NASA ADVANCED SUPERCOMPUTING (NAS) FACILITY TIMELINE

www.nasa.gov
www.nas.nasa.gov

1987 • *von Karman*
Convex 3820 • 8 CPUs
Peak Performance: 1.92 GFLOPS

1987 • *Pierre*
Thinking Machines CM-2 • 1,024 CPUs
Peak Performance: 14.34 GFLOPS

1985 • *Navier – 1987 • Stokes (pictured)*
Cray 2 • 4 CPUs
Peak Performance: 1.95 GFLOPS

1985
Cray X-MP 12 • 1 CPUs
Peak Performance: 0.21 GFLOPS

1990 • *LaGrange*
Intel iPSC/860 • 128 CPUs
Peak Performance: 7.68 GFLOPS

1988 • *Reynolds*
Cray Y-MP • 8 CPUs
Peak Performance: 2.54 GFLOPS

1994 • *Babbage*
IBM SP-2 • 128 CPUs
Peak Performance: 34.05 GFLOPS

1993 • *Boltzmann*
Thinking Machines CM5 • 128 CPUs
Peak Performance: 16.38 GFLOPS

1993 • *Grace*
Intel Paragon • 208 CPUs
Peak Performance: 15.60 GFLOPS

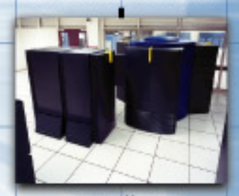

1993 • *von Neumann*
Cray C90 • 16 CPUs
Peak Performance: 15.36 GFLOPS

1993 • *Eagle*
Cray C90 • 8 CPUs
Peak Performance: 7.68 GFLOPS

1996 • *Newton*
Cray J90 • 36 CPUs
Peak Performance: 7.20 GFLOPS

1998 • *Steger*
SGI Origin 2800 • 256 CPUs • 300 MHz
Peak Performance: 128.00 GFLOPS

1999 • *Lomax*
SGI Origin 2800 • 512 CPUs • 400 MHz
Peak Performance: 409.60 GFLOPS

1997 • *Turing*
SGI Origin 2000 • 24 CPUs • 195 MHz
Peak Performance: 9.36 GFLOPS

2001 • *Chapman*
SGI Origin 3800 • 1,024 CPUs • 600 MHz
Peak Performance: 1.23 TFLOPS

2001 • *Bright*
Cray SV1ex • 32 CPUs • 500 MHz
Peak Performance: 64.00 GFLOPS

2004 • *Columbia*
SGI Altix 3700 • 10,240 CPUs • 1.5 GHz
Peak Performance: 61.44 TFLOPS

2004
Cray X1 • 16 CPUs • 800 MHz
Peak Performance: 51.20 GFLOPS

2003 • *Kalpana*
SGI Altix 3000 • 512 CPUs • 1.3GHz
Peak Performance: 2.66 TFLOPS

1984 1985 1986 1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005

Between 1985 and 2004, the NAS Facility increased its peak computing speed by 260,000 times

1 gigaflop = 1 billion floating point operations/second • 1 teraflop = 1 trillion floating point operations/second

NETWORK ROADMAP TIMELINE

1984

- First major supercomputer center to adopt TCP/IP across all systems.

1985

- Connected to Ames network via Ethernet, 10 Mb/s. (LAN)
- Created NASA's first TCP/IP WAN (NASnet).
- HyperChannel link installed between workstation and Cray 2. (LAN)

1986

- Long Haul Communications Prototype (LHCP) connects Ames to Langley, Lewis via satellite (224 Kb/s).
- Participated in development of BARRnet.

1988

- T1 link (1.544 Mb/s) installed to Langley.
- NASnet increased to 30 sites with 700 users.

1989

- NASnet circuits transitioned from switched to dedicated 56 Kb/s.

1990

- NAS proves IP routing technology through Research Internet Gateway (RIG) project.

1991

- Installed high speed Ultranet LAN and FDDI ring.
- Transition from NASnet (bridged) to AERONet (routed network).

1992

- AERONet 1 becomes operational, providing direct access for NAS remote users, using largest OPSF network.
- Completed first Switched Multimegabit Data Service (SMDS) test demonstrating variable bandwidth.
- Implementation of the router discovery (RFC1256), a software program that allows a host to dynamically choose its next-hop router for default routing, independently of the actual routing protocol used among the routers.

1993

- Distributed computing team established from NAS staff and users. (LAN)
- HIPPI deployed on LAN.
- Frame Relay deployed into AERONet.
- SMDS employed into AERONet.
- I of NEWT demonstrates remote access windtunnel concept.
- Supercomputing Conference demonstrated new technologies with the world's highest performance supercomputing network.
- Patent on HNMS – hierarchical network management system, a distributed, scalable IP network monitor and manager. Used an x-window interface to graphically display and control medium to large-sized IP networks. Based on SNMP.

1994

- Upgraded AERONet backbone to T3.
- First cross country ATM trial conducted over DS-3 circuits.
- Multiple File Transfer Protocol (MFTP) enhanced FTP that opens multiple streams for transferring data, to overcome the bandwidth delay product (1994).
- Demonstrated satellite-based megabit applications using the ACTS satellite and associated ground terminals.

1995

- Deployed first federal native ATM network cross-country over DS-3 (NREN), Phase 1.
- Deployed NREN Phase II Native ATM over OC-3.
- Asynchronous Remote Copy Program (ARCP). A reliable, restartable file transfer mechanism. ARCP guarantees the successful completion of a file transfer.
- Demonstrated 155 Mbps ATM interconnects across NREN.

1996

- WAN consolidation within NASA began and AERONet transitioned to NISN.
- Demonstrated the interoperability between independently managed NREN networks that are based on ATM technology supplied by multiple vendors.
- NREN joins the federal Next Generation Internet (NGI) initiative.

1997

- Collaborated with Paramount studios to demonstrated the first uncompressed D1 video streaming over WAN (OC12) at Supercomputing Conference.
- Deployed OC-3 155 Mbps network across NREN WAN.
- Established NREN peering at the NGIX-Chicago Exchange point.
- Demonstrated real-time, echocardiography Over a Wide Area Network (WAN).

1998

- Deployed ATM OC-12 (155Mb/s) backbone to DARWINnet.
- Demonstrated real-time high-definition video over an ATM WAN, with remote access to windtunnels at SC98.
- Established the NGIX-West Exchange Point at NASA Ames.
- Transferred ATM-based WAN architecture to NISN.
- Transferred NGIX peering technology to NISN.
- Supported the remote operation of the Nomad Rover between the US and Chile.

1999

- Participated in series of aviation safety projects, RTSS, SMA, AEN, PDARS.
- Deployed native multicast across NREN.
- Deployed OC-12 622 Mbps network across NREN WAN.
- Utilized Uni-Directional Link Routing (UDLR), across satellite links, to support the virtual reality based distributed Virtual Collaborative Clinic application.

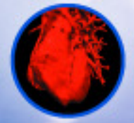

2000

- Set up high-speed interconnect (622 Mbps) for DAO between Ames and Goddard, utilizing ATM links and the NREN backbone.
- Began collaboration in the Secure Advanced Federated Environment (SAFE) project.
- Developed PCMon, a passive network monitoring tool.
- Demonstrated High Definition TV streams across NREN and Internet2.
- At the NREN Gigabit Networking workshop, demonstrated the Digital Sky Virtual Observatory.

2001

- Expanded NAS LAN beyond Ames boundaries to Research Park.
- Deployed high-speed Gigabit Ethernet with Jumbo-frames.
- Provided multicast functionality into the NAS LAN.
- Provided dynamic IP address allocation (DHCP) into NAS LAN for visiting scientists.
- The first Multicast Peering point, Multicast was established at Ames.
- Prototype Quality of Service (QoS) mechanisms in a WAN environment between two NREN sites.

2002

- Deployed 802.11b wireless technology into the NAS LAN.
- Co-developed with NAS security a Wireless Firewall Gateway, featured on the cover article of Government Computer News (GCN).
- Completed HIPPI translation to GigE, except CRAY Sv1 (no capability).
- Demonstrated MPEG2 video conferencing between Ames and NASA Headquarters across the NREN WAN.
- Utilized Multi-Path Label Switching (MPLS) to partition NREN WAN in conjunction with the NASA Prototyping Network (NPN).
- Supported remote collaboration among aerospace design teams, utilizing the Virtual Flight Rapid Integrated Test Environment via WAN-based Multicast technologies.

2003

- QoS enables High Definition Television (HDTV) meetings between Ames and JPL.
- NREN partners with Earth Sciences Enterprise for Aura Spacecraft Checkout, utilizing QoS.
- Networking and grid technologies merged to enable Earth Science ground-truthing experiment.

2004

- Established 1 Gbps link between Ames and Jet Propulsion Laboratory (JPL).
- Established 1 & 10 Gbps networks between the nodes of the Columbia supercomputer.

2005

- Established 1 Gbps link between NASA Ames and NASA Goddard.

National Aeronautics and Space Administration
www.nasa.gov

AERONAUTICS RESEARCH MISSION DIRECTORATE

The goal of the Aeronautics Research Mission Directorate is to pursue research and technology development that: Protects air travelers, the public, and the environment; increases mobility and enhances national security; and pioneers revolutionary aeronautical concepts for science and exploration. The objective of the Aeronautics Research Mission Directorate is to pioneer and validate high-value technologies that enable new exploration and discovery, and improve quality of life through practical applications.

National Aeronautics and Space Administration
www.nasa.gov

NASA ENGINEERING AND SAFETY CENTER (NESC)

The NASA Engineering and Safety Center (NESC) is an independent organization, charted in the wake of the Space Shuttle Columbia accident to serve as an agency-wide technical resource focused on engineering excellence. The objective of the NESC is to improve safety by performing in-depth independent engineering assessments, testing, and analyses to uncover technical vulnerabilities and to determine appropriate preventative and corrective actions for problems, trends, or issues within NASA's programs, projects, and institutions.

National Aeronautics and Space Administration
www.nasa.gov

EXPLORATION SYSTEMS MISSION DIRECTORATE

The Exploration Systems Mission Directorate aims to develop a constellation of new capabilities, supporting technologies, and foundational research that enables sustained and affordable human and robotic exploration. Key research areas include: development of a Crew Exploration Vehicle for astronaut travel in space; health and safety assurance of crew on long-duration space flights; and evaluation of plans for a new capability to service, repair, and eventually de-orbit the Hubble Space Telescope.

SCIENCE MISSION DIRECTORATE

NASA's Science Mission Directorate carries out scientific exploration of Earth, the moon, Mars, and beyond; charts the best route of discovery; and reaps the benefits of Earth and space exploration for society. It seeks to: establish an understanding of Earth and other planets and their evolution; bring the lessons learned from our study of Earth to the exploration of the solar system; and assure the discoveries made here will enhance our work there.

National Aeronautics and Space Administration
www.nasa.gov

SPACE OPERATIONS MISSION DIRECTORATE

The Space Operations Mission Directorate provides agency leadership and management, including top-level requirements development, policy, and programmatic oversight of NASA space operations related to human exploration in and beyond low-Earth orbit. 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.

National Aeronautics and Space Administration
www.nasa.gov

SPACE OPERATIONS MISSION DIRECTORATE

PROPULSION SUBSYSTEMS: ADVANCED TURBOPUMPS & FLOWLINERS

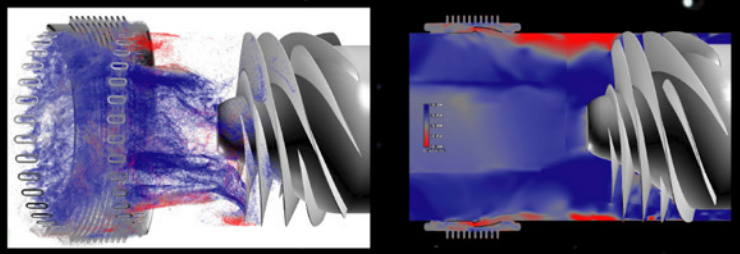

Illustration of unsteady interaction between the backflow and the flow in the bellows cavity, considered one of the major contributors to high-frequency cyclic loading.

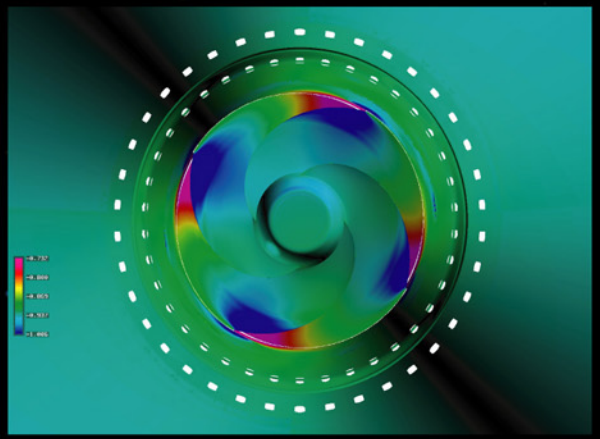

Strong backflow causing high-frequency pressure oscillations.

NASA researchers are working to provide a computational framework for design and analysis of the entire fuel supply system of a liquid rocket engine, including high-fidelity unsteady turbopump flow analysis.

This effort will decrease design costs, improve performance and reliability, and provide developers with information such as transient flow phenomena at startup, impact of non-uniform flows, and impact on the structure. The computational framework has been used to investigate the root cause of fuel-line cracks, and will enable technologies for future engine designs.

Various computational models have also been developed, and time-accurate computations have been carried out to characterize various aspects of the flow field surrounding the flowliner.

The low-pressure fuel turbopump inducer creates backflow regions that extend both upstream and downstream of the flowliner regions. The upstream propagation of the pressure field poses high-frequency loading on the flowliners. This unsteady interaction between the backflow and the flow in the bellows cavity is considered one of the major contributors to the high-frequency cyclic loading.

National Aeronautics and Space Administration
www.nasa.gov

SPACE OPERATIONS MISSION DIRECTORATE

SIMULATION OF SPACE SHUTTLE ASCENT AND DEBRIS

Plot of several trajectories depicting the path of debris shed from the shuttle's external tank, generated using the Ames-developed debris-transport analysis software. (Stuart Rogers, NASA Ames Research Center)

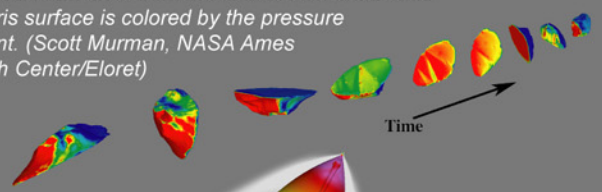

Time sequence of computed six-degree-of-freedom trajectory of an actual foam divot from the shuttle's external tank. The debris surface is colored by the pressure coefficient. (Scott Murman, NASA Ames Research Center/Eloret)

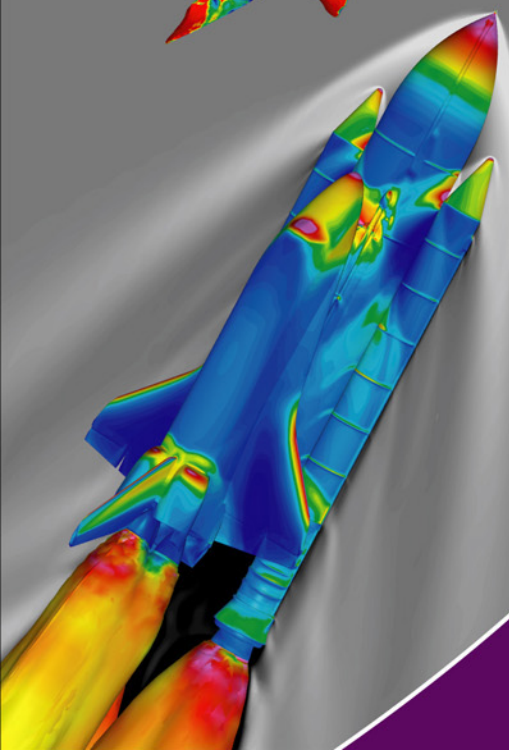

Flowfield around the Space Shuttle Launch Vehicle traveling at Mach 2.5 during ascent, as simulated by NASA computational fluid dynamics software run on the Columbia supercomputer. (Darby Vicker, NASA Johnson Space Center)

Drawing from more than 400 computational fluid dynamics simulations of Space Shuttle Launch Vehicle ascent aerodynamics, the Ames-developed debris-transport analysis software package has become the primary engineering tool used to perform analyses of all potential debris, and is critical to space shuttle Return to Flight efforts.

During development of the debris-transport analysis software, unsteady six-degree-of-freedom moving-body simulations were used to create drag and cross-range models for the different shapes of ice and foam. These models were validated with experimental results, and then incorporated into the software. During the Discovery mission, the software was used to compute trajectories of the foam debris seen in the launch video and the photos taken by the crew.

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

CODE PORTING AND SCALING EXPERIENCES ON COLUMBIA

Data Parallel Line Relaxation Code (DPLR) - a space shuttle aero-thermal analysis code which is critical for NASA's Return-to-Flight (RTF) mission and is being used in the design of the Crew Exploration Vehicle.

OVERFLOW2 - a computational fluid dynamics code used in a large-scale static aeroelastic stability derivative computation needed in the design of aerospace vehicles.

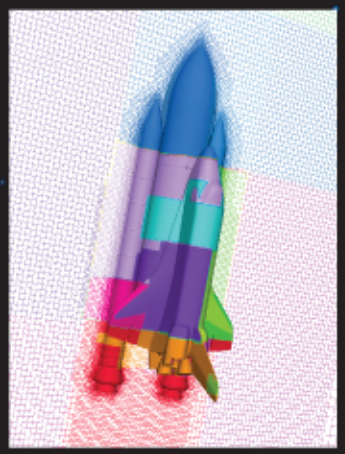

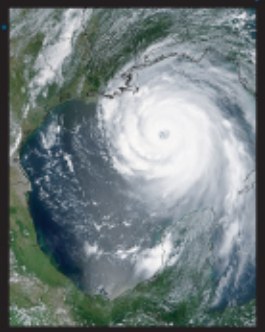

finite-volume General Circulation Model (fvGCM) - an ultra-high resolution code used in hurricane predictions (pictured here is hurricane Katrina on August 28, 2005).

The support staff at the NASA Advanced Supercomputing Division at Ames Research Center in California provides full-service code porting and scaling assistance to Columbia supercomputer users.

Code *porting* assistance involves some or all of the following:

- Choosing compiler flags or options to preserve accuracy of results
- Modifying or creating makefile and run scripts
- Replacing system-dependent function calls
- Experimenting with different compiler versions and/or libraries
- Downloading, building, and installing third-party libraries/software from the public domain
- Debugging code

Code *scaling* assistance involves some or all of the following:

- Choosing compiler options and/or versions for aggressive optimization without sacrificing accuracy
- Profiling code to measure the relative amount of time spent in the various routines
- Modifying or restructuring code to improve cache utilization or memory access
- Identifying and eliminating performance bottlenecks such as memory contention, unnecessary barrier synchronization, and restructuring or reducing I/O
- Converting serial code to parallel code

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

AN APPLICATION-BASED PERFORMANCE CHARACTERIZATION OF THE COLUMBIA SUPERCOMPUTER

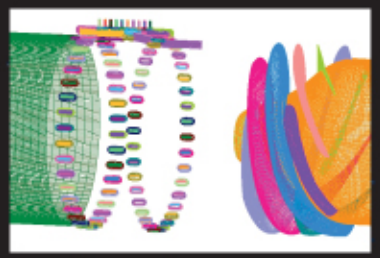

Incompressible Navier-Stokes Solver in General 3D Coordinates (INS3D) simulates 3D viscous fluid flow over arbitrary surfaces, including unsteady flow and moving boundaries. Left: INS3D surface grids for the low pressure fuel pump inducer and the flowliner on the Space Shuttle. (Cetin Kiris and William Chan, NASA Ames Research Center)

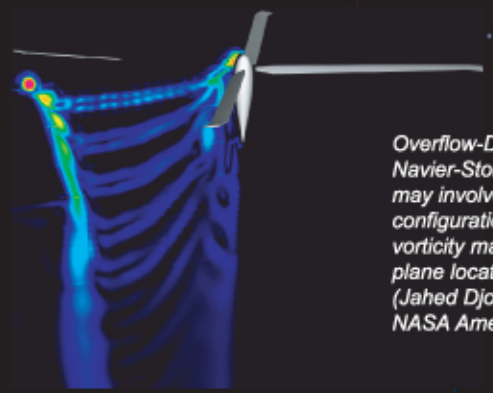

Overflow-D, a general purpose Navier-Stokes solver for problems that may involve relative motion between configuration components, computed the vorticity magnitude contours on a cutting plane located 45° behind the rotor blade. (Jahed Djomehri and Roger Strawn, NASA Ames Research Center)

Comparison of NAS Parallel Benchmarks Multi-zone (NPB-MZ) performance under three different networks: in-node, NUMAlink4 (XPM), and InfiniBand (IB). (Henry Jin, NASA Ames Research Center)

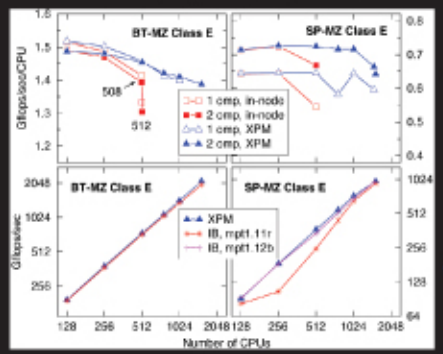

Columbia Supercomputer

- Comprised of 20 Altix nodes with 512 processors each
- Includes a 4-node shared memory system (2,048 processors)

Performance Characteristics Evaluated

- Floating-point performance
- Memory bandwidth
- Message passing latency and bandwidth
- Compilers

Tools Used

- High Performance Computing (HPC) Challenge Benchmarks
- NASA Advanced Supercomputing Parallel Benchmarks (NPB)
- Three scientific applications of interest to NASA

Results: 4-node shared memory system shows promise for multi-node application scaling to 2,048 processors

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

HIGH FREQUENCY COLUMBIA MONITORING TOOL – "NODEMON"

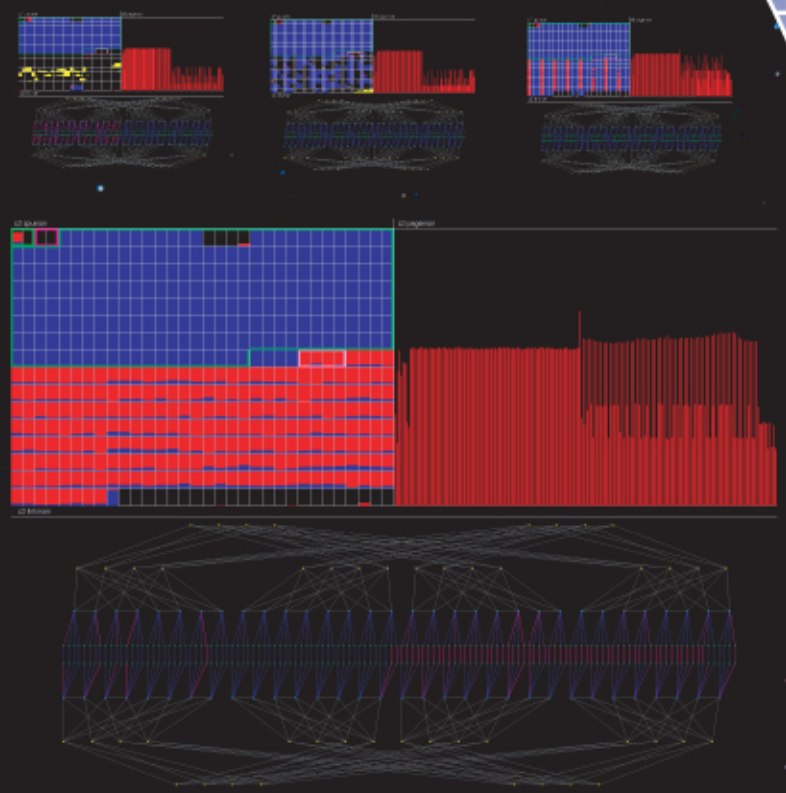

A series of four screen shots of the nodemon tool running on the hyperwall. Upper left of each screen: Activity for all 512 CPUs on 1 of Columbia's 20 nodes. Upper right: Memory usage for each 2-CPU node; top range is 4GB. Bottom: Dual-plane fat tree NUMAlink interconnect topology, with activity shown in magenta between the 2 nodes within the green squares in a C-brick, and between the nodes and the routers represented by cyan squares; level 1 and level 2 metarouters are indicated by the yellow squares.

cpumon
Black = idle
Blue = user time
Green = nice
Red = system time
Cyan = IO wait
Yellow = soft interrupts
Orange = hard interrupts

linkmon
Magenta = active
Blue = inactive
Gray = unreadable

Nodemon is a visualization tool for monitoring system resource utilization on the Columbia supercomputer, a 10,240-processor system at NASA Ames Research Center.

- Monitoring and visualization components are developed using the "growler" distributed-object, modular plug-in architecture
- Collects high frequency statistics, typically at 20 hertz
- Runs one growler-nodemon daemon on each single system image (SSI) Columbia node
- Maintains ability to run viewers on any Linux workstation

cpumon and *cpusetmon* represent central processing unit (CPU) activity information, collected from /proc/stat and the cpuset kernel interface, organized in grid form (one cell per CPU).

- cpusets (CPU allocations) are represented by colored outlines
- Cell colors indicate amount of activity

pagemon represents memory distribution information across processor nodes, collected from /proc/discontig or /sys/devices, organized in bar graph form (one bar per 4 gigabyte region).

linkmon represents numalink traffic statistics, collected via the Performance Copilot (PCP) library, in a topologically correct graph.

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

HIGH SPEED NETWORKING

Target configuration for the NASA Research and Engineering Network in fiscal year 2006.

NASA's large science and engineering challenges are too large to be solved by investigators in a single location. Many of these challenges that utilize modeling and simulation, require investigators, computing resources, and data distributed across the globe. As a result, networking systems will be required to support geographically diverse engineers and remotely located scientists with reliable access to high-end computing, simulation, and massive data distribution. The build of Columbia, a 10,240 processor supercomputer comprised of 20 512-processor SGI Altix nodes, coupled with NASA's unique applications represents a specific resource requiring high speed networking between NASA centers and its partners.

Columbia Local Area Network (LAN)

- 10 Gigabit Ethernet connections to each of Columbia's 20 single system image (SSI) nodes
- Defense in Depth (strategy to delay rather than prevent the advance of an attack) security mechanisms

NASA Research and Engineering Network (NREN) Wide Area Network (WAN)

- 1 to 10 Gigabit Ethernet connections between many NASA centers
- High performance peering to gain access to NASA external modeling and simulation partners

www.nren.nasa.gov

INTEGRATED COLUMBIA SUPERCOMPUTING ENVIRONMENT

As the main computational resource for the entire agency, and as a national leadership-class system, the Columbia supercomputer and the accompanying integrated support environment provide scientists and engineers with the computational resources, simulation tools, and computer science expertise vital to carrying out work critical to NASA mission success, and to making new scientific discoveries for the benefit of humankind.

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

ULTRASCALE REACTIVE ATOMISTIC SIMULATIONS

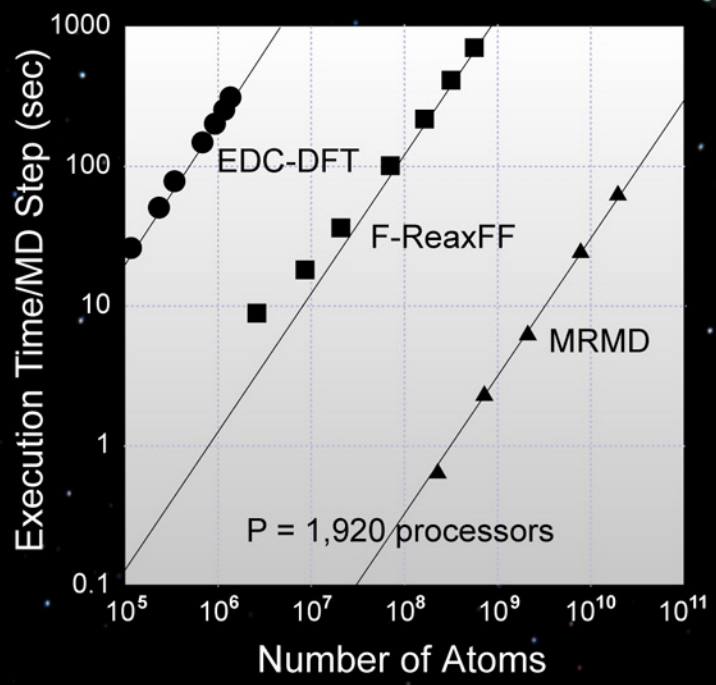

Design-space diagram for reactive and nonreactive molecular dynamics (MD) simulations on 1,920 of Columbia's Itanium2 processors. This figure shows the total execution time per MD step as a function of the number of atoms, N, for three linear-scaling algorithms: quantum-mechanical MD based on the embedded divide-and-conquer density functional theory (EDC-DFT, circles); fast reactive force-field MD (F-ReaxFF, squares); and space-time multi-resolution MD (MRMD, triangles). Lines show ideal O(N) scaling.

Recently, unprecedented scale of chemically reactive atomistic simulations has been achieved on NASA's Columbia supercomputer. These quantum-mechanically accurate, well-validated simulations have broad application to areas that will benefit society, including the understanding and inhibition of stress corrosion cracking, combustion of nano-energetic materials, and micrometeorite impact damage to aerospace vehicles.

Scientists at the University of Southern California, the California Institute of Technology, and NASA Ames Research Center have collaborated to parallelize two newly designed linear-scaling molecular dynamics (MD) algorithms: first principles-based fast reactive force-field (F-ReaxFF) MD; and quantum mechanical MD based on an embedded divide-and-conquer (EDC) density functional theory (DFT) on adaptive multigrids, both based on a tunable hierarchical cellular decomposition framework to expose maximum data localities. Benchmark tests on 1,920 of Columbia's Intel Itanium2 processors have achieved 0.56 billion-atom F-ReaxFF MD and 1.4 million-atom (0.12 trillion grid points) EDC-DFT MD simulations, with the isogranular parallel efficiency as high as 0.953, in addition to 18.9 billion-atom nonreactive space-time multi-resolution MD simulations.

National Aeronautics and Space Administration
www.nasa.gov

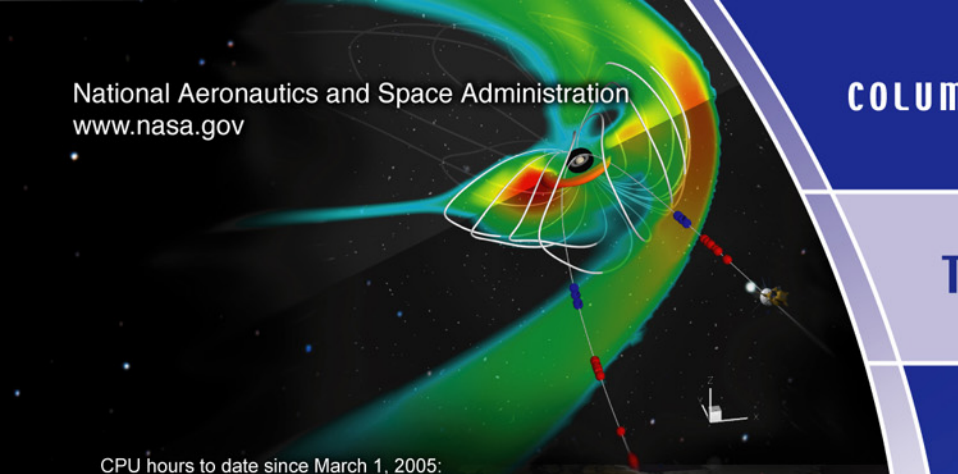

CPU hours to date since March 1, 2005:
Total production: 24.6 M total CPU hours delivered

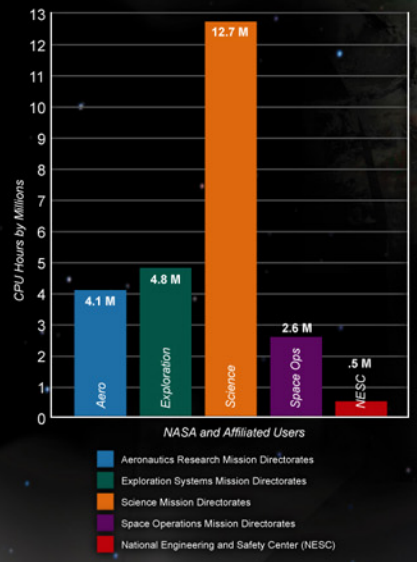

"Presently, Columbia is the only system that can help us meet our goal of achieving faster than real-time performance for modeling and predicting conditions in the near-Earth space environment. Using 256 processors, we can now simulate 600 seconds of physical time in 326 seconds using normal grid resolutions."

– Tamas Gombosi, University of Michigan

THE COLUMBIA SUPERCOMPUTER: YEAR ONE

During its first year of operation, the 10,240-processor SGI Altix supercomputer, Columbia was used heavily to support work critical to the NASA mission, and enabled science and engineering results not previously attainable. Since March 1, 2005, the system has delivered more than 20 million central processing unit hours to NASA and its affiliated users, simultaneously supporting all four of NASA's mission directorates and the NASA Engineering and Safety Center. Some of the work being conducted on Columbia includes:

- Application of engineering analysis tools for assessing viability of space shuttle heat shield survival
- "Real-time" simulation global climate events including tropical storm events in both the Atlantic and Pacific
- Risk assessment of proposed crew exploration vehicle designs during reentry
- Development of cutting-edge nanoseonsors with the capability of ultra-rapid DNA sequencing
- Risk assessment of atomic propulsion devices for future missions to the outer reaches of the solar system

National Aeronautics and Space Administration
www.nasa.gov

COLUMBIA SUPERCOMPUTER

COLUMBIA PERIMETER PROTECTION SYSTEM

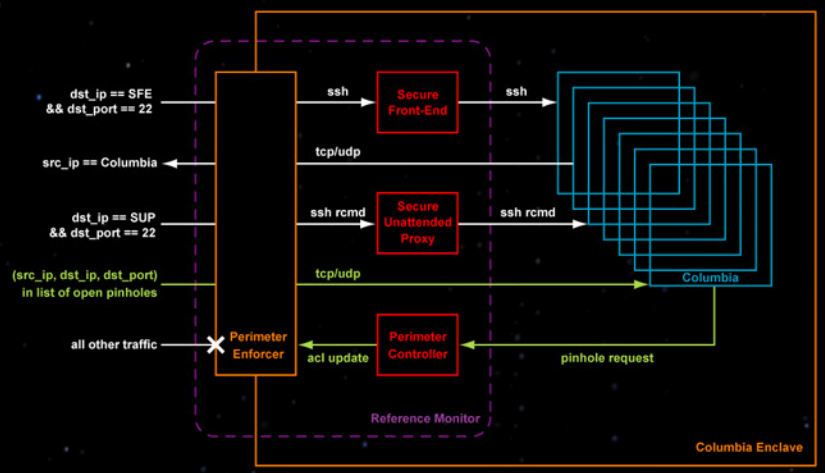

Columbia Perimeter Protection System

The Columbia perimeter protection system is based on the fundamental security principles of a reference monitor and least privilege. Its four key components are:

- A Secure Front-End that mediates all interactive access using two-factor authentication
- A Secure Unattended Proxy that mediates unattended remote commands based on time- and scope-limited private keys obtained using two-factor authentication
- A Perimeter Enforcer that mediates all network traffic and enforces the designated network security policy at line rate
- A Perimeter Controller that processes requests for inbound perimeter access from internal services and opens the perimeter as appropriate by dynamically updating the configuration of the Perimeter Enforcer

Together, these components provide a strong perimeter protection system that:

- Implements a least privilege network security policy
- Allows sites with similar policies to maintain communication
- Supports the highest possible network bandwidths and modern protocols for file transfer and grid computing