



National Aeronautics and Space Administration

World-Class Supercomputer System, “Columbia” Revolutionizes NASA Science and Engineering Capabilities

*HPC Online Conference and Expo
October 6, 2005*

Dr. Walt Brooks
Chief, NASA Advanced Supercomputing Division
NASA Ames Research Center
Moffett Field, CA 94035
<http://www.nas.nasa.gov>





Success Through Partnerships

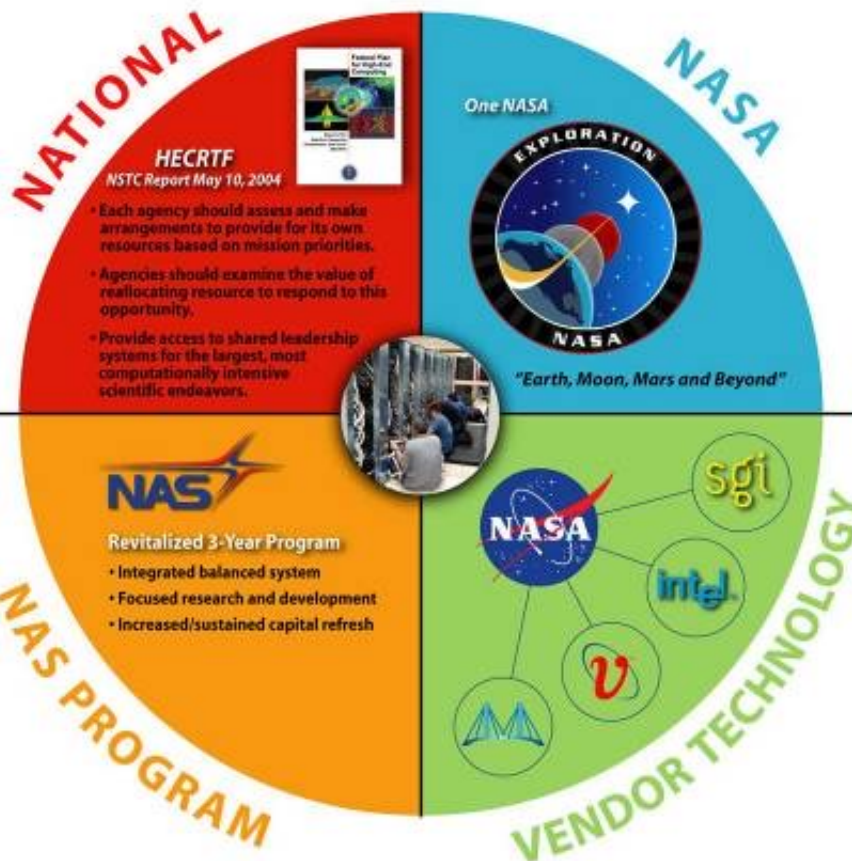
"LEADERSHIP" SYSTEMS



MISSION CRITICAL APPLICATIONS



NASA's Mission Directorates



INTEGRATED ENVIRONMENTS

Next Generation Codes & Algorithms	Rapid Deployment
<ul style="list-style-type: none"> OVERFLOW - STS 107 Honorable Mention, 1998 INS3D - Pump Analysis NASA Software of the Year, 1994 CART3D - STS 107 NASA Software of the Year, 2002 	<ul style="list-style-type: none"> - Systems - Networks - Archives
Data Analysis & Visualization	Computer Science Tools & Expertise
<ul style="list-style-type: none"> - Hyperwall visualization systems - Process large data sets - Interactive real-time visualization 	<ul style="list-style-type: none"> - Program development tools - Automatic multilevel parallelization - Scaling, porting, and optimization

SUPERCOMPUTING





Success Through Partnerships: Leadership Systems as National Resources

"LEADERSHIP" SYSTEMS



SGI Altix




IBM Blue Gene



National LambdaRail




Cray X1



AMD Opteron

NATIONAL



HECRTF


NSTC Report May 10, 2004

- Each agency should assess and make arrangements to provide for its own resources based on mission priorities.
- Agencies should examine the value of reallocating resource to respond to this opportunity.
- Provide access to shared leadership systems for the largest, most computationally intensive scientific endeavors.



Success Through Partnerships: NASA's Mission Critical Applications

MISSION CRITICAL APPLICATIONS



NASA's Mission Directorates




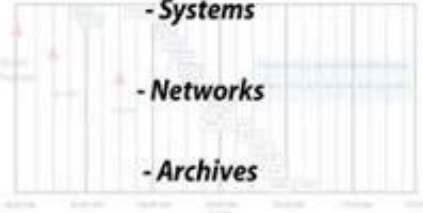




Success Through Partnerships: Vendor Technology and Supercomputing





Success Through Partnerships: Integrated Simulation Environments

INTEGRATED ENVIRONMENTS

Next Generation Codes & Algorithms	Rapid Deployment
 OVERFLOW - STS 107 <i>Honorable Mention, 1998</i>	 <ul style="list-style-type: none">- Systems- Networks- Archives
 INS3D - Pump Analysis NASA Software of the Year, 1994	
 CART3D - STS 107 NASA Software of the Year, 2002	
Data Analysis & Visualization	Computer Science Tools & Expertise
 <ul style="list-style-type: none">- Hyperwall visualization systems- Process large data sets- Interactive real-time visualization	 <ul style="list-style-type: none">- Program development tools- Automatic multilevel parallelization- Scaling, porting, and optimization



Revitalized 3-Year Program

- Integrated balanced system
- Focused research and development
- Increased/sustained capital refresh

NAS PROGRAM




Integrated Support for High-Performance Modeling and Simulation

The Columbia supercomputer provides a peak of 62 teraflops to the user community. The NASA Advanced Supercomputing (NAS) facility, where the Columbia system is housed, provides integrated support for the full life-cycle of high-end science and engineering applications.

NASA Scientists and Engineers

Scientists and engineers set up computational problems, choosing effective codes and resources to solve NASA's complex mission problems.



NASA'S MISSION DIRECTORATES

Performance Optimization



NAS software experts utilize tools to parallelize and optimize codes, dramatically increasing simulation performance while decreasing turnaround time.

Data Analysis and Visualization




NAS experts apply advanced data analysis and visualization techniques to help scientists explore and understand large data sets.

Supercomputers, Storage and Networks



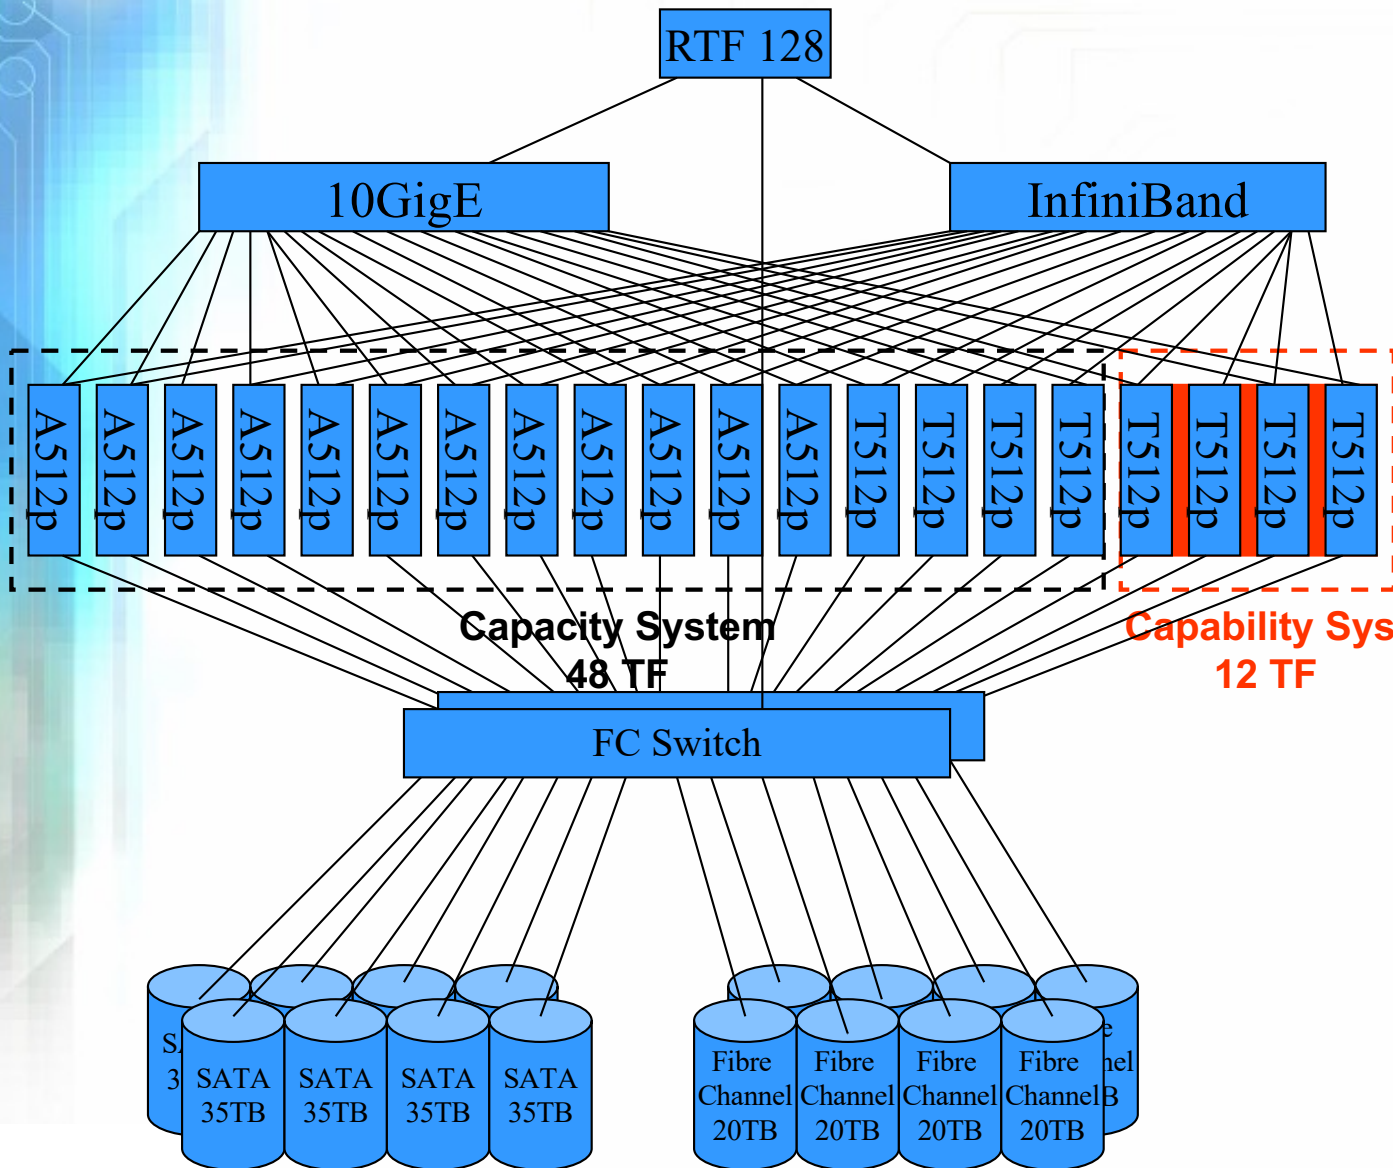
The NAS supercomputer environment (hardware, software, network and storage) is used to execute the optimized code to solve NASA's large computational problems.



November 2003
The Basic Building Block - World's First SGI 512
supercomputer with Intel Itanium 2 processors.
(Named "Kalpana" in honor of NASA Ames scientist and
Columbia astronaut, Dr. Kalpana Chawla)



Columbia Configuration



Front End

- 128p Altix 3700 (RTF)

Networking

- 10GigE Switch 32-port
- 10GigE Cards (1 Per 512p)
- InfiniBand Switch (288port)
- InfiniBand Cards (6 per 512p)
- Altix 3700 2BX 2048 Numalink Kits

Compute Node (Single System Image)

- Altix 3700 (A) 12x512p
- Altix 3700 BX2 (T) 8x512p

Storage Area Network


- Brocade Switch 2x128port

Storage (440 TB)

- FC RAID 8x20 TB (8 Racks)
- SATARAID 8x35TB (8 Racks)

PROJECT COLUMBIA

SUPPORTING NASA'S MISSION DIRECTORATES





Project Columbia

Dedicated October 2004





60 TERAFLUPS in 120 DAYS

NASA's factor of 10 Supercomputer:

- Developed 10X faster than previous systems (4 months versus 40 months).
- 10X less Expensive (\$50M versus \$500M).
- 10X more Productive (600 users versus 60 users).

- **Program**

- Beginning May 18, 2004, obtain all of the necessary approvals and procure the system by June 18, 2004.

- **Physical Plant**

- Make all of the necessary power and cooling changes to run *Columbia*.
- Reconfigure and retrofit decommissioned water cooling loop.

- **Production**

- SGI build and deliver 19 Altix 512s in less than 4.5 months, including the first Altix 3700BX2.

- **Integration**

- Assemble and test 20 x 512p with GigE and Infiniband connectivity.

- **Continuous production**

- Continue NASA science and engineering in support of NASA Missions.


- **Provide a national capability**

- Build and utilize the first shared-memory 2048.



Columbia16 Marks First 3700BX2

- Density of CPUs is twice earlier generation.
- Necessary racks for 512 reduced from 16 to 11.
- Supports shared memory across 2,048 processors.
- Utilizes chilled water-cooling in the doors to reduce load on existing cooling units on the floor.
- Brings “Earth Simulator” class computing to 30x30 room.





60 TERAFLUPS in 120 DAYS!

NASA's Factor of 10 Supercomputer:

- Developed 10X faster than previous systems (4 months versus 40 months).
- 10X less Expensive (\$50M versus \$500M).
- 10X more Productive (close to 700 to 800 users versus previous 60 to 100 users).



NASA Vision and Mission

Vision:

- To improve life here,
- To extend life to there,
- To find life beyond.

Mission:

- To understand and protect our home planet,
- To explore the universe and search for life,
- To inspire the next generation of explorers,
... as only NASA can.

President's Information Technology Advisory Committee quote:


“Information Technology will be one of the key factors driving progress in the 21st century - it will transform the way we live, learn, work, and play.

Advances in computing and communications technology will create a new infrastructure for business, scientific research, and social interaction.”






Debris Transport Analysis






Debris Code Analysis Options

LH ₂ Flange Foam	
Mass	= 0.023 lb _m
r_{foam}	= 2.34 lb/ft ³
$V_{\text{pop-off}}$	= 113 ft/sec



Deterministic
Zero Lift Trajectory + Range of Initial Velocities



Probabilistic
Zero Lift Trajectory + Crossrange Cone



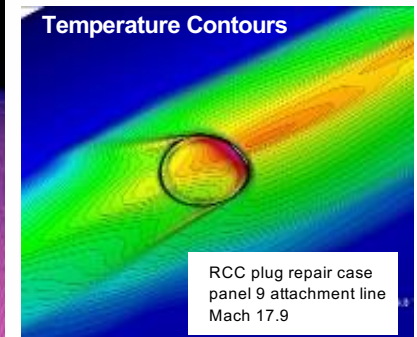

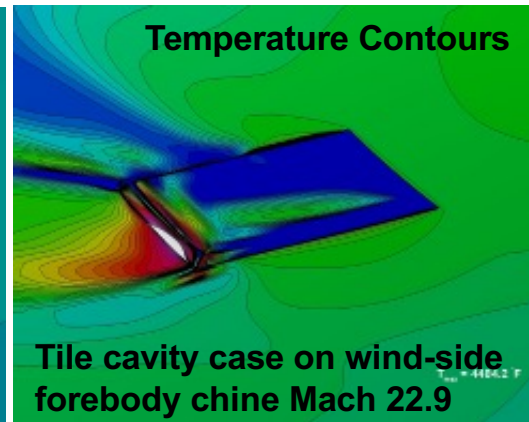


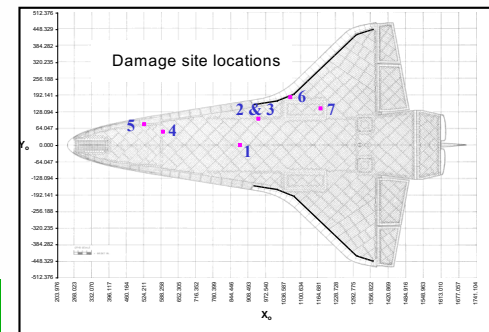
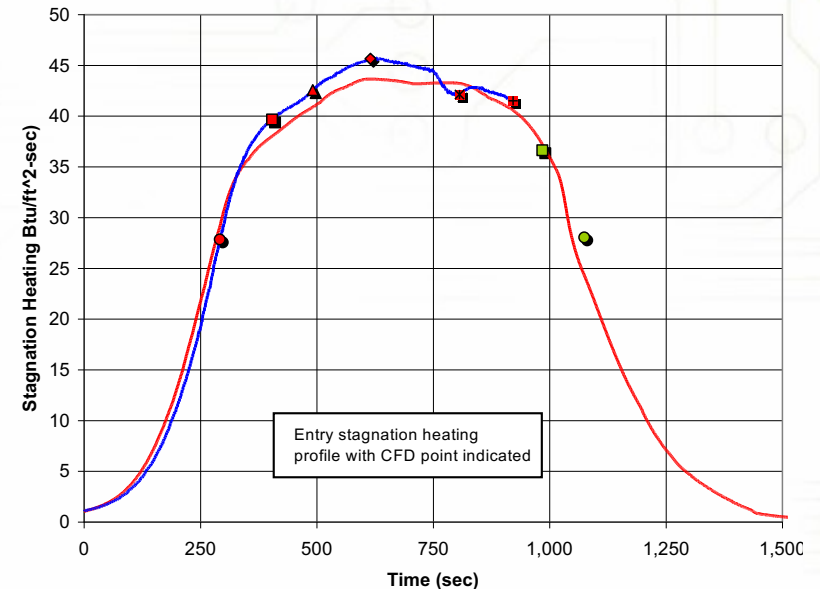
Rapid Aerothermal Analysis Demonstration

• Objectives:

- Demonstrate a new rapid aerothermal CFD analysis capability
- The new capability shall permit near real-time analysis of observed Orbiter damage during flight
- The capability would provide an alternate high-fidelity evaluation of local heating bump factors calculated from engineering codes
- Required capability: 10 damage sites at 10 trajectory points (> 100 solutions) assessed in under 24 hours
- Capability prior to STS-107 accident ~ 1 solution in 3 weeks

• Results: **Success!!!**


- 8 (512 processor) nodes of the *Columbia* supercomputer were used for 24 hours
- 12 different computational meshes were generated for 10 different damage/repair sites
- 10 trajectory points calculated for each damage/repair site
- More than 100 high-fidelity Navier-Stokes calculations performed
- Some unsteadiness observed in some tile cavity CFD solutions
- Some pre- and post processing improvements needed





Improved Hurricane Track Prediction

fvGCM Code Simulations - Hurricane Francis 08/04




(Resolution: 1/4th of a degree)

PI: Bowen Shen, Goddard Space Flight Center



Higher Resolution Hurricane Track Prediction fvGCM Code Simulations - Hurricane Francis 09/04

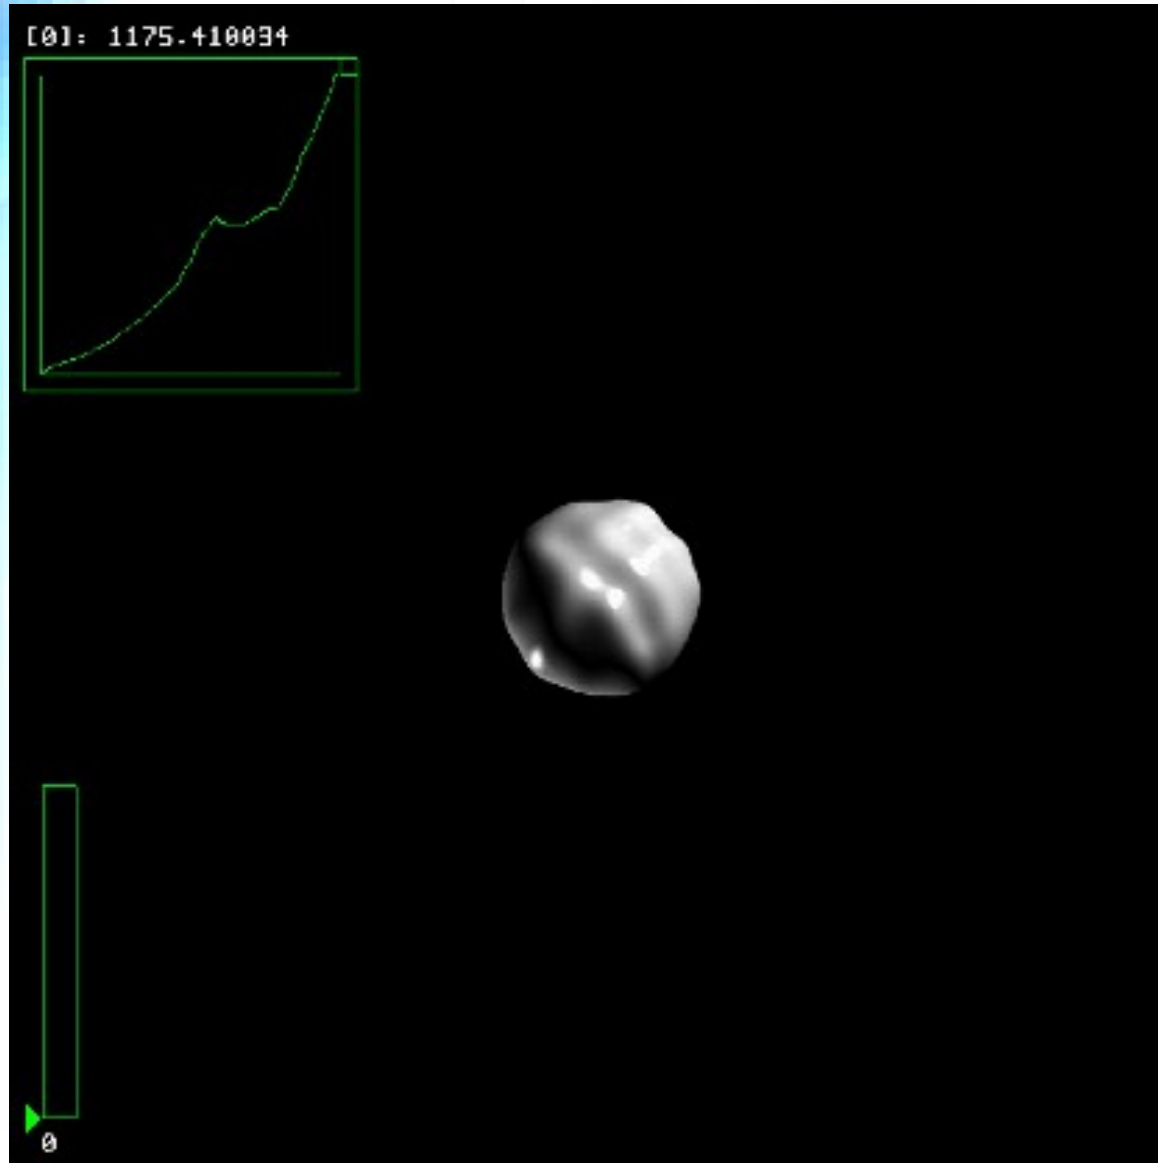


(Resolution: 1/12th of a degree)


PI: Bowen Shen, Goddard Space Flight Center



Type Ia Supernova Detonation Simulation



Approximate Relative CPU Performance




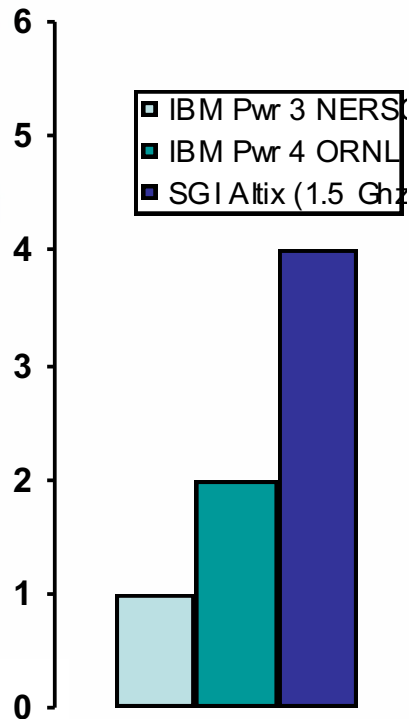
PI: Woosley, Zingale, etal, UC Santa Cruz



Large-Scale Structure of the Universe

Simulation of large-scale dark matter distribution throughout the universe. Problem domain 120M parsec cube with 256-cubed particles. Small bright knots ~ 10 Milky Way masses.

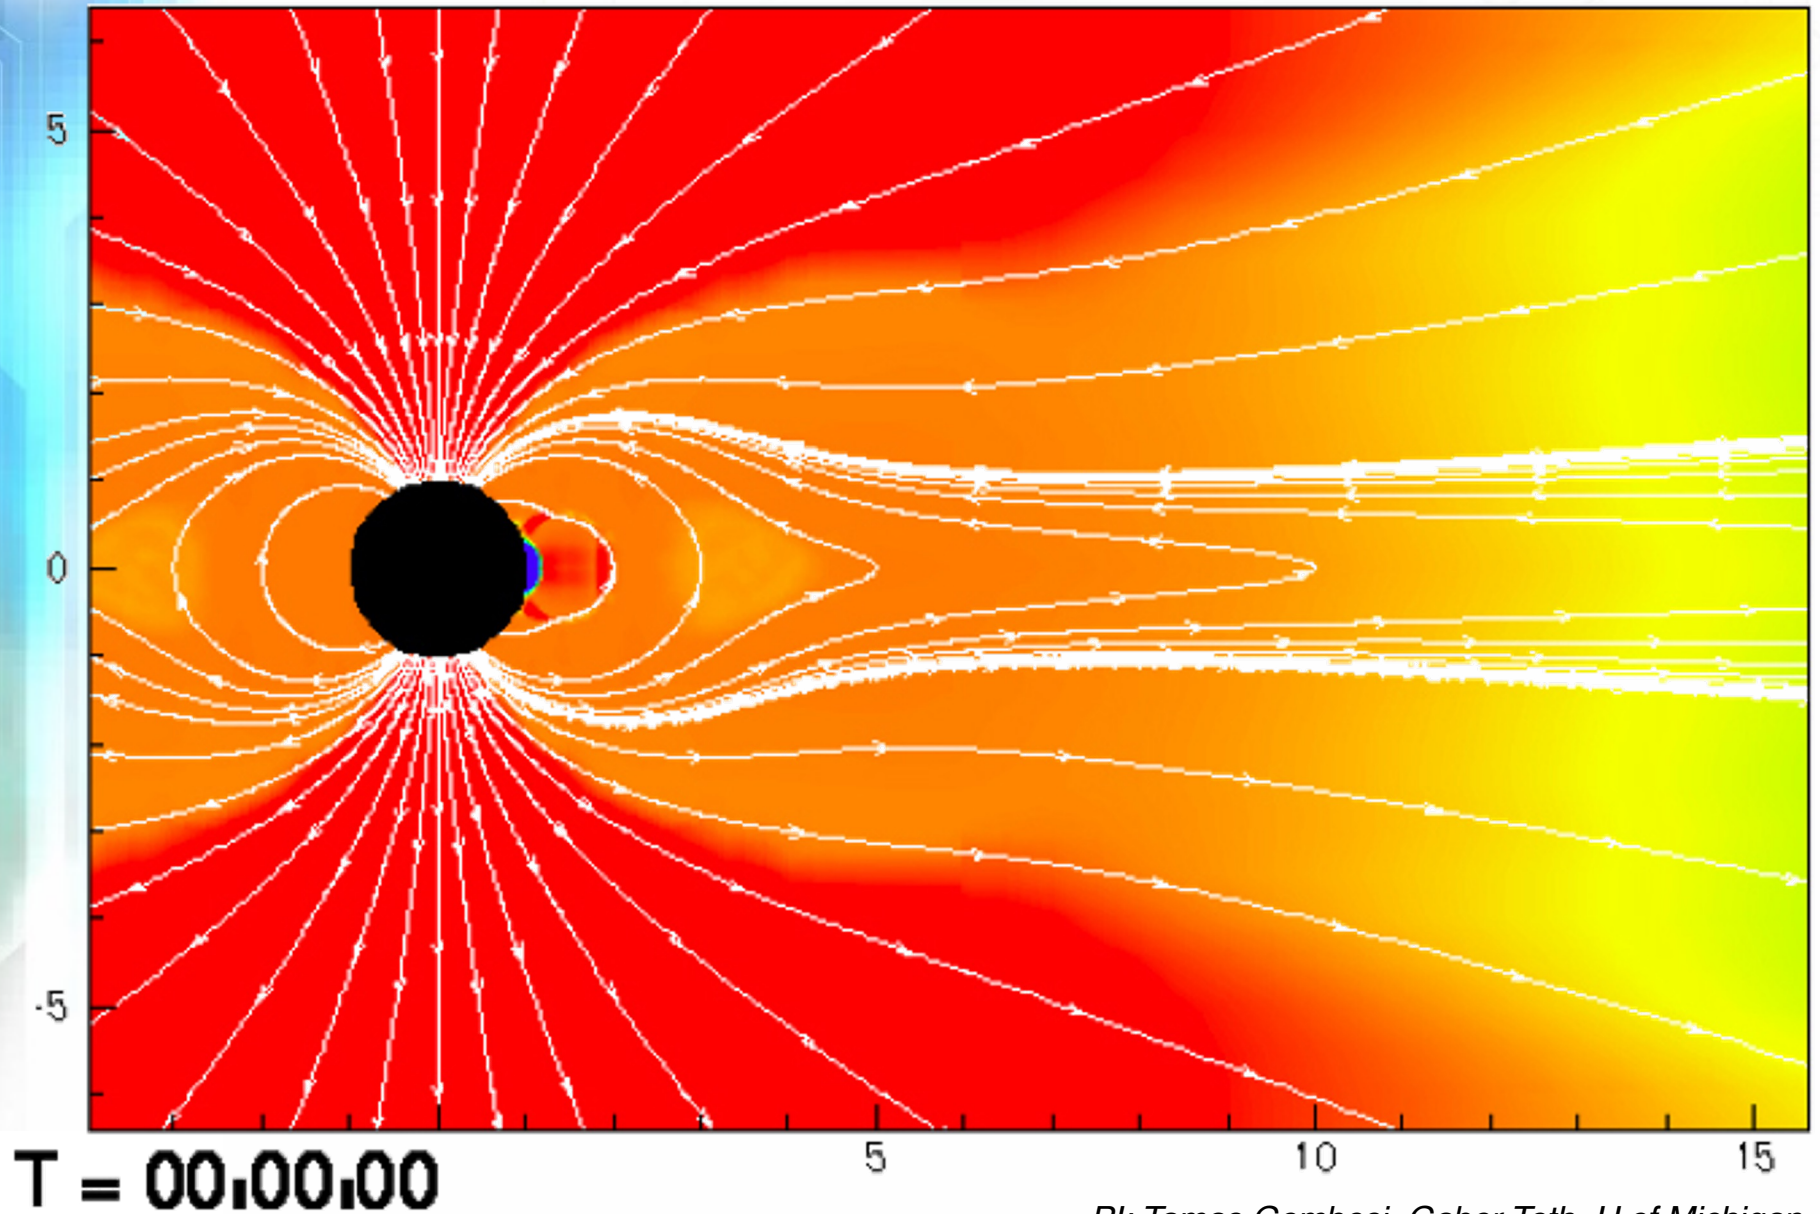
Approximate Relative CPU Performance



PI: Primak, Allgood, etal, UC Santa Cruz



Space Weather: Radiation Storms



PI: Tamas Gombosi, Gabor Toth, U of Michigan



hyperwall-1: Parameter Studies





So Where Are We?

- **The Science**

- Production CFD codes executing 100x C90 numbers of just a few years ago.
- Throughput 100x (or more) above that of a few years ago.
- Earth/Space Science codes executing 2-4x faster than last year's best efforts, 100x throughput over last year's efforts.

- **The Systems (1997 - present)**

- New expanded shared memory architectures: First 256, 512, and 1024 CPU Origin systems. First 256p, 512p Altix SSI systems.
- First 2048p NUMAlinked 512p Altix cluster.

- **The Future?**

- Expanded Altix SSI to 4096?
- Expanded Altix NUMAlinked clusters to 16Kp?
- Serious upgrades to CPUs

