NASA

NAS

NUMERICAL   AERODYNAMIC SIMULATION

NAS

NAS

NASA ADVANCED SUPERCOMPUTING

NAS Division 2008

25

Years of Innovation

LETTER FROM THE DIRECTOR OF AMES

Greetings,

I'm delighted to present this special chronicle celebrating NAS' 25th anniversary. Created as the Agency's bold initiative in simulation-based aerospace vehicle design and stewardship, NAS has earned an international reputation as a pioneer in development and application of high-performance computing technologies, providing its diverse customers with world-class aerospace modeling and simulation expertise, and state-of-the-art supercomputing services.

Within these pages, you'll find an overview of NAS' 25-year history, pictorial highlights from the organization's legacy of accomplishments in supercomputing and its contributions to the Space Agency's exciting missions, a snapshot of the division's current supercomputing environment, plus a glimpse at what's in store for NAS in the future. We also honor key contributors and partners in NAS' success story—the "superstars" who had the energy and ingenuity to turn NAS from concept to reality, anticipate trends, and develop new technologies to lead the division to new heights of innovation and mission impact.

NAS has achieved some amazing things over the past several decades, drawing from the synergy between computer science and modeling and simulation. From applying computational fluid dynamics to improving the design and safety of the Space Shuttle Main Engine, to adapting Shuttle technology for a life-saving heart assist device, to leadership in single-system image supercomputing and concurrent visualization that enabled near-real-time hurricane track and intensity forecasts—NAS' high-end computing resources and modeling expertise have been, and will continue to be, a critical part of mission success throughout the Agency.

The NAS team's sustained dedication to excellence and mission success over the years has made the organization a unique asset to both NASA and the nation's high-end computing, aerospace, and scientific communities. I hope you enjoy this overview of our 25-year heritage, as well as the exciting vision of what lies ahead at NAS.

S. Pete Worden
Director, NASA Ames Research Center

NAS OVERVIEW
a Look at NAS' 25 Years of Innovation and a Glimpse at the Future

In the mid-1970s, a group of Ames aerospace researchers began to study a highly innovative concept: NASA could transform U.S. aerospace R&D from the costly and time-consuming wind tunnel-based process to simulation-centric design and engineering by executing emerging computational fluid dynamics (CFD) models on supercomputers at least 1,000 times more powerful than those commercially available at the time. In 1976, Ames Center Director Dr. Hans Mark tasked a group led by Dr. F. Ronald Bailey to explore this concept, leading to formation of the Numerical Aerodynamic Simulator (NAS) Projects Office in 1979.

At the same time, a user interface group was formed consisting of CFD leaders from industry, government, and academia to help guide requirements for the NAS concept and provide feedback on evolving computer feasibility studies. At the conclusion of these activities in 1982, NASA management changed the NAS approach from a focus on purchasing a specially developed supercomputer to an on-going Numerical Aerodynamic Simulation Program to provide leading-edge computational capabilities based on an innovative network-centric environment. The NAS Program plan for implementing this new approach was signed on February 8, 1983.

Grand Opening

As the NAS Program got underway, its first supercomputers were installed in Ames' Central Computing Facility, starting with a Cray X-MP-12 in 1984. However, a key component of the Program plan was to create the NAS facility: a state-of-the-art supercomputing center as well as a multi-disciplinary innovation environment, bringing together CFD experts, computer scientists, visualization specialists, and network and storage engineers under one roof. Groundbreaking for the NAS facility took place on March 14, 1985. In 1986, NAS transitioned from a projects office to a full-fledged division at Ames. In January 1987, NAS staff and equipment were relocated to the new facility, which was dedicated on March 9, 1987. At the grand opening, excitement levels were high about what had been accomplished so far, and about what was yet to come.

Pioneering Achievements

From its earliest years, NAS has achieved many firsts in computing and in enabling NASA's challenging science and engineering missions, including deployment of

the first Cray-2 in 1985, which attained unprecedented performance on real CFD codes. NAS was also the first facility to put the UNIX operating system on supercomputers, and the first to implement TCP/IP networking in a supercomputing environment, linking to users across the country. In 1984, NAS became the first organization to connect supercomputers and workstations together to distribute computation and visualization (what is now known as the client-server model). In another innovation, NAS developed the first UNIX-based hierarchical mass storage system, NAStore. NAS was the first customer for Silicon Graphics Inc. (SGI), beginning an enduring partnership that remains important today, and that led to the successful NASA-SGI-Intel collaboration on the Columbia supercomputer in 2004.

NAS has also been a leader in visualization, benchmarking standards, and job management. The graphics tools, PLOT3D and FAST, both had vast user bases and won major software awards from NASA. The NAS Parallel Benchmarks (NPB), which became the industry standard for objective evaluation of parallel computing architectures, are still widely used today. And since its first release in 1994, the Portable Batch System (PBS) has been commercialized and adopted broadly in the supercomputing community.

Throughout its history, NAS has also been a pioneer and committed partner in both high-fidelity modeling tool development and applications. Redesign of the Shuttle main engine, failure of the Challenger O-ring, wing stall and aerodynamic loads for fighter airplanes, thrust loss in vertical take-off jets—all demonstrated the potential of CFD supported by faster supercomputing to help lives and millions of dollars. Many of the award-winning codes that supported these analyses were also developed at NAS, including INS3D, OVERFLOW, and Cart3D. These codes have continued to evolve through the years, and remain workhorses for NASA missions today.

New Directions

With a wealth of information technology expertise and a visionary spirit, NAS helped pioneer new technologies, including computational nanotechnology and grid computing, starting in the mid-1990s. NAS even set up one of the earliest websites in the world in 1993. In the late 1990s, NASA's Shuttle flights began to seem routine, and aeronautics research and modeling lost momentum in the Agency, leading to diminished investment in supercomputing capabilities and CFD research. Nevertheless, NAS managed to sustain its international leadership in single-system image supercomputing with funding from an information technology R&D project and NASA's Earth Science Enterprise. In 1995, NAS changed its name to the Numerical Aerospace Simulation Division, and in 2001, NAS became the NASA Advanced Supercomputing Division—the name it retains today.

Rediscovering the Mission

Starting in 2001, NASA conducted a series of major engineering studies, including the X-37 vehicle atmospheric reentry and the Shuttle's hydrogen fuel line flowliner cracking, each of which required most of NAS' computational resources for three to six months, significantly delaying other important simulations. Following the Columbia tragedy in February 2003, the Agency again turned to NAS CFD experts for foam debris transport analysis, further straining NAS computational resources in an effort to determine the physical cause of the accident. The Earth Science Enterprise stepped in once more to support the purchase of Kalpana, NAS' first SGI Altix system, which enabled the even larger computational modeling effort required to improve the Shuttle design for a successful Return to Flight.

By June 2004, the importance of supercomputing-enabled high-fidelity modeling and simulation for aerospace and other missions was clear. Purchase of the Columbia supercomputer was approved by NASA and Congress in record time, and the system returned supercomputing leadership to the U.S. after two and a half years with the Japanese Earth Simulator having the highest computational performance. More importantly, NASA was able to pursue its new Space Exploration mission and simultaneously conduct leadership research in Earth science, astrophysics, and aerospace modeling. In October 2005, NASA's long-term commitment to supercomputing at NAS was assured with the formation of the High End Computing Columbia project as the first asset in the Strategic Capabilities Assets Program (SCAP). NAS remains widely recognized as the premier supercomputing facility for the Agency, and vital to each of NASA's mission directorates. In 2007, the Aeronautics Research and Exploration Systems Mission Directorates procured additional resources and placed them at NAS to augment their SCAP allocations. NAS facilities and staff are playing key roles in Shuttle operations, designing the Ares and Orion space exploration vehicles, advancing Earth and space science understanding, and conducting aeronautics research in all flight regimes.

Future

While reflecting on NAS' 25-year legacy is inspiring, current NAS and Agency leadership is carrying on the tradition of the visionaries of a quarter century ago. The newly-installed Hyperwall-2 graphics and data analysis system—currently the world's highest resolution display at a quarter-billion pixels—will allow supercomputer users to view and explore their computational results in unprecedented detail and speed. Visualization will become more tightly integrated into the traditional environment of computing, storage, and networks to provide a more powerful tool to scientists and engineers. The NAS facility is being significantly upgraded and expanded in preparation for installation of new generations of leadership-class supercomputers over the coming months and years to address NASA's increasingly challenging modeling and simulation requirements. Emerging innovative architectures and systems will be strategically leveraged to provide the most effective supercomputing platforms and environments. Advanced algorithms and applications will be developed in lockstep to exploit these petascale systems. Partnerships with other NASA centers, government laboratories, the computer industry, and academia are being aggressively developed to continue Ames and the NAS facility's proud reputation of being the "go-to" place for supercomputing and advanced simulations. Discussions are underway regarding bleeding-edge concepts such as 10-petaflop systems and 1-exabyte data archives at NAS, connected to the science and engineering community via 1 terabit-per-second links, enabling near-real-time aerospace design and deeper understanding of our planet and the universe.

At NAS, we honor our past most sincerely by continually pursuing more impressive achievements in the future. And as these triumphs attract a new generation of innovators and leaders to NAS, soon enough their aspirations and achievements will exceed our own.

A visionary team of managers, engineers, and CFD researchers develops the NAS Program plan, which leads to the beginning of a NASA New Start Program in October. One of the primary goals included in the plan is to provide a national computational capability as a necessary element in ensuring continued leadership in CFD and related disciplines.
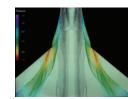
25 Silicon Graphics Inc. (SGI) IRIS 1000 graphics terminals arrive, which soon make a significant impact on post-processing and visualization of CFD results. This first SGI purchase marks the beginning of a successful partnership that remains important today.

NAS Program officially goes into production phase on July 21 at 5 a.m. with Cray-2 links open to 27 remote computer locations across the country, catching nationwide news media attention.
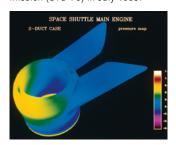
90,000 square-foot NAS facility opens its doors on March 1, 1987 and is dedicated by NASA Administrator, James C. Fletcher eight days later.

NAS purchases its first hardware—two DEC VAX 11/750s named Orville and Wilbur, linked by Ethernet and hyperchannel networks. Engineers design a network-centric environment with a UNIX operating system and TCP/IP protocol.

First 4-processor Cray-2 system, Navier (1.95 Gflops peak) is delivered, which later helps NAS achieve its initial performance goal of a sustained 250 Mflops rate—the world's most powerful supercomputer to date.

Taking advantage of the Cray-2's large memory capacity and speedy calculations, General Dynamics researchers apply CFD methods to obtain Euler and Navier-Stokes solutions for a complete, realistic F-16 fighter configuration.

A second Cray-2, Stokes (same performance characteristics as Navier) arrives, doubling NAS' computing capacity.

1983 ... 1984 ... 1985 ... 1986 ... 1987 ... 1988

In collaboration with Rocketdyne, and supported by NASA Marshall's Development Engine Project Office, researchers generate CFD solutions for the Space Shuttle Main Engine (SSME) hot gas manifold flow, using newly developed INS3D software. This first known application of CFD to rocket propulsion systems development led to replacement of the original three-duct hot gas manifold design. The more reliable, better performing two-duct design made its first flight on Shuttle Discovery's 20th mission (STS-70) in July 1995.

Construction contract for NAS facility is awarded on March 14, marking the groundbreaking of the building.

Development of the Graphics Animation System (GAS) begins, providing useful animation of particle traces so they appear to move through a flow field. Example applications include pressure distributions both inside the SSME, vortex flows over the wing surface of F-16 aircraft, and over the high-speed National Aerospace Plane.

The Remote Interactive Particle-tracer (RIP) application is developed—a unique program to combine capabilities of SGI IRIS graphics workstations and the Cray-2 to reduce demands on computational resources. The workstations perform excellent 3D graphics but lack the Cray-2's memory and speed to process large datasets—RIP exploits the strengths of both systems.
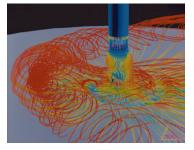
Researchers run CFD calculations on the Cray-2s to investigate V/STOL (vertical and short take-off and landing) flows to better understand the interaction of propulsive jets with V/STOL aircraft and the ground, among other phenomena.

A single-processor Cray X-MP-12 (210.5 Mflops peak performance) arrives at building 233a at NASA Ames (where NAS computer hardware was housed prior to construction of the NAS facility). This is the first customer installation of a UNIX-based Cray supercomputer, and the first to provide client-server interface with UNIX workstations.

NASnet, a high-speed, long-haul communications subsystem using TCP/IP, connects users across the U.S. to supercomputing resources at NAS. This is the first instance of a supercomputer being connected to a TCP/IP network without using a front-end system.

NASA Langley researchers use NAS systems to understand and analyze events in the Shuttle Challenger disaster, and in particular, failure of the O-ring; this work results in the largest 3D structural simulation in the world to date.

NAS partners with the Defense Advanced Research Projects Agency (DARPA) and the University of California, Berkeley in the RAID (Redundant Array of Inexpensive Disks) project, which was completed in 1992, resulting in the RAID storage technology standard in use today.

TIMELINE (1988–1992)

NAS installs two StorageTek 4400 cartridge tape robots (each with a capacity of about 1.1 TB), reducing tape retrieval time from 4 minutes to 15 seconds. In 1990, this new hardware becomes a part of NAStore, the first UNIX-based hierarchical mass storage system.
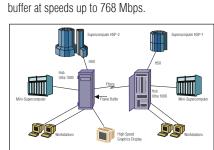
A 16,000-processor Connection Machine model CM-2, Pierre (14.34 Gflops peak), is installed, initiating NAS computer science research in the area of massively parallel computing.

Original NAS Parallel Benchmarks (NPB) are developed in response to the program's increasing involvement with massively parallel architectures. The synthetic application benchmarks mimic computation and data movement of large-scale CFD applications, providing objective evaluation of parallel computing architectures. NPBs become an industry standard and are still widely used today.

Installation of NAS-developed AEROnet is underway, the first high-speed wide-area network (WAN) connecting supercomputing resources to remote customer sites. This new network provides increased bandwidth to remote users, keeping NAS on the forefront of networking technology while implementing open systems standards.

NAS supports the U.S. Navy and McDonnell Aircraft Company in demonstrating the effectiveness of applied CFD in meeting immediate engineering design objectives—to reduce the number of expensive, time-consuming wind tunnel tests required. This effort focuses on developing methods and demonstrating the ability of CFD to predict aerodynamic loads on the F/A-18 vehicle.

1988 · · · 1989 · · · 1990 · · · 1991 · · · 1992 · · ·

NAS deploys UltraNet technology, taking a giant step forward in network speed. Initial tests confirm that UltraNet can sustain data movement from a Cray-2 to a network frame buffer at speeds up to 768 Mbps.

Reynolds, an 8-processor Cray Y-MP (2.54 Gflops peak) arrives to replace a Cray-2 now used by the Department of Defense for classified computing at NAS. Reynolds meets NAS' goal of a sustained 1 Gflops performance rate for CFD applications.
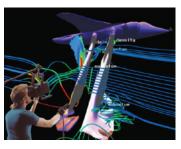
NAS develops the Virtual Windtunnel (VWT), applying virtual reality technology to visualization of CFD simulation results. This highly interactive 3D tool allows users to explore and analyze complex structures found in time-varying fluid flow simulations—for example, momentum and density characteristics on a Harrier Jump Jet model.

Research scientists successfully run incompressible Navier-Stokes simulations of isotropic turbulence—the first application code to demonstrate the effectiveness of Lagrange, the highly parallel Intel iPSC/860 system (128 processors, 7.68 Gflops peak) for very large CFD simulations—the largest having a mesh size of about 16 million grid points.

A 4-processor Convex 3240 system with a peak performance of 200 Mflops is installed.

NAS makes extensive modifications to GRIDGEN, including the addition of four new modules. The multi-program software package, originally developed by Wright Patterson Air Force Base and General Dynamics, is used at NAS for modeling complex-structured aerodynamic grid configurations.
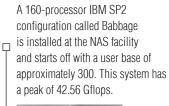
The OVERFLOW code is released outside of NASA. This is the first general-purpose NASA CFD code for overset (Chimera) grid systems—resulting in a production-level CFD capability for complex geometries, and initial moving-body simulations.
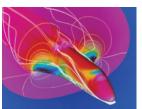
NAS and the National Energy Research Supercomputer Center (NERSC) of Lawrence Livermore National Laboratory release the first phase (alpha test release) of the Portable Batch System (PBS), providing improved batch job scheduling and resource management for massively parallel computers, clustered workstations, and traditional supercomputers. PBS is later commercialized in 1998.
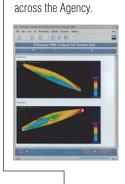
The Flow Analysis Software Toolkit (FAST) team wins NASA Software of the Year Award for their work on the highly modular software environment for visualizing and analyzing CFD scalar and vector data. FAST builds on the earlier PLOT3D program.
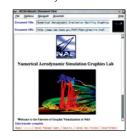
NAS networking group successfully connects the internal NASA AEROnet wide-area network and DARWINnet, a secure local-area network at Ames, to facilitate remote access to wind tunnel and other aircraft design tests across the Agency.

To prepare for a January 2000 test flight over the Pacific, NASA Langley researchers use a NAS Cray system to simulate the scramjet-powered Hyper-X vehicle's unprecedented horizontal separation from its Pegasus booster rocket—a risky Mach-7 maneuver never before attempted.

A 16-processor Cray C90 (15.36 Gflops peak) arrives at the NAS facility, just hours after removal of the Cray-2. This new system, named von Neumann, allows NAS to run many more multitasking jobs around the clock, some as large as 7,200 MB.

A 160-processor IBM SP2 configuration called Babbage is installed at the NAS facility and starts off with a user base of approximately 300. This system has a peak of 42.56 Gflops.

A 64-processor SGI Origin 2000 testbed system named Turing—the first of its kind to be run as a single system—is installed. This joint purchase by the Ames Information Technology Program and NASA's Data Assimilation Office is part of a collaborative plan for systems software development and code optimization.

1993 · · · · · · · · · · **1994** · · · · · · · · · · **1995** · · · · · · · · · · **1996** · · · **1997** · · · · · · · · · · **1998**

NAS' first website goes live—Internet "surfers" can now visit the NAS facility via hyperlinks on the World Wide Web, a wide-area hypermedia information retrieval system.

A 1992 NAS-produced video covering the Division's services and projects wins a nationally recognized Silver Telly Award— the highest honor in the video industry for non-network and cable films.
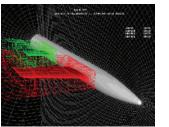
The INS3D code, designed to solve incompressible Navier-Stokes equations in 3D for both steady-state and time-varying flows, wins NASA Software of the Year Award.

Ames computational molecular nanotechnology group is created, securing $1 million for the first-ever nanotechnology program at a federally funded laboratory.
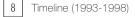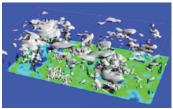
NAS embarks on the ambitious Information Power Grid (IPG) project to seamlessly integrate computing systems, data storage, networks, and analysis software across the Agency. Among the IPG's primary goals is enabling researchers to spend more time conducting research and less time dealing with computing resources required.

NASA Ames computational molecular nanotechnology group wins the Richard P. Feynman Prize in Nanotechnology for Theoretical Work.

NAS researchers are developing CAPO, an automated tool that helps speed up and simplify the tedious process of parallelizing NASA's large serial codes for distributed-memory machines. Among the success stories, NASA Goddard uses CAPO to parallelize one of their codes, enabling them to run larger cloud simulations much faster than before.

The PLOT3D team is awarded the fourth largest prize ever given by the NASA Space Act Program for development of their software program, which has revolutionized scientific visualization and analysis of 3D CFD solutions.

NAS and NASA Langley are co-creating the NASA Metacenter, the Agency's first successful endeavor to dynamically distribute real-user production workloads across supercomputing resources at geographically distant locations. The project involves two IBM-SP2s—one at NASA Langley, the other at NAS.
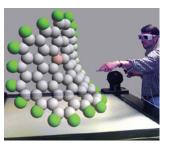
Taking advantage of the multiprocessor architecture of the SGI Origin 2000s at both NAS and the National Center for Supercomputing Applications, NAS researchers are exploring the potential of using nanotubes for tiny electromechanical sensors and semiconductor devices.

Using NAS' 512-processor SGI Origin 2800, Lomax, NASA Langley researchers run CFD simulations of the NASA/Boeing X-37 experimental launch vehicle—which flies at speeds over Mach 20—to analyze thermodynamic conditions of the vehicle's reentry through the atmosphere.

NAS' data analysis group develops Collaborative Virtual Mechanosynthesis, an application for distributive, collaborative, remote computational steering of atoms and interaction with their structures.

The NASA Research and Education Network (NREN) continues to serve as a global leader in deploying new multicast routing and forwarding protocols used for sending information across the Internet.
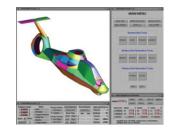

The life-saving DeBakey Ventricular Assist Device (VAD), a collaborative effort between researchers at NASA Ames, NASA Johnson, and MicroMed Technology, Inc., is named NASA's Commercial Invention of the Year. By this date, the device—improved using SSME technology and CFD, coupled with NAS' HEC resources—has been implanted in 115 heart patients in Europe with no reported failures.
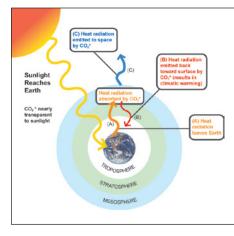

NAS aerospace engineers develop the AeroDB software system to automate CFD processes. With its database-centric approach, AeroDB provides users with a convenient central warehouse for investigating simulation results in real time. Since its launch in September, it has been used during several Shuttle missions to conduct rapid-turnaround engineering analyses.

The Cart3D team wins the 2002 NASA Software of the Year Award. The software package, developed by researchers from Ames and the Courant Institute at New York University, streamlines the design and analysis process for complex configurations such as aerospace vehicles.

1999 2000 .... 2001 .... 2002 ....

Multi-Level Parallelism (MLP), a new method for application parallelization, is developed by NAS computer scientists to take advantage of shared-memory, multiprocessor computing architectures such as the SGI Origin 2000 housed at NAS.

Version 1.0 of the NASA-developed Chimera Grid Tools (CGT) software package is released under a beta testing agreement, and used in a wide variety of disciplines including aerospace, marine, automotive, environmental, and sports. This comprehensive tool for overset grid generation has since been licensed to more than 500 user organizations at different NASA centers, government labs, industries, and universities.
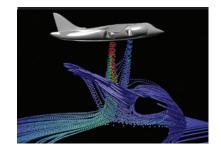

Computational chemistry is being used to calculate the global warming potential of industrial compounds being released into Earth's atmosphere—exploring phenomena impossible to capture in laboratory experiments due to the instability of crucial compounds at Earth's surface.
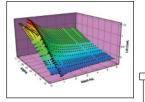

A 1,024-processor SGI Origin 3800, Chapman (819.2 Gflops peak), is installed.

NAS visualization scientists develop two new algorithms, GEL and Batchvis, to help discern unsteady flow features not readily distinguishable or reproducible with traditional software packages. These algorithms are being applied to the YAV-8B Harrier jet to visualize the vehicle's exhaust flows and resulting ground vortices during vertical take-offs—phenomena that can lead to rapid loss of engine thrust and flight failure.
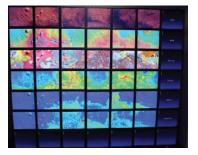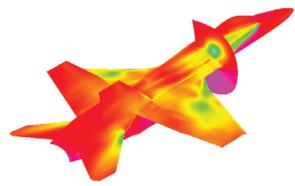

NAS visualization experts develop an innovative visualization system called the Hyperwall, including a graphics cluster and LCD panel wall, which enables scientists to view complex datasets and series of computational results across a dynamic, linked array of 49 displays.

Since the Shuttle Columbia accident in February 2003, NAS resources have largely been dedicated to analysis of the accident's cause. These time-critical analyses underscore the Agency's requirement for expanded supercomputing power to complete human spaceflight analyses much more quickly, while still providing sufficient resources to continue NASA's broad modeling and simulation activities.

The MAP'05 team (under the auspices of the Modeling, Analysis, and Prediction Program), implements two high-resolution versions of the Goddard Earth Observing System (GEOS), a global atmospheric model that enables "real-time" prediction of Atlantic tropical cyclones and other meteorological events. These models run 4 times daily for over 6 months on one of Columbia's 512-processor nodes to produce hurricane forecasts for immediate integration into the national ensemble forecast.

The finite-volume General Circulation Model (fvGCM) runs on Columbia throughout the 2004 hurricane season, validating its real-time, numerical weather prediction capabilities to improve hurricane tracking and intensity forecasts.

NASA-developed HiMAP (High-fidelity Multidisciplinary Analysis Process) software is applied in a joint project with the U.S. Navy to understand phenomena involved in abrupt wing stall of an F-18 E/F aircraft. To date, calculations include more than 17 million grid points involving 104 processors on Lomax.

The first two SGI Altix 512-processor systems of the Columbia supercomputer (63 Tflops peak) are installed, with all 10,240 processors operational by October. Columbia captures the second highest spot on the TOP500 list of the fastest supercomputers in the world, with a LINPACK rating of 51.9 Tflops.

The Simulation Assisted Risk Assessment (SARA) team gets into full swing with Constellation-related work, incorporating CFD analyses performed on Columbia into a probabilistic risk assessment model for Ares I crew launch vehicle (CLV) ascent. Ares I will launch Orion, the crew exploration vehicle (CEV) that will take over NASA human spaceflight missions after the Space Shuttle is retired in 2010.

2003   **2004**   **2005**

The first SGI Altix system, a 512-processor single-system image supercomputer, is installed. The new system, named Kalpana after astronaut and former NAS colleague Kalpana Chawla, is being used for a joint effort between NASA Headquarters, the Jet Propulsion Laboratory, Ames, and Goddard to deliver high-resolution ocean analyses for the ECCO (Estimating the Circulation and Climate of the Ocean) model. The work dramatically accelerates development of highly accurate analyses of global ocean circulation and sea-ice dynamics.

NAS scientists are extensively involved in the Agency's Return to Flight (RTF) activities, employing state-of-the-art CFD codes to simulate steady and unsteady flow fields around the Shuttle during ascent. NAS' findings help make crucial design changes to the Shuttle for future flights.

The Integrated Modeling and Simulation project's Entry Descent and Landing team greatly streamlines efficiency of vehicle design work for the Agency's new Constellation Program, created in May 2004. By coupling engineering methods with high-fidelity aerodynamic and aerothermal simulations run on Columbia, engineers are able to analyze six potential design shapes for the Orion crew exploration vehicle within a three-month period—saving years of design work and demonstrating the value of using high-fidelity simulations early in the design cycle.

The High End Computing Revitalization Task Force (HECRTF) is established in March to fuel improvements in U.S. HEC capacity, capability, and accessibility. NAS is heavily involved, looking for ways to help the U.S. gain a competitive, cost-effective edge over Japan's Earth Simulator, which has ranked highest on the TOP500 list of fastest supercomputers for three years running.
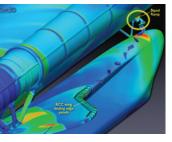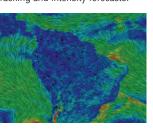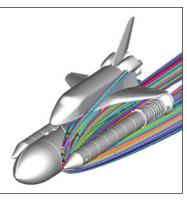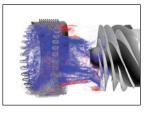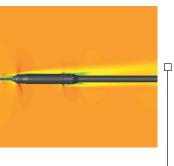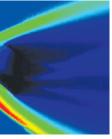
NAS CFD experts are performing high-fidelity computations on Columbia to analyze the Shuttle orbiter's liquid hydrogen feedline flowliner. Several computational models are being used to characterize unsteady flow features, which have been identified as major contributors to cyclic loading that damages the flowliner over time.

NREN completes contract negotiations with National Lambda Rail (NLR) and Level3 Communications, Inc., enabling them to establish 10 Gigabit-per-second wide-area networking between Ames and other NASA field centers. Such robust connection speed is of paramount importance for optimal utilization of Columbia, and for transferring massive datasets required to solve today's challenging problems.
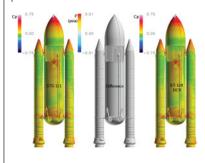
NAS visualization experts assist NASA Goddard researchers in modeling the spectacular merger of two black holes using a 2,048-processor component of Columbia. The Goddard simulation represents one of the largest single calculations ever run on Columbia.

The Division releases the NAS Technology Refresh (NTR) request for information in September, initiating the process of enhancing NASA's high-end computing capacity to keep pace with the Agency's growing HEC requirements. NAS' ultimate goal is to increase total computing capacity at least four-fold over Columbia's existing size by the end of FY09.

Aerospace engineers from NASA Johnson and Ames are relying heavily on Columbia to support redesign efforts for the Shuttle's external tank, consuming more than 600,000 processor-hours on the task to date.
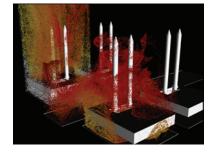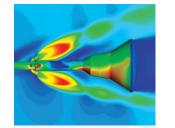

In support of NASA Kennedy ground operations, NAS simulation and modeling experts are conducting time-accurate simulations of the launch pad with double and single solid rocket booster configurations (to mimic the design of the Ares I-X) to analyze effects of ignition over-pressurization on the vehicle.

Two new systems arrive as part of the NTR effort—Schirra, a 640-core IBM POWER5+ (4.8 Tflops peak) and a 2,048-core SGI Altix 4700 (13 Tflops peak, Columbia node 22), which is the largest SSI computer to date. NAS staff members are evaluating the two architectures to determine their potential to help meet the Agency's future computing needs.

NAS visualization team members enhance their concurrent visualization framework to extract complicated vortex core isosurfaces directly on Columbia, shedding new light on phenomena associated with V-22 tiltrotor flight. Without concurrent data extraction, storing and interactively exploring the full data from every simulation timestep (43 terabytes total) would not be possible.

2006 ·········· 2007 ·············· 2008

NAS modeling and simulation experts are using Columbia to determine whether the Vehicle Assembly Building (VAB) at NASA Kennedy, designed for the Shuttle Program, is properly equipped to safely handle storage of the significantly greater number of solid rocket booster segments required for next-generation launch vehicles.

The CEV Aerosciences Project (CAP) Aerodynamics team begins using NASA's Cart3D and OVERFLOW codes on Columbia to generate aerodynamic performance databases for each major design iteration of the Orion launch abort vehicle (LAV). These results enable the team to examine critical performance factors such as stability, control, and separation dynamics for the vehicle, and to recommend specific modifications to the design.

The Army Research Laboratory selects the NAS facility as its new operations base for the Army High Performance Computing Research Center (AHPCRC). NAS now houses a Cray X1E and a Cray XT3, plus other supporting equipment. NASA Ames, Stanford University, High Performance Technologies, Inc., and several minority university partners serve as the consortium for AHPCRC.
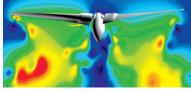

To meet increasing computational demands, the Agency's Aeronautics Research and Exploration Systems Mission Directorates fund the purchase of additional computing resources at NAS. The ESMD systems are two 1,024-core SGI Altix 4700s with the latest Montvale Itanium 2 processors, which become nodes 23 and 24 of Columbia (now 88.9 Tflops). The ARMD system is a 4,096-core SGI Altix ICE cluster (43.5 Tflops peak) with Xeon quad core processors.

After months of performance testing and analysis, NAS announces the outcome of its large technology refresh effort to carry the Agency through the next several years of high-end computing—a 20,480-core SGI ICE cluster system with a peak performance of 245 Tflops. The system will be installed in summer 2008, initially augmenting and eventually, replacing Columbia.

NAS visualization, applications, and facilities teams work together to complete construction of Hyperwall-2, a follow-on version of the Hyperwall visualization system that enables scientists to display high-fidelity simulation results in a single view. This new version comprises 128 state-of-the-art graphics processing units and 1,024 cores, providing 74 teraflops peak processing power and a total of 245 million pixels—the highest resolution display in the world today.

NAS SUPERSTARS

Among the large cast of important players deserving applause for their contributions to the reputation and success of NAS over the past quarter-century, these 25 individuals are spotlighted for their vision, guidance, and influence on the organization's achievements in supercomputing and computational science for NASA and the nation.

David Bailey Senior Scientist 1984–1998: Expert in high-performance scientific computing who led development of the NAS Parallel Benchmarks that are still widely used to evaluate sustained performance of highly parallel supercomputers.

Ron Bailey Founding NAS Division Chief 1982–1990: Computational fluid dynamics (CFD) pioneer and leader of the team that developed the NAS concept and laid its technical foundation. Anticipated and promoted the transition from vector to parallel computing, and remains an important contributor to NAS' vision today.

Bill Ballhaus Ames Center Director 1984–1989: Among the researchers who created the field of CFD at NASA Ames. Earlier, as Director of Astronautics when the NAS Program plan was approved, strongly supported the NAS vision.

Bob Bishop SGI CEO 2004–2007: Key partner in NASA-industry collaboration to build Columbia. Enabled radical alteration of SGI production to turn out 5–10 times more Altix systems to meet the aggressive Columbia schedule. Contributed intellectual and fiscal support to create this powerful national resource.

Bruce Blaylock NAS Chief Engineer, 1994–1999: Important early contributor to shaping of NAS Program. Earlier, as NAS Systems Development Branch Chief, provided clear vision and staff support for software development, and introduced rigorous design methodologies, tools, and environments.

Loren Bright Director of Ames Research Support 1969–1979: Early advocate of supercomputing at Ames who led the ILLIAC IV project. Teamed with Chapman to champion and support early NAS technical studies and program planning.

Walt Brooks NAS Division Chief 2003–2005: Created the NASA-industry partnership and gained Congressional approval for Columbia, which reestablished NAS as a world-class supercomputing facility and created critical simulation capabilities in support of Agency missions.

Dean Chapman Ames Director of Astronautics 1974–1980: Visionary who established the first research group at Ames to develop and combine CFD with high-speed computers to lead NASA into the future.

Dan Clancy Ames Director of Information Sciences 2003–2004: Recognized Columbia's potential for NASA, played a critical role in supporting and advocating the project at the Agency's highest levels. Worked closely with Brooks to make HEC a hallmark of Ames technology.

Dale Compton Ames Center Director 1989–1994: Dedicated supporter of NAS Program and facility, and a strong advocate for aeronautics user community involvement to help define requirements for NAS. Earlier, a strong advocate for NAS project development and implementation.

Dave Cooper NAS Division Chief 1991–1995: NAS advocate extraordinaire who tirelessly campaigned for NAS recognition and funding. Influenced the transition from vector to parallel computing, and contributed significantly to the advancement of U.S. supercomputing.
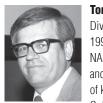
Bill Feiereisen NAS Division Chief 1998–2002: Gave NAS its start in single-system image (SSI) architectures, increasing NAS' computing capability ten-fold, and was a key player in Ames' foray into grid technologies. Previously, Program Manager for NASA's HPCC Program.

Randy Graves NASA Director of Aerodynamics 1984–1989: As the first Headquarters Program Manager for NAS, understood and promoted its importance as part of the U.S. supercomputing effort. Helped shape national HPC policies and was influential in starting the HPCC Program of the 1990s.

Scott Hubbard Ames Center Director 2002–2006: Recognized Columbia's importance to supporting Agency missions; was key to bringing all pieces together to reach the NASA-industry agreement and gain Headquarters' approval.

Paul Kutler Ames Applied Computational Aerodynamics Branch Chief 1980–1984: CFD research scientist and manager who developed and promoted creation of aerospace-related software to run on supercomputers. Dedicated much of his NASA career to advocacy, planning, and utilization of NAS capabilities.

Tom Lasinski NAS Deputy Division Chief (acting) 1995–1999: Established and led NAS Applied Research Branch and fostered development of key technology projects. Collaborated with aeronautics scientists at Ames to create the NAS prototype that tightly integrated workstations with supercomputers.

Harvard Lomax Ames Computational Fluid Dynamics Branch Chief 1970–1992: Pioneer in emergence of CFD, and intellectual powerhouse in combining CFD with supercomputers. With Kutler, Peterson, and Chapman, showcased NAS' CFD accomplishments to U.S. aerospace companies.

Hans Mark Ames Center Director 1969–1977: "Godfather" of numerical simulation at Ames, appointed task team to study concept leading to NAS. Managed Ames' research and applications efforts in aeronautics, space science, life science, and technology, and strongly advocated CFD benefits in these areas.

Harry McDonald Ames Center Director 1996–2002: Supported pioneering work in SSI architectures, known internationally for his work in computer applications. Responsible for defining and executing Ames' role as a Center of Excellence for Information Technologies.

Paul Otellini Intel President and Chief Executive Officer 2005–present: Envisioned the partnership between NAS and SGI to deliver the Intel Itanium 2 processor-based Columbia supercomputer to accelerate NASA's design and engineering research.

Dave Parry SGI Senior Vice President and Product General Manager 2002–present: Chief architect in hardware engineering of SGI's Altix servers that enabled Columbia's implementation. His vision and leadership moved the Altix to Open Linux and provided scaling to 512 and 2,048 processors, enabling Columbia's growth.

Vic Peterson Ames Deputy Director 1989–1994: Earlier, as Chief, Thermo- and Gas-Dynamics Division and then Director of Aerophysics, an originator (with Bailey, Chapman, and Mark) of initiative to develop the NAS capability. Devoted much time to advocating the importance of supercomputing to maintain U.S. leadership in aerospace.

Paul Rubbert Boeing Commercial Airplane Group's Chief of Aerodynamics Research 1984–1992: Championed NAS from concept through maturity and expounded NAS as a pathfinding organization, leading to duplication of NAS model at other sites. Provided long-term support, experience, and vision for evolution of CFD and supercomputing.

Sy Syvertson Ames Center Director 1978–1984: Signatory on the NAS Program plan, who led aerodynamics and wind tunnels at Ames into the supersonic and hypersonic regimes. His earlier research on lifting bodies helped define the basic shape of the Apollo capsules, and the Space Shuttle orbiter.

Steven Wheat Intel Senior Director of HPC Platforms 1995–present: Innovator of Intel's flagship supercomputing project, who worked directly with NAS management to develop the initial framework, key objectives, and parameters for Columbia.

Special appreciation goes to our stakeholders—the many managers and leaders throughout NASA and up through the halls of Congress, whose sponsorship, commitment, and financial support have allowed NAS' successful 25-year run.

NAS TODAY
an Integrated Supercomputing Environment Supporting NASA Missions 24x7

The NAS concept was a revolutionary idea for its time: A balanced supercomputing capability consisting of processors, networks, storage, and visualization tools, all speaking the same language. Over the last quarter-century, NAS has continued to evolve as a leader in high-end computing (HEC). Our supercomputers are still integrated with high-speed networks and mass data storage systems, but are now augmented by customizable support in user services, software performance optimization, high-fidelity modeling and simulation, and scientific visualization. Together, they provide our users with unparalleled support throughout the entire life-cycle of their projects.

Production Supercomputing
Our expert HEC team manages all aspects of the NAS production supercomputing environment to ensure users get the secure, reliable resources they expect. The current environment includes three supercomputers, as well as two secure front-end systems requiring two-factor authentication, and two secure unattended proxy systems for remote operations. The Columbia supercomputer—recently upgraded to 14,336 cores and 88.9 Tflops of peak processing power—continues to enable striking advances in addressing NASA's real-world science and engineering challenges. In addition, NAS houses two smaller but still very powerful testbed systems, RTJones (43.5 Tflops) and Schirra (4.8 Tflops).

Networking Technologies
NAS' high-speed networking technologies and services are essential to the performance of the HEC systems and their high impact for our customers. Using high-capacity network connections and our local and wide-area network expertise, scientists' massive data transfers (some exceeding 1 terabyte per hour) travel seamlessly between local and remote systems. Our network engineers have also implemented innovative transfer strategies and protocols to vastly reduce time-to-solution for users.

Scientific Visualization
The NAS visualization team develops and applies tools and techniques customized for NASA science and engineering problems. Our tools help users view and interact with their results—right from their desktops—to quickly pinpoint important details in complex datasets. Working closely with NASA users, our visualization experts capture, parallelize, and render the enormous amounts of data needed to produce high-resolution, 3D visualizations and videos. NAS' new Hyperwall-2 visualization system—the highest known resolution display in the world at a quarter-billion pixels—is designed to provide users with a way to visualize massive datasets directly from the supercomputing systems.

High-Fidelity Modeling and Simulation
Many of NASA's missions present unique aerospace design challenges that no off-the-shelf software can address. NAS scientists specializing in physics-based modeling and simulation have years of experience in developing world-class CFD software packages, custom tools, and methods. Our CFD experts provide critical services to NASA teams—often dramatically reducing time-to-solution and substantially increasing model resolution. Recently, the team has been involved in simulations to assess design risks for proposed crew exploration and launch vehicles.

User Services
Our user services team works 24x7 to assure users can make effective, productive use of the NAS HEC systems. This team's experience in anticipating and heading off problems and quickly solving challenges as they arise is greatly appreciated by users. During Space Shuttle missions, the team marshals all HEC components and NAS staff to ensure essential support is in place for vital analysis tasks. This enables engineers to rapidly provide mission managers with critical data to clear Shuttles for landing. User services staff monitor all systems, networks, job scheduling, and resource allocations to ensure a stable, seamless computing environment.

Performance Optimization
The NAS application optimization team specializes in enhancing performance of NASA's complex codes so researchers can do more science and engineering in less time by utilizing our HEC systems more effectively. Optimization services range from answering basic questions to partnering with users for in-depth code performance optimization. The team also evaluates tools and technologies best suited for the NAS environment and gives feedback to outside tool developers.

Data Storage
Our HEC customers often require vast amounts of data storage. With about 25 petabytes (PB) of tertiary storage and over 2 PB of disk storage on the floor, NAS' mass storage system allows users to archive and retrieve important results quickly, reliably, and securely. Our storage specialists create custom filesystems to temporarily store large amounts of data for special user projects, and provide individual training to help users efficiently manage and move such large amounts of data.

For more details on NAS systems and services, please visit http://www.nas.nasa.gov/Resources/resources.html on the web.

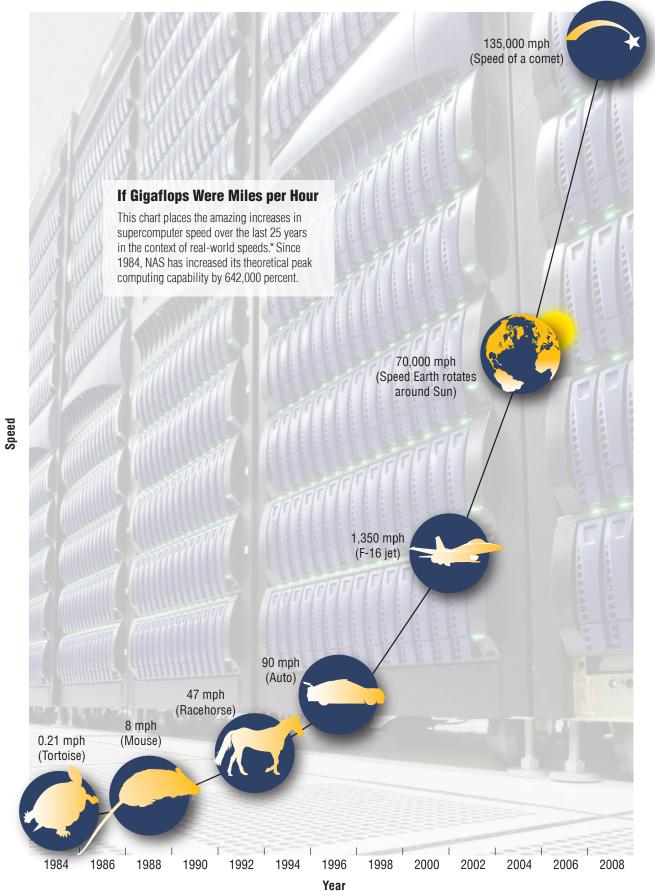Theoretical peak performance of NAS' first supercomputer, the Cray X-MP (1984), in Teraflops **:** 0.00021

Minimum number of Apple Mac Mini computers it would take to exceed performance of the Cray X-MP **:** 1

Theoretical peak performance of Columbia (2008), in Teraflops **:** 88.9

Number of Mac Minis it would take to equal peak performance of Columbia **:** 329,185

Estimated total number of people who have ever lived on Earth, as of 2002 **:** 106 billion

Amount of time, in seconds, it would take the Cray X-MP to add all their heights together **:** 505

Amount of time, in seconds, it would take Columbia to add all their heights together **:** 0.0012

Amount of time in years it would take a person to complete 1 hour of Columbia calculations, if they performed 1 calculation per second for 8 hours per day for 365 days per year **:** 30 billion

Approximate number of books housed in the U.S. Library of Congress collection **:** 32 million

Number of complete copies of the U.S. Library of Congress collection that would fit on NAS' current mass storage system **:** 1,000+

Current amount of unique data, in Petabytes, stored on NAS mass storage systems **:** 3.2

Approximate number of trees it would take to make the paper needed to print all of this data **:** 160 million

Total amount of memory (RAM), in Gigabytes, available on the Cray X-MP supercomputer **:** 0.016

Number of iPod songs that could be stored in the Cray X-MP's memory **:** 4

Total amount of memory (RAM), in Gigabytes, currently available on the Columbia supercomputer **:** 28,672

Number of iPod songs that could be stored in Columbia's memory **:** 6,880,320

Speed of network connections, in bits per second, to some of NAS' remote users in 1983 **:** 300

Number of years it would take to download a feature-length movie using the original network connection **:** 3.97

Speed of NAS' current high-bandwidth wide-area network, in Gigabits per second **:** 10

Number of seconds it would take to download a feature length movie now **:** 3.76

Hyperwall-2 (2008) graphics processing capability, in texture elements/second **:** 4.71 trillion

Number of Microsoft Xbox game consoles required to equal this graphics capability **:** 589

If Gigaflops Were Miles per Hour

This chart places the amazing increases in supercomputer speed over the last 25 years in the context of real-world speeds.* Since 1984, NAS has increased its theoretical peak computing capability by 642,000 percent.

Speed

135,000 mph
(Speed of a comet)

70,000 mph
(Speed Earth rotates around Sun)

1,350 mph
(F-16 jet)

90 mph
(Auto)

47 mph
(Racehorse)

8 mph
(Mouse)

0.21 mph
(Tortoise)

Year

1984 1986 1988 1990 1992 1994 1996 1998 2000 2002 2004 2006 2008

*The numbers in this chart represent the aggregate theoretical peak of the supercomputers operational at NAS on each date. For the purposes of this presentation, 1 Gflops = 1 mph