



# **NASA Supercomputing: Impact on Present and Future NASA Missions**

**DIESA Symposium**

May 5&6, 2006


**Dr. Walter Brooks**

*NASA Advanced Supercomputing Division  
NASA Ames Research Center, Moffett Field, California*


<http://www.nas.nasa.gov>




# NASA's Engineering and Science Applications



**Current Status**




**Applications Drivers**




**System Balance**

**Pathway to Petaflops**



**Summary**





## Columbia – World Class Supercomputing

- **Provide supercomputing & storage environment to maximize NASA user productivity and achievement**
- **Fast build: Order to full ops in 120 days; ~10X faster than similar systems; dedicated Oct. 2004**
  - Unique partnership w/ industry (SGI, Intel, Voltaire)
  - Built from trusted components (Altix)
- **Immediate impact: Full production and increased capacity during build**
- **Leadership capability: 4th fastest supercomputer in world: 62 Tflops peak**
  - Linpack runs at 51.9 Tflops
- **Effective architecture: Easier application scaling for high-fidelity, shorter time-to-solution, higher throughput**
  - 20 x 512p/1TB shared memory nodes
  - Some apps scaling to 2048p and above
- **Supporting all NASA Mission**
  - >160 projects; >900 accounts; ~150 simultaneous logins
  - Users from across and outside NASA
  - 24x7 support; high Quality of Service
- **Deep and Broad Mission Impact!**



Systems: SGI Altix 3700 and 3700-BX2  
Processors: 10,240 Intel Itanium 2  
Global Shared Memory: 20 Terabytes

Front-End: SGI Altix 3700 (64 proc.)  
Online Storage: 440 Terabytes RAID  
Offline Storage: 6 Petabytes STK Silo

Internode Comm: Infiniband  
Hi-Speed Data Transfer: 10 Gigabit Ethernet  
2048p subcluster: NUMalink4 interconnect



# HEC in NASA

- Overall: NAS will remain the Agency lead for HEC
  - Pathfind capabilities (compute, data, communication, application scaling, data analysis, vis) to support emerging computational requirements.
- NASA has funded NAS at the Agency level, and may add NCCS to that in the future
  - Recognizing that HEC is an essential tool for missions that seek to advance human knowledge, technology, and achievement.
- The U.S. Space Exploration Vision will increasingly leverage HEC
  - Enables the development of new space vehicles and subsystems, materials, understanding of space environment effects on physiology, and more.



## HEC in NASA(cont'd)

- Multi-Center: NASA's two primary HPC centers are in the initial stages of developing a multi-center HPC environment.
- Features include
  - a shared information environment
  - single sign-on
  - workload sharing
  - backup for critical data and functions.
- NLCS: We plan to continue offering ~5% of the cycles on NASA's leadership systems to enable system insight and to support leadership science and technology advancement by other agencies, industry, and academia




## Aerospace Analysis and Design

### Relevance to NASA missions:

- High-fidelity CFD techniques are developed and applied to many NASA aerospace analysis & design problems:
  - Full Space Shuttle Launch Vehicle (SSLV) configuration including orbiter, external tank, solid rocket boosters, and fore / aft attached hardware
  - Exploration vehicles (e.g. CEV, CRV, HLV)
  - Solid Rocket Booster (SRB) blast within Vehicle Assembly Building (VAB)
  - Flame trench analysis for Launch Pad

### Numerical Methods of Research:

- Two high-performance aerodynamic simulation packages were used on Columbia:
  - *Overflow*: To perform analysis at the most important flight conditions and drive the high-fidelity design optimization procedure
  - *Cart3D*: To validate the new design over a broad range of flight conditions
- Combination of *Overflow* & *Cart3D* enables high-fidelity characterization of aerospace vehicle design performance over entire flight envelope



*Pressure contours around full SSLV configuration*

**POC:** Michael Aftosmis, NASA Ames Research Center  
(650) 604-4499, Michael.J.Aftosmis@nasa.gov



# Aerospace Analysis and Design with Petascale Computing

## What can be accomplished with petascale computing?

- Get closer to reaching the full potential of what high-fidelity numerical dynamic simulation techniques have to offer: deliver more optimal designs (larger parametric analysis) and accelerate design cycle (reduced time-to-solution)
  - Aerospace vehicles could potentially be “flown” through the database by guidance and control system designers to:
    - Explore issues with stability and control
    - Determine vehicle’s suitability for various NASA mission profiles

## What are the architecture and algorithm bottlenecks?


- Bandwidth is the biggest problem facing CFD solvers today
  - Important also when synchronizing calculation over large number of processors
- Current high-fidelity NASA CFD codes are processor-speed-bound on Columbia
  - Runs utilizing many hundreds of processors typically use small fraction of available memory
- Significantly higher grid resolution required to improve accuracy of CFD methods
- Investment in scalable solution techniques must be made to replace current block tri- and penta-diagonal solvers used in NASA’s production codes
  - Program models which emphasize good scalability for message passing architectures need to be developed
  - Fundamental change in solution strategy and algorithm are required to make use of petascale systems (multiple cores, high clock rate, huge concurrency)




## Propulsion Subsystem Analysis

### Relevance to NASA missions:

- High-fidelity unsteady flow simulation techniques for design and analysis of propulsion systems for NASA missions (Shuttle - retrofit, CEV - new design)
  - Liquid rocket engine flowliner analysis for Space Shuttle Main Engine (SSME)
  - Reduces cost of space flight—make design decisions early to improve efficiency, performance, and reliability based on computational models of propulsion systems



*Instantaneous surface pressure contours on inducer and flowliner (flow unsteadiness cause flowliner cracks)*



*Flowliner overset grid system: 264 grid blocks of various sizes, total 66 M grid points*

QuickTime™ and a YUV420 codec decompressor are needed to see this picture.

*Computed particle traces colored by axial velocity direction (Blue: Forward; Red: Backward)*


**POC:** Cetin Kiris, NASA Ames Research Center, (650) 604-4485, Cetin.Kiris@nasa.gov



# Propulsion Subsystem Analysis with Petascale Computing

## What can be accomplished with petascale computing?

- Increase the fidelity of the current propulsion subsystem analysis to full-scale, multi-component, multidisciplinary propulsion applications
  - Extend single feedline and 1-SSME capability to 3-SSME and simulate Shuttle flight condition
  - Analyze engine integration for CEV to predict flow-induced vibration
- Model propulsion systems of new / existing launch vehicles to reach full flight rationale:
  - Modeling turbulent combustion in solid rocket boosters
  - Cavitating hydrodynamic pumps in the space shuttle main engine



## What are the architecture and algorithm bottlenecks?

- Improvements to numerical and domain decomposition algorithms for overset grid systems so as to enhance scalability while retaining robustness and convergence
- Scalability bottlenecks inherent in current H/W archs, particularly if built out to 100K procs
  - Require faster interconnection technology
- Development of multi-phase and multi-fluid flow models
  - Very small time scales, highly nonlinear and unsteady, require more physical time steps: implies longer runtimes and lead to intractable problems
- Grid generation (moving & stationary), dynamic adaptation (multiple length and time scales), numerical stability




## fvGCM Columbia Results

*Simulations of  
Hurricane  
Frances and  
Typhoon  
Songda  
simultaneously  
in the Atlantic  
and Pacific  
Oceans (1/8 deg  
resolution)*

QuickTime™ and a  
YUV420 codec decompressor  
are needed to see this picture.

QuickTime™ and a  
YUV420 codec decompressor  
are needed to see this picture.



*Hurricane Ivan's track and  
intensity as forecasted by  
fvGCM 5 days before landfall  
(solid black line), the official  
National Hurricane Center  
(NHC) forecast (dashed blue  
line), and the NHC observed  
positions (solid blue line)*

QuickTime™ and a  
YUV420 codec decompressor  
are needed to see this picture.



# Weather Modeling with Petascale Computing

## What can be accomplished with petascale computing?

- Reliable longer-duration weather forecasts
- Global non-hydrostatic Earth modeling system, including eddy-resolving oceans, cloud-resolving atmosphere, and land models, coupled with chemical and biological components

## What are the architecture and algorithm bottlenecks?

- As resolution increases significantly, horizontal scales become smaller and hydrostatic assumption therefore no longer valid
  - Non-hydrostatic dynamics must be included before trying 4km resolution runs
  - New schemes to better represent physical processes at 1-10km resolutions may be needed
  - Assumptions of physical parameterizations needed at coarse resolutions but may be invalid at resolutions of 10 km and less (e.g. clouds)
- New grid systems (e.g., geodesic grid) are needed due to inefficiencies of (non-uniform) latitude-longitude grid at ultra-high resolutions
- Earth modeling systems typically both computation- and memory-intensive; thus, faster processor (multi-core) and larger cache / local memory would be beneficial




## COLUMBIA SYSTEM FOLLOW ON

We are pursuing the Columbia follow-on as a 4-step process:

- **Market Research**: Market-wide research and briefings from all interested HPC system vendors on their technology roadmaps.
- **Assessment**: Broad assessment of the performance, usability, reliability, and other characteristics of the most promising HPC system architectures.
- **Testbeds**: Procurement and comprehensive testing of a few large testbed HPC systems representing the most promising HPC system architectures.
- **Partnership**: Design, in partnership with typically 1 or 2 selected HPC system vendors, of the next-generation HPC system(s) to be acquired.

*These systems will be selected considering the full cost of implementation and delivered performance for large-scale NASA applications, with a goal of tracking commodity price-performance improvements.*



## Status on Columbia Replacement/Upgrade

- We are at the transition between steps 1 and 2 above.
- The partnership decision is expected early in 2007
- Full-scale system(s) to arrive at NAS starting mid-2007.
- NAS will continue to have requirements to support
  - leadership scale computations
  - a heavy workload of capacity computations,
  - time-critical, mission-critical analyses
    - Shuttle RTF and stewardship, CEV/CLV/etc. design analyses, hurricane forecasting,...



## Storage Systems


- The file system is evolving from direct connect to disks, to a shared file system.
- Our disk storage has risen to ~1PB on the floor, and we will upgrade our archive system to the latest tape technology (500GB drives and tapes) this summer, for at least a 5X increase in bandwidth to tape
- We will have to continue to aggressively enhance the storage and archive system to keep up with the data growth rate.
- Projections for 2010
  - Store 125 TB/day
  - Retrieve 25 TB/day
  - 25 PB archive RAID cache
  - 200 Gb/s across the network, a 200 PB tape media archive
  - 1 billion files and directories.





## What are the networking requirements for the National Leadership Computing System (NLCS) applications?

*Of the four NLCS Awards, two have massive data requirements. These requirements can only be met via a high-speed net such as NREN.*

- Dr. Greg Holland, NCAR, PI
  - .95 M processor hours is about ~8 hours/day every day for a year/512p or ~8 hours/day for three months for 2048p
  - Sample Simulation Parameters for 2048p:
    - 10008 \* 2295 \* 50 = 1148 million gridpoints
    - = 561,000 gridpoints / CPU
    - 1.7 simulated days / wall-clock day
    - 2.1 TB / wall-clock-day
  - 20-88 GBytes per hour of running time
    - 20 GB/hour => 44 Mbps
    - **88 GB/hour => 196 Mbps**
  - ~38-150 TB total NAS => NREN/Abilene =>



- Dr. Marcus Day, LBNL, PI
  - 1.8 M processor hours => ~8 hours/day every day for a year for 1024p or six months for 2048p
  - Every 1/2 hour during 1024p runs
    - 35 GB solution data file to archival
    - 5 GB plot file to archival, analysis, and visualization
  - Additional resolution => +10 GB/hr
  - Either 10-20 GB/hour or 80-90 GB/hour (100 GB/hour) DoE NERSC depending on storage and bandwidth
    - 20 GB/hour => 44 Mbps
    - **90 GB/hour => 200 Mbps**
  - ~40-45 TB NAS => NREN/ESnet => NERSC



- The Columbia local storage system requires that this data be disposed of at the rate of a minimum of 50 MB/sec (400 Mbits/sec) to offload the data from the local filesystems.
- Before the expected upgrade (July '06), the archival system "lou" can sustain 10-15 MB/sec on 4 streams simultaneously, for a maximum ingest rate of 60 MB/sec. Lou is already heavily loaded.
- Therefore, these jobs currently will require access to offsite storage via the WAN in order to provide the required archival rate. This storage will be accessed via NREN, ESnet, and Internet2/Abilene.
- After the lou upgrade, the archival system will support a sustained 1 Gbytes/sec ingest rate.

"A Cloud-resolving Tropical Simulation for Studying Scale-Interactions and Hurricane Variability," led by Dr. Greg Holland, National Center for Atmospheric Research. This project was awarded 0.95 million processor-hours to simulate weather at high, cloud-resolving resolution to determine how moist convection impacts natural climate change, such as hurricane frequency and intensity. This will provide objective information on future hurricane impacts, leading to much improved understanding of how climate both influences, and is influenced by, human activities.

"Interaction of Turbulence and Chemistry in Lean Premixed Combustion," led by Dr. Marcus Day, Department of Energy's Lawrence Berkeley National Laboratory. This project was awarded 1.8 million processor-hours to investigate the dynamics of ultra-lean, turbulent, premixed burners incorporating detailed chemistry and transport for a variety of fuels. The resulting knowledge can accelerate the development of new gaseous burners, which have the potential to dramatically reduce emissions in transportation systems, heat generation, and power generation.

# NASA Research and Engineering Network



	Dedicated 10 GigE (NLR)
	Dedicated 1 GigE (NLR)
	Shared 1 GigE (NLR)
	OC-3 (Legacy Qwest)

NREN Site	
Peering Points	
1 GE	
10 GE	

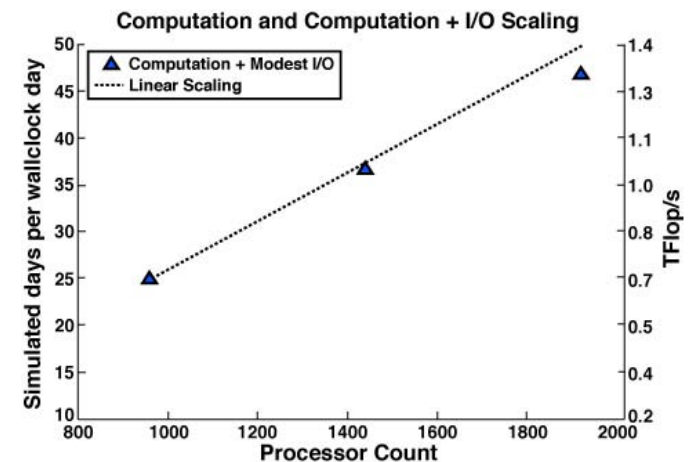
FY'06 Target Configuration



## Visualization

- Concurrent visualization is an increasingly important tool for large data computations.
  - "Concurrent visualization enabled MAP 05 project scientists at Goddard, for the first time, to see simulation results at every time step, which is not possible with traditional batch and post-processing techniques (not possible to write every time step to disk).
- Visualization pipeline dramatically increases temporal resolution, enabling scientists to see high-frequency dynamics in their models not otherwise observable
- Hyperwall II
  - It may be as large as 16x8, with each screen 1600x1200, which would give a total size of 23'x9', and about 250M pixels.
  - This would enable scientists to more quickly analyze and explore much larger and more complex datasets than the current 7x7 hyperwall

QuickTime™ and a YUV420 codec decompressor are needed to see this picture.





# Colliding Black Hole Galaxies

QuickTime™ and a  
TIFF (Uncompressed) decompre  
are needed to see this picture.

- This visualization shows what Einstein envisioned. Researchers crunched Einstein's theory of general relativity on the Columbia supercomputer at the NASA Ames Research Center to create a three-dimensional simulation of merging black holes.
- This was the largest astrophysical calculation ever performed on a NASA supercomputer.
- The simulation provides the foundation to explore the universe in an entirely new way, through the detection of gravitational waves.
  - (7.4 Mb - no audio). Credit: Henze, [NASA+ High resolution image 1 + High resolution image 2+ High resolution image 3+ High resolution image 4+ High resolution image 5](#)



## Summary

- Current Columbia system is heavily utilized in all of NASA's Science and Engineering
- Maintaining systems balance is difficult as users scale up to fully utilize Columbia
- National Leadership Computing System Awards are providing “bleeding edge “ users and simulations that give us insight into the problems our next system will face
- We are well into the process of technology evaluation in anticipation of a major system upgrade in 2007