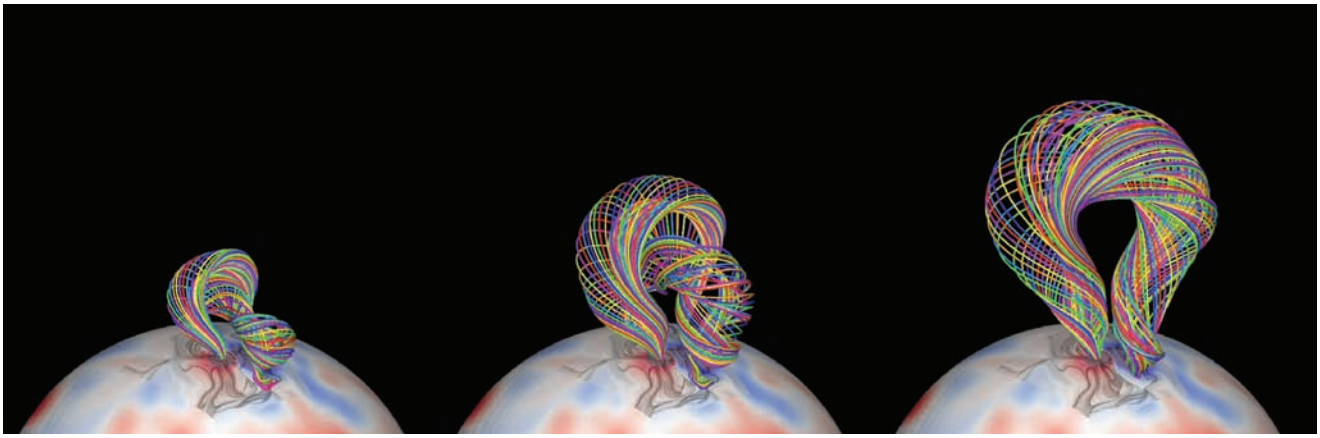
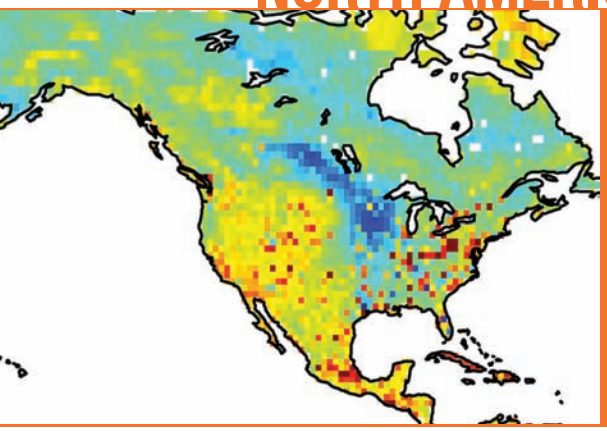


Figure 1: A numerical simulation of the eruption of a coronal mass ejection from the Sun on May 12, 1997. The magnetic field lines show a flux rope that eventually becomes an interplanetary magnetic cloud, which was observed at Earth.



Figure 2: The August 1, 2008 total solar eclipse corona as predicted by a magnetohydrodynamic (MHD) model and as observed from Bor Udzuur, Mongolia.

HIGH-RESOLUTION WIND FIELDS FOR CONSTRAINING NORTH AMERICAN FLUXES OF CARBON DIOXIDE



SCIENCE MISSION DIRECTORATE

ANNA M. MICHALAK
The University of Michigan
(734) 763-9664
amichala@umich.edu
http://www-personal.umich.edu/~amichala/Research/

◀ Close-up of Figure 2.

Project Description: Our work is part of an ongoing project entitled *Constraining North American Fluxes of Carbon Dioxide and Inferring Their Spatiotemporal Covariances through Assimilation of Remote Sensing and Atmospheric Data in a Geostatistical Framework*, funded through NASA's North American Carbon Program (NACP). The overall goal is to use remotely sensed and atmospheric data in a geostatistical inverse modeling framework to quantify North American sources and sinks of carbon dioxide (CO_2) with unprecedented spatial and temporal resolution. Using the Columbia supercomputer at the NASA Advanced Supercomputing (NAS) facility, we are producing high-resolution meteorological fields with the Weather Research and Forecasting (WRF) atmospheric model, and using these to drive the Stochastic Time-Inverted Lagrangian Transport (STILT) atmospheric transport model and calculate the sensitivity of atmospheric CO_2 measurements to surface fluxes.

This project will address the NACP's stated goals of (i) developing quantitative scientific knowledge, robust observations, and models to determine emissions and uptake of CO_2 and the factors regulating these processes, and (ii) developing the scientific basis to implement full carbon accounting, including natural and anthropogenic fluxes of CO_2 on regional and continental scales. We will achieve our goal of quantifying CO_2 surface fluxes without relying on *a priori* flux estimates, while rigorously quantifying the magnitude and spatiotemporal covariance of the various components of model, measurement, and flux errors. In addition, we will evaluate the sensitivity of the inferred fluxes to available remotely sensed environmental datasets, providing the process-based understanding needed to improve bottom-up inventories and biospheric models, thereby enabling more accurate flux accounting.

Relevance of Work to NASA: The NACP implementation plan calls for the development, demonstration, and evaluation of inverse techniques for estimating monthly CO_2 exchange on a 100-km grid from atmospheric observations. We address this need by merging available atmospheric measurements of CO_2 with NASA remote-sensing data to estimate fluxes of CO_2 directly at high spatiotemporal resolutions, without relying on prior estimates of flux distributions. The project also presents a significant opportunity to use NASA's spatially distributed satellite observations of the Earth to improve our understanding of the North American carbon budget.

Computational Approach: In order to resolve mesoscale circulation, cloud venting, and other detailed atmospheric phenomena that affect CO_2 transport, the WRF model is nested down to high resolution (2 km) over target regions surrounding tall towers with continuous, long-term, hourly monitoring of CO_2 concentrations in Wisconsin, Maine, and Texas. The 2-km domain is the innermost grid; it is embedded within a 10-km grid domain over the Eastern United States, and an outermost 40-km grid domain covering Mexico, the U.S., and Canada. Two-way nesting allows the dynamics simulated within the innermost nest to propagate outward into coarser nests. The WRF model is forced at the lateral boundaries by meteorological fields from the National Centers for Environmental Prediction (NCEP) North American Regional Reanalysis, which has already assimilated an extensive array of atmospheric observations from radiosondes, satellites, and ground stations. These fields are also used for analysis nudging of the outermost domain. The nested WRF windfields contain much more fine-scale information than the 2-degree (or coarser) products used by typical inverse modeling studies, or readily available regional analysis fields typically at 40-km resolution. The WRF fields produced from this study will be freely available to the carbon community.

Results: We began using the NAS Columbia system actively in spring 2007. Since that time, we have generated nested meteorological fields for 2004, 2005, 2006, and 2007. In spring 2008, we began experimenting with the STILT model on Columbia to quantify the sensitivity of available measurements of atmospheric CO_2 to terrestrial fluxes at a 1-degree

resolution for the North American continent (Figures 1 and 2). We are now using the system to run STILT for all years for which meteorological fields have been generated.

Role of High-End Computing: Columbia has enabled us to create an unprecedented dataset of meteorological fields for 2004 to 2007 using the WRF model, a well-established community mesoscale modeling framework that supports multiple dynamical cores, numerical schemes, and physical parameterization packages. The software framework has built-in support for parallel architectures. Specifically, we use the Advanced Research WRF (ARW) supported by the National Center for Atmospheric Research. The ARW is a finite-difference code using 3rd-order Runge-Kutta time-stepping. Physics packages are available for the treatment of soil, surface, boundary layer, moist physics, and radiative processes.

Future: Our work focuses on quantifying the sensitivity of available continuous measurements of atmospheric CO_2 to North American CO_2 fluxes at fine spatiotemporal resolutions. We expect this work to continue for the upcoming year. Through ongoing collaborations, we plan ultimately to create a 5-year record of meteorology, atmospheric transport, and estimated fluxes.

Co-Investigators

- Adam I. Hirsch, University of Colorado and NOAA Earth System Research Laboratory
- Thomas Nehrkorn, Atmospheric and Environmental Research, Inc.
- John C. Lin, University of Waterloo

Publications

[1] Mueller, K.L., Gourdji, S.M., and Michalak, A.M., "Global Monthly Averaged CO_2 Fluxes Recovered Using a Geostatistical Inverse Modeling Approach: 1. Results Using Atmospheric Measurements," *Journal of Geophysical Research*, 113, D21114, doi:10.1029/2007JD009734, 2008.

[2] Gourdji, S.M., Mueller, K.L., Schaefer, K., and Michalak, A.M., "Global Monthly Averaged CO_2 Fluxes Recovered Using a Geostatistical Inverse Modeling Approach: 2. Results Including Auxiliary Environmental Data," *Journal of Geophysical Research*, 113, D21115, doi:10.1029/2007JD009733, 2008.

[3] Michalak, A.M., "Technical Note: Adapting a Fixed-Lag Kalman Smoother to a Geostatistical Atmospheric Inversion Framework," *Atmospheric Chemistry and Physics*, 8, 6789–6799, 2008.

Figure 1: Sensitivity of measurements available in June 2004 to CO_2 sources and sinks throughout the North American continent.

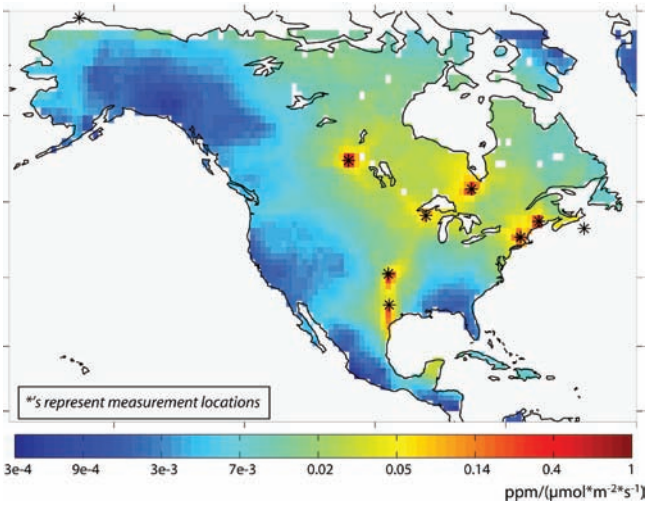
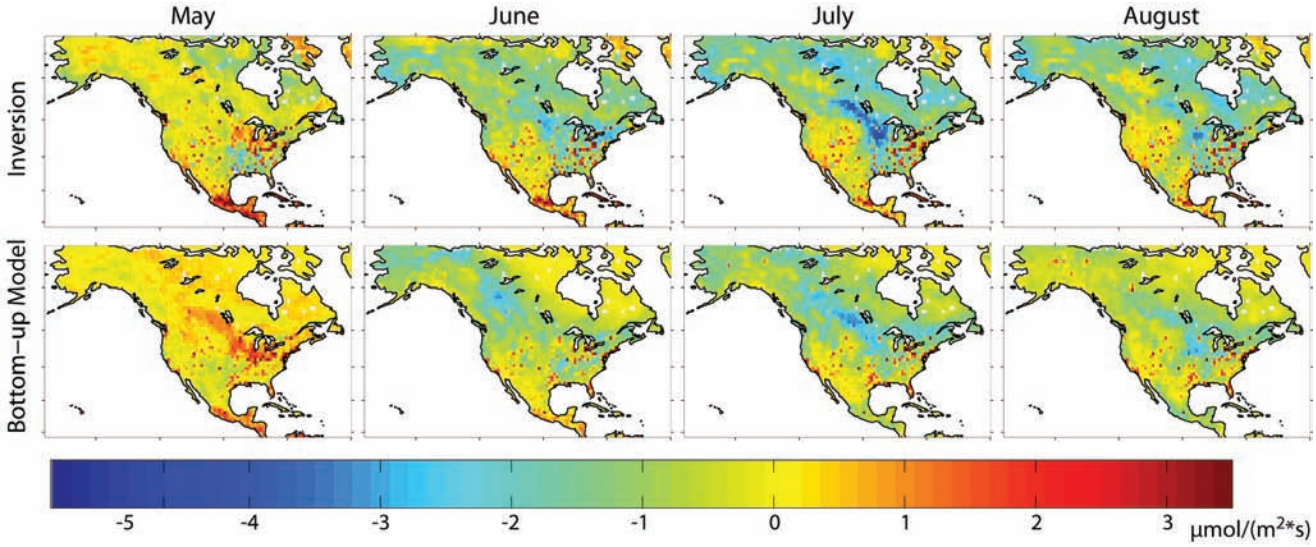
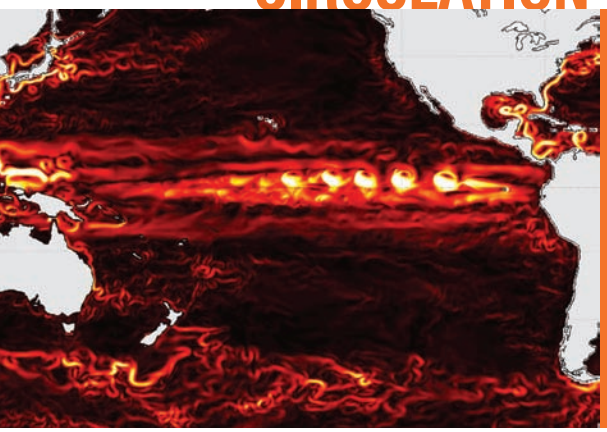


Figure 2: Top row: monthly grid-scale CO_2 flux estimates for 2004 resulting from the North American geostatistical inversion. In addition to atmospheric measurements and transport information, investigators incorporated several Moderate Resolution Imaging Spectroradiometer (MODIS)-derived data products, climate parameters, and fossil fuel inventories into the inversion. Bottom row: for comparison, grid-scale flux estimates derived from bottom-up or mechanistic models (the Carnegie-Ames-Stanford Approach [CASA] biospheric model, fire emissions, and fossil fuel inventories) are also presented. Because the geostatistical inversion does not make use of bottom-up models, inversion results can also serve as an evaluation tool for such models.



NON-BOUSSINESQ OCEAN GENERAL-CIRCULATION MODEL AND GRACE APPLICATIONS



◀ Close-up of Figure 1.

SCIENCE MISSION DIRECTORATE

Y. TONY SONG
NASA Jet Propulsion Laboratory
(818) 393-4876
Tony.Song@jpl.nasa.gov
http://science.jpl.nasa.gov

Project Description: With funding from NASA's Modeling, Analysis, and Prediction (MAP) Program, we are modeling global ocean-ice circulation and validating the results with data from the Gravity Recovery and Climate Experiment (GRACE) satellites. Our goal is to develop an advanced mass-conserving (Non-Boussinesq) ocean general-circulation model (OGCM), allowing use of satellite data for a better understanding of the ocean's climate. The scientific objectives are threefold:

- Compare GRACE, Earth rotation, and other geodetic observations with mass-conserving models from the NASA Jet Propulsion Laboratory (JPL) to improve NASA's next-generation assimilation system.
- Quantify the dynamic balance of wind-stress curl and bottom-pressure torque by using wind data from NASA's Quick Scatterometer (QuikSCAT), bottom-pressure data from GRACE, and simulation products from NASA's Global Modeling and Assimilation Office.
- Study ocean-solid-Earth interactions and global sea-level changes by applying geodetic observations to the climate model system.

The global ocean-ice model is designed to better represent ocean mass variation and the effect of topography on ice-shelf flows. Our coupled ocean-ice system includes a dynamic-thermodynamic sea-ice component based on an elastic-viscous-plastic rheology. The ocean and ice models communicate by the exchange of heat, fresh water, and momentum at the ocean-ice interaction layer. The water depth of the model is divided into 30 terrain-following, stretched-pressure levels from shallow coast to the deep ocean [1]. Topography data is from ETOPO2 with modification from the Navy's DBDB2 bathymetry to give special care to strait geometry when generating the model grid. The model is spun up for 30 years with an annual mean forcing to reach an approximate steady state, and then driven by daily National Centers for Environmental Prediction (NCEP)/National Center for Atmospheric Research (NCAR) reanalysis forcing from 1970 to the present. With the northern pole displaced towards Russia, the mass-conserving ocean and ice model has included both polar regions. A snapshot of the model ice concentration and ocean circulation is shown in Figure 1.

Relevance of Work to NASA: This project will provide an improved picture of ocean bottom pressure and its dynamic link to sea-surface height and geodetic observations. This will provide a direct comparison with GRACE and Earth rotation observations as well as insights into the determination and interpretation of global sea-level change. The derived ocean and ice data will be made available on-line (http://earthquake-tsunami.jpl.nasa.gov). The synergistic applications of the satellite data should provide insights to addressing NASA's research strategy questions: "How is the global ocean circulation varying on interannual, decadal, and longer time scales?," "How can climate variations induce changes in global ocean circulation?," and "How is global sea-level affected by climate change?"

Computational Approach: The global ocean and ice model is modular Fortran 90 and Fortran 95 code. It uses C preprocessing to activate the physical and numerical options. We have established several coding standards to facilitate model readability, maintenance, and portability. All the state model variables are dynamically allocated and passed as arguments to the computational routines via de-referenced pointer structures. All private or scratch arrays are automatic; their size is determined when the procedure is entered. This code structure facilitates computations over nested and composed grids.

Results: We have completed three tasks: (i) compared model ocean-bottom-pressure results with GRACE observations (Figure 2); (ii) carried out coupled earthquake-OGCM simulations for the 2004 Indian Ocean tsunami [2, 3]; and (iii) coupled a sea-ice model into the OGCM [4–7].

Role of High-End Computing: The parallel framework is coarse-grained, with both shared- and distributed-memory paradigms coexisting in the same code. The code has extensive pre- and post-processing software for data preparation, analysis, plotting, and visualization. All model input and output is via NetCDF, which facilitates the interchange of data between computers, user communities, and other independent analysis software. The speed and parallel capabilities of NASA's High-End Computing facilities are crucial to the project.

Future: We plan to couple and develop an ice-sheet model with JPL's OGCM for a better understanding of ocean-ice interactions, ice-shelf/ice-sheet/glacier melting, and their impact on oceanic climate and sea-level rise. We also plan to consider both polar regions in assessing mass exchanges between the cryosphere and oceans with data from GRACE and ICESat (Ice, Cloud, and land Elevation Satellite)—which measure ice-sheet and glacier mass balance and sea-ice thickness change.

Co-Investigators

- Richard Gross, Victor Zlotnicki, NASA Jet Propulsion Laboratory
- C.K. Shum, Ohio State University
- Dale Haidvogel, Rutgers University

Publications

[1] Song, Y.T., and Hou, T.Y., "Parametric Vertical Coordinate Formulation for Multiscale, Boussinesq, and Non-Boussinesq Ocean Modeling," *Ocean Modelling*, Vol. 11, Nos. 3–4, pp. 298–332, Doi:10.1016/j.ocemod.2005.01.001, 2006.

[2] Song, Y.T., Fu, L.-L., Zlotnicki, V., Ji, C., Hjorleifsdottir, V., Shum, C.K., and Yi, Y., "The Role of Horizontal Impulses of the Faulting Continental Slope in Generating the 26 December 2004 Tsunami," *Ocean Modelling*, Vol. 20, No. 4, pp. 362–379, doi:10.1016/j.ocemod.2007.10.007, 2008.

[3] Song, Y.T., "Detecting Tsunami Genesis and Scales Directly from Coastal GPS Stations," *Geophysical Research Letters*, Vol. 34, L19602, doi:10.1029/2007GL031681, 2007.

[4] Zheng, Q., Susanto, R.D., Ho, C.R., Song, Y.T., and Xu, Q., "Statistical and Dynamical Analyses of Generation Mechanisms of Solitary Internal Waves in the Northern South China Sea," *Journal of Geophysical Research-Oceans*, Vol. 112, No. C3, C03021, doi:10.1029/2006JC003551, 2007.

[5] Song, Y.T., and Zlotnicki, V., "The Subpolar Ocean-Bottom-Pressure Oscillation and its Links to ENSO," *International Journal of Remote Sensing*, Vol. 29, No. 21, pp. 6091–6107, 2008.

[6] Zlotnicki, V., Wahr, J., Fukumori, I., and Song, Y.T., "The Antarctic Circumpolar Current: Seasonal Transport Variability During 2002–2005," *Journal of Physical Oceanography*, Vol. 37, doi:10.1175/JPO3009.1 2006.

[7] Song, Y.T., "Estimation of Interbasin Transport using Ocean Bottom Pressure: Theory and Model for Asian Marginal Seas," *Journal of Geophysical Research*, Vol. 111, C11S19, doi:10.1029/2005JC003189, 2006.

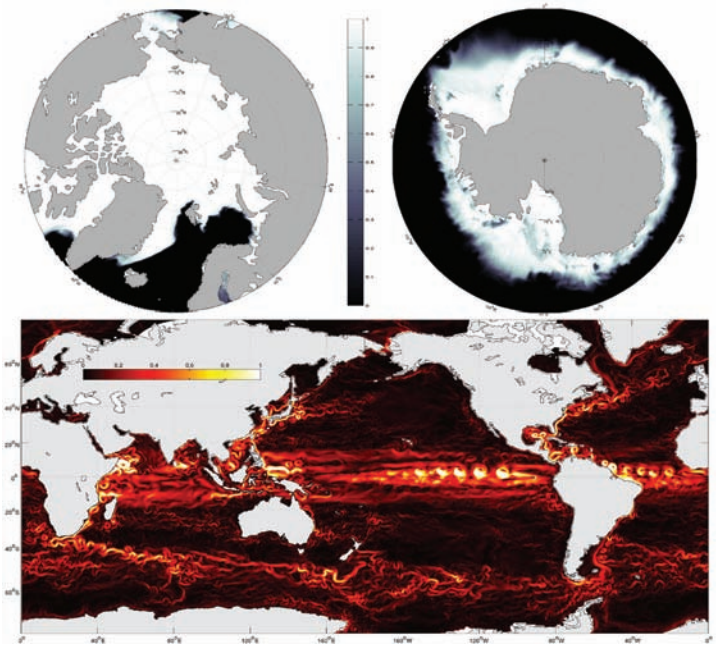


Figure 1: Coupled ocean-ice model for the global ocean. The top panels show simulated ice concentration in the Arctic and Antarctic regions. The bottom panel shows speeds of simulated ocean currents. In this projection, the model's north pole is shifted towards Russia to avoid the computational singularity.

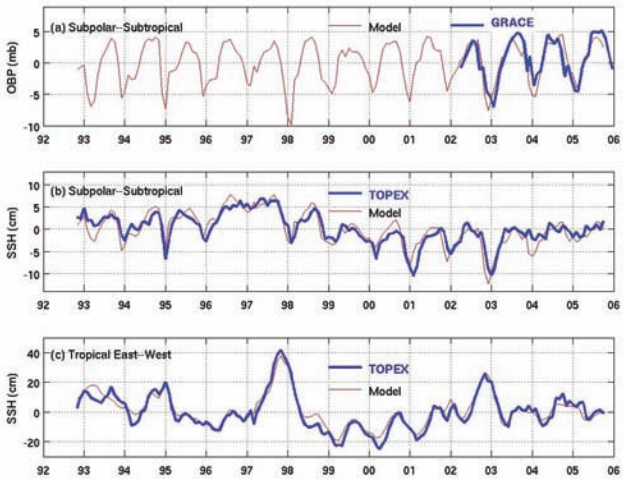
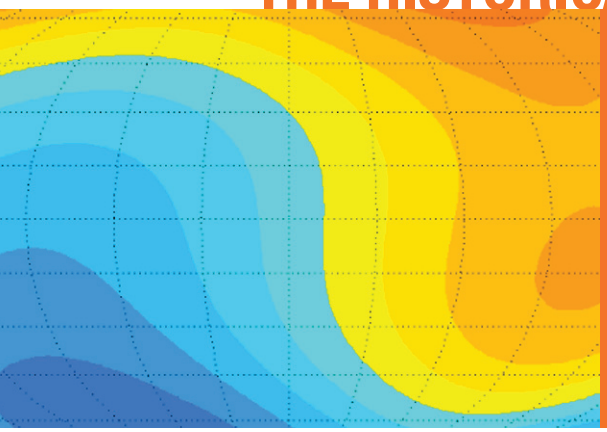


Figure 2: Comparison of the ocean model with observed data from 1993 to 2006. The top panel shows simulated ocean-bottom-pressure variability compared with Gravity Recovery and Climate Experiment (GRACE) observations of the ocean in the subpolar-subtropical regions. The middle and bottom panels show simulated sea-surface heights compared with TOPEX/Poseidon observations of the Pacific Ocean in the subpolar-subtropical and tropical east-west regions.

NUMERICAL SIMULATION OF THE HISTORICAL MARTIAN DYNAMO



SCIENCE MISSION DIRECTORATE

WEIJIA KUANG
NASA Goddard Space Flight Center
(301) 614-6108
Weijia.Kuang-1@nasa.gov

◀ Detail of Figure 1.

Project Description: The Mars Global Surveyor's detection of a Martian magnetic anomaly suggests that Mars once possessed an active, global, internal magnetic field that was generated and maintained by convective flow in an electrically conducting fluid core (dynamo). It also appears that the dynamo action lasted for only several hundred million years after the accretion of the planet.

We aim to understand the Martian dynamo and its termination by means of large-scale numerical simulations with the *m*MoSST (Message Passing Interface (MPI)-based, Modular, Scalable, Self-consistent and Three-dimensional) core dynamics model developed at NASA Goddard Space Flight Center. Our simulations focus on: (1) subcritical dynamos near the end of the Martian dynamo era; (2) the effects of Martian interior structures on the subcritical dynamo states; and (3) the geophysical implications of the subcritical dynamos on Martian magnetic anomalies and the evolution of Mars.

Our research objective is to answer fundamental science questions related to the observed magnetic anomaly: How much energy was needed to sustain the Martian dynamo? When and how was it terminated? What information or constraint could Martian magnetism provide to the evolution of Mars?

Relevance of Work to NASA: By studying Mars' magnetic properties, the research will improve knowledge of the planet's interior and its evolution history. It may identify geophysical mechanisms for the termination of the Martian dynamo, provide better interpretation of past observations and current missions to Mars, and potentially support scientific goals for future missions. It directly addresses NASA Planetary Science Research Objective 3C.1, "Learn how the Sun's family of planets and minor bodies originated and evolved," of the NASA Strategic Sub-Goal 3C, "Advance scientific knowledge of the origin and history of the solar system, the potential for life elsewhere, and the hazards and resources." It also contributes to Research Objectives 3C.2 and 3C.3 via investigation of Mars magnetic field evolution. Funding for this research comes through the NASA Mars Fundamental Research Program and the NASA Earth Surface and Interior Program.

Computational Approach: Our *m*MoSST core dynamics model uses a hybrid, spectral/finite-difference algorithm to solve a chaotic magnetohydrodynamic system in a rapidly rotating spherical shell. The code is written in Fortran 90/95 with modular structures. MPI libraries perform communication among distributed processors. We carry out most of our simulation runs with 128 processors for grid sizes up to 128×128×128. More resources are required for higher-resolution simulations.

Results: The research has shown, first, that the Martian dynamo can be subcritical, i.e., that the energy to sustain the dynamo can be lower than the critical energy necessary to excite the dynamo. Also, a subcritical dynamo tends to reverse more frequently than supercritical dynamos, resulting in a mean dipole field aligning closer to the equator (Figure 1). Furthermore, existence of a subcritical dynamo is not affected by the dimension of the inner core, though smaller inner cores would lead to smaller subcritical domains. Finally, a subcritical Martian dynamo could be terminated very quickly over a short period (fewer than 1 million years) by a small perturbation (e.g., less than 1% perturbation to the heat flow across the core-mantle boundary), and once terminated, it could not be reactivated even if the core's original geophysical state were restored.

Role of High-End Computing: We run all of our simulations on the NASA Advanced Supercomputing (NAS) facility's Columbia supercomputer. This system provides over 1 million processor-hours each year for our project, and we could not achieve our project goals without it. For example, simulations of subcritical dynamos require model domains at least

of the order 100×100×100 for accurate determination of the critical points for onset and termination of the dynamo. A single simulation run needs approximately 100 processors for 240 wall-clock hours. The massive storage system at NAS is also necessary for archiving all numerical solutions (on the order of terabytes) for post-simulation analysis and other geophysical applications.

Future: Recent results from several research groups (including ours) suggest that the Martian dynamo might have been terminated in the Later Heavy Bombardment (LHB) period, which occurred between 3.8 and 4.1 billion years ago. The giant impacts during LHB could have provided sufficient perturbation to shut down a subcritical Martian dynamo. To further investigate this scenario, one must understand the properties of Martian dynamos in a heterogeneous thermodynamic environment resulting from giant impacts on Mars. This will be a focus for our future simulations. Another research activity will be geomagnetic data assimilation, in which we will use surface geomagnetic observations and our numerical geodynamo model to better understand the dynamics inside the Earth's core and its impact on changes of the Earth over long periods.

Co-Investigators

- Weiyuan Jiang, University of Maryland, Baltimore County

Publications

[1] Kuang, W., Jiang, W., and Wang, T., "Sudden Termination of Martian Dynamo?: Implications from Subcritical Dynamo Simulations," 35, L14204, *Geophysical Research Letters*, doi:10.1029/2008GL034183, 2008.

[2] Jiang, W., and Kuang, W., "An MPI-Based MoSST Core Dynamics Model," *Physics of the Earth and Planetary Interiors*, Vol. 170, Issues 1–2, September 2008, pp. 46–51, doi:10.1016/j.pepi.2008.07.020, 2008.

[3] Kuang, W., Tangborn, A., Jiang, W., Liu, D., Sun, Z., Bloxham, J., and Wei, Z., "MoSST DAS: The First Generation Geomagnetic Data Assimilation Framework," *Communications in Computational Physics*, Vol. 3, pp. 85–108, 2008.

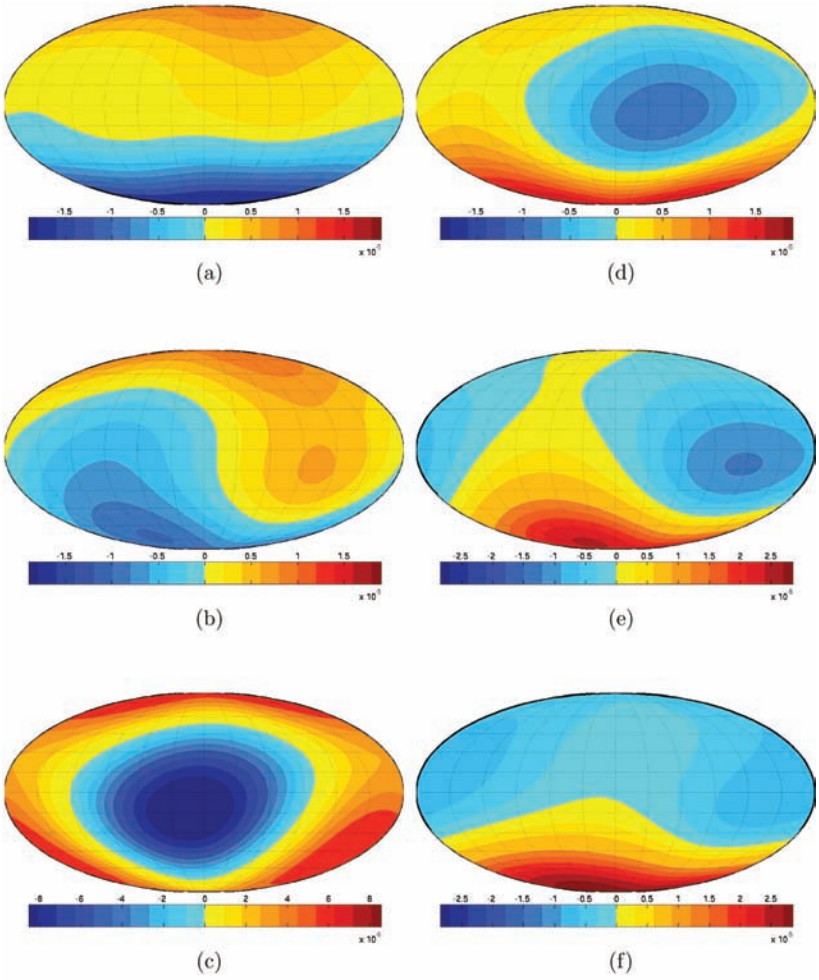
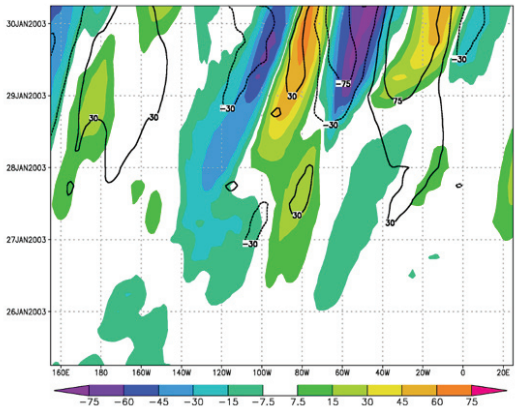


Figure 1: Snapshots of the radial component of the magnetic field at the surface of Mars during a reversal process. The polarity of the field in (a) reverses by the end of the process, shown in (f).

OBSERVING SYSTEM EXPERIMENTS: EVALUATING AND ENHANCING THE IMPACT OF SATELLITE OBSERVATIONS



SCIENCE MISSION DIRECTORATE

ORESTE REALE
NASA Goddard Space Flight Center/University of Maryland, Baltimore County
(301) 614-6254
Oreste.Reale-1@nasa.gov

◀ **Figure 1:** Hovmoeller diagram (time upward) of the latitude-averaged (40N–80N) Advanced Infrared Spectrometer (AIRS)-induced impact on the 500 hectopascal (hPa) geopotential height between longitudes 160E and 20E. The shading represents the AIRS vs. Control difference (in meters). The solid lines represent the verifying National Centers for Environmental Prediction (NCEP) analyses minus the Control [1].

Project Description: This project aims to evaluate and optimize the use of satellite data, in particular the Advanced Infrared Spectrometer (AIRS), in global data assimilation and modeling. An Observing System Experiment (OSE) assesses the impact of an observational instrument by producing two or more data assimilation runs, one of which (the Control run) omits data from the instrument under study. From the resulting analyses, we initialize corresponding forecasts and evaluate them against the National Centers for Environmental Prediction (NCEP) operational analyses. A synoptic and dynamic evaluation then shows how the additional data propagate from the initial conditions and are amplified by the model's dynamics.

Relevance of Work to NASA: This work aims to improve modeling and prediction of weather and climate, and ultimately to increase our understanding of the Earth's atmosphere. Primary funding comes from NASA's Modeling, Analysis, and Prediction (MAP) Program. This work also contributes to a multi-agency team, the Joint Observing System Simulation Experiment (OSSE), led by M. Masutani at NCEP.

Computational Approach: The main tool for our work is the Goddard Earth Observing System Model, Version 5 (GEOS-5), provided by the Global Modeling and Assimilation Office (GMAO)—in particular its Data Assimilation System (DAS) and forecasting system. We also develop diagnostics to assess the impact of different instruments or datasets. We run the GEOS-5 DAS at a resolution of ½ degree longitude and latitude, with 72 vertical levels; forecasts have a resolution of ½ or ¼ degree.

Results: Our research has demonstrated the impact of quality-controlled AIRS observations under partly cloudy conditions. The improved coverage leads to a substantially different thermal structure in boreal winter conditions, particularly in the high latitudes. This improves forecast skill by enhancing the representation of the jet stream and baroclinic activity [1]. Figure 1 shows the impact of AIRS: between longitudes 120W and 20E, the AIRS 5-day forecast goes in the same direction as the NCEP analyses. This impact arises out of a substantially different representation of the mid-low tropospheric thermal structure in a data-poor region (northeastern Siberia and polar regions), which is captured in the AIRS analyses [1].

With the GMAO's L.P. Riishojgaard and E. Liu, we are also comparing the impact of AIRS temperature retrievals with that of clear-sky radiances. Despite an overall consensus that radiances are better, we have found that rejecting all data from cloud-contaminated areas, as done in the clear-sky radiance approach, severely reduces the spatial coverage and undermines the AIRS impact on the DAS and forecasting system.

In modeling tropical cyclogenetic processes, we have found that using AIRS data under partly cloudy conditions leads to better-defined tropical storms and improved GEOS-5 track forecasts. We performed two sets of experiments. One set centered on April–May 2008, when Tropical Cyclone Nargis hit Myanmar. The other centered on August–September 2006, overlapping with the NASA African Monsoon Multidisciplinary Analysis (NAMMA) observing campaign. The first of these studies aims to improve analysis and prediction of cyclones over the northern Indian Ocean. We are finding that AIRS strongly affects GEOS-5 analyses, leading to a deeper and better-located storm center. Clear-sky radiances have an intermediate impact, possibly due to less extensive coverage. As shown in Figure 2, AIRS temperature retrievals under partly cloudy conditions deepen Nargis' center, displaying a well-defined low and closed circulation not seen in the Control run. The corresponding forecasts compare well with the observed storm track.

In the second study, AIRS has improved the representation of African Easterly Waves, their interaction with Saharan air, and

the organization of developing waves in closed circulations. We simulated the genesis of Hurricane Helene—observed at the end of the NAMMA campaign—using the GEOS-5 DAS with and without AIRS retrievals under partly cloudy conditions. The AIRS data significantly deepened the center of the simulated storm when the observed system was still an intensifying tropical storm. Forecasts initialized from these improved analyses led to more accurate storm tracks.

Finally, within the Joint OSSE team we contributed to assessing the suitability of the new Nature Run produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) [2].

Role of High-End Computing: NASA High-End Computing (HEC) resources are crucial for this work. In the last 2 years, we have run about 70 month-long assimilation experiments at different resolutions and corresponding 5-day forecasts on the Explore supercomputer at the NASA Center for Computational Sciences and the Columbia supercomputer at the NASA Advanced Supercomputing facility. The mass storage allows us to continue analyzing model results with diagnostic tools that we have developed within the HEC environment.

Future: We are further investigating the impact of AIRS in the analysis and prediction of Tropical Cyclone Nargis. In collaboration with Riishojgaard and Liu, we are also comparing the retrieval and radiance approach from a global point of view, with rigorous and comprehensive tests in different conditions (boreal winter and boreal summer). We continue to study the impact of AIRS in modeling Atlantic cyclone development and the impact of model resolution on tropical cyclone representation, with forecasts at ½- and ¼-degree resolution.

In the future, we plan to assimilate retrievals of moisture, and temperature retrievals above clouds, for improved representation of the outflow above tropical cyclones. We will also evaluate the AIRS cloud-cleared radiances (that is, radiances that would be observed if the scene were clear), and will start assessing AIRS version 6 when available.

We will continue to benefit from collaboration with J. Susskind, Riishojgaard, and the GMAO, which provides GEOS-5 and frequent scientific and technical updates.

Co-Investigators

- William K. Lau, NASA Goddard Space Flight Center

Publications

[1] Reale, O., Susskind, J., Rosenberg, R., Brin, E., Liu, E., Riishojgaard, L.P., Terry, J., Jusem, J. C., "Improving Forecast Skill by Assimilation of Quality-Controlled AIRS Temperature Retrievals under Partially Cloudy Conditions," *Geophysical Research Letters*, Vol. 35, L08809, 2008.

[2] Reale, O., Terry, J., Masutani, M., Andersson, E., Riishojgaard, L.P., Jusem, J.C., "Preliminary Evaluation of the European Centre for Medium-Range Weather Forecasts (ECMWF) Nature Run over the Tropical Atlantic and African Monsoon Region," *Geophysical Research Letters*, Vol. 34, L22810, 2007.

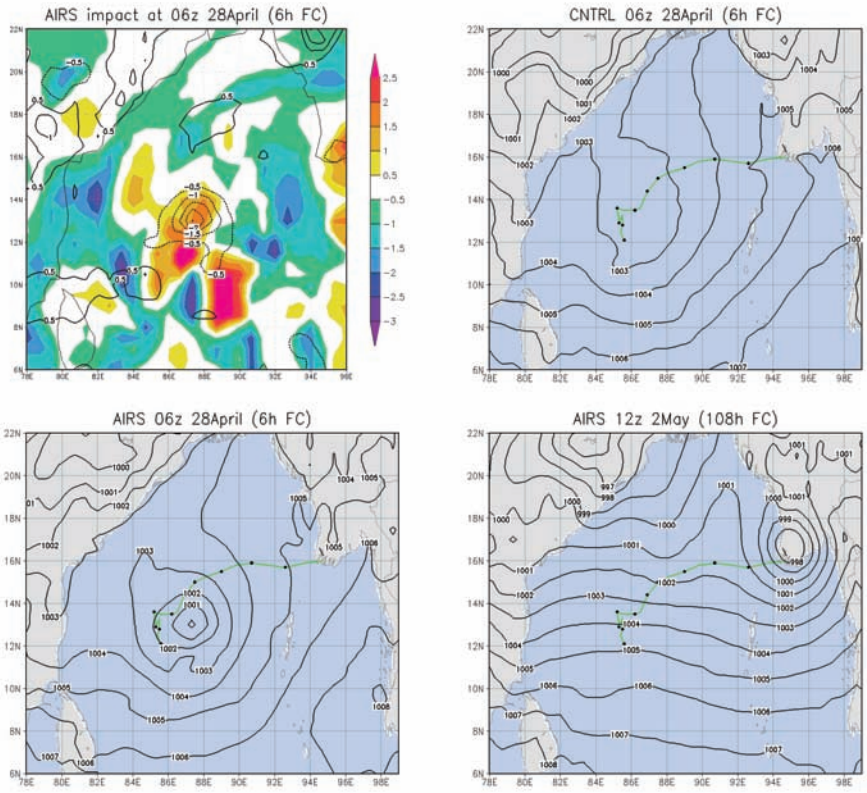
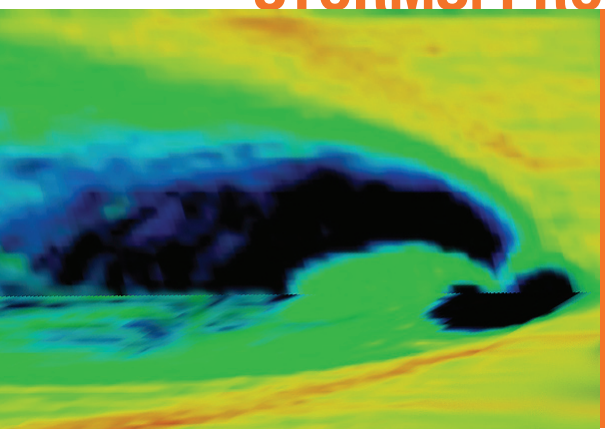


Figure 2: Impact of AIRS on the ½-degree Goddard Earth Observing System Model, Version 5 (GEOS-5) forecast for Tropical Cyclone Nargis. Upper left: Differences (AIRS minus Control) in 6-hour forecasts of 200 hPa temperature (°C, shaded) and sea-level pressure (hPa, solid line). Lower left: The 6-hour sea-level pressure forecast from the AIRS run shows a well-defined low close to the observed storm track (green solid line). Lower right: The corresponding 108-hour forecast for 2 May 2008 (landfall time) compares very well with the observed track. Upper right: The 6-hour sea-level pressure forecast from the Control run shows no detectable cyclone.

PLASMA REDISTRIBUTION DURING GEOSPACE STORMS: PROCESSES AND CONSEQUENCES



SCIENCE MISSION DIRECTORATE

THOMAS E. MOORE
NASA Goddard Space Flight Center
(301) 286-5236
thomas.e.moore@nasa.gov
http://gpl.gsfc.nasa.gov/public/traj/

◀ Close-up of Figure 2.

Project Description: This effort supports global modeling of the coupled solar wind, ionosphere, and magnetosphere through theoretical analysis and assimilation of observational results into empirical-statistical models. We are establishing: 1) how the flow of energy during geospace storms alters ionospheric plasma expansion into the magnetosphere and 2) how this expansion influences the dynamics and coupling of the solar wind, magnetosphere, and ionosphere to create these storms. Emphasis is on the inner magnetosphere's plasma and geomagnetic field conditions (including distortion by the ring current) and the evolution of ionospheric conductance, temperature, and densities. We aim to identify the principal features of plasma redistribution and to assess their impacts quantitatively over the range of storm conditions driven by the solar wind and the interplanetary magnetic field.

By comparing our model results with data from the Heliophysics Great Observatory, we validate the dynamic local response of source regions to solar wind influences and the simulated characteristics of the magnetospheric circulation. Resulting improvements in our simulation results lead toward enhanced global circulation models of geospace and its response to the dynamic heliosphere. We use data from the Advanced Composition Explorer (ACE), Wind, Geotail, and other missions to establish external drivers of the magnetospheric system. We also use data from Polar, the Fast Auroral SnapshoT (FAST), the Defense Meteorological Satellites Program (DMSP), and other low-Earth orbit missions to establish the spatial distribution and rate of ionospheric expansion; and data from the Polar, Imager for Magnetopause-to-Aurora Global Exploration (IMAGE), Cluster, and Department of Energy geosynchronous missions to make contact with observed magnetospheric responses.

Relevance of Work to NASA: Large-scale redistribution and restructuring of the ionosphere by storm-induced currents and electric fields produce massive ion plasma flows into the magnetosphere. Consequences include an enhanced polar wind, a heavy-ion auroral wind, and convective entrainment of the eroding plasmaspheric plumes. Entrained ionospheric plasmas populate the plasma sheet and ring current, modify magnetospheric convection and current systems, and thereby couple back into ionospheric plasma electrodynamics. We are taking major steps toward quantifying the effects of storm-time ionospheric restructuring on the magnetosphere and the dynamic development of this feedback—which are essential to forecasting near-Earth space weather. Funding for this research comes from NASA's Living With a Star Targeted Research and Development Program.

Computational Approach: This work uses a large number of theoretical and empirical models (Figure 1). We compute the outer magnetosphere and its interaction with the solar wind by integrating the equations of motion of millions of non-interacting charged particles in specified electric, magnetic, and gravitational fields. We make no drift approximations and integrate the full gyro-motion of the ions using adjustable resolution in space and time, based on a 4th-order Runge-Kutta scheme developed by collaborator Dominique Delcourt of the Centre d'Étude des Environnements Terrestre et Planétaires (CETP) in France.

We simulate the inner magnetosphere using the Comprehensive Ring Current Model (CRCM), which includes gyration and bounce-averaged particle motions, pitch-angle distributions, and charge exchange collisions. It also includes electrodynamic interaction with the ionosphere, leading to a self-consistent, stress-driven solution for the ring current pressure distribution. Outer magnetosphere results from the single-particle simulations supply the outer boundary condition.

Results: We successfully introduced the plasmasphere into the global circulation and outer magnetosphere configuration (Figure 2). We used the CRCM's plasmaspheric simulation

to specify a source of cold ionospheric plasma at the outer boundary of the plasmasphere, where individual ions enter the outer magnetospheric circulation. The ions can return to the plasmasphere after substantial acceleration by the solar wind interaction, supplying energized ions back into the inner magnetosphere. We have implemented these capabilities using iterations of the outer and inner magnetospheric simulations [3].

Role of High-End Computing: Our primary computing platform is the Discover system at the NASA Center for Computational Sciences (NCCS). High-end computing is central to our work, owing to the large number of particle trajectories we need for bulk parameter calculations and kinetic work. To date, the particles are non-interacting, for simpler computation; yet the results have been well received and have motivated others in the global simulation community to develop multi-fluid codes that incorporate ionospheric plasmas as a dynamic element of the outer magnetosphere.

Future: Our next goal is to develop "virtual spacecraft" that can simulate the measurements of any particular spacecraft. This capability will facilitate comparing our results with specific spacecraft data. Having done this for fixed locations in space, our main work is to implement specified orbit parameters to enable the simulation to sample the magnetosphere as the spacecraft would do.

Another goal is to simulate particle trajectories using the new generation of global, multi-fluid simulations containing ionospheric plasma outflows in the circulation. Our first case will study only solar wind (ignoring ionospheric outflow) to test the multi-fluid code and prepare code interfaces for more interesting work that accounts for ionospheric outflows. This work will be of mainly technical interest, but it may lead to kinetic studies of solar wind ions.

Co-Investigators

- Mei-Ching H. Fok, NASA Goddard Space Flight Center

Publications

[1] Fok, M.-C., Moore, T.E., Brandt, P.C., Delcourt, D.C., Slinker, S.P., and Fedder, J.A., "Impulsive Enhancements of Oxygen Ions During Substorms," *Journal of Geophysical Research*, 111, A10222, doi:10.1029/2006JA011839, 2006.

[2] Moore, T.E., Fok, M.-C., Delcourt, D.C., Slinker, S.P., and Fedder, J.A., "Global Aspects of Solar Wind-Ionosphere Interactions," *Journal of Atmospheric and Solar-Terrestrial Physics*, 69, 265, doi:10.1016/j.jastp.2006.08.009, 2007.

[3] Moore, T.E., Fok, M.-C., Delcourt, D.C., Slinker, S.P., and Fedder, J.A., "Plasma Plume Circulation and Impact in an MHD Substorm," *Journal of Geophysical Research*, 113, A06219, doi:10.1029/2008JA013050, 2008.

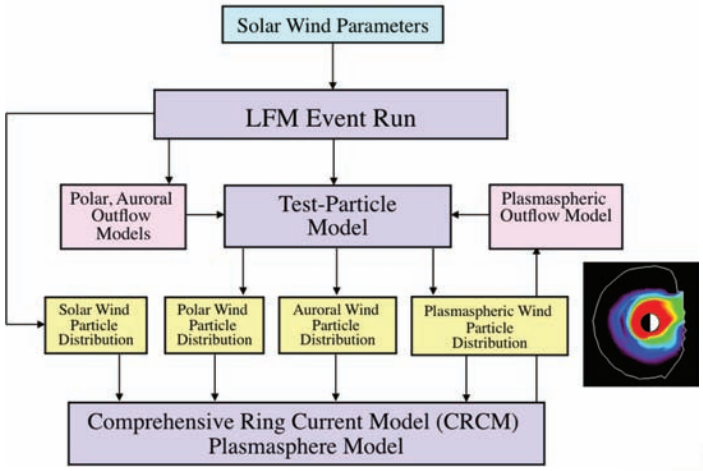


Figure 1: A simplified flowchart of a Global Ion Kinetic Model, which simulates plasma redistribution during solar storms. Solar wind parameters are inputs to the Lyon-Fedder-Mobarry (LFM) global fluid simulation, which computes the distribution of solar wind fluid in the magnetosphere. An empirical model specifies ionospheric outflows to generate test particle weightings and initial condition ranges and then launches millions of particles with random properties selected from those ranges. Resultant plasma properties serve as the boundary condition for the Comprehensive Ring Current Model (CRCM) of the inner magnetosphere. CRCM serves as a boundary condition for launching plasmaspheric ions back into the outer magnetosphere, where they circulate. These ions may return to the inner magnetosphere, where CRCM processes them.

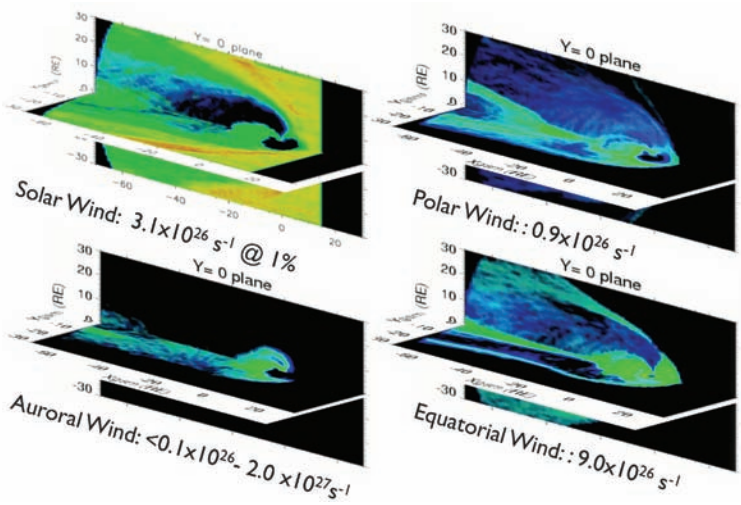
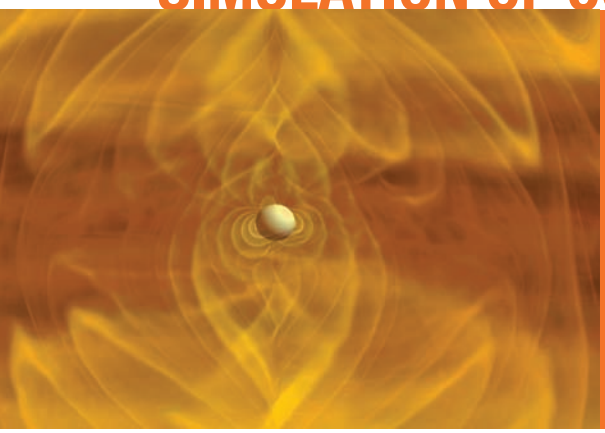


Figure 2: Distributions of multiple magnetospheric plasma populations as simulated using our global ion kinetic methods. Each case—Solar Wind, Polar Wind, Auroral Wind, and Equatorial (or Plasmaspheric) Wind—shows the density at a time representative of moderate magnetospheric convection flow in two planes. Labels under each plot indicate the fluence of ions into the magnetosphere.

SIMULATION OF COALESCING BLACK HOLE BINARIES



SCIENCE MISSION DIRECTORATE

JAMES VAN METER
NASA Goddard Space Flight Center
(301) 286-8241
james.r.vanmeter@nasa.gov
http://astrophysics.gsfc.nasa.gov/gravity/

◀ Detail of Figure 1.

Project Description: With support from NASA's Astrophysics Theory and Fundamental Astrophysics Program, we are simulating black hole binaries, in which black holes inspiral towards each other and merge into a single black hole remnant (Figure 1). Every stage of this process generates copious gravitational radiation; indeed, coalescing black holes are considered the most promising source of observable gravitational waves. Our simulations solve Einstein's field equations for spacetime using a finite-differencing code. The primary purpose of this work is to supply waveform templates for analyzing data from the planned Laser Interferometer Space Antenna (LISA) mission. Such templates are essential for performing matched filtering of the data and discerning signals from noise.

Relevance of Work to NASA: Understanding the dynamics and gravitational radiation of coalescing black holes is critical for the success of LISA and for meeting NASA science goals:

"Understand how the first stars and galaxies formed, and how they changed over time into the objects observed in the present universe": Calculating gravitational radiative recoil from asymmetric mergers helps to predict the probability that black holes will escape their host galaxies. By supplying waveform templates for LISA data analysis, we will aid both LISA and the ground-based Laser Interferometer Gravitational-Wave Observatory (LIGO) in pinpointing the positions and characteristics of the coalescing supermassive black holes that are believed to populate galactic cores, where they play central roles (literally) in galaxy evolution.

"Understand the origin and destiny of the universe, phenomena near black holes, and the nature of gravity": Black hole mergers may eventually serve as "standard candles," supplanting supernovae as the most accurate measure of the universe-expanding "dark energy." Accurately identifying and locating such mergers will require waveform templates for matched filtering of the data. Further, verification of our computed waveforms by laser interferometers will test Einstein's theory of general relativity to unprecedented precision and confirm the prediction of gravitational waves for the first time.

Computational Approach: We use a homegrown, finite-differencing code to solve Einstein's field equations. For adaptive mesh refinement (AMR) and parallelization, we currently use the PARAMESH software but are transitioning to an alternative AMR package called Carpet. Differencing and interpolation is currently 5th-order accurate. A 4th-order Runge-Kutta algorithm performs time-integration.

Results: We have made substantial scientific progress, including achievement of several milestones. We were the first to accurately compute the radiative recoil from an unequal mass binary and the first to simulate a merger preceded by as many as seven orbits. We have made important contributions to ongoing analytical efforts to model radiative recoil from binaries of arbitrary mass ratio and spins. We were the first to demonstrate consistency of numerical waveforms with post-Newtonian predictions. And we have contributed to waveform analysis in various other ways, including refinement of an effective-one-body, post-Newtonian model.

Our progress has benefited from significant technological developments. We developed coordinate conditions tailored to moving black holes and adapted these conditions for unequal mass binaries to facilitate sufficient numerical accuracy around the smaller black hole. Implementation of dissipation and constraint-damping has enabled stable and accurate evolutions of arbitrary duration. We have moved from a 2nd-order accurate to a fully 5th-order accurate evolution code and from 2nd-order to 6th-order accurate radiation extraction. We also have invented a novel algorithm for computing the spin of a black hole.

Role of High-End Computing: Our computational grid must have sufficient extent and resolution to simulate a vast region of physical space accurately, which includes the black hole sources, a "wave zone" where radiation is measured, and an outer boundary sufficiently removed to minimize reflections into the wave zone. Our simulations have often employed on the order of 10 million grid points, which have been just barely adequate so far—we desire an increase by at least an

order of magnitude. At each grid point we store on the order of 100 double-precision variables. Consequently, running on a large number of parallel processors—at least 100 and ideally more than 1,000—is essential. Our calculations have benefited from such resources at both of NASA's High-End Computing Program facilities.

Future: First, we aim to simulate a mass ratio of greater than 10:1—which has never been done before—to verify analytic waveform models. Second we intend to simulate a variety of mass ratio and spin configurations, including exotic "transitional precession" events involving both large spin and mass ratio. Such exotic events may result in interesting exceptions to the relatively simple wave-forms simulated to date. Finally, we plan to model accretion disks around black holes, first with collisionless particles and then with a newly acquired hydrodynamic code. This research will determine whether electromagnetic correlates to gravitational radiation might be observable.

Co-Investigators

- John Baker, Joan Centrella, NASA Goddard Space Flight Center

Publications

[1] Baker, J., Boggs, W., Centrella, J., Kelly, B., McWilliams, S., and van Meter, J., "Mergers of Nonspinning Black Hole Binaries: Gravitational Radiation Characteristics," *Physical Review*, D78, 044048, 2008.

[2] Baker, J., Boggs, W., Centrella, J., Kelly, B., McWilliams, S., and van Meter, J., "Modeling Kicks from the Merger of Generic Black Hole Binaries," *Astrophysical Journal Letters*, 682, L29, 2008.

[3] Schnittman, J., Buonanno, A., van Meter, J., Baker, J., Boggs, W., Centrella, J., Kelly, B., and McWilliams, S., "Anatomy of the Binary Black Hole Recoil: A Multipolar Analysis," *Physical Review*, D77, 044031, 2008.

[4] Pan, Y., Buonanno, A., Baker, J., Centrella, J., Kelly, B., McWilliams, S., Pretorius, F., and van Meter, J., "Data-Analysis Driven Comparison of Analytic and Numerical Coalescing Binary Waveforms: Nonspinning Case," *Physical Review*, D77, 024014, 2008.

[5] Choi, D., Kelly, B., Boggs, W., Baker, J., Centrella, J., and van Meter, J., "Recoiling from a Kick in the Head-On Collision of Spinning Black Holes," *Physical Review*, D76, 104026, 2007.

[6] Buonanno, A., Pan, Y., Baker, J., Centrella, J., Kelly, B., McWilliams, S., and van Meter, J., "Approaching Faithful Templates for Nonspinning Binary Black Holes Using the Effective-One-Body Approach," *Physical Review*, D76, 104049, 2007.

[7] Baker, J., van Meter, J., McWilliams, S., Centrella, J., and Kelly, B., "Consistency of Post-Newtonian Waveforms with Numerical Relativity," *Physical Review Letters,* 99, 181101, 2007.

[8] Baker, J., Boggs, W., Centrella, J., Kelly, B., McWilliams, S., Miller, M., and van Meter, J., "Modeling Kicks from the Merger of Nonprecessing Black Hole Binaries," *Astrophysical Journal*, 668, 1140, 2007.

[9] Baker, J., McWilliams, S., van Meter, J., Centrella, J., Choi, D., Kelly, B., and Koppitz, M., "Binary Black Hole Late Inspiral: Simulations for Gravitational Wave Observations," *Physical Review*, D75, 124024, 2007.

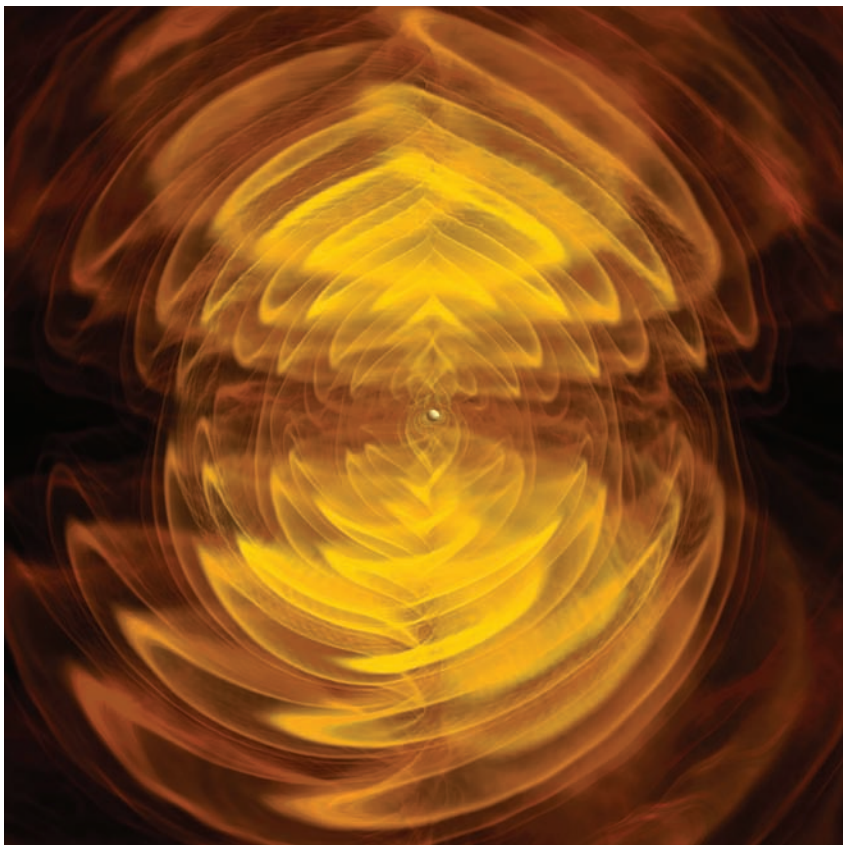


Figure 1: A simulation of X-polarized gravitational radiation from the merger of two black holes. The spherical shape in the center represents the horizon of the merged remnant. For a typical supermassive black hole of 1 million solar masses, this picture is on the order of 1 billion kilometers across.

SOLAR SURFACE MAGNETO-CONVECTION



SCIENCE MISSION DIRECTORATE

ROBERT STEIN
Michigan State University
(517) 884-5613
stein@pa.msu.edu
http://www.pa.msu.edu/~steinr/research.html

◀ **Figure 1:** Vertical cut through a 48-megameter (Mm)-wide simulation domain showing vertical velocity (red upward, blue downward) and streamlines near the solar surface (top of frame). Diverging upflows sweep downflows toward each other at the boundaries of the larger, deeper-lying upflows.

Project Description: The overall goal of this project is to understand the magneto-dynamics of the solar surface by making realistic simulations of surface magneto-convection and quantitatively comparing them with solar observations. The specific objectives are to investigate the nature of supergranulation and the evolution of the magnetic network; study the emergence of magnetic flux and its role in controlling the structure of the solar surface; and test and validate local helioseismology methods.

We perform simulations of: (i) supergranule-scale, quiet-Sun, solar-surface, hydrodynamic convection; (ii) supergranule-scale, quiet-Sun, solar-surface magneto-convection; (iii) very-high-resolution, granule-scale, solar-surface magneto-convection; and (iv) active-region evolution.

Relevance of Work to NASA: Funding for this research comes from NASA's Living With a Star (LWS) and Solar and Heliospheric Physics Programs. We use data from our simulations to interpret solar observations from the current Solar and Heliospheric Observatory (SOHO)/Michelson Doppler Imager and Hinode missions, and will do this in the future with the Solar Dynamics Observatory. We also use simulations to validate and improve observational analysis procedures, especially inversion methods in local helioseismology. Results from the magneto-convection simulations enhance our understanding of the solar magnetic field and our predictions of magnetic flux emergence before it is visible at the solar surface—an important goal of the LWS Program.

Computational Approach: Using a staggered 3D mesh, we solve the equations for mass, momentum, and internal energy in conservative form as well as the magnetic induction equation for fully compressible flow. The code uses finite differences, with 6th-order derivative operators and 5th-order interpolation operators. Time integration employs a 3rd-order, low-memory, Runge-Kutta scheme. We stabilize the code by diffusion in the momentum, energy, and induction equations. The grid is uniform in horizontal directions and non-uniform in the vertical (stratified) direction. Horizontal boundary conditions are periodic, while top and bottom boundary conditions are transmitting. By loading ghost zones at the top and bottom boundaries, we can use the same spatial derivative scheme throughout the domain.

We use a tabular equation of state, which includes local thermodynamic equilibrium (LTE) ionization of the abundant elements as well as hydrogen molecule formation, to obtain the pressure and temperature as a function of log density and internal energy per unit mass. We calculate the radiative heating/cooling by solving the radiation transfer equation in both continua and lines, assuming LTE. Using a multi-group method drastically reduces the number of wavelengths for which the transfer equation must be solved.

Results:

- *Simulation of supergranule-scale, hydrodynamic convection:* We study domains 48 and 96 megameters (Mm) wide and 20 Mm deep (Figures 1–3). The simulations cover only 10% of the geometric depth of the solar convection zone but 50% of its pressure-scale heights. They include all of the hydrogen and most of the helium ionization zones. The internal (ionization) energy flux is the largest contributor to the convective flux for temperatures less than 40,000 Kelvin; the thermal energy flux is the largest contributor at higher temperatures. The horizontal velocity spectrum is a power law, and the horizontal size of the dominant convective cells increases with depth. Convection arises from buoyancy work, which is largest close to the surface but significant over the entire domain. Two thirds of the area is upflowing fluid, except very close to the surface.
- *Application of simulation results to validate local helioseismic methods:* The simulations have a spectrum of resonant modes that agrees well (although sparser) with solar observations. We have performed time-distance and ring diagram analyses using the simulated photospheric velocities and have made comparisons with the simulation flowfield. We have found

that horizontal velocities can be determined down to several megameters below the surface, but vertical velocities cannot be accurately determined.

- *Magneto-convection simulations:* We have begun simulating both granule (6 Mm wide by 3 Mm deep) and supergranule scales (24 Mm wide by 20 Mm deep).

Role of High-End Computing: Without NASA's High-End Computing resources, we could not perform either the supergranule-scale or the high-resolution, granule-scale, magneto-convection simulations. We need to run the simulations for tens of thousands of time-steps on grids up to $1{,}000^2$ by 500. This project needs a large number of processors to obtain results in a reasonable amount of time. We get the greatest throughput with 125 processors on the Columbia system at the NASA Advanced Supercomputing (NAS) facility, although we could run efficiently with up to 1,000 processors. The NAS visualization group has been extremely helpful to us in interpreting our results.

Future: We intend to perform very-high-resolution (6 km horizontal and 5 to 14 km vertical), granule-scale (6 Mm wide by 3 Mm deep), magneto-convection simulations. They will have both initial vertical and horizontal field advected by inflows at the bottom into the computational domain. Supergranule-scale (48 Mm and 96 Mm wide by 20 Mm deep), magneto-convection simulations will have the initial horizontal field varying as a square root of density and horizontal flux advected into the domain by inflows at the bottom.

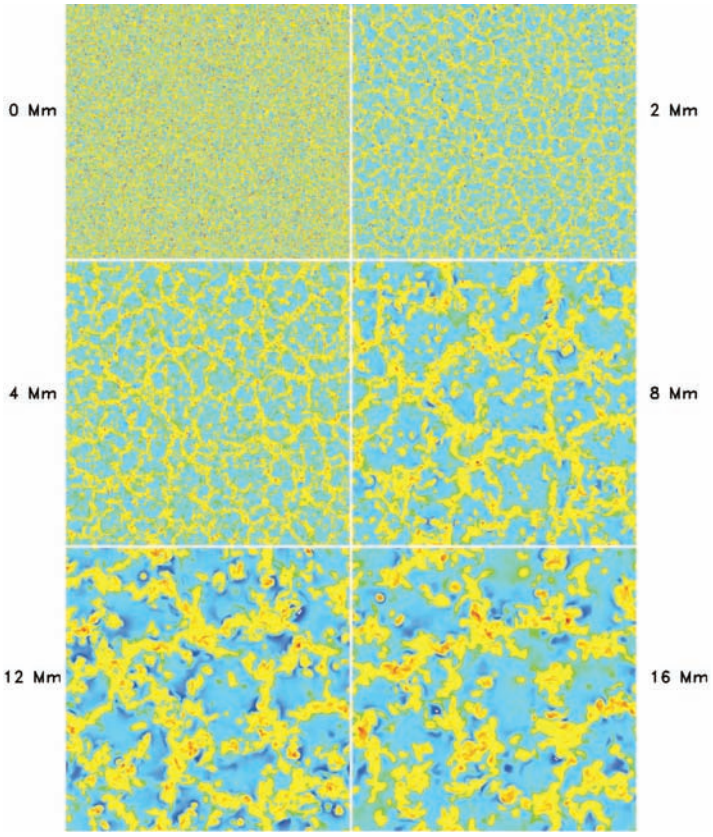


Figure 2: Horizontal slices at the solar surface (0) and 2, 4, 8, 12, and 16 Mm below the surface, showing vertical velocity. Red and yellow show downflows. Blue and green show upflows. The dominant horizontal scale of the convection increases monotonically with increasing depth.

Co-Investigators

- Åke Nordlund, Copenhagen University
- Dali Georgobiani, Werner Schaffenberger, Michigan State University
- David Benson, Kettering University

Publications

[1] Braun, D.C., Birch, A.C., Benson, D., Stein, R.F., and Nordlund, Å., "Helioseismic Holography of Simulated Solar Convection and Prospects for the Detection of Small-Scale Subsurface Flows," *Astrophysical Journal*, 669, 1395, 2007.

[2] Zhao, J., Georgobiani, D., Kosovichev, A.G., Benson, D., Stein, R.F., and Nordlund, Å., "Validation of Time-Distance Helioseismology by Use of Realistic Simulations of Solar Convection," *Astrophysical Journal*, 659, 848, 2007.

[3] Georgobiani, D., Zhao, J., Kosovichev, A.G., Benson, D., Stein, R.F., and Nordlund, Å., "Local Helioseismology and Correlation Tracking Analysis of Surface Structures in Realistic Simulations of Solar Convection," *Astrophysical Journal*, 657, 1157, 2007.

[4] Stein, R.F., Benson, D., Georgobiani, D., Nordlund, Å., and Schaffenberger, W., "Surface Convection," *Unsolved Problems in Astrophysics*, AIP Conference Proceedings 948, 111–115, 2007.

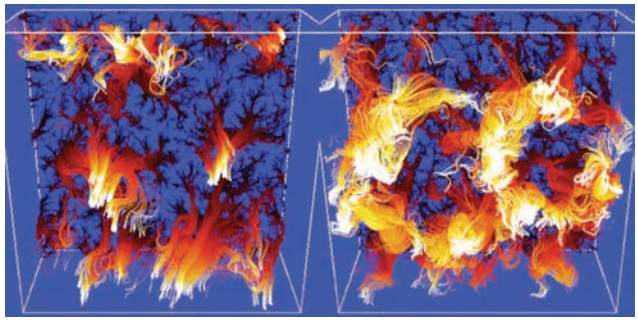
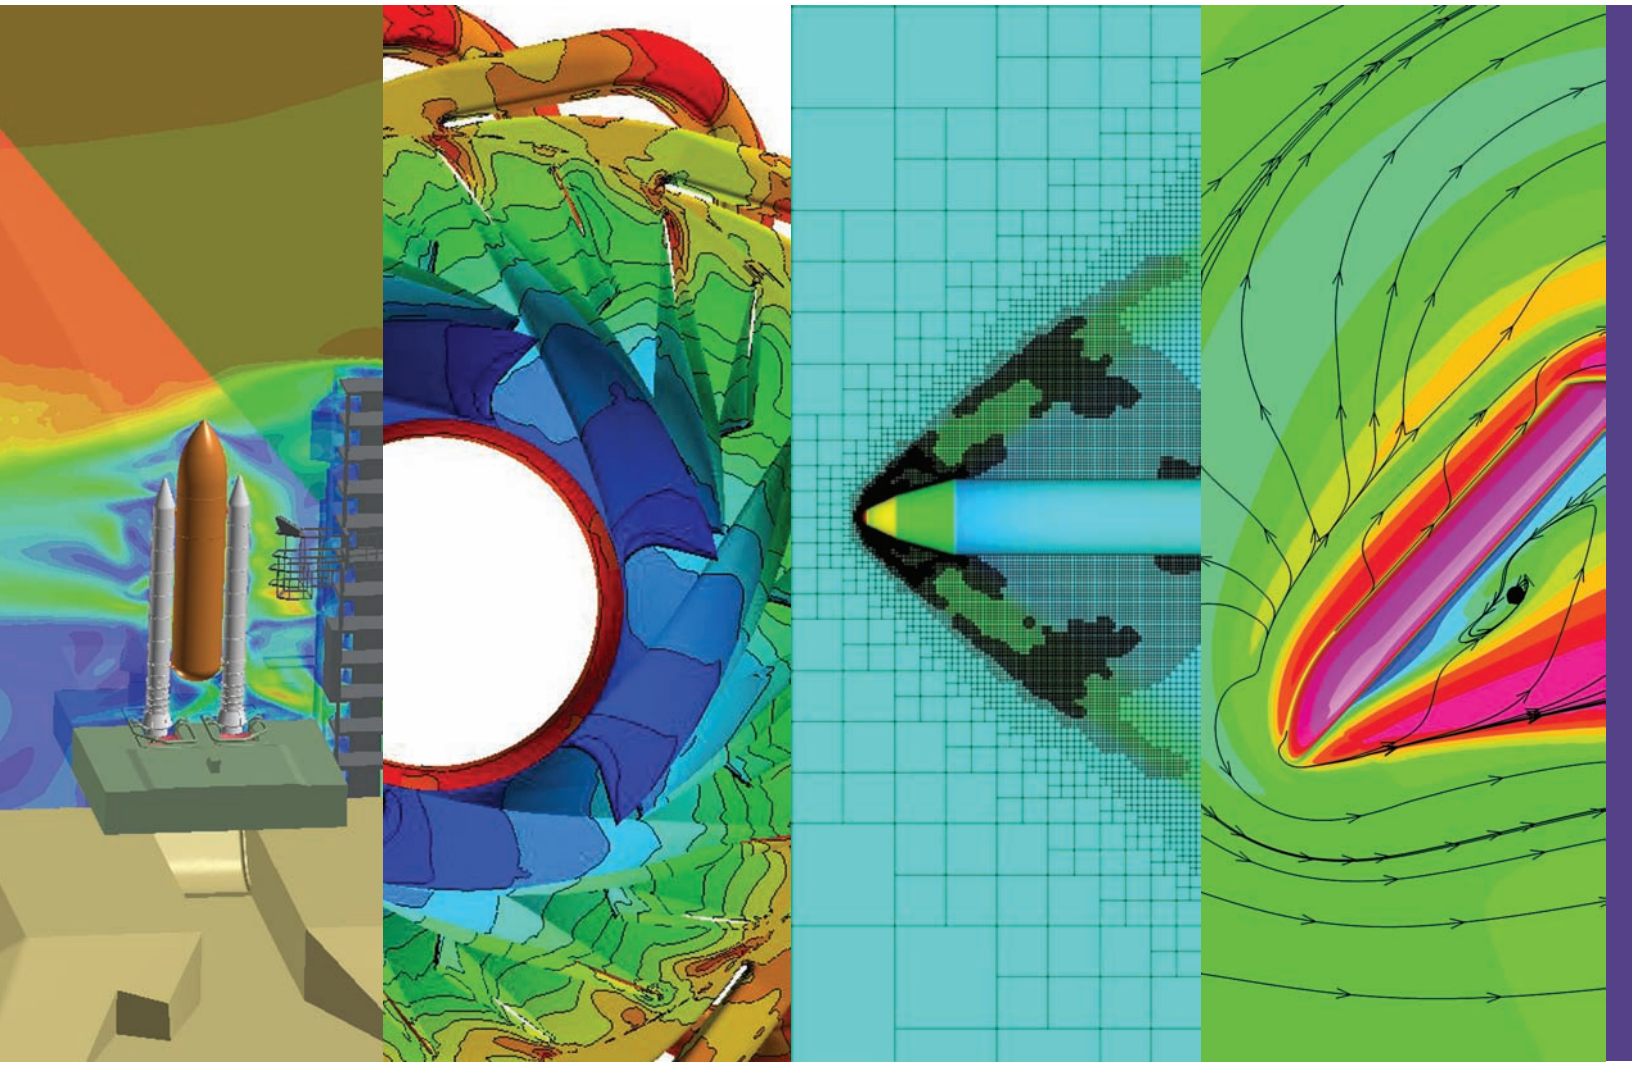


Figure 3: Fluid streamlines in a 48-Mm-wide simulation. In the left box, fluid moving up to the solar surface (background) originates from a small area in the upflow cells at the bottom (foreground). In the right box, fluid moving down from the surface (background) collects in the downflow boundaries of the large, supergranulation cells at the bottom (foreground).



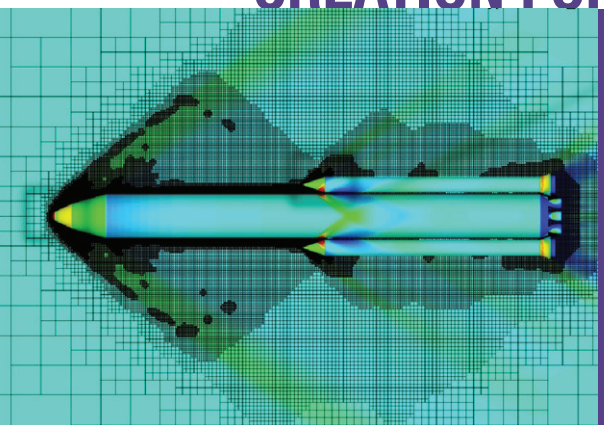
MISSION DIRECTORATE

SPACE OPERATIONS

The Space Operations Mission Directorate provides the Agency with leadership and management of NASA space operations related to human exploration in and beyond low-Earth orbit. Space Operations also directs low-level requirements development, policy, and programmatic oversight. Current exploration activities in low-Earth orbit are the Space Shuttle and International Space Station programs. The directorate is similarly responsible for Agency leadership and management of NASA space operations related to Launch Services, Space Transportation, and Space Communications in support of both human and robotic exploration programs.

WILLIAM H. GERSTENMAIER
Associate Administrator
http://www.nasa.gov/directorates/somd/home/index.html

AUTOMATED AERO-DATABASE CREATION FOR LAUNCH VEHICLES



SPACE OPERATIONS MISSION DIRECTORATE

MICHAEL J. AFTOSMIS
NASA Ames Research Center
(650) 604-4499
Michael.Aftosmis@nasa.gov
http://people.nas.nasa.gov/~aftosmis/cart3d/

◀ Detail of Figure 1.

Project Description: Computational fluid dynamics (CFD) is routinely used to analyze aerospace vehicle performance using high-fidelity methods at only a handful of critical design points. Recent progress in automated methods for numerical simulation of vehicle aerodynamics, however, now enables complete simulations to be performed with little to no human intervention. This progress coincides with the unprecedented increase in NASA's high-performance computing capacity afforded by the Agency's Pleiades and Columbia superclusters. These two concurrent developments place NASA in a unique position to explore the viability of developing fully automated aerodynamic performance databases for new aerospace vehicles. Such databases describe the aerodynamic performance of new vehicles throughout their entire flight envelope, and enable vehicles to be "flown" through the databases, to quantify vehicle performance for any candidate mission profile. These parametric studies vary both flight conditions as well as all permissible control surface deflections.

The aim of this project is to develop and deploy a prototype system for rapid generation of aerodynamic performance databases, and to use the system on real-world problems faced by NASA's Exploration Systems and Space Operations mission directorates. The project focuses on developing efficient, automated tools for generating aerodynamic data; and on establishing accurate, formal, and quantitative estimates of the uncertainty of this data. Such error estimates can be fed back into the simulations to improve their accuracy to customer-driven tolerances, yielding aerodynamic performance databases of certifiable quality.

These aerodynamic databases give a much broader picture of aerodynamic performance and can be used both in the preliminary design of new vehicles and for providing insight into the detailed aerodynamics of existing vehicles. The database can be used with six-degree-of-freedom (6-DOF) trajectory simulations coupled to guidance and control (G&C) systems. "Flying" a design through an aerodynamic database in faster-than-real time enables performance estimates for prototypes to be rapidly evaluated, and supports G&C system design. The broad utility of such simulation-based aero-performance data makes the accuracy of this data of paramount importance—underscored by NASA's development of an Agency-wide Standard for Models and Simulation, which outlines formal standards for quantitative estimation of errors in modeling and simulation data.

Relevance of Work to NASA: NASA's Exploration Systems Mission Directorate is currently faced with the challenge of developing the Orion Crew Exploration Vehicle (CEV), the Ares I Crew Launch Vehicle (CLV), and related exploration architecture systems to provide the nation's access to space after the planned retirement of the Space Shuttle. During vehicle development, NASA continues to fly and modify its fleet of Space Shuttle Launch Vehicles operated by the Space Operations Mission Directorate. With these challenges ahead, our need for advanced simulation technology, such as a system for automated aero-database creation, has never been greater.

Computational Approach: A typical CFD aerodynamic database contains on the order of 10^4–10^6 simulations, depending on the problem requirements. We employ Cart3D, a massively parallel, automated aerodynamic simulation package that scales linearly to thousands of processors. Automated tools drive this package to manipulate the geometry for control surface deflections; produce surface and volume grids; and manage the massive numbers of simulations that populate such large datasets. Error estimation tools shadow each simulation and automatically refine the computational grid to control numerical error and provide quantitative error bounds on the calculations. Further automation is used to harvest meaningful results and capture them in a performance database.

Results: This prototype database generation system has been developed largely in direct support of both the Space Shuttle and Constellation programs. Figures 1 and 2 show snapshots extracted from error-controlled simulations in performance databases of two different vehicles. Figure 1 shows supersonic aerodynamics on an early design of an Ares-based heavy-lift launch vehicle. The Ares simulations are part of a design effort to look for potential gains in aerodynamic performance through modifications to the vehicle's shape. This figure includes an overlay of the computational mesh, which the error-estimation module has refined to minimize errors in force coefficient predictions. Figure 2 shows simulations from large parametric studies—spanning multiple Mach numbers, angles-of-attack, and thrust settings—used to study the control effectiveness of forward-mounted jets on the Orion Launch Abort Vehicle (LAV).

The error-estimation module developed under this project has driven mesh refinement to resolve the intricate shock structures that occur when the (abort control motor) plumes (see Figure 2) first emerge from the vehicle, and to modify the shape of the plumes as they evolve downstream, affecting force and moments on the overall vehicle. The simulation tools developed under this project help assess such mechanisms with greater fidelity than ever before.

Role of High-End Computing: This unprecedented simulation capability is contingent upon high-end computing. Each simulation in a database typically has 15–50 million degrees-of-freedom, and a performance databases usually consist of 5–100 thousand such simulations. Calculations supporting error analysis and flowfield sensitivity approximately double this load. Using Cart3D and the prototype automation system, a dedicated node of the Pleiades or Columbia supercomputers can perform 10,000–20,000 simulations per hour. Given the low cost per processor-hour on these systems, this is by far the cheapest method available to obtain high-quality aerodynamic data.

Future: As NASA continues to develop vehicles for the Constellation Program, simulation requirements continue to grow. Not only is the number of designs to analyze increasing, but error-estimation and validation due-diligence mandate that we re-create wind tunnel test databases, as well. As Constellation evolves, these tools offer NASA an unprecedented ability to "fly" candidate designs through various mission profiles to gain insight into vehicle performance and carry out trade studies.

Co-Investigators

- Marsha J. Berger, Courant Institute, New York University
- Marian Nemec, ELORET Corp.

Publications

[1] Aftosmis, M.J. and Rogers, S. E., "Effects of jet-interaction on pitch control of a launch abort vehicle," *AIAA Paper 2008-1281*, Jan. 2008.

[2] Biswas, R., Aftosmis, M.J., Kiris, C., and Shen, B.-W., "2007: Petascale Computing: Impact on Future NASA Missions," *Petascale Computing: Algorithms and Applications* (D. Bader, ed.), Chapman and Hall /CRC Press, pp. 29–46, Dec. 2007.

[3] Mavriplis, D.J., Aftosmis, M.J., and Berger, M.J., "High resolution aerospace applications using the NASA Columbia supercomputer," *International Journal of High Performance Computing Applications*, 21(1), pp. 106–126, Jan. 2007.

[4] Nemec, M., and Aftosmis, M.J., "Adjoint error-estimation and adaptive refinement for embedded-boundary Cartesian meshes," *AIAA Paper 2007-4187*, 18th AIAA CFD Conference, Miami, FL, Jun. 2007.

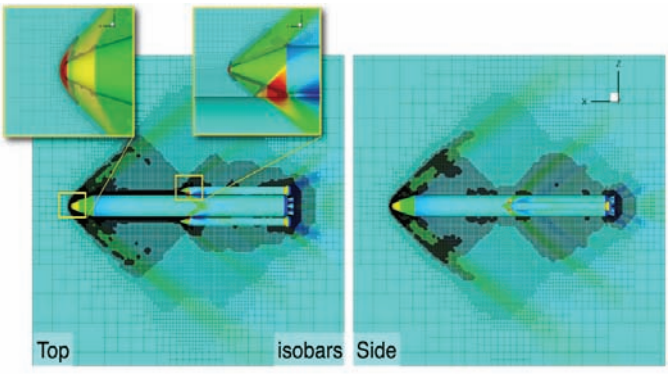


Figure 1: Top and side views of an Ares-based heavy-lift launch vehicle at supersonic conditions.

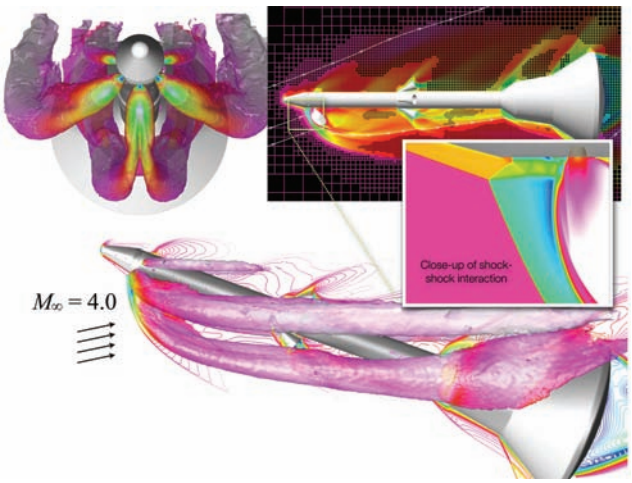
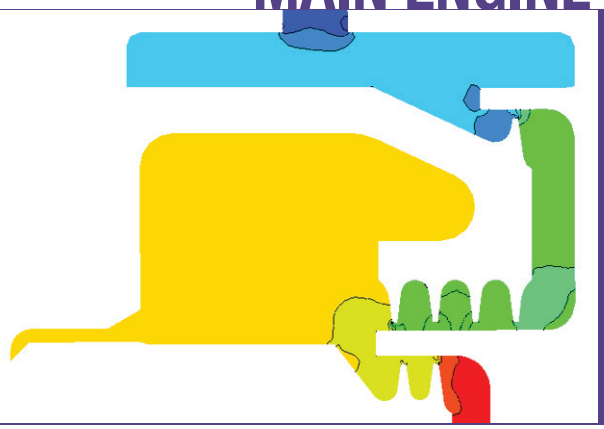


Figure 2: This set of figures shows front and side views of the Orion Launch Abort Vehicle at Mach 4, with the abort control motors firing. Cart3D was used with the adjoint-based error-estimation and mesh adaptation module to control numerical error in predicted aerodynamic forces and capture the many scales of the flow.

CFD ANALYSIS OF SHUTTLE MAIN ENGINE TURBOPUMP SEAL CRACKS



◀ Close-up of Figure 1.

SPACE OPERATIONS MISSION DIRECTORATE

KRIS MCDOUGAL
Marshall Space Flight Center
(256) 544-0659
Kristopher.McDougal@nasa.gov

DAN DORNEY
Marshall Space Flight Center
(256) 544-5200
Daniel.J.Dorney@nasa.gov

Project Description: The primary goal of this effort was to utilize computational fluid dynamics (CFD) models to characterize the flow drivers that were suspected as the root cause cracking mechanism on the Space Shuttle Main Engine (SSME) High Pressure Oxidizer Turbopump (HPOTP) Knife Edge Seals (KES) (Figure 1). The CFD analysis also provided detailed load characterization at discrete engine operating conditions.

The SSME HPOTP uses knife-edge seals to control leakage flows that pass between turbine rotating and stationary parts. The leakage flow on the backside of the HPOTP third-stage turbine is controlled by a series of seals. After hot-fire testing, two of these seals were found cracked or broken on several turbopumps. In this project, CFD models have been developed to provide insight into the root cause of these cracks.

Relevance of Work to NASA: This project provided analyses that contributed to identifying the root cause of the seal cracks. Furthermore, the analysis offered guidance for a new design, and was extended to provide dynamic fluid loads. Starting with shuttle mission STS-114, this new design has been incorporated in high-pressure oxidizer turbopumps that have successfully flown on the Space Shuttle.

Computational Approach: The code used in the numerical KES simulations was Loci-CHEM. Loci-CHEM (version 2) is a finite-volume flow solver (with combustion kinetics) for generalized grids. Developed at Mississippi State University, in part through NASA and National Science Foundation funding, Loci-CHEM uses high-resolution approximate Riemann solvers to simulate finite-rate, chemically reacting, viscous turbulent flows. Several turbulence models are available, including the Spalart-Allmaras one-equation model and a family of three-equation turbulence models. Loci-CHEM is comprised entirely of C and C++ coding and is supported on all popular Unix variants and compilers. Parallelism is supplied by the Loci15 framework, which exploits multi-threaded and Message Passing Interface libraries.

Results: Pre-test three-dimensional (3D) steady and unsteady CFD predictions showed good-to-excellent agreement with the ensemble of steady and unsteady pressure data measured in a cold flow KES air-flow test model. The two-dimensional (2D) unsteady CFD predictions were conservative such that the predicted unsteadiness was greater than that observed in the experiments. Therefore, we have concluded that the 2D unsteady CFD tooth-loading predictions used in the structural response analyses of the redesigned seals (Figure 2) are conservative with respect to the acoustic loadings in the engine, and that the redesigned KES are sufficient.

Role of High-End Computing: Two-dimensional CFD grids containing on the order of 3 million grid points, as well as 3D grids containing on the order of 13 million grid points were run to model the SSME HPOTP KES for investigation. A steady simulation was first run for each condition, followed by an unsteady simulation initiated from the steady solution. Computations of this magnitude require intensive levels of processing. The simulations were run on 20–160 processors of SGI Altix Linux clusters located at NASA's Marshall Space Flight Center and Ames Research Center. Without NASA high-end computing (HEC) resources, this analysis could not have supported the timely redesign of the KES.

Future: While the SSME HPOTP KES project has been closed out with no immediate need for additional modeling, many CFD analyses and models are currently benefitting from the vast computing capabilities of the HEC Program. Examples include CFD analyses of the cracked STS-126 Flow Control Valve to support STS-119's launch, as well as J-2X turbomachinery CFD analyses which is on the J-2X engine design critical path.

Publications

[1] Dorney, D., "Comparison of Pre-Test Predictions and Experimental Data for Knife-Edge Seals," Joint Army Navy NASA Air Force (JANNAF) Meeting, 54th Propulsion Meeting, May 2007.

[2] Dorney, D., "Investigation of the Flow in a Series of Knife-Edge Seals," JANNAF, 53rd Propulsion Meeting, Dec. 2005.

[3] Hawkins, L., "SSME High Pressure Oxidizer Turbopump Find 072 Inner," and, 011 Outer Turbine Outlet Duct Seal *INCREMENTAL VERIFICATION COMPLETE REPORT (FINAL) VRS-0648-2*, Contract NAS8-01140, March 2007.

[4] McDougal, K., "NASA MSFC Knife Edge Seal Air Rig Post Test Analysis *Report Test Number P2499*, March 2006.

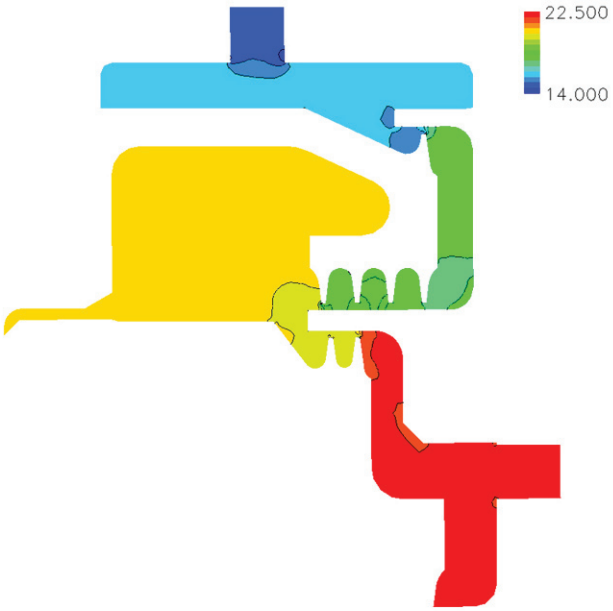


Figure 1: Cross-section of a representative design case of Space Shuttle Main Engine (SSME) Knife Edge Seal time-averaged static pressures at design conditions.

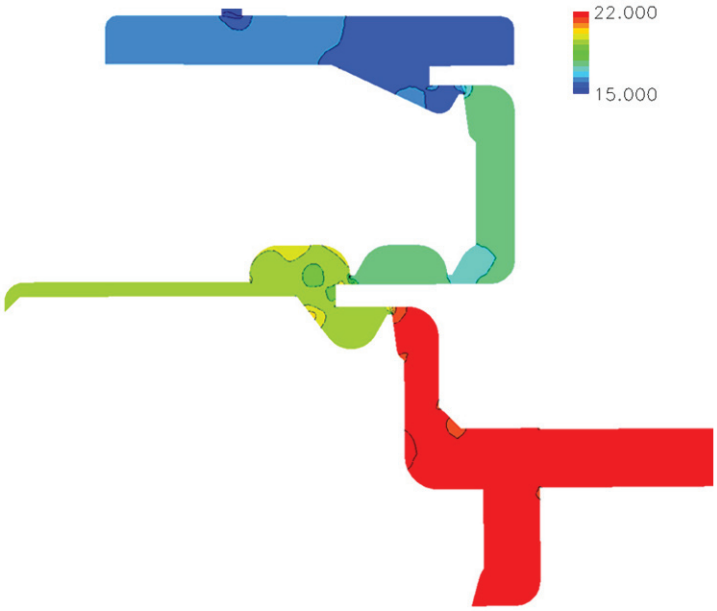
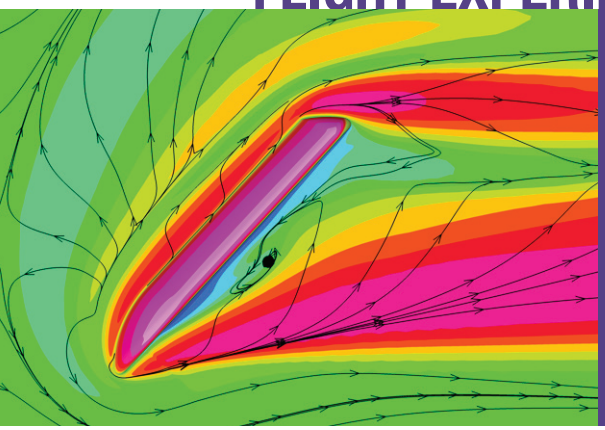


Figure 2: Cross-section of the SSME Knife Edge Redesign time-averaged static pressures at redesign conditions.

NUMERICAL ANALYSIS OF BOUNDARY LAYER TRANSITION FLIGHT EXPERIMENTS ON THE SPACE SHUTTLE



◀ Close-up of Figure 2.

SPACE OPERATIONS MISSION DIRECTORATE

CHUN TANG
NASA Ames Research Center
(650) 604-3480
Chun.Y.Tang@nasa.gov

WILLIAM WOOD
NASA Langley Research Center
(757) 864-8355
William.A.Wood@nasa.gov

Project Description: The objective of this work is to study the behavior of early flow transition due to the presence of discrete roughness elements. Since transition from laminar to turbulent flow can result in significant increases in heating on the thermal tiles of the Space Shuttle during reentry to Earth, a better understanding of this flow phenomenon is useful in estimating the impact that protuberances (such as protruding gap fillers and thermal blankets) may have on the safe return of the Orbiter. In addition, knowledge gained from the numerical simulations and experimental data will lead to better designs of thermal protection systems (TPS) for future spacecraft such as the Orion crew exploration vehicle and planetary probes.

A boundary layer transition (BLT) flight experiment, scheduled for launch in March 2009 onboard shuttle mission STS-119, consists of a 0.25-inch tall by 4-inch long protuberance placed on the windward surface of the Orbiter along with pressure sensors and calorimeters downstream from the protrusion. The goal of the experiment is to measure any elevated heating caused by flow transition. Numerical analysis of the BLT protrusion is used to predict peak heating rates on the surface of the protuberance and the surrounding acreage tiles during the Orbiter's reentry. These solutions are used to select the appropriate TPS material and shape of the protuberance.

Relevance of Work to NASA: This work is highly relevant to the safety of NASA's shuttle fleet and crews, since local protuberances such as protruding gap fillers may cause early transition from laminar to turbulent flow and result in higher heating on the Orbiter's heat tiles. The ability to accurately predict the aerothermal effects of these protrusions is important in real-time safety assessments during a shuttle mission. A better understanding of flow transition is also important in optimizing the design of future spacecraft for space exploration. Additional BLT flight experiments will include testing of Orion TPS materials in the turbulent zone created by the protuberance. Data and computer simulations from these experiments improves our understanding of the ablation process in a turbulent hypersonic flow environment.

Computational Approach: High-fidelity computational fluid dynamics (CFD) Navier-Stokes codes developed at NASA Ames and NASA Langley Research Centers are used to predict the aerothermal effects of a protuberance on the Space Shuttle. Instead of running the CFD solver on the entire vehicle, a "local" rapid analysis process (developed for real-time analysis of damages sustained by the shuttle during a mission) is used to evaluate the flow around the BLT protrusion. This approach uses existing solutions from the Return-to-Flight (RTF) database (for example, see Figure 1) to provide boundary conditions for a local simulation. The heating environment (as seen in Figure 2) is then fed into a thermal analysis model to predict the internal and bondline temperatures of the protrusion and surrounding heat tiles. This information is used to assess possible risks associated with the BLT flight experiment.

Results:

- The original BLT design plus seven geometry revisions were analyzed at flight conditions using two CFD codes, Data-Parallel Line Relaxation (DPLR, developed at NASA Ames) and LAURA (developed at NASA Langley).
- Arc jet experiments and corresponding arc jet CFD simulations were completed at NASA Johnson Space Center.
- BLT Flight Experiment #1 launched aboard STS-119 (March 2009).
- Additional laminar and turbulent CFD simulations on taller protuberances were computed in preparation for future BLT flight experiments.

Role of High-End Computing: NASA's Columbia supercomputer provided the necessary resources to quickly compute the protuberance simulations. For example, each protrusion calculation took approximately 8 hours of wall-time using 70

processors on Columbia. This rapid turnaround allowed us to run numerous simulations to optimize the protuberance shape for minimum surface heating. In addition, personnel at the NASA Advanced Supercomputing facility are optimizing the DPLR code for maximum performance. These improvements should result in better turnaround time for obtaining high-fidelity CFD solutions, which will be useful for real-time risk assessment during a shuttle mission.

Future: Data gathered from the first flight experiment will be used to assess accuracy of the numerical predictions. It is expected that post-flight CFD simulations will be needed to match the actual freestream conditions of the experiment. Two additional BLT flight experiments are planned for future Orbiter missions. These experiments will involve taller protuberances to induce earlier flow transition and the testing of different TPS materials to investigate the behavior of catalycity in a turbulent flow environment. Thus, more CFD simulations at flight and arc jet conditions will be needed to support these future flight experiments. Once again, supercomputers will play a critical role in providing the resources needed to quickly compute and post-process these calculations.

Co-Investigators

- Kerry Trumble, NASA Ames Research Center
- Victor Lessard, Genex Systems
- Charles Campbell, NASA Johnson Space Center

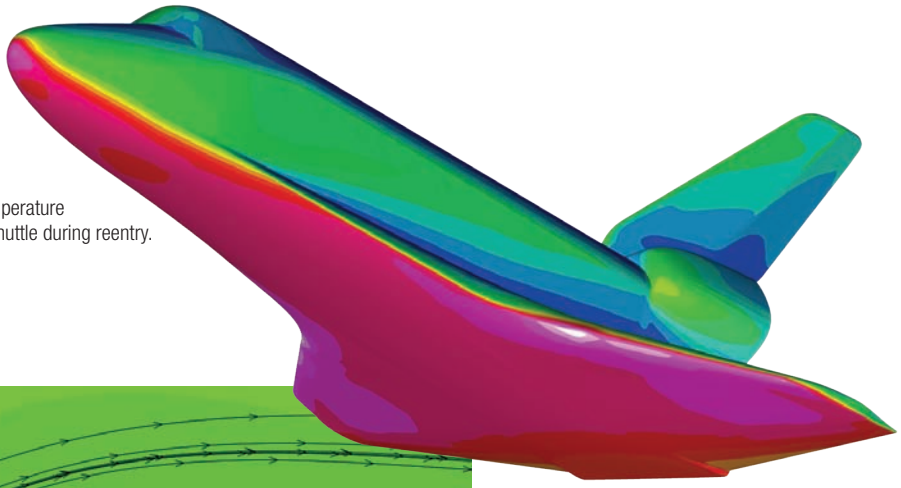


Figure 1: Computed temperature contours on the Space Shuttle during reentry.

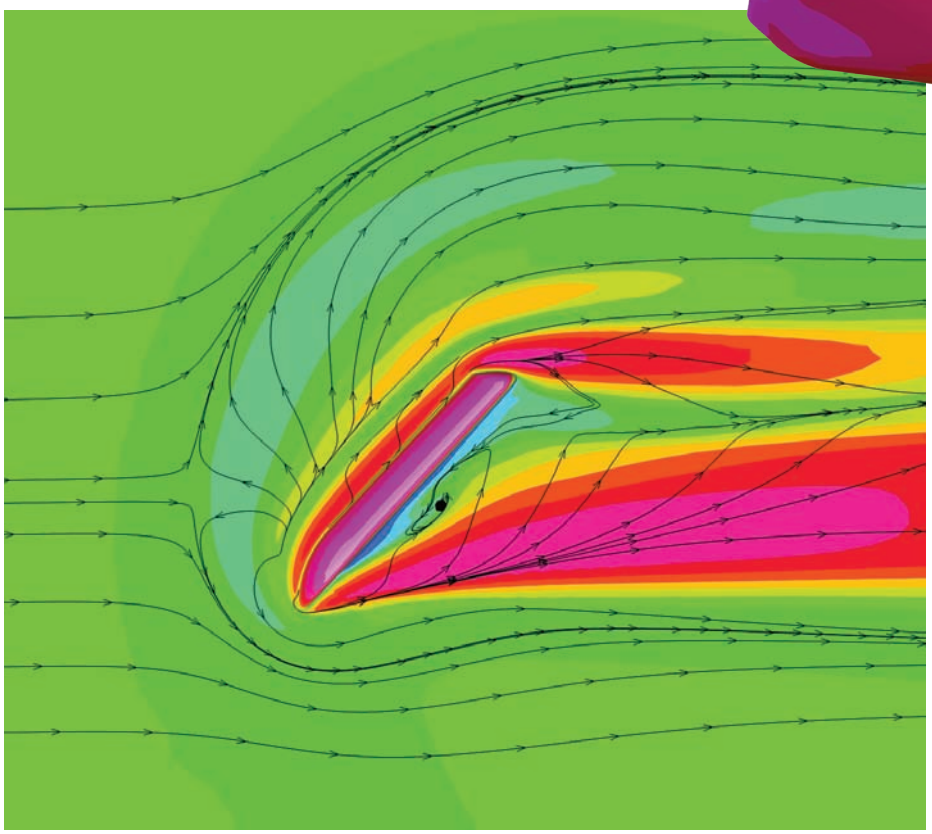


Figure 2: Computed streamlines and temperature contours on the boundary layer transition protuberance during shuttle reentry.

SPACE SHUTTLE ASCENT AERODYNAMICS AND DEBRIS TRANSPORT ANALYSES



SPACE OPERATIONS MISSION DIRECTORATE

REYNALDO J. GOMEZ III
NASA Johnson Space Center
(281) 483-6108
Reynaldo.J.Gomez@nasa.gov

◀ **Figure 1:** Surface pressure and shockwaves for the Space Shuttle at Mach 1.4.

Project Description: The Space Shuttle Program relies on detailed computational simulations to assess the wide range of aerodynamic and aerothermodynamic environments encountered during a shuttle mission. From redesigns to eliminate potential ice and foam debris, to inflight assessments used to clear the shuttle for entry and landing, the Columbia supercomputer plays a key role in NASA's Space Operations Mission Directorate. Probabilistic debris risk assessments use ascent computational fluid dynamics (CFD) flowfields along with probabilistic aerodynamic models to prioritize redesigns of the external tank and inspections on the Space Shuttle Orbiter.

This project supports Space Shuttle aerodynamic, aerothermodynamic, and debris transport assessments through the use of high-fidelity, unsteady CFD simulations. Redesigns of the shuttle's external tank have progressed so rapidly and involved such small details of the geometry that wind tunnel testing would be difficult to schedule and instrument. Computational models, anchored to previous wind tunnel tests, have been used to produce the various types of data required to assess these redesigns. Additionally, debris transport assessments use these flowfields to predict debris trajectories, and are key inputs for the probabilistic risk assessments used to prioritize external tank redesigns.

Relevance of Work to NASA:

- The last major outer mold line change to the Space Shuttle external tank was the replacement of four of the LO_2 feedline brackets with titanium (Ti), covered by a minimal amount of thermal protection system foam. The lower thermal conductivity of the Ti brackets enabled the shuttle program to reduce the amount of potential foam and ice debris that could be shed by the brackets, and increased clearance between the brackets and external tank. Detailed simulations of this configuration change, run on the Columbia supercomputer, were a key part of the redesign assessments and provided insight into flowfield details that would be difficult or impossible to extract from a wind tunnel test. Surface grids and a representative flowfield around the launch vehicle with these modifications are shown in Figures 1 and 2.
- Maximum allowable iceball maps, created using in-house debris transport tools, have also been assessed using Columbia. The primary result of these analyses is a map of the allowable ice-ball size that can be shed from the external tank without exceeding the impact-damage threshold for a given component. The result of this study has been incorporated as a launch commit criteria in the document "*Space Shuttle Ice/Debris Inspection Criteria.*"

Computational Approach: NASA's overset CFD flow solver, OVERFLOW, is the primary tool used to produce ascent airloads on the Space Shuttle. Typically, 64–128 processors are used in parallel to produce a set of solutions along an ascent or entry trajectory. NASA's Cart3D moving body Euler code is also used to produce probabilistic aerodynamic models for a wide range of debris shapes.

Results:

- *Iceball Allowables Debris Transport Analyses:* The results of these analyses are used during the prelaunch countdown to ensure that the shuttle is not launched with a potentially dangerous ice accumulation on the external tank [1].
- *OVERFLOW Simulations:* Improvements include the extension of transonic, slender-body performance to high-speed, complex configurations, and significant robustness enhancements for lower-speed flows with high-speed jets. The increased robustness has paid significant dividends for analyzing the shuttle and next-generation Orion space vehicle designs. The improved CFD code has also been applied to Space Shuttle Orbiter calculations. In a case at Mach 4, the modified code achieved a two-orders-of-magnitude reduction in processor time to converge the solution on Columbia [2].

- *External Tank ET-128 Ascent Airloads, Aerothermal, and Debris Transport Simulations:* This project used the OVERFLOW code to simulate the flowfield around the last major outer mold line change to the external tank. This design change replaced the aluminum LO_2 feedline brackets with Ti, decreasing the potential for ice and foam debris from these components, and increasing safety of the vehicle and crew. The detailed pressure distributions extracted from these solutions were used to assess changes to the aerodynamic and aerothermodynamic loads on the modified components. Further simulations are currently being used to simulate ascent debris environments and address inflight issues [3].

Role of High-End Computing: Producing an assessment database for a geometry as complex as the Space Shuttle Launch Vehicle is a computationally intensive task. Each unsteady ascent simulation of the new design required approximately 7,600 processor-hours and 46 gigabytes of disk space for the time varying results. Approximately 625,000 processor-hours were used to produce 80 solutions of the new design at a range of ascent flight conditions. The complex nature of the flowfield around this complex geometry requires the use of unsteady simulations significantly more computationally expensive than simpler steady-state solutions. Each of these time-accurate simulations is time averaged to create ascent flowfield environments.

Future: Currently, NASA is planning to retire the shuttle in 2010. Modifications to the external tanks and procedural changes have slowed production, and we have been using CFD tools to assess a number of reproducibility enhancements intended to reduce the time required to manufacture external tanks without affecting launch safety.

Co-Investigators

- Stuart Rogers, Scott Murman, and Edward Tejnil, all of NASA Ames Research Center
- Darby Vicker, NASA Johnson Space Center

Publications

[1] Rogers, S. E., "Hemisphere Ice Ball Debris-Transport Analysis," NAS Technical Report," NAS-07-004, June 2007.

[2] Murman, S.M., "Dynamic Stability Analysis using CFD Methods," in Experimental Determination of Dynamic Stability Parameters, von Karman Institute for Fluid Dynamics, February 2008.

[3] Tejnil, E., Rogers, S. E., and Gomez, R. J., "OVERFLOW CFD Analysis of SSLV with ET-128 Design Changes,"June 2008.

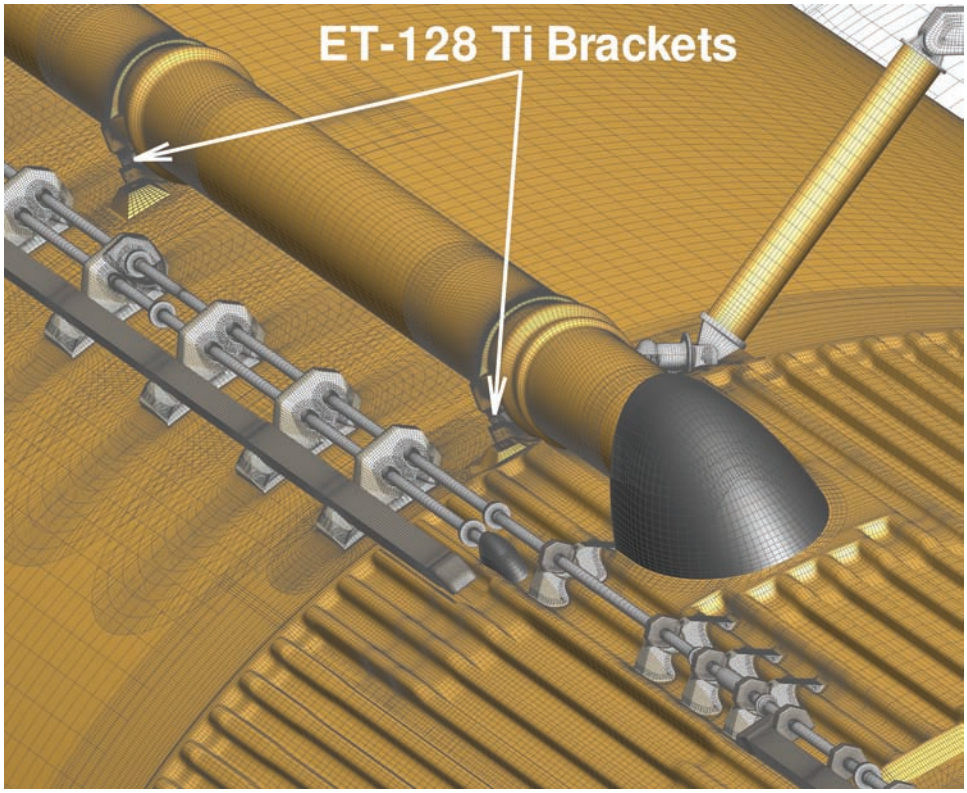
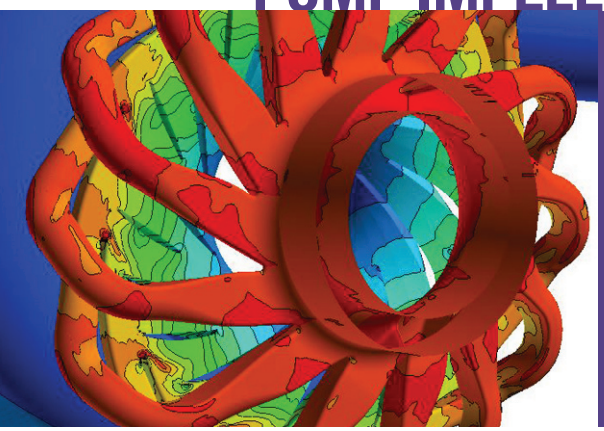


Figure 2: Surface grids for the ET-128 external tank redesign.

SSME HIGH PRESSURE FUEL PUMP IMPELLER CRACK INVESTIGATION



SPACE OPERATIONS MISSION DIRECTORATE

PRESTON SCHMAUCH
NASA Marshall Space Flight Center
(256) 544-1218
Preston.B.Schmauch@nasa.gov

◀ Close-up of Figure 1.

Project Description: In 2007, a liquid dye penetrant test revealed a crack on a Space Shuttle Main Engine (SSME) High Pressure Fuel Pump (HPFP) impeller mid-blade. In order to more fully understand the events leading to the impeller crack initiation, as well as the dynamic environment that caused the crack to grow, various unsteady computational fluid dynamics (CFD) simulations were performed at SSME power levels experienced by the cracked impeller. These simulations calculated unsteady blade loading and associated frequencies that were used to determine expected part life. The simulations were also run with multi-phase flow to determine at what power levels cavitation was likely to exist near the crack site. Additionally, running conditions from prior water flow testing on the HPFP were compared with unsteady CFD water testing simulation results to validate and correlate the computational model with the engine running conditions. Numerous simulations that varied the inlet guide vane trailing-edge thickness were also run to determine whether vortex shedding of the inlet guide vane could have been a contributing factor.

Relevance of Work to NASA: With safety as one of NASA's primary goals, it is imperative to fully analyze damage to any critical parts of the SSME. These simulations provided key information to the impeller crack investigation team so they could make educated decisions regarding the cause of the crack and the next course of action. Data used to support those decisions included the number of engine starts since crack initiation, mean and alternating stresses near the crack location, and expected margin of safety at that location. Those decisions led to the new requirements placed on the impeller related to part life and intervals between inspections. This process helps to ensure the safety of our astronauts, NASA employees and contractors, and helps to improve future designs.

Computational Approach: Phantom, a NASA code that has been anchored and validated for supersonic turbines, was used for the unsteady CFD simulations. Phantom uses three-dimensional, unsteady Navier-Stokes equations as the governing equations. The Baldwin-Lomax turbulence model is used for turbulence closure. For this simulation, H grid topology with moving grids was employed to simulate blade motion. In order to accurately resolve flow features for the impeller it was necessary to simulate the inlet scroll, inlet guide vane, impeller, and crossover diffuser (see Figures 1 and 2). Merlin, a code similar to Phantom but with two-phase flow capabilities, was used to simulate the unsteady CFD with cavitation effects.

Results: Simulation results have been used to help understand the original robustness of the impeller design and factors involved in crack initiation and growth. These simulations helped the investigation team to arrive at three conclusions based on the unsteady CFD solutions. First, all unsteady rotor-stator flow drivers appear to affect all flow channels, and a rotating stall or a planar disturbance driver will be experienced by each blade. Second, results indicated that cavitation would be present on the mid-length blade. Third, further investigation into the vibration phenomena seen in the CFD simulation and from experimental data is recommended to gain a better understanding of the loading impacts. It was determined that the unsteady rotor-stator loads seen in the CFD were the source for the cracking. The reasons that only one blade on one pump has experienced cracking—especially since other flight engines have higher cumulative run times than the affected engine—are still unknown. All simulations pertaining to this investigation have been completed.

Role of High-End Computing: Quick simulation turnaround time in this study allowed multiple suspected causes of the crack to be investigated. Suspected crack initiators were simulated, analyzed, and determined to be either a contributor or non-factor. HEC resources significantly reduced overall simulation time and allowed the team to focus more on pertinent areas of investigation. The processor power, processor time, and high-speed file transfers were important factors in getting the work done in a timely manner. The simulations run for

this investigation totaled over 300,000 processor-hours on the Columbia supercomputer.

Future: Future work utilizing HEC resources includes investigations into vibratory loading effects on the impeller blade. Due to the high mass flow variation from passage to passage of the impeller, there is a corresponding velocity and incidence change at the leading edge of the blade. The resulting loading effect could be a factor in crack initiation. Other areas of interest involve adding secondary two-phase leakage paths to the current simulations for a better model of the SSME HPFP flow environments.

Co-Investigators

- Daniel Dorney, NASA Marshall Space Flight Center

Publications

[1] Marco B., Zabo R., Dorney D., and Zoladz T., "The Effect of Acoustic Disturbances on the Operations of the Space Shuttle Main Engine Fuel Flowmeter," 43rd AIAA Joint Propulsion Conference, *AIAA 2007-5534*, July 2007.

[2] Dorney D., Griffin L., Marcu B., and Williams M., "Unsteady Flow Interactions between the LH2 Feedline and SSME LPFP Inducer," 42nd AIAA Joint Propulsion Conference, *AIAA 2006-5073*, July 2006.

[3] Marcu B., Tran K., Dorney D., and Schmauch P., "Turbine Design and Analysis for the J-2X Engine Turbopumps," 44th AIAA Joint Propulsion Conference, *AIAA 2008-4660*, July 2008.

[4] Dorney D., Griffin L., and Schmauch P.,"Unsteady Flow Simulations for the J-2X Turbopumps," 54th JANNAF Propulsion Meeting, May 2007.

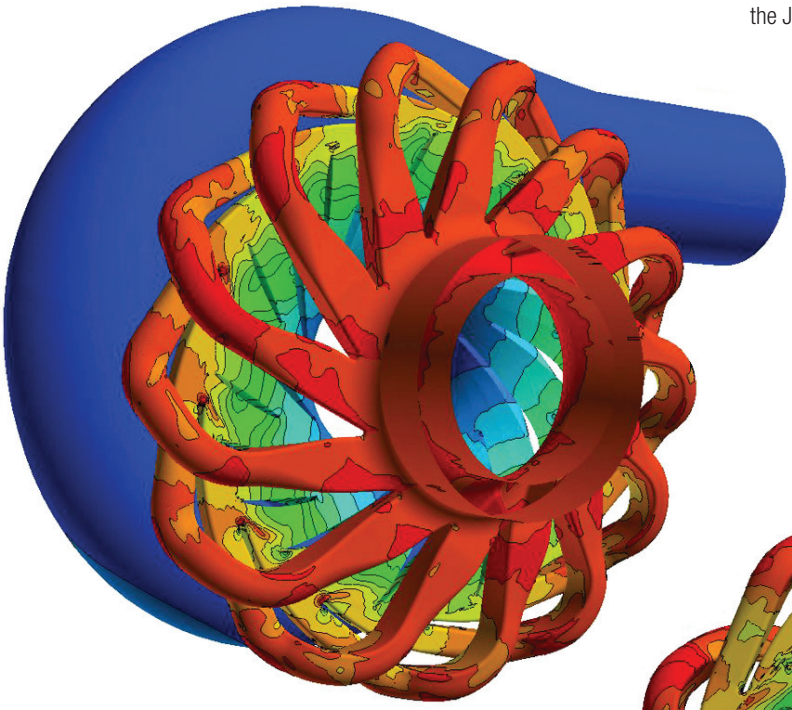


Figure 1: Space Shuttle Main Engine (SSME) High Pressure Fuel Pump static pressure. View is looking upstream from the first crossover diffuser exit.

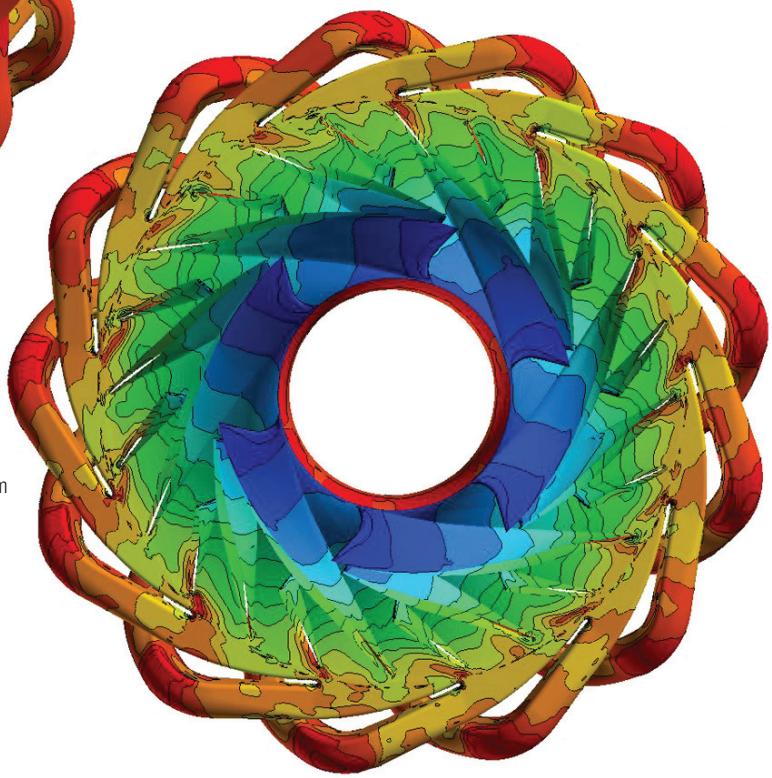
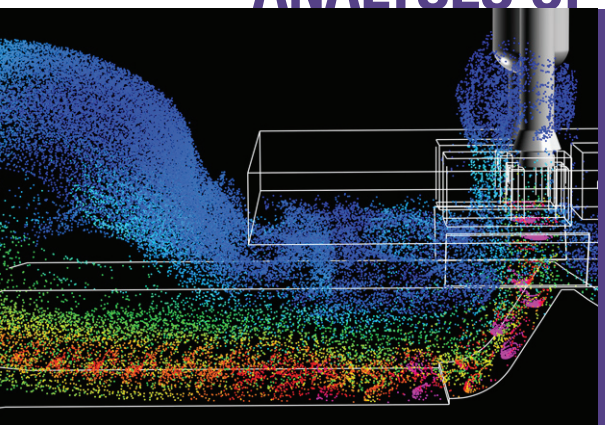


Figure 2: SSME High Pressure Fuel Pump static pressure. View is looking downstream from the impeller inlet.

TIME-ACCURATE COMPUTATIONAL ANALYSES OF THE LAUNCH PAD FLAME TRENCH



SPACE OPERATIONS MISSION DIRECTORATE

CETIN KIRIS
NASA Ames Research Center
(650) 604-4485
Cetin.C.Kiris@nasa.gov

◀ Detail of Figure 1.

Project Description: During ignition of rocket propulsion systems such as that of the Space Shuttle, interaction of the exhaust plume with the flame trench below the launch pad produces a series of strong pressure waves that travel back through the inlet of the trench, where they may affect the stability of the launch vehicle during takeoff. For this project, we have performed time-accurate computational fluid dynamics (CFD) simulations to characterize ignition overpressure (IOP) phenomena during liftoff of new and existing launch vehicles. In preparation for the launch of the Ares I-X flight test vehicle in 2009, we developed a computational model to investigate the feasibility of using the shuttle mobile launch platform (MLP) for next-generation Ares launch vehicles.

To validate the geometric model and CFD procedure, we first simulated IOP waves generated during ignition of shuttle mission STS-1, which did not use a water suppression system, and compared results with measured data (Figure 1). The model was then modified to study the effects of different MLP configurations on IOP waves for the Ares I-X vehicle, represented by a single, modified solid rocket booster (SRB) positioned over the left flame trench hole. Ares I-X ignition simulations were performed for three alternate configurations of the existing MLP (which has two exhaust holes with deflectors to direct flow into the flame trench): both holes open with both deflectors present (Figure 2); both holes open with right deflector removed; and right hole closed with right deflector removed.

The flame trench computational model has also subsequently been used to support repair efforts at NASA Kennedy Space Center after the launch of STS-124 damaged a large section of the trench wall. Time-accurate flow data and pressure values were provided for 21 points on the flame trench wall and at six inches out from each point.

Relevance of Work to NASA: This work has provided valuable insight into launch environments for NASA's current and future space vehicles. CFD analyses provide an efficient, cost-effective means of reassessing ground operations infrastructures to determine and plan modifications required for the Ares launch vehicles. This study has demonstrated the value of these efforts by showing that the MLP exhaust holes will not need to be modified for the Ares I-X vehicle, saving millions of dollars in engineering and construction expenses. Additionally, computed flame trench pressure values helped determine structural requirements for the critical repair of flame trench damage that occurred during the launch of STS-124.

Computational Approach: Using a grid generation script based on the Chimera Grid Tools library, structured viscous overset grid systems were built to model the different launch site configurations, including the flame trench, MLP, plume deflectors, support structures, launch vehicles, and surrounding terrain. The overset grid and scripting approach facilitated modifications to the grids to model different options for the launch vehicle, MLP, and deflectors. The Space Shuttle grid system contained 129 zones and 92 million grid points, and grid systems for the single SRB with various MLP options contained 92 to 120 grids, and 73 to 87 million grid points.

The NASA-developed CFD code OVERFLOW, an implicit, structured, overset, Reynolds-Averaged Navier-Stokes solver, was used to simulate the exhaust plume interaction with the flame trench. To capture the correct plume behavior with a single species model, the SRB chamber conditions were modified to obtain correct nozzle exit thrust, temperature, and Mach number conditions critical to simulation accuracy. Impulsive start conditions were used for STS-1 and Ares I-X simulations to assess the magnitude of the IOP waves, while a transient ramp-up of SRB chamber conditions was applied for the later wall damage simulations.

Results: For the initial validation study, good correlation was found between the predicted CFD results and the recorded flight data for STS-1. Peak pressure levels agreed closely, and qualitative behavior also compared well once the acoustic noise was removed from the flight data. Results of the Ares I-X MLP exhaust hole studies showed that, in each configuration, the peak IOP waves are reflected from the SRB's own exhaust hole, and that blocking the right hole and removing the deflector has no effect in reducing the IOP on the test vehicle. Comparison of the predicted pressure peaks with the STS-1 data showed that similar IOP behavior is predicted for all three of the alternate configurations, and that none of them will generate significantly larger peak pressures than those experienced by STS-1. With additional damping provided by the launch pad's current water suppression system, which was excluded from the STS-1 model, the shuttle MLP should provide an adequate launch platform for the Ares I-X vehicle without requiring costly modifications to the right-side hole.

Results of the STS-124 flame trench wall damage study show that the initial pressure wave reaches each of the 21 examined points according to the distance from the main deflector. Peak pressure magnitudes were found to be highest near the flame trench floor and lower near the top of the trench. These computed pressures were compared to STS-4 flight data and showed good correlation, although some differences were observed due to the presence of the water suppression system used for the STS-4 launch.

Role of High-End Computing: The HEC resources supported by the NASA Advanced Supercomputing facility were essential to performing the large-scale, time-accurate CFD simulations needed for these studies. Using 128 processors on the Columbia supercomputer, computation of two seconds of ignition conditions for the STS-1 and Ares I-X configurations required several weeks of runtime. For the time-critical wall damage support study, unsteady flame trench simulations were carried out to 1.15 seconds in four and a half days using 504 processors on Columbia.

Future: Significantly more modeling and simulation of ground operations infrastructures and environments will be required to prepare for the Ares I Crew Launch Vehicle and the Ares V Cargo Launch Vehicle in the coming years. In addition to the launch pad configurations, the Vehicle Assembly Building at NASA Kennedy is also being modeled to evaluate its potential future use for the Ares vehicles.

Co-Investigators

- Jeffrey Housman, Daniel Guy Schauerhamer, Marshall Gusman, all of ELORET Corp.
- William Chan, Dochan Kwak, both of NASA Ames Research Center

Publications

[1] Kiris, C., Chan, W., Kwak, D., Housman, J., "Time-Accurate Computational Analysis of the Flame Trench," Fifth International Conference on Computational Fluid Dynamics, Seoul, Korea, July 2008.

[2] Kiris, C., Schauerhamer, D.G., Housman, J., Gusman, M., Chan, W., Kwak, D., "Time-Accurate Computational Analysis of the Flame Trench Applications," 21st International Conference on Parallel Computational Fluid Dynamics, May 2009.

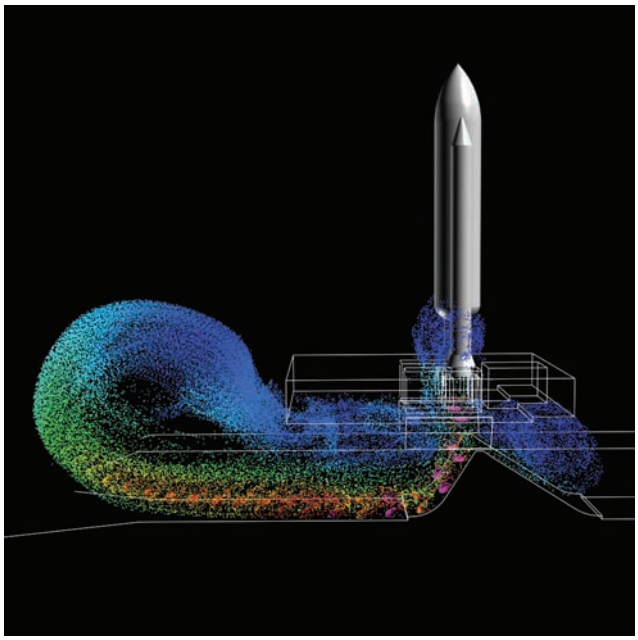


Figure 1: Visualization of shuttle flame trench CFD simulation showing instantaneous particle traces colored by Mach number.

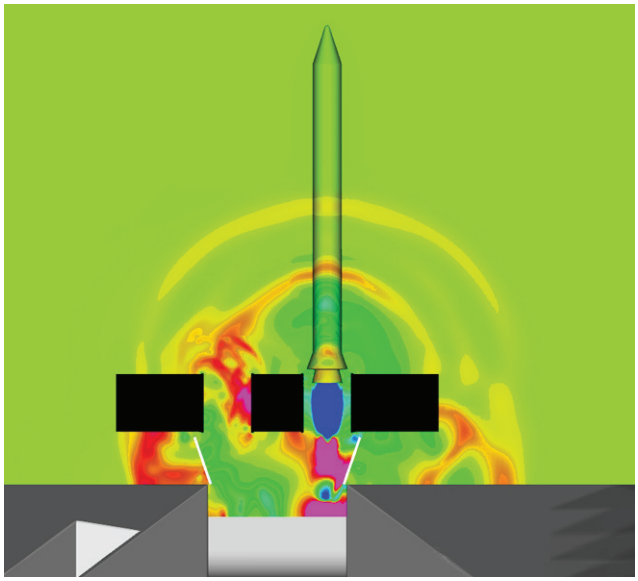
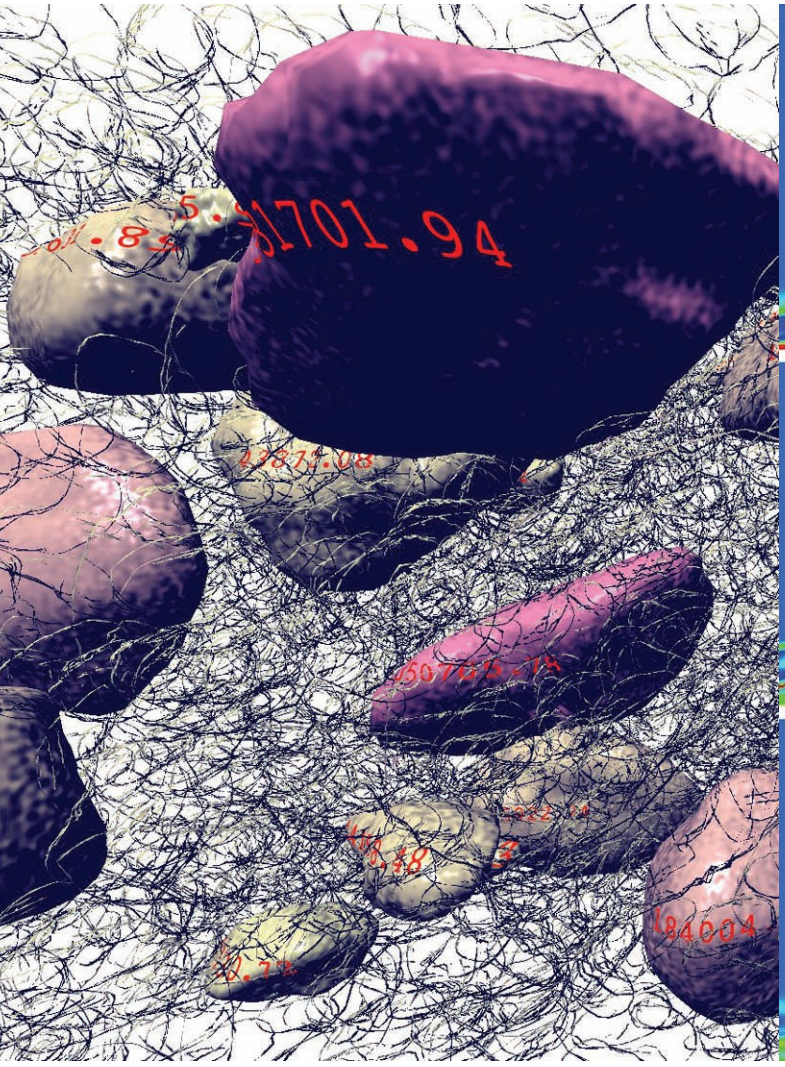
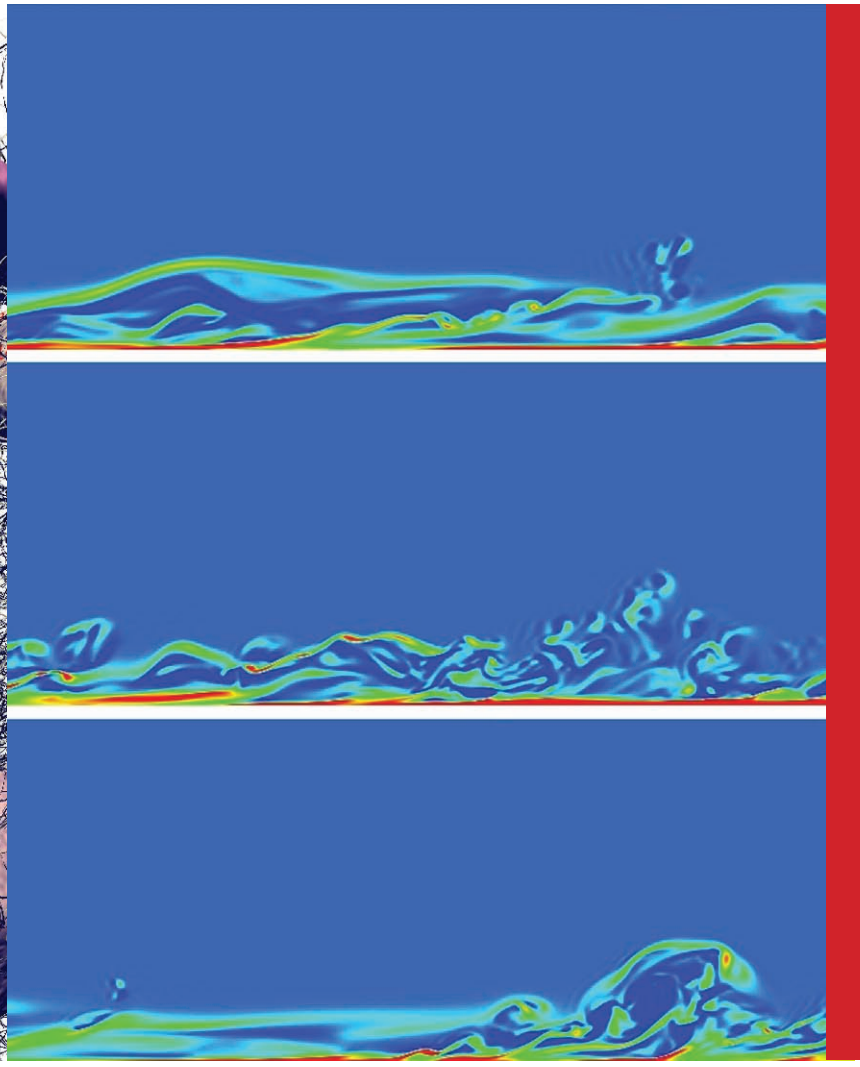


Figure 2: Instantaneous ignition overpressure waves for a single solid rocket booster ignition on the shuttle mobile launch platform configured with right hole open and right deflector present.

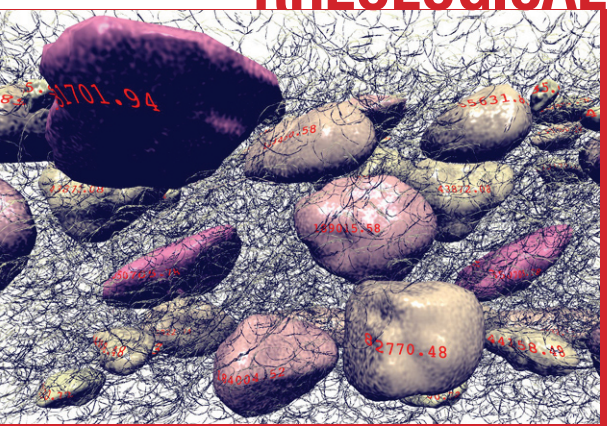




NATIONAL LEADERSHIP COMPUTING SYSTEM

NASA's National Leadership Computing System (NLCS) initiative provides access to the Agency's largest supercomputers to selected non-NASA researchers doing cutting-edge, computationally intensive science and engineering of national interest. NLCS demonstrates the Agency's support for important national priorities, and its commitment to continued U.S. leadership in high-end scientific and technical computing and computational modeling. By inviting industry and academia participation, NASA can help advance U.S. technology and education, and assist U.S. competitiveness. In return for NLCS awards, much of the resulting knowledge will be made publicly available.

MODELING THE RHEOLOGICAL PROPERTIES OF CONCRETE



NATIONAL LEADERSHIP COMPUTING SYSTEM

WILLIAM GEORGE
National Institute of Standards and Technology
(301) 975-4943
wgeorge@nist.gov

◀ **Figure 1:** Snapshot of a flowing suspension of rocks in a mortar matrix (concrete). The mortar matrix is not visible in this image. Graphical processing unit programming was used within the application to place the stress value on the surface of each rock. Rocks with a stress value below a specified threshold are shown only as silhouettes.

Project Description: Modeling and predicting physical properties of concrete remains a great challenge, as it involves phenomena taking place over many length and time scales. For example, concrete is typically composed of cement (micrometer scale with interactions at the nanometer scale), sand (millimeter scale), and aggregates (centimeter scale). Even at each representative length scale, there can be considerable variation and other factors to consider, such as the addition of chemical admixtures in the cement paste, sand diameter variation of over a factor of 100, or shape variation of aggregates (rounded or crushed).

The rheological properties of complex fluids, specifically viscosity and yield stress (the applied stress when flow begins), play an important role in a wide variety of technological and environmental processes. Understanding how a fluid yields under stress is a subject of great interest and remains an outstanding problem in the field of fluid physics.

In this project, we study the flow of dense suspensions composed of rigid bodies having a wide range of sizes and shapes, under a variety of flow conditions (shear and around obstacles). Our goal is to advance the general understanding of the flow properties of these complex fluids. While this research is applicable to many areas, the focus of our study is to understand, and ultimately to predict, the rheological properties of cement-based materials.

Relevance of Work to NASA: This research focuses on advancing our understanding of fundamental mechanisms that control the flow of suspensions. It is important to understand how these materials start to flow and to tune their physiochemical properties in order to control their movement as the material is applied. This knowledge will be useful in the development of materials and techniques for the building and repair of structures under various conditions such as low- and high-gravity environments. Further, suspensions are utilized in a wide variety of technological processes and, because this study is largely parametric in nature, results are transferable to other fields.

Computational Approach: Recently, a new computational fluid dynamics (CFD) method called Dissipative Particle Dynamics (DPD) has been developed, which holds promise for modeling complex fluids. Indeed, DPD may have some advantages over other CFD methods because DPD can naturally accommodate many boundary conditions and does not require meshing (or re-meshing) of the computational domain. On the surface, DPD looks very much like a molecular dynamics algorithm where particles, subject to inter-atomic forces, move according to Newton's laws. However, the particles in DPD are not atomistic, but a mesoscopic-scale (between the macroscopic- and atomic- scale ranges) representation of the fluid. We have adopted this approach due to its potential for modeling rigid bodies with a wide variety of shapes. Custom visualization software is used to explore the results of the simulations in detail. Representative snapshots of this software in use are shown in Figures 1 and 2.

Results: Results from recent simulations have advanced our understanding of suspensions in two particular areas: yield stress and the formation of stress chains in dense suspensions. Our studies have helped provide fundamental insights into the physical mechanisms that control yield stress. These simulations indicate that current theoretical models are inadequate for describing the yielding behavior of dense colloidal suspensions [1, 3]. We have also successfully modeled and visualized the formation of stress chains in dense suspensions. These stress chains are accompanied by giant stress fluctuations and have been seen in recent physical experiments [2].

Role of High-End Computing: Our ability to simulate dense suspensions has been enabled by access to NASA's computational resources. Simulations of the size and duration needed

to reveal the characteristics of these dense suspensions are only possible on large parallel supercomputers such as Columbia. We have typically utilized 500–1,000 processors for each simulation—some of which still required over 100 hours of continuous compute time to complete. Results from these simulations, some exceeding 20 gigabytes, have been transferred from the NASA Advanced Supercomputing (NAS) facility in California to the National Institute of Standards and Technologies (NIST) in Maryland via a high-speed Internet2 connection facilitated by NAS network administrators in collaboration with their NIST counterparts.

Future: We are in the process of completing an analysis of the localized stress transmission in dense colloidal suspensions, which is expected to provide greater insight into the fundamental mechanisms that control yield stress. The next goal is to study the dependence of yield stress and viscosity on the strength of inter-particle interactions. We will also explore the role of aggregate shape and size distribution on rheological properties of suspensions, comparing results with currently available experimental data. We are interested in studying the structural rearrangements that occur in very dense colloidal suspensions. These systems serve as idealized models of molecular glasses, which are difficult to visualize because the length and time scales are too small. Finally, we have recently modified our code to model suspensions with a non-Newtonian fluid matrix in order to expand the types of suspensions that we can study. Such suspensions are very realistic and representative of many building materials, such as mortars and concrete. We intend to apply for additional compute time on NASA's computational resources for advancing this research and to expand the research to include the study of the chemical properties of cement as it cures.

Co-Investigators

- Nicos S. Martys, Edward J. Garboczi, John G. Hagedorn, Judith E. Terrill, all of the National Institute of Standards and Technology

Publications

[1] Martys, N., Lootens, D., George, W., Satterfield, S., and Hebraud, P., "Spatial-Temporal Correlations at the Onset of Flow in Concentrated Suspensions," The XVth International Congress On Rheology, Ed. A. Co, L.G. Leal, R. H. Colby, A.J. Giacomin, *AIP Conf Proc. Vol. 1027*, Monterey, CA, pp. 207–209, 2008.

[2] Lootens, D., Martys, N., George, W., Satterfield, S., and Hebraud, P., "Stress Chains Formation under Shear of Concentrated Suspension," The XVth International Congress On Rheology, Ed. A. Co, L. G. Leal, R. H. Colby, A. J. Giacomin, *AIP Conf Proc. Vol. 1027*, Monterey, CA, pp. 677–679, 2008.

[3] Martys, N., Lootens, D., George, W., and Hebraud, P., "Spatial-Temporal Correlations of Colloids in Start-up Flows," (submitted to *Physical Review Letters*).

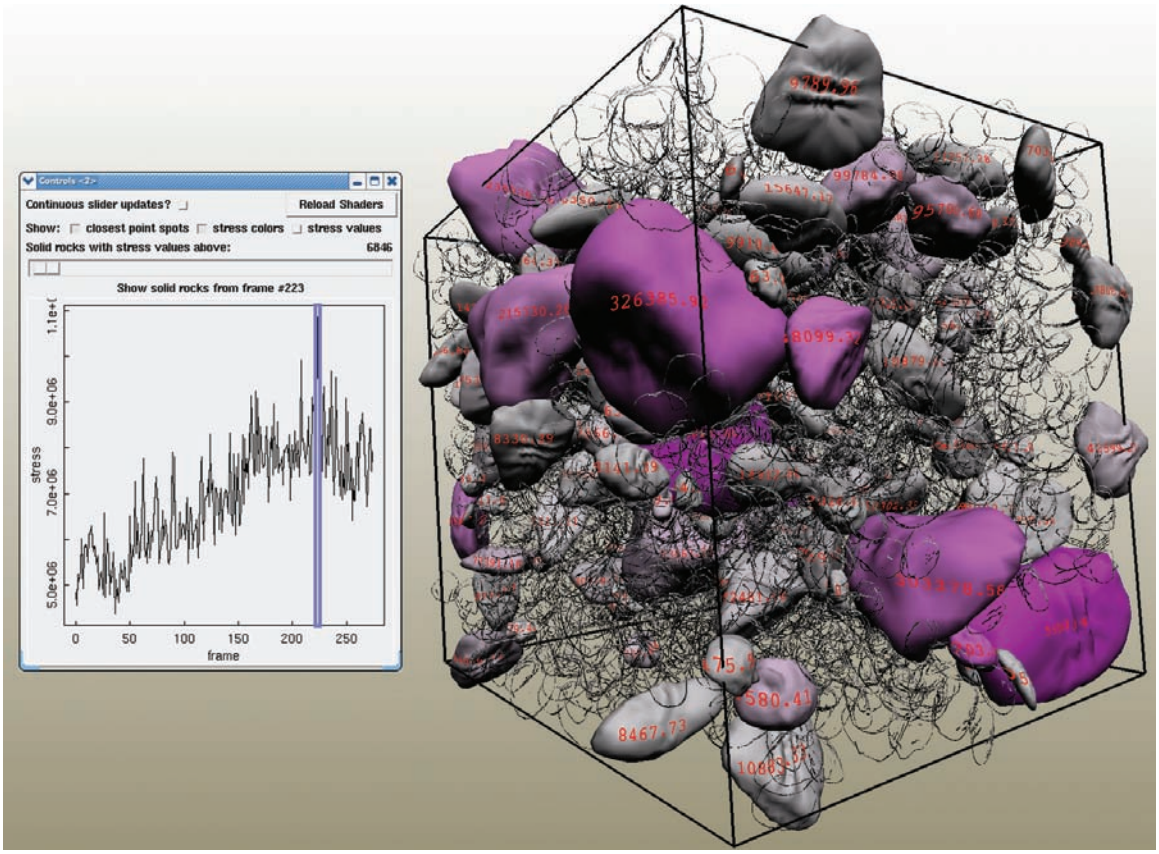
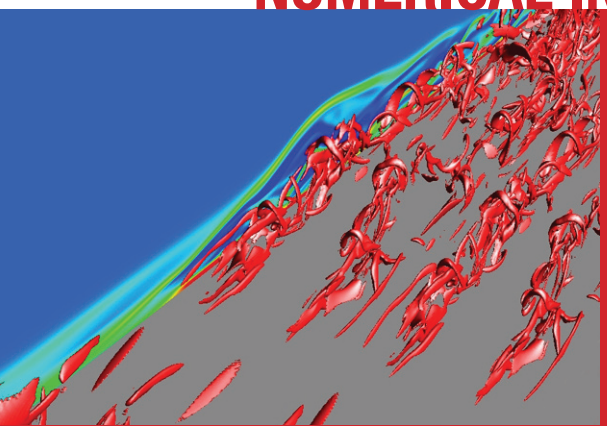


Figure 2: Graph of the average system stress, that is, applied force per unit area over time, is shown to the left of a visualization of the dense suspension. Interactive controls allow for the detailed exploration of the simulation results. This custom exploration software is supported on both desktop systems and an immersive visualization environment.

TRANSITION IN HIGH-SPEED BOUNDARY LAYERS: NUMERICAL INVESTIGATIONS USING DNS



NATIONAL LEADERSHIP COMPUTING SYSTEM

FRANK HUSMEIER
Greenberg Traurig, LLP
(520) 621-4424
husmeierf@gtlaw.com
http://cfdlab.web.arizona.edu

◀ Detail of Figure 1.

Project Description: Flight vehicles, cruising faster than the speed of sound, experience high heating rates at their surface. In a high-density environment, these aero-thermal loads are even further increased due to the transition process of a laminar high-speed boundary layer to turbulence. In the past, engineers used a relatively conservative approach for the design of thermal protection systems (TPS). In this approach, the turbulent boundary layer was assumed to be present over the entire TPS. For the design of future high-speed vehicles, however, design margins will need to be reduced to enhance payload capabilities. To reach this goal, the transition process of a high-speed boundary layer must be better understood to provide the design community with accurate physical models for prediction of the transition point.

The transition process of a laminar high-speed boundary layer to turbulence is studied using direct numerical simulations (DNS). In this approach, the Navier-Stokes equations governing the flow are solved directly using very efficient and accurate numerical methods that scale well with large parallel computers such as NASA's Columbia supercomputer. Understanding the transition physics is mandatory for finding advanced methods to: control the transition process; provide a numerical database for an in-depth validation of engineering predictions (for example, turbulence models) of the various transition stages; and provide the design community with accurate physical models for prediction of the transition point.

Relevance of Work to NASA: This project is partially funded by NASA's Aeronautics Research Mission Directorate and is an important part of NASA's quest to understand the transition process of high-speed boundary layers. Reliable transition predictions are critically important for the design and safe operation of any high-speed, advanced flight vehicle. Examples include the Hyper-X Program as well as the Constellation Program with both the Orion Crew Exploration Vehicle and the Ares Crew Launch Vehicle. The large increase in wall heat transfer due to transition to turbulence in high-speed boundary layers is one of the major difficulties in the design and operation of high-speed flight vehicles. In addition to these aero-thermal loads, transition to turbulence has a large impact on the aerodynamic performance and flight characteristics of these vehicles as skin friction is drastically increased. This fact is especially important for flight vehicles that have only a lifting body, such as the X-43. Accurate transition prediction can, in some cases, increase its range or, in others, result in a reduction of size and weight with comparable performance.

Computational Approach: The Navier-Stokes equations are integrated in time by employing a fourth-order Runge-Kutta scheme. The spatial derivatives are discretized using fourth-order split-finite differences in the streamwise and wall-normal direction. Assuming that the spanwise direction is periodic, integration variables are transformed into spectral space using Fast Fourier Transforms. Furthermore, the spanwise discretization is pseudo-spectral; that is, all nonlinear terms in the governing equations are computed in physical space [1, 2]. The use of an explicit time integration scheme and a standard finite difference method for discretization of the spatial derivatives allows the simulation code to scale very well on large parallel computers such as Columbia.

Results: The project was divided in two subprojects. One part focused on the transition process of a hypersonic cone and flat-plate boundary layer at Mach 8, whereas the second subproject concentrated on the transition process of a supersonic flat-plate boundary layer at Mach 3. For the first subproject, the following milestones were achieved:

- Two strong transition mechanisms (so-called "oblique breakdown" and "oblique fundamental resonance") were identified, which may lead to a fully developed turbulent boundary layer at Mach 8.
- For these mechanisms, highly resolved direct numerical simulations were performed—the first simulations of this kind at hypersonic speeds.

- The influence of the nose tip radius of a cone on the transition process for a hypersonic boundary layer was studied intensively.

For the second project, key results are summarized as follows:

- The entire path of laminar-turbulent transition for a supersonic boundary layer via oblique breakdown was studied.
- It was demonstrated that oblique breakdown indeed breaks down to turbulence since a turbulent stage was reached near the outflow.
- This is one of the first highly resolved direct numerical simulations that captures the entire transition path for a supersonic, flat-plate boundary layer.

Role of High-End Computing: Numerical study of high-speed boundary-layer transition is crucial for understanding the transition process and underlying physical mechanisms since experimental investigations in high-speed wind tunnels are very difficult to perform. Only a few experimental studies that provide high-quality datasets exist. Hence, accurate simulations of the entire transition process—from the linear stage to breakdown to turbulence, including the turbulent regime—are mandatory for understanding high-speed boundary-layer transition. These simulations pose a formidable computational challenge, requiring grid sizes approaching half a billion grid points and more. Even with the current high-end supercomputing systems (for example, Columbia), such calculations still take several weeks to complete on 256 processors.

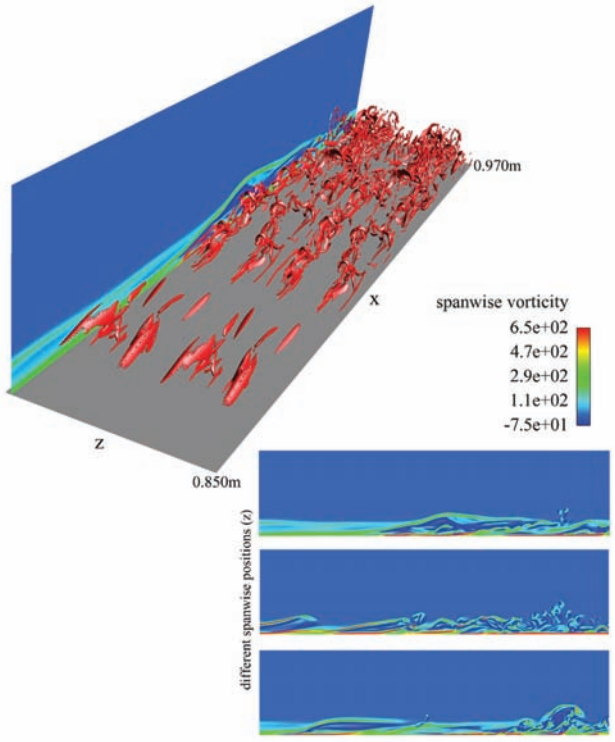


Figure 1: Transition to turbulence initiated by "oblique breakdown" for a supersonic flow over a flat plate. Ma=3, T=103.6 K, Re=2.181E6 1/m, f=6.36 kHz, forcing location x=0.452 m. Isosurfaces of Q=30000 [3].

Future: We are planning to extend our computational efforts by using our recently developed high-order accurate Navier-Stokes code to investigate boundary-layer transition for a cone at Mach 3.5 (in close cooperation with experimental efforts at NASA Langley Research Center). This code will enable us to submit our future simulations on NASA's Pleiades supercomputer since it utilizes Message Passing Interface for its parallelization routines. Furthermore, we will be able to efficiently allocate a larger number of processors for these simulations (beyond 1,000). To meet this challenge, we will work closely with NASA's high-end computing experts to optimize our new Navier-Stokes code.

Co-Investigators

- Christian Mayer, Andreas Laible, The University of Arizona

Publications

[1] Husmeier, F. and Fasel, H., "Numerical Investigations of Hypersonic Boundary Layer Transition over Circular Cones," *AIAA Paper 2007-3843*, 2007.

[2] Laible, A., Mayer, C., and Fasel, H., "Numerical Investigation of Supersonic Transition for a Circular Cone at Mach 3.5," *AIAA 2008-4397,* 2008.

[3] Mayer, C.S.J., von Terzi, D.A., and Fasel, H., "DNS of Complete Transition to Turbulence Via Oblique Breakdown at Mach 3," *AIAA Paper 2008-4398*, 2008.

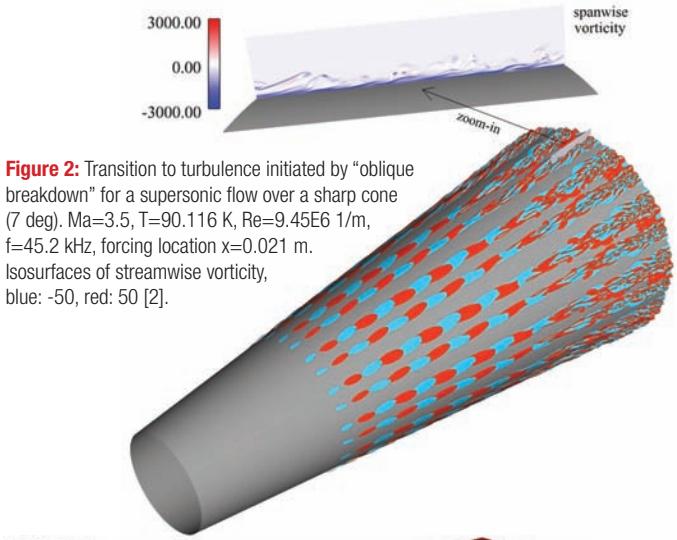


Figure 2: Transition to turbulence initiated by "oblique breakdown" for a supersonic flow over a sharp cone (7 deg). Ma=3.5, T=90.116 K, Re=9.45E6 1/m, f=45.2 kHz, forcing location x=0.021 m. Isosurfaces of streamwise vorticity, blue: -50, red: 50 [2].

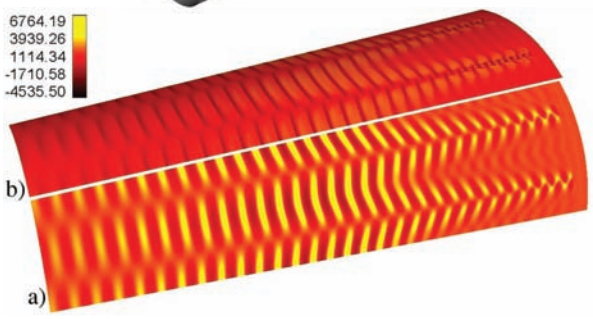
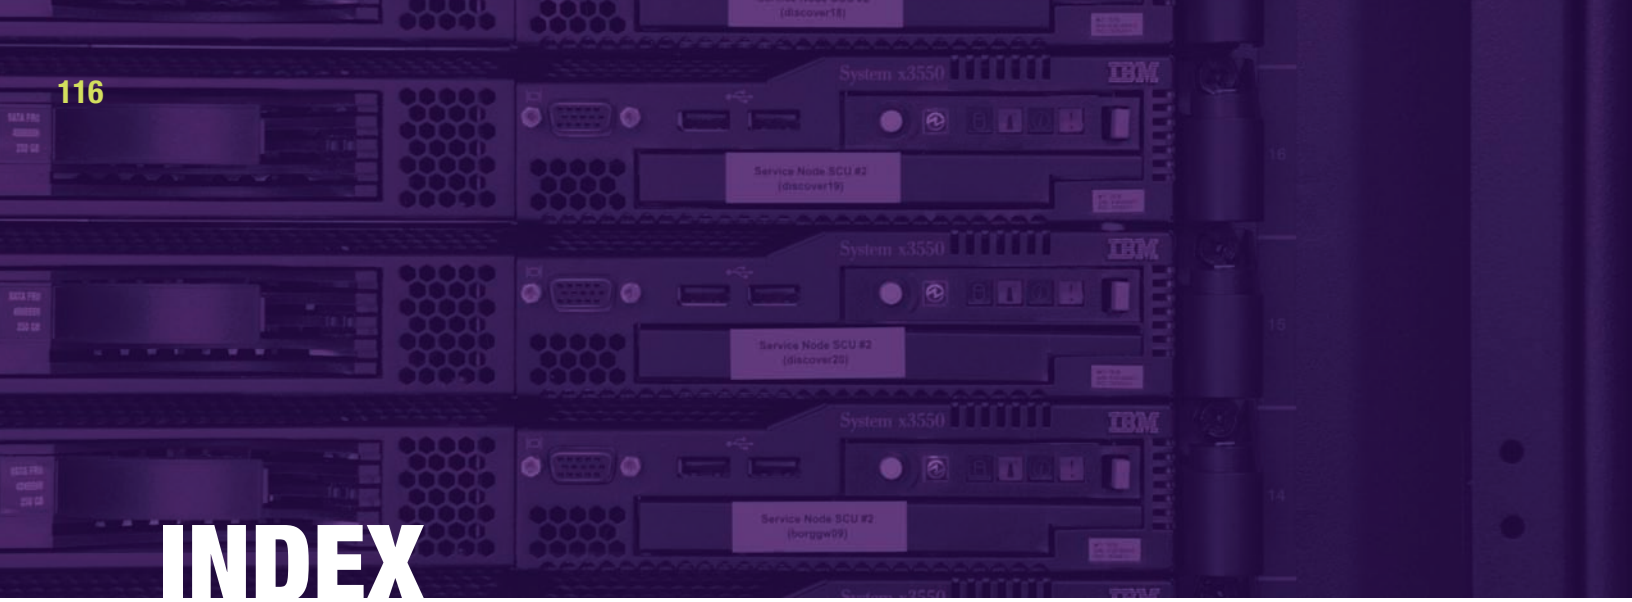


Figure 3: Transition to turbulence initiated by "oblique breakdown" for a hypersonic flow over a sharp cone (7 deg). Ma=8, T=53.35 K, Re=3333333, f=88 kHz. Wall-normal heat flux at a) the wall, and b) boundary layer edge [1].



INDEX

Abdol-Hamid, Khaled S., 42
Aftosmis, Michael J., 98
Ahmad, Jasim, 22

Bacmeister, Julio, 72
Baker, John, 92
Bakhle, Milind, 24
Balaji, Venkatramani, 68
Balakumar, Ponnampalam, 28
Barnhardt, Mike, 50
Bartels, Robert, 48
Bauer, Steven, 44
Benson, David, 94
Berger, Marsha J., 98
Bian, Huisheng, 76
Biedron, Robert, 48
Bosilovich, Michael, 72
Brown, James, 60
Buning, Pieter, 62

Campbell, Charles, 102
Carbon, Duane, 30
Cen, Renyue, 70
Centrella, Joan, 92
Chaban, Galina, 30
Chaderjian, Neal M., 22
Chan, William, 54, 108
Chang, I-Chung, 22
Chen, Junye, 72
Chin, Mian, 76, 78
Chwalowski, Pawel, 48, 62
Colarco, Peter, 74
Compton, William, 44

da Silva, Arlindo, 74

Dalton, Jeff, 24
Davis, Philip, 56
Deere, Karen, 42, 44
Diehl, Thomas, 76
Dorney, Daniel, 40, 52, 100, 106
Dyakonov, Artem, 60

Edquist, Karl, 60
Envia, Edmane, 32

Fok, Mei-Ching H., 90

Garboczi, Edward J., 112
Gelaro, Ron, 72
George, William, 112
Georgobiani, Dali, 94
Gökçen, Tahir, 50
Gomez, Reynaldo J., 104
Gross, Richard, 84
Guruswamy, Guru, 22
Gusman, Marshall, 108

Hagedorn, John G., 112
Hah, Chunill, 26
Haidvogel, Dale, 84
Hathaway, Michael, 24
Hawke-Wong, Veronica, 54
Hirsch, Adam I., 82
Holst, Terry, 22
Housman, Jeffrey, 108
Hui-Chun Liu, Emily, 72
Hunter, Craig, 44
Huo, Winifred, 30
Husmeier, Frank, 114

Jaffe, Richard, 30
Jiang, Weiyuan, 86
Jin, Haoqiang H., 26

Kandil, Osama A., 28
Kao, David, 22
Kara, Kursat, 28
Kim, Gi-Kong, 72, 74
Kim, Kyu-Myong, 78
Kiris, Cetin, 108
Kirk, Benjamin, 46
Klopfer, Goetz H., 54
Klypin, Anatoly, 66
Krist, Steven, 44
Kuang, Weijia, 86
Kucsera, Tom, 76
Kwak, Dochan, 108

Laible, Andreas, 114
Lau, William K., 78, 88
Lessard, Victor, 102
Lin, John C., 82
Linker, Jon A., 80
Lionello, Roberto, 80
Liu, Yen, 30
Lytle, John, 24

McDougal, Kris, 100

Magee, Todd, 36
Magin, Thierry, 30
Martys, Nicos S., 112
Massey, Steve, 48
Matsui, Toshihisa, 78
Mayer, Christian, 114
Michalak, Anna M., 82
Mikic, Zoran, 80
Mineck, Ray, 48
Mizuno, Yosuke, 64
Mok, Yung, 80
Moore, Thomas E., 90
Murman, Scott, 104

Nehrkorn, Thomas, 82
Nemec, Marian, 98
Nielsen, Eric, 74
Niemiec, Jacek, 64
Nishikawa, Ken-ichi, 64
Nordlund, Åke, 94

Olejniczak, Joseph, 46
Onufer, Jeffrey, 54

Pandya, Shishir, 54
Pao, Jenn Louh, 34
Pao, S. Paul, 42
Prabhu, Dinesh, 30, 50, 60
Primack, Joel, 66
Pulliam, Thomas, 22

Reale, Oreste, 88
Reddy, T. S., 24
Rienecker, Michelle, 72, 74
Riley, Pete, 80
Rogers, Stuart, 46, 104
Romander, Ethan, 22
Rosati, Anthony, 68

Schaffenberger, Werner, 94
Schauerhamer, Daniel Guy, 108
Schmauch, Preston, 40, 52, 106
Schubert, Siegfried, 72
Schwenke, David, 30
Shi, Jainn J. (Roger), 78
Shum, C.K., 84
Song, Y. Tony, 84
Srivistava, Rakesh, 24
Suarez, Max, 72, 74
Suresh, Ambady, 24
Stein, Robert, 94

Takacs, Larry, 74
Tan, Qian, 76
Tang, Chun, 60, 102
Tao, Wei-Kuo, 78
Tejnil, Edward, 104
Terrill, Judith E., 112
Thompson, Richard, 46
Titov, Viacheslav, 80
To, Wai Ming, 24
Todling, Ricardo, 72, 74
Trumble, Kerry, 102

Van Meter, James, 92
Van Norman, John W., 62
Van Zante, Dale, 32
Vicker, Darby, 104

Watson, Michael, 64
White, Todd, 50, 60
Wood, William, 102
Wray, Alan, 30
Wright, Michael J., 50, 60

Yu, Hongbin, 76

Zhang, Shaoqing, 68
Zlotnicki, Victor, 84

For further information about NASA's High-End Computing Program, please contact:

Dr. Tsengdar Lee
Scientific Computing Portfolio Manager
NASA Science Mission Directorate
NASA Headquarters
(202) 358-0860
Tsengdar.J.Lee@nasa.gov

Dr. Rupak Biswas
High-End Computing Capability (HECC) Project Manager
NASA Advanced Supercomputing Division
NASA Ames Research Center
(650) 604-4411
Rupak.Biswas@nasa.gov

Dr. W. Phillip Webster
NASA Center for Computational Sciences (NCCS) Project Manager
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center
(301) 286-9535
Phil.Webster@nasa.gov



Many colleagues at the HEC Program facilities contributed long hours and dedicated effort to producing this 2007–2008 High-End Computing Report. We are especially grateful to the technical publications teams at NAS and NCCS, whose hard work and expertise made this report possible. Thanks also to the principal investigators who contributed science and engineering highlights, and to reviewers at both facilities and within all mission directorates.

Executive Editor: Bryan Biegel
Writers/Editors: Holly Amundson, Jarrett Cohen, Jill Dunbar, John Evans, Harper Pryor, Lorien Wheeler
Graphic Designer: Sue Kim
Editorial and Design Consultants: Nick Bonifas, John Hardman, Marco Librero
Photographers: Dominic Hart, Pat Izzo, Sue Kim, Marco Librero

Trade names or manufacturers' names are used in this report for identification only. This usage does not constitute an official endorsement, either expressed or implied, by the National Aeronautics and Space Administration.

Images shown on back cover (from left to right):

- IBM DS4700 Fibre Channel drives provide online disk storage for the Discover Linux cluster at the NASA Center for Computational Sciences.
- Close-up of a front panel from the NASA Advanced Supercomputing facility's 14,336-processor core Columbia supercomputer, which enables timely computational results for many of NASA's scientific and engineering projects.
- The 51,200-processor core Pleiades supercomputer at the NASA Advanced Supercomputing facility is connected via an overhead trellis system containing over 20 miles of InfiniBand network cables.
- Each cabinet of the 6,656-processor core Discover Linux cluster at the NASA Center for Computational Sciences tightly couples processors via a high-speed InfiniBand network.


004:29
RUNNING
004:25
RUNNING