

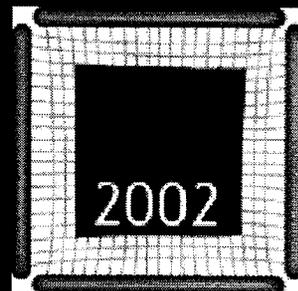
An Overview of Recent Supercomputing Applications at JPL

Chris Catherasoo, *Jet Propulsion Laboratory*

ABSTRACT: *The Jet Propulsion Laboratory (JPL) is responsible for much of NASA's unmanned space exploration activities including astronomy, astrophysics, earth observations and planetary explorations. The Supercomputing and Visualization Facility at JPL provides the various projects with a state-of-the-art supercomputing and visualization environment for modelling and data analysis. Our facilities include a 16-processor Cray SV1-1A, an SGI Origin2000 (with 128 processors and 6 graphics pipes), a StorageTek tape silo and a Powerwall room with seven projectors. Examples of recent work that has been done on the JPL supercomputers include Martian rover navigation simulations, Shuttle Radar Topography Mission (SRTM) data analysis, ocean modeling and spacecraft engine and heat-shield simulations. This presentation will cover some of these recent activities.*

An Overview of Recent Supercomputing Applications at JPL

Chris Catherasoo
20 May 2002

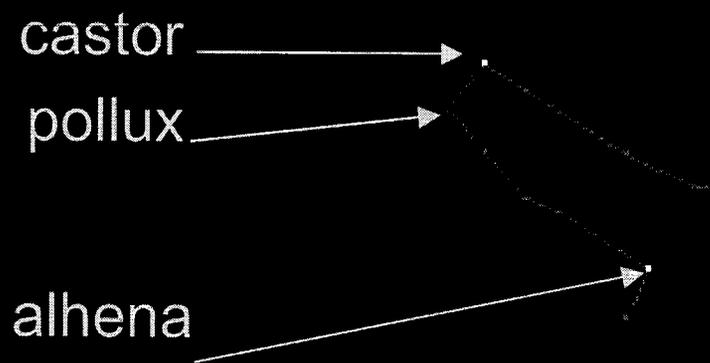




Outline

- Overview
 - Supercomputing overview
- Resources
 - Supercomputers, visualization facility
- Applications
- Conclusion

Overview

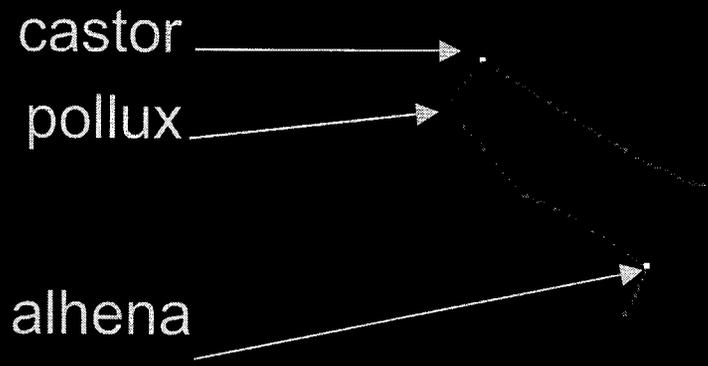




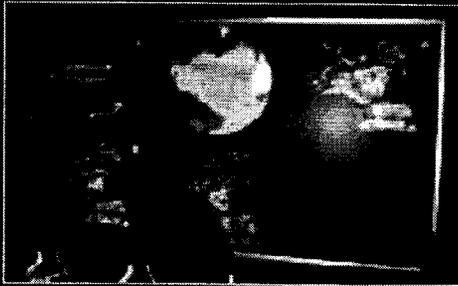
JPL Supercomputing Overview

- The JPL Supercomputing and Visualization Facility (SVF) provides a state-of-the-art, high-performance computing environment for solving computationally-intensive problems.
- It supports all local scientists and engineers, and their collaborators at other institutions.
- These capabilities enable the development of expertise in supercomputing applications in support of JPL's mission, for the NASA Offices of Space Science, Earth Science and Aeronautics.

Resources

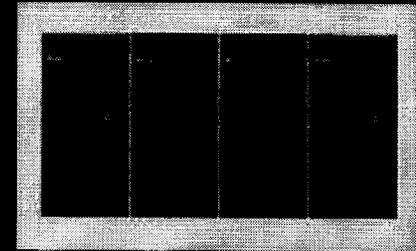


Supercomputing and Visualization Resources



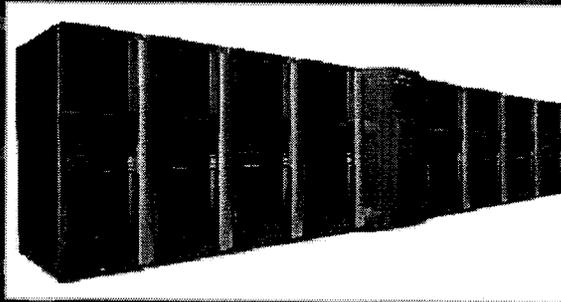
PowerWall

- 6 DLV rear projectors
- 1 DLP rear projector
- 7.8 MegaPixel display
- Wireless switching control



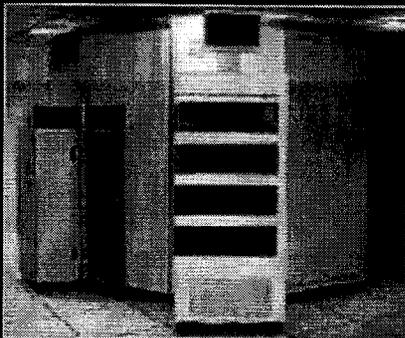
SGI Dual Power Onyx

- 12 R10000 processors
- 4 GB main memory
- 4 Infinite Reality graphics pipes
- 3.2 TB FC RAIDS disks



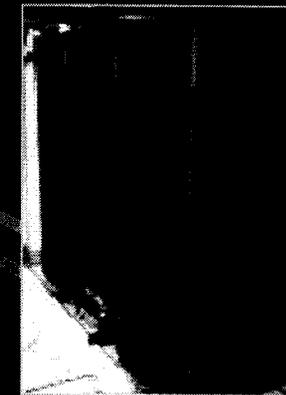
SGI Origin2000 Reality Monster

- 128 R12000 300 MHz processors
- Peak speed: 76.8 GFLOPS
- 64 GB main memory
- 72 GB local disk space
- 6 Infinite Reality II graphics pipes
- 3.2 TB FC RAIDS disks



StorageTek Silo

- 6000 tape capacity
- 60 GB per cartridge
- Maximum capacity: 360 TB
- Average access time: 11 seconds



Cray SV1-1A

- 16 300 MHz CPU Parallel Vector Processor
- 1200 MFLOPS/CPU
- Peak Speed: 19.2 GFLOPS
- 8 GB main memory
- 480 GB local disk space

SGI Origin2000 –

Machine Characteristics

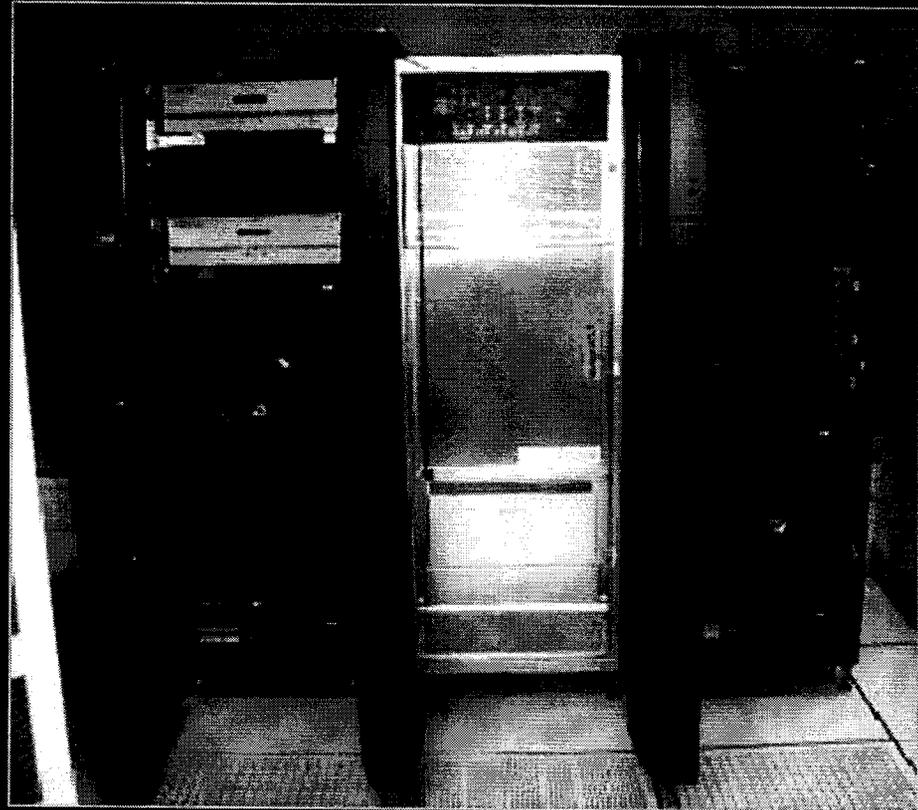
- Cache-coherent nonuniform memory access (ccNUMA)
- 128 R12000 processors
 - 300 MHz
 - 600 MFLOPS/processor
- 76.8 GFLOPS peak speed
- 64 GB main memory
- 72 GB local disk
- 3.2 TB FibreChannel RAID
- 6 Infinite Reality II graphics engines
 - Maximum of two channels per graphics engine
- Networking
 - 4 HiPPI
 - 5 100BaseT
 - 4 ATM OC12
 - 4 1000BaseT



Cray SV1-1A –

Machine Characteristics

- Parallel vector processor (PVP)
- 16 CPUs
 - 300 MHz
 - 1200 MFLOPS/CPU
- 19.2 GFLOPS peak speed
- 8 GBytes memory (1 GWord)
 - CMOS DRAM memory
 - 70-ns access speed
- 480 GBytes disk(s)
 - 10 MB/sec access speed per channel





Powerwall Visualization Laboratory

Characteristics

- Located in supercomputing center
- Rear-projection system with a 144" x 84" screen
- Six Christie Digital 1280 LCD DLV projectors
 - A single high-resolution tiled image displayed at 3840 x 2048 pixels
 - Six independent images displayed at a standard resolution of 1280 x 1024 pixels each
- One Digital Projection 6sx DLP projector
 - Resolution of 1280 x 1024 pixels
 - 5000 lumens
- Powered by the two SGI PowerOnyx machines or the SGI Origin2000



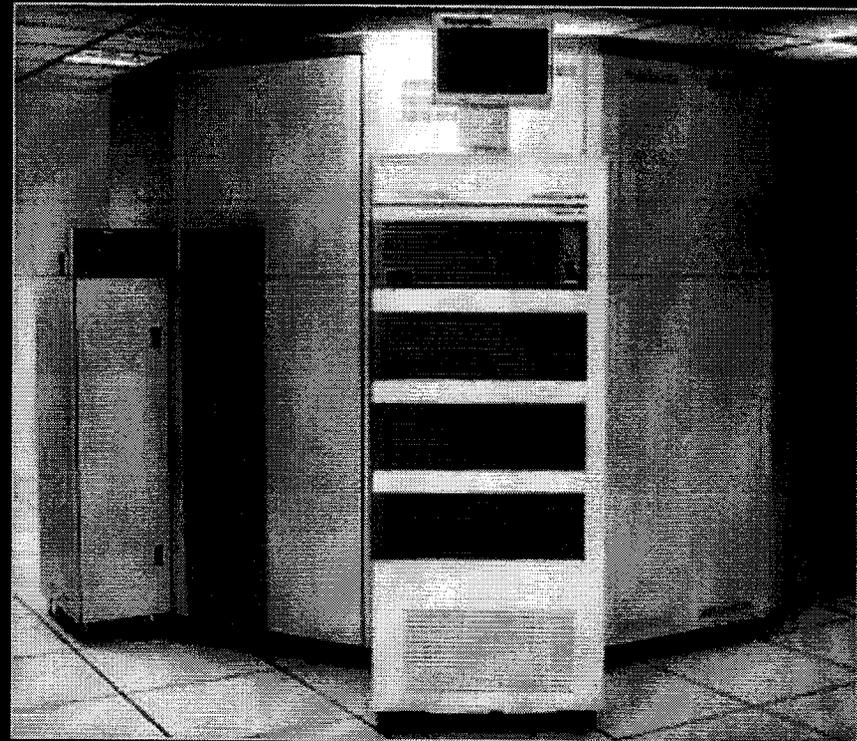


Storage Technology 9310 Automatic Cartridge System

NASA

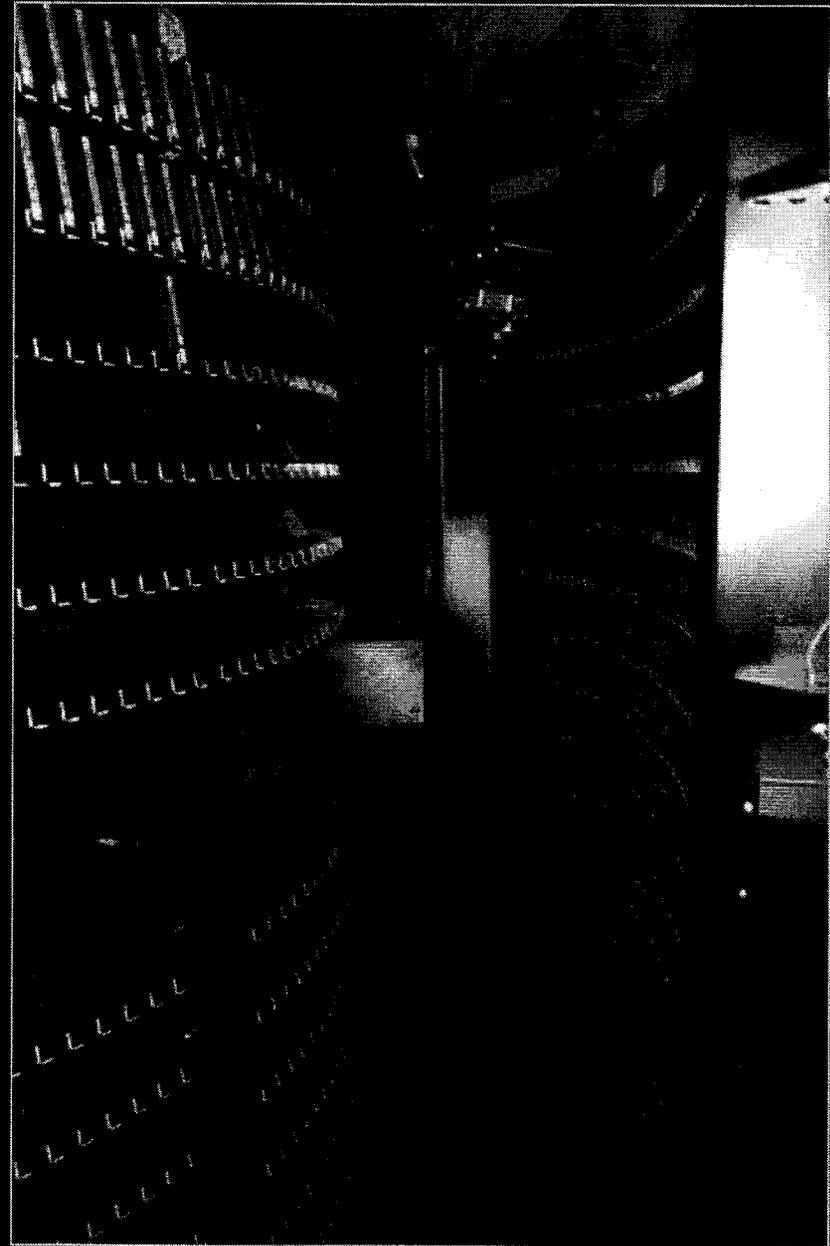
Silo Characteristics

- 6000 tape capacity
- 60, 50 or 20 Gigabytes per cartridge
- Maximum capacity: 360 Terabytes
- Average access time: 11 seconds
- Three 9840 fast access drives
 - Sustained transfer rate: 10 MB/sec per drive
 - Tape rewind speed: 16 sec
- Three helical-scan Redwood tape drives
 - D-3 helical-scan cartridges
 - Sustained transfer rate: 11 MB/sec per drive
 - Tape rewind speed: 89 sec
- Uses automatic data archival software

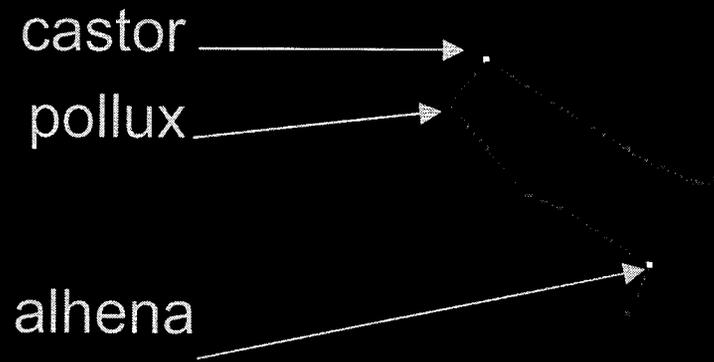


Data Migration Facility

- Automatic data storage archival software
- 472 GByte disk cache
- Available from Nebula and Alhena
- Stage files in */silo/username*
- System will migrate */silo* automatically
- Virtual storage facility
- */u1 - /u5* backed up nightly



Applications





Applications

- Physical sciences
 - Earth sciences
 - Planetary sciences
 - Fluid dynamics modeling
 - Strength of materials modeling
 - Spacecraft design
 - Radar design
 - Visualization
-
- Visit our success story slides at:

<http://sc.jpl.nasa.gov/>



Powerwall Visualization Laboratory

Characteristics

- Located in supercomputing center
- Rear-projection system with a 144" x 84" screen
- Six Christie Digital 1280 LCD DLV projectors
 - A single high-resolution tiled image displayed at 3840 x 2048 pixels
 - Six independent images displayed at a standard resolution of 1280 x 1024 pixels each
- One Digital Projection 6sx DLP projector
 - Resolution of 1280 x 1024 pixels
 - 5000 lumens
- Powered by the two SGI PowerOnyx machines or the SGI Origin2000



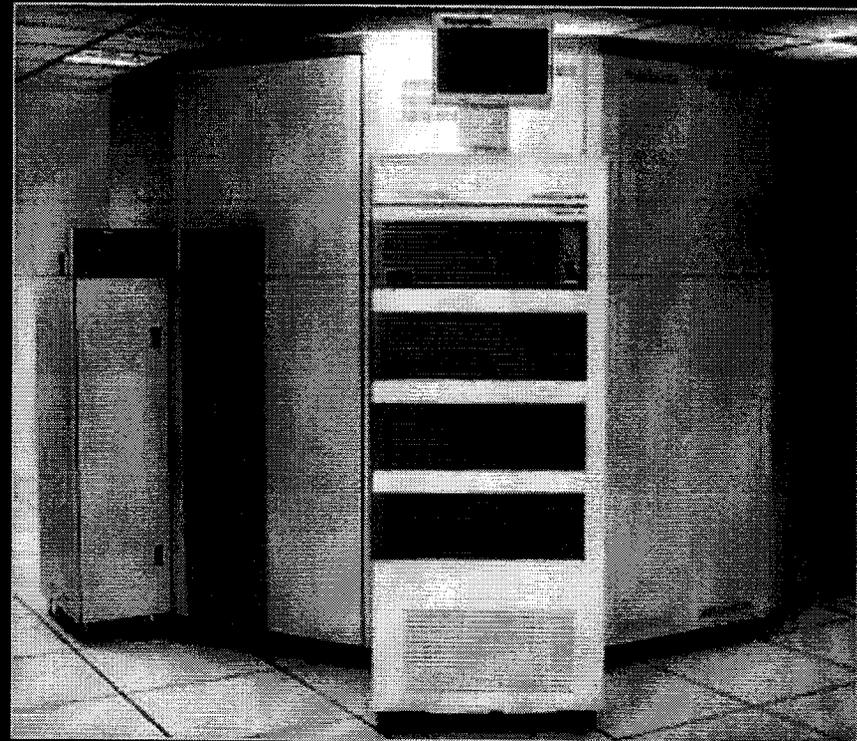


Storage Technology 9310 Automatic Cartridge System

NASA

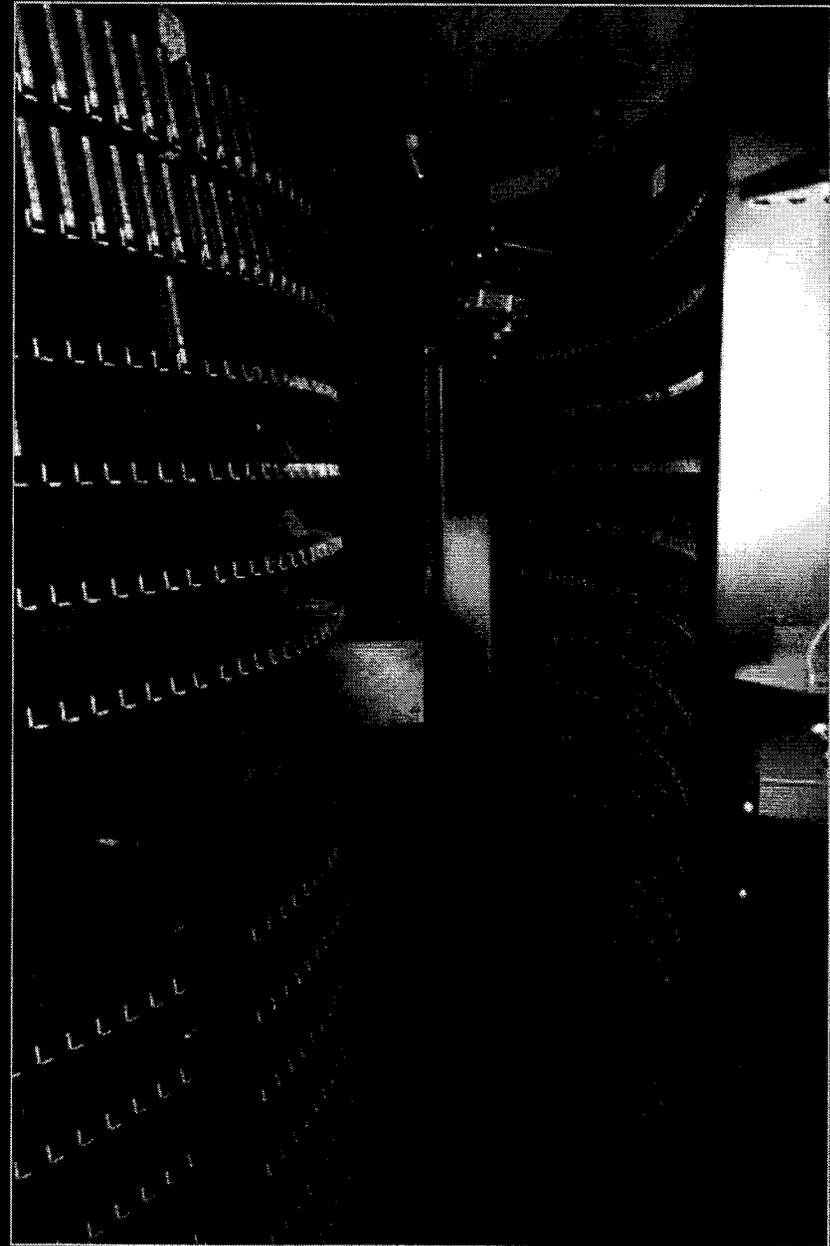
Silo Characteristics

- 6000 tape capacity
- 60, 50 or 20 Gigabytes per cartridge
- Maximum capacity: 360 Terabytes
- Average access time: 11 seconds
- Three 9840 fast access drives
 - Sustained transfer rate: 10 MB/sec per drive
 - Tape rewind speed: 16 sec
- Three helical-scan Redwood tape drives
 - D-3 helical-scan cartridges
 - Sustained transfer rate: 11 MB/sec per drive
 - Tape rewind speed: 89 sec
- Uses automatic data archival software

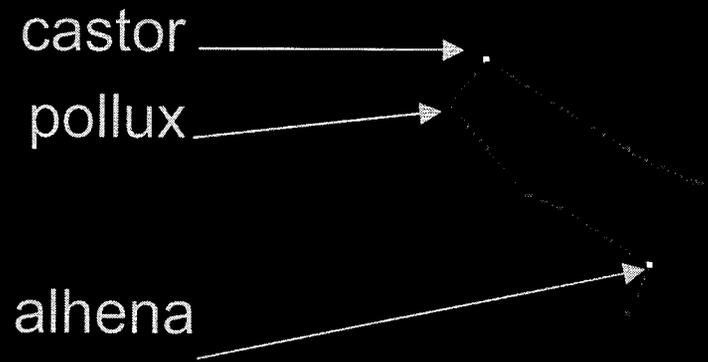


Data Migration Facility

- Automatic data storage archival software
- 472 GByte disk cache
- Available from Nebula and Alhena
- Stage files in */silo/username*
- System will migrate */silo* automatically
- Virtual storage facility
- */u1 - /u5* backed up nightly



Applications





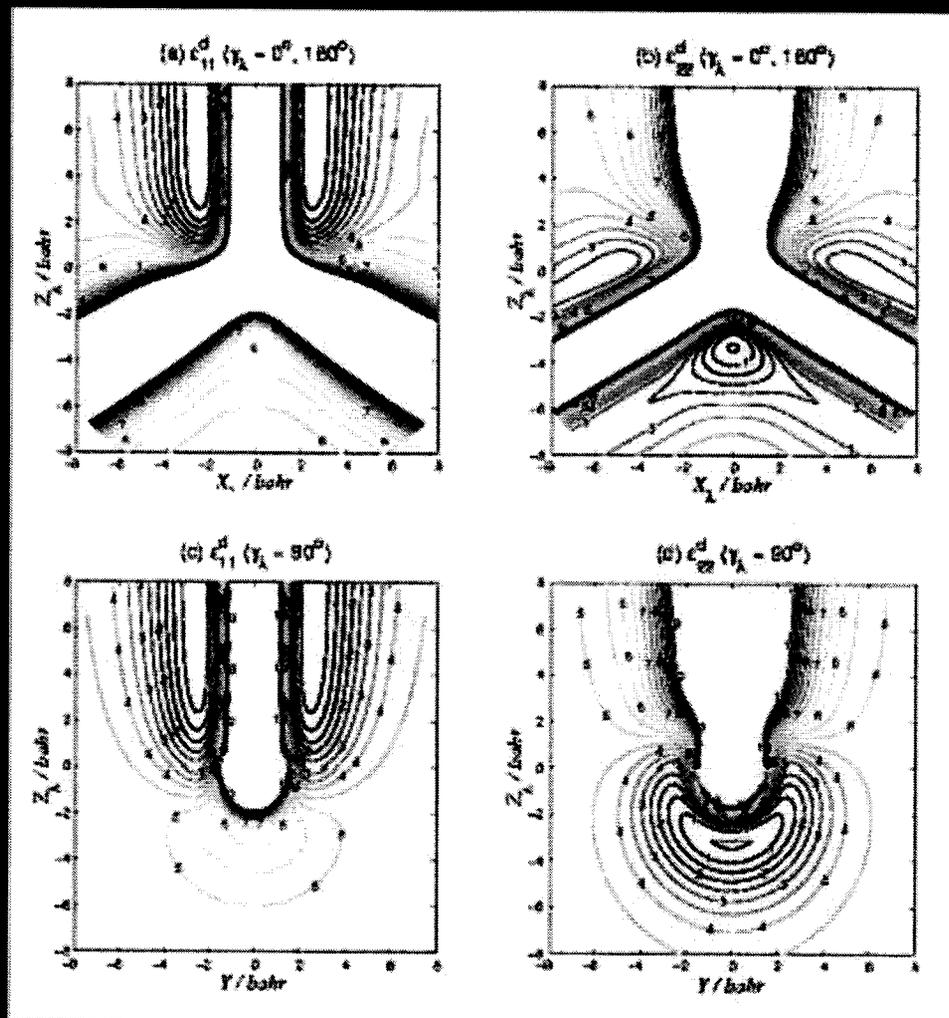
Applications

- Physical sciences
 - Earth sciences
 - Planetary sciences
 - Fluid dynamics modeling
 - Strength of materials modeling
 - Spacecraft design
 - Radar design
 - Visualization
-
- Visit our success story slides at:

<http://sc.jpl.nasa.gov/>

Optimal Diabatization of Multiple Electronic States for Quantum Reaction Dynamics

Ravinder Abrol and Aron Kuppermann, Caltech



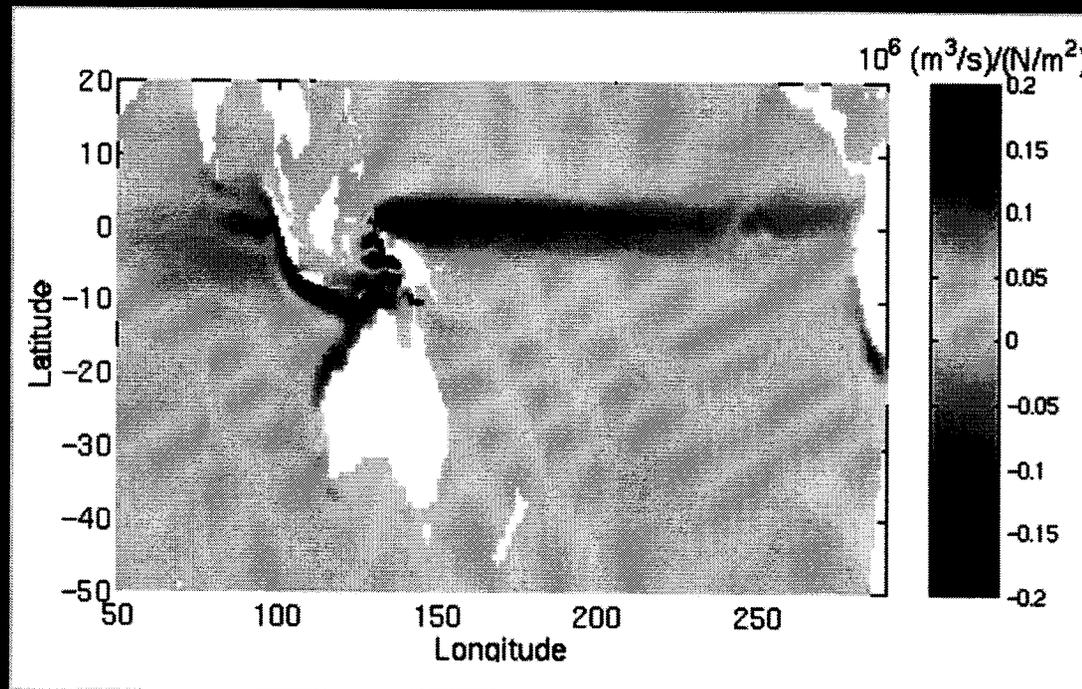
Diagonal diabatic energy contours for the ground state (ϵ^d_{11}) and the first excited state (ϵ^d_{22}) of the H atom to the H_2 molecule reaction. Figures (a) and (b) correspond to a collinear approach of the H atom to the H_2 molecule, while (c) and (d) correspond to a perpendicular approach.

- This study obtains an optimal diabaticization of the two lowest electronic states of the H atom to H_2 molecule reaction.
- It uses an adiabatic-to-diabatic transformation to eliminate the contribution of the longitudinal part of the nuclear motion Schrödinger equation, leaving the contribution of the transverse part.
- Finding the optimal solution involved solving a three-dimensional Poisson equation in the triatomic system's three internal nuclear coordinates (ρ , θ and ϕ_λ), using both Dirichlet and Neumann boundary conditions.
- The number of points used in the ρ , θ and ϕ_λ directions was 513, 257 and 65 respectively, for a total of about 8.6 million grid points.
- The calculations were run on the JPL Cray SV1-1A, and used the MUDPACK library. The time per run was about 30 minutes.
- The study provides insight into the dynamical importance of excited states, and conical intersections in small-molecule atmospheric reactions.
- This work was partially supported by NSF.

Sensitivity of Ocean Currents to Wind Forcing

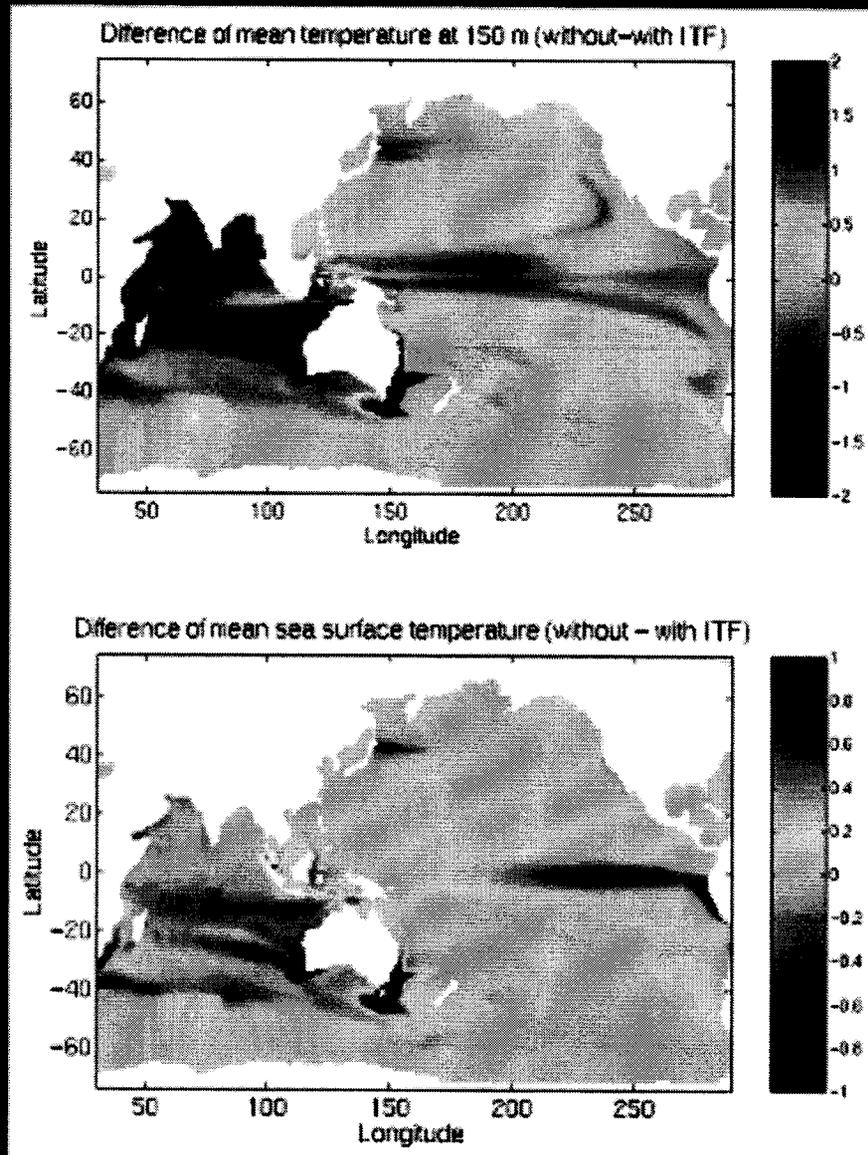
T. Lee, B. Cheng and R. Giering, JPL

- This is a study on how interannual winds change the oceanic transport of the Indonesian throughflow (ITF). Understanding the forcing mechanism of the interannual transport variation of the ITF can improve our knowledge of its role on the coupled ocean-atmosphere system.
- The figure shows the sensitivity of annual-mean ITF transport to annual-mean zonal winds over the western equatorial Pacific, the eastern Indian Ocean, especially southwest of Sumatra and Java, and off the Australian and South American coasts. These are indicated by the blue and red regions.
- These sensitivities, which can be explained through oceanic wave physics, help provide a systematic view of local and remote wind forcing of the ITF.
- The computations were conducted on 64 CPUs of the JPL SGI Origin2000 supercomputer.



Ocean Modelling of the Indonesian Throughflow

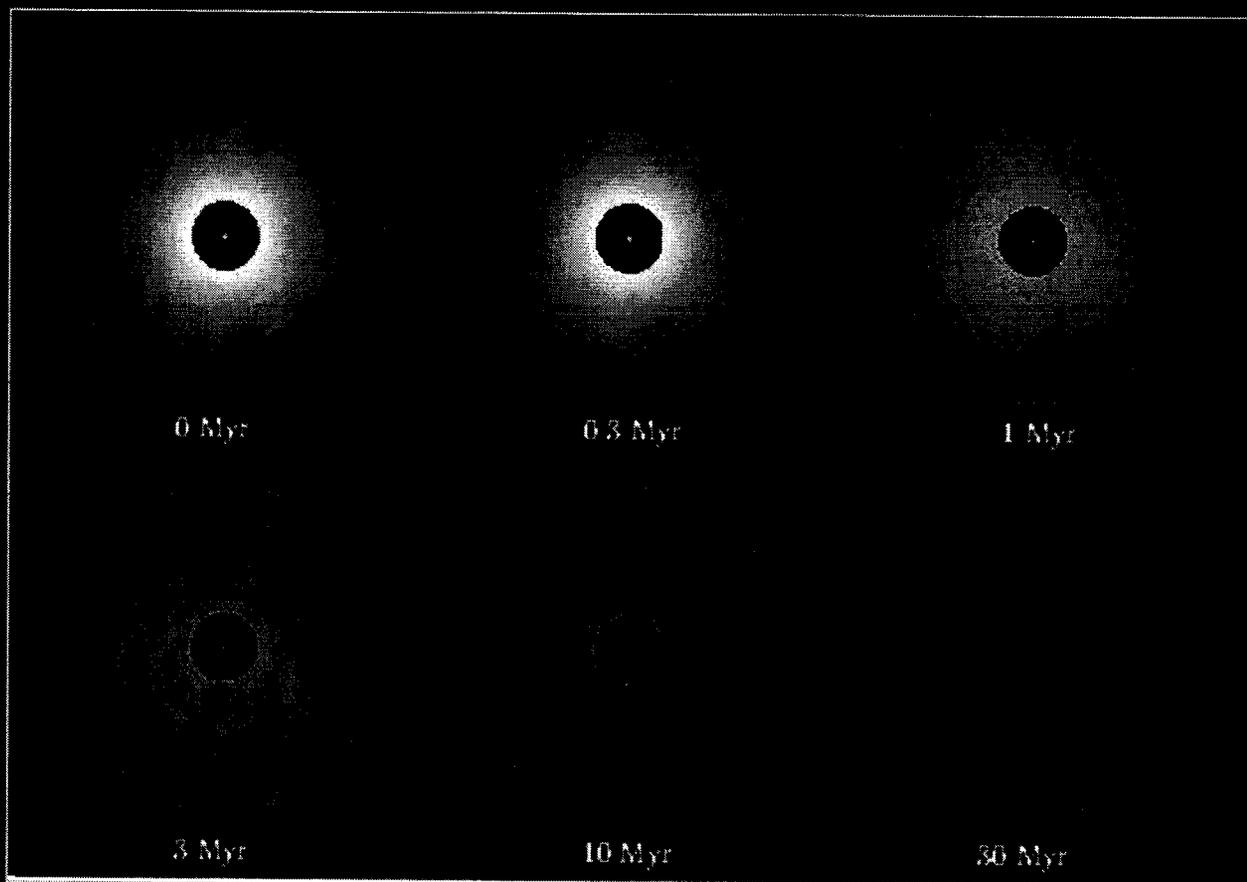
T. Lee, I. Fukumori, D. Menemenlis, Z. Xing and L.-L. Fu, JPL



- A global ocean general circulation model is used to investigate the effects of the Indonesian throughflow (ITF) on global ocean circulation and on climate.
- The ITF is the only low-latitude connection between major oceans.
- Two numerical experiments were performed, driven by surface forcing for the period 1980-1997, with the Indonesian passages open and closed, respectively.
- The figures show the time-mean differences in temperature without and with the ITF. Blocking the ITF warms the Pacific Ocean and cools the Indian Ocean.
- These changes affect not only ocean circulation, but heat transfer between the ocean and the atmosphere, and thus the coupled ocean-atmosphere system.
- In particular, the change in sea surface temperature in the eastern equatorial Pacific (lower panel) has important implications to El Niño formation.
- The computations were performed on 64 CPUs of the JPL SGI Origin2000 super-computer.

Cascades in Planetary Debris Disks, Stellar Flybys

Benjamin Bromley, University of Utah and Scott Kenyon, Smithsonian Astrophysical Observatory



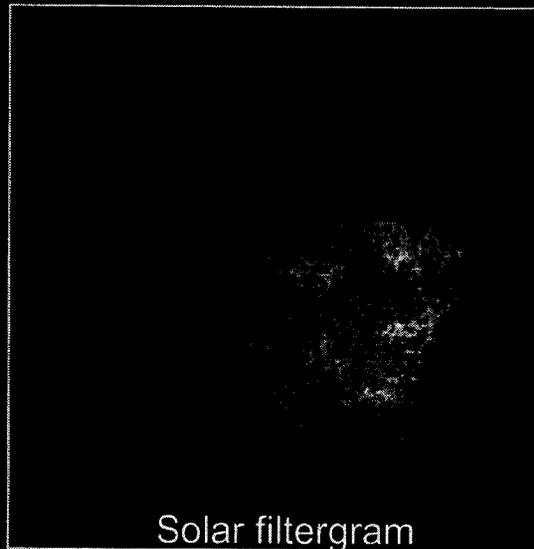
- Each disk shows the surface brightness (luminosity per surface area) in the simulation at six different times (in millions of years).
- Note that the central surface brightness of the disk decays with time.
- Over 12,000 CPU-hours have been used on the SGI Origin2000, via a parallel code that uses 8 to 16 CPUs.
- This work was sponsored by NASA's Office of Space Sciences.

A new multi-annulus planetesimal accretion code is used to investigate the evolution of a planetesimal disk following a close encounter with a passing star. Calculations include fragmentation, gas and Poynting-Robertson drag and velocity evolution from dynamic friction and viscous stirring. We assume the stellar encounter increases planetesimal velocities to the shattering velocity, initiating a collisional cascade in the disk. During the early stages of our calculations, erosive collisions damp particle velocities and produce substantial amounts of dust. For a wide range of initial conditions, the time evolution of the dust luminosity follows a simple relation: $L_d / L_* = L_0 / [\alpha + (t / t_d)^\beta]$. Less massive disks produce smaller dust luminosities and dampen on longer timescales. Since stellar encounters with field stars are rare, results imply that stellar flybys cannot explain collisional cascades in debris disk systems with stellar ages of 100 million years or longer.

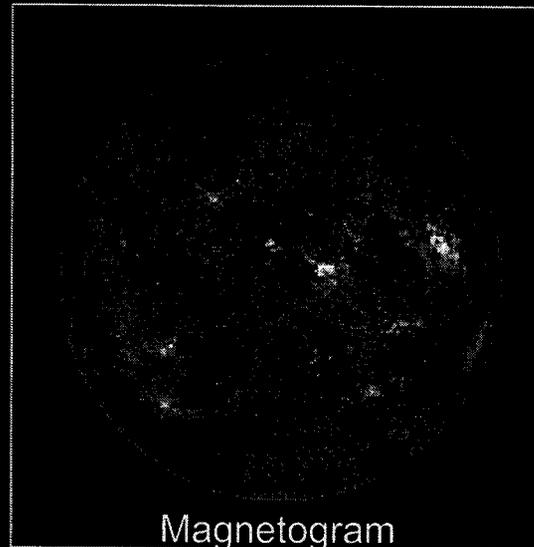
Helioseismology of Solar Structure and Dynamics

Edward Rhodes Jr., JPL, and Perry Rose, USC

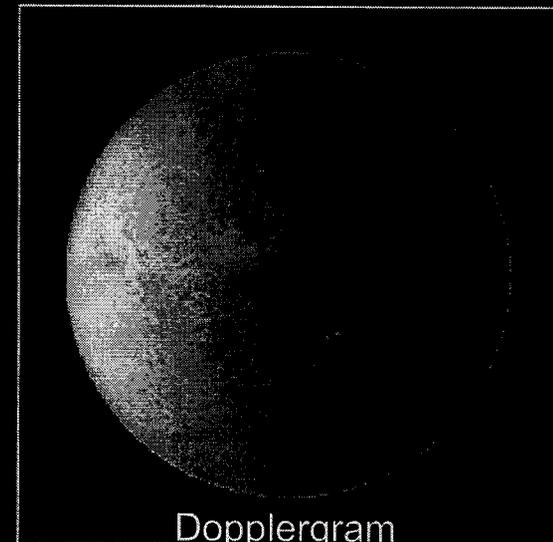
- The study of solar earthquakes is called helioseismology. At the 60 foot solar tower telescope at Mount Wilson, California operated by Prof. Edward Rhodes Jr. for NASA, the internal structure and dynamics of the Sun are studied by monitoring the wave-like motion observed in the solar atmosphere. The tower collects electronic high-resolution full-disk images of the Sun taken in sodium light. Digital techniques are then used to obtain dopplergrams that reveal the radial motion of the solar atmosphere.
- By studying the Sun, which is constantly vibrating, helioseismology tells us about the complicated differential rotation of the Sun, and more precisely reveals the hydrogen and helium abundances in the Sun. Eventually, helioseismology may be able to predict where and when sunspots may occur.
- Most of the analysis for this study was performed on supercomputers at USC and at JPL. Numerous complex programs, that were both memory and I/O intensive, were used to assemble time-varying pictures of the Sun's oscillations and vibrations.



Solar filtergram



Magnetogram



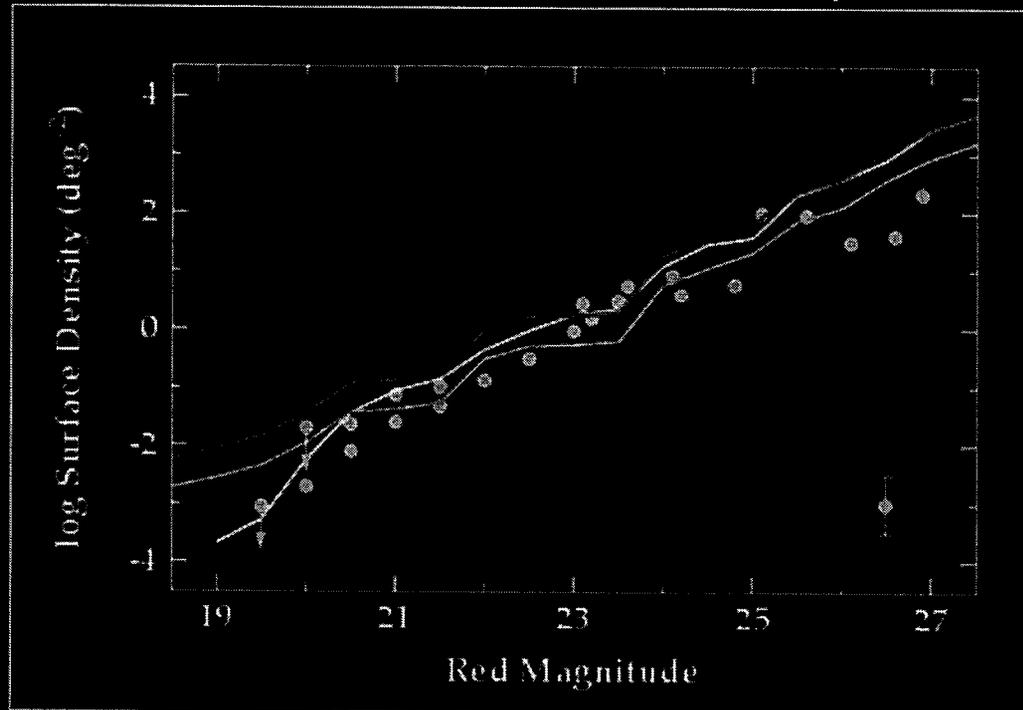
Dopplergram

The left image was taken July 2000 near the peak of high solar activity. The middle image shows the corresponding chromosphere of the Sun (the dark and light areas are the negative and positive magnetic regions respectively). The right image, created at a slightly different time than the other two, reveals the inward and outward motions of the super granulation on the solar surface. When observed over time, these dopplergrams produce an oscillatory motion similar to the motion of earthquake waves.

Planet Formation in the Outer Solar System

Scott Kenyon, Smithsonian Astrophysical Observatory

- In the past decade, observations have revealed several hundred Kuiper Belt Objects (KBOs), with radii of 50 km to 500 km, that lie in the ecliptic plane, at distances of 35-50 AU from the Sun.
- Planets grow by the collision and merger of small bodies called planetesimals. This process is simulated by dividing the outer solar system into concentric annuli, and then following the collisional evolution of bodies with radii of 10 cm up to 3000 km for 5 Gyr.
- The picture below is a multi-annulus model that compares the results of our calculations with observations of KBOs. The yellow points represent the surface density of KBOs projected on the sky as a function of the red magnitude 'R'. KBOs with $R \sim 20$ to 24 can be observed with a 2 m to 4 m class telescope; fainter KBOs with $R \sim 25$ to 27 require one of the Keck telescopes.
- The lines plot the number of KBOs predicted by our model. The white line is the result for 50 Myr, when the first Pluto forms at 40AU.



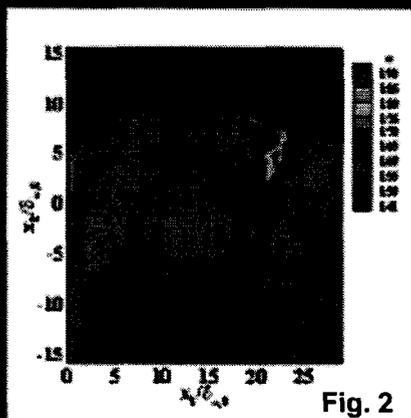
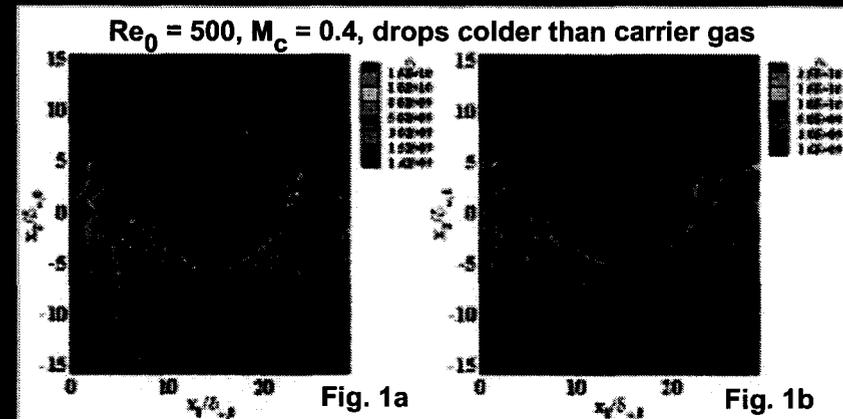
The white line is the result for 50 Myr, when the first Pluto forms at 40AU. At 2 Gyr (magenta line), the largest objects are 10 times the size of Pluto, but the number of smaller objects is unchanged at $R=21$ and fainter. After 2 Gyr (cyan line), collisions deplete the population of large objects, and the number of KBOs declines closer to the levels observed today.

- This work was carried out on the SGI Origin2000 supercomputer, and used well over 3000 CPU-hours. The MPI-parallel code was run on 8 to 16 CPUs.

Direct Numerical Simulations of Mixing Layers Laden with Evaporating Drops of a Liquid Composed of Many Chemical Species

Patrick LeClercq and Josette Bellan, JPL

- This simulation uses the compressible Navier-Stokes, species and energy equations, with source terms that originate from the drop-gas dynamic and thermodynamic interactions.
- It is a unique model for drops-in-mixing-layer simulations, which treats the composition of the liquid drops and evolving gas from the evaporation in a probabilistic manner.
- Comparison of results from a simulation where the liquid is represented by a single chemical species (Fig. 1a), to the more realistic simulation using a model accounting for multiple species (Fig. 1b), shows that the multi-species droplets survive longer, and therefore interact longer with the flow than those of a single species. This gives the flow a more convoluted structure.
- The molar mass of the gas is stratified for the case of the multi-species liquid drop (Fig. 2), but it is inherently constant for the single species case.

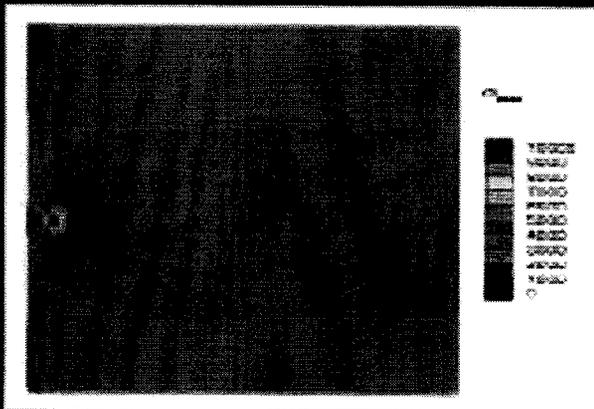


- This work is of interest to spray simulations for combustion applications (e.g. furnaces and engines), and other spray-nozzle applications. It will enable improvements to spray design and control, so as to target the spatial distribution of specific species released through evaporation.
- The three-dimensional simulations were performed using 256 x 288 x 160 grid nodes, and it is one of the largest simulations in the field.
- These simulations were conducted on 32 CPUs of the SGI Origin2000, and used 3200 CPU-hours.
- This research was sponsored by the ACS Petroleum Research Fund, and by the DOE Hydrogen Program.

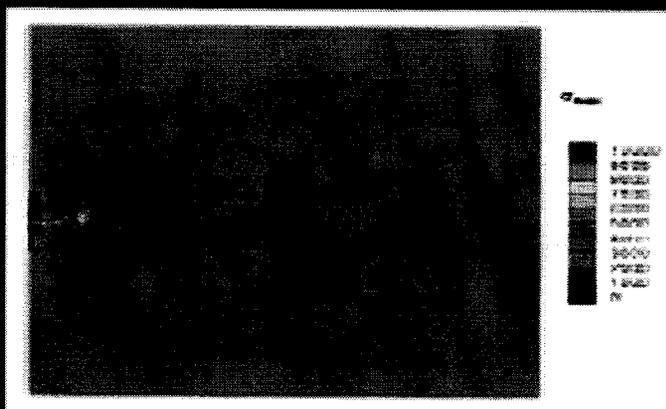
Micromechanical Modeling of Dynamic Failure of Materials

Min Zhou and Jun Zhai, Georgia Institute of Technology

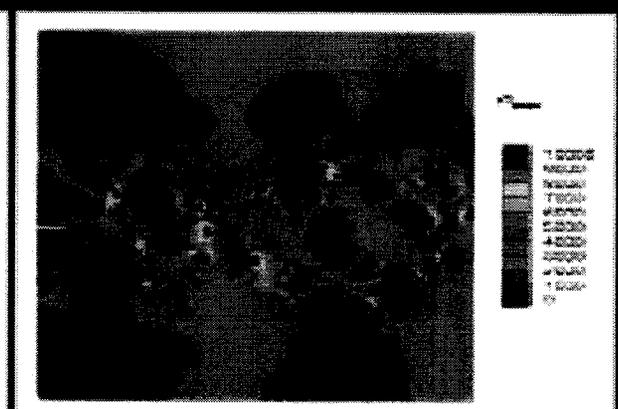
- This research numerically analyzes dynamic crack propagation in microstructures of ceramic composites, using a micromechanical model that accounts for arbitrary microstructures and arbitrary fracture patterns:
 - The analysis considers arbitrary phase distributions in the actual microstructures of alumina / titanium diboride ($\text{Al}_2\text{O}_3/\text{TiB}_2$) composites.
 - The simulations concern the effects of microstructural morphologies on fracture.
- The approach uses both a constitutive law for the bulk solid constituents and a constitutive law for the fracture surfaces. The model is based on the cohesive surface formulation of Xu and Needleman, and represents a phenomenological characterization for atomic forces on potential crack/microcrack surfaces.
- Fractures evolve as an outcome of bulk material response, interfacial behavior and applied loading. This approach provides a unified and self-consistent treatment of mixed mode fracture.
- The bonding strength of interfaces between the phases is found to significantly influence the failure characteristics.
- The overall local crack speed, defined as the time rate of change of arc length along zig-zagging crack paths, is found to reach the intersonic range (i.e., greater than shear-wave speeds and smaller than longitudinal-wave speeds).
- The model allows the energy release rate to be evaluated easily, and the results demonstrate that large TiB_2 reinforcements significantly impede crack propagation and increase the fracture resistance of the composites.
- This research is important for the development of failure-resistant ceramic composites through microstructural design.
- Run on the Cray SV1-1A, a typical case takes about 8 CPU-hrs, with some cases requiring as much as 20 CPU-hrs.



Dynamic crack growth in alumina:
crack branching after short propagation.



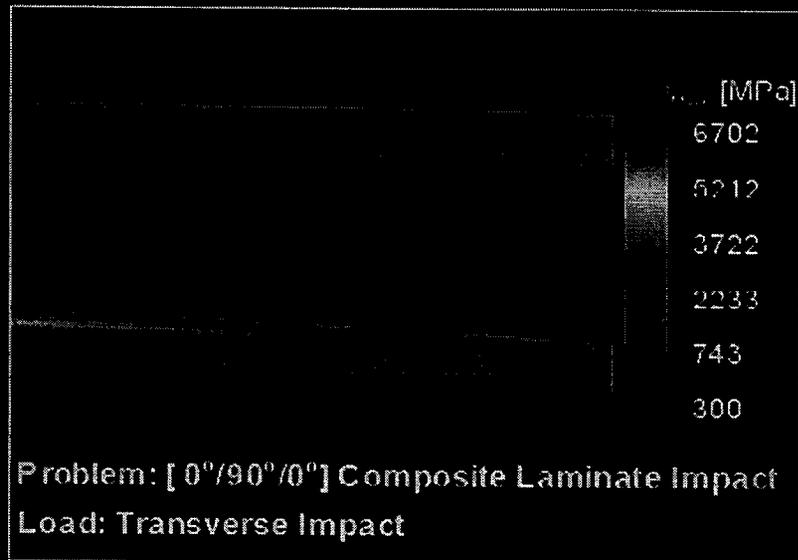
Dynamic crack growth in ceramic -
composite with strong interface:
crack deflection and branching.



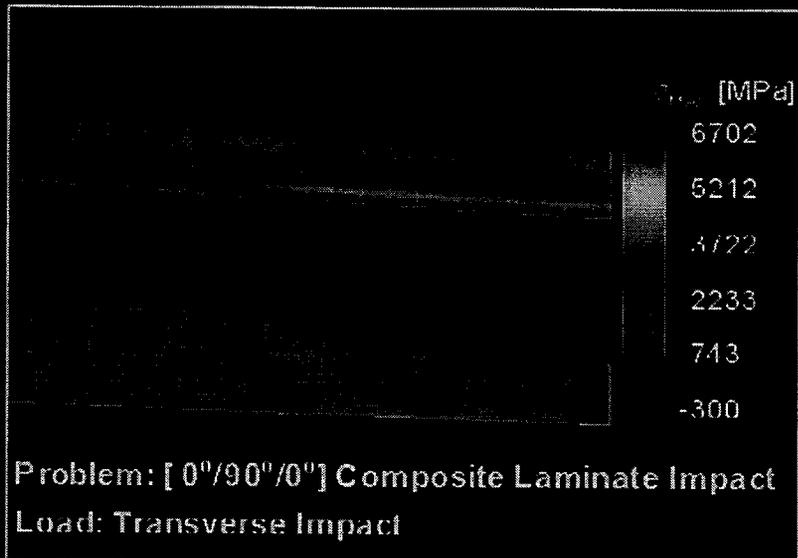
Dynamic crack growth in ceramic -
composite weak interface: microcrack
nucleation and coalescence along interfaces.

Impact Damage in Composite Laminates

Karel Minnaar and Min Zhou, Georgia Institute of Technology



Fracture in a laminated composite

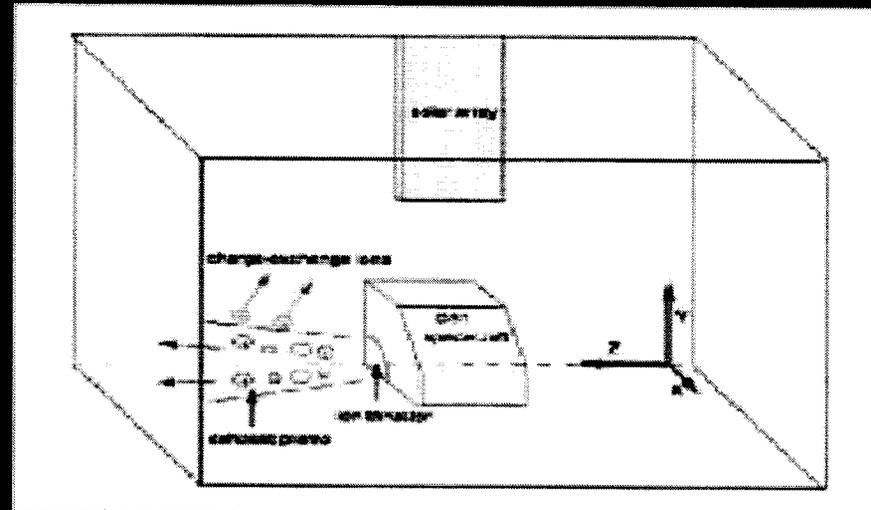
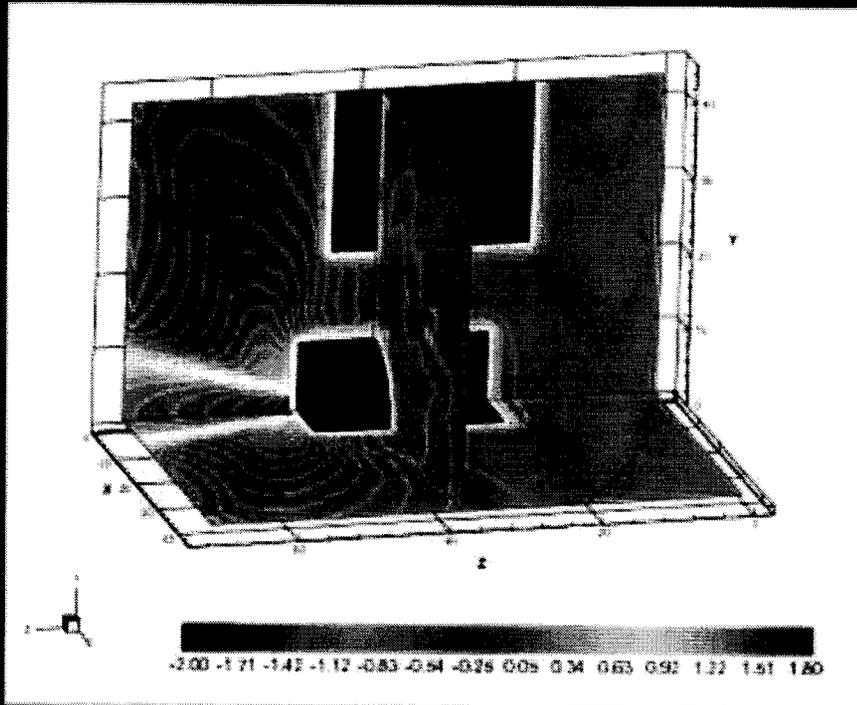


Delamination

- This work consists of a two-dimensional dynamic finite-element model with cohesive surfaces embedded in the transverse ply and between the laminate layers.
- Cohesive finite elements explicitly model two failure modes, matrix cracking and delamination, common in composite laminates.
- Complex fracture patterns evolve, from the initial impact load, to the extents that are shown here; cracks start in the transverse ply and propagate to the ply interfaces, and delamination starts at the points where the matrix cracks reach the ply interfaces.
- This technique allow researchers to understand the influence of loading mode, loading rate, bonding strength and material lay-up on the failure behavior of composite materials.
- Further, it offers an exciting capability to model fracture as a natural outcome of the applied load and boundary conditions.
- The Cray SV1-1A supercomputer was used for these computations. Each case that was studied required an average of 8 CPU-hrs, with some cases taking up to 20 CPU-hrs.

3-D Particle Simulations of Ion Thruster Plume for Deep Space 1

Joseph Wang, JPL and Virginia Tech



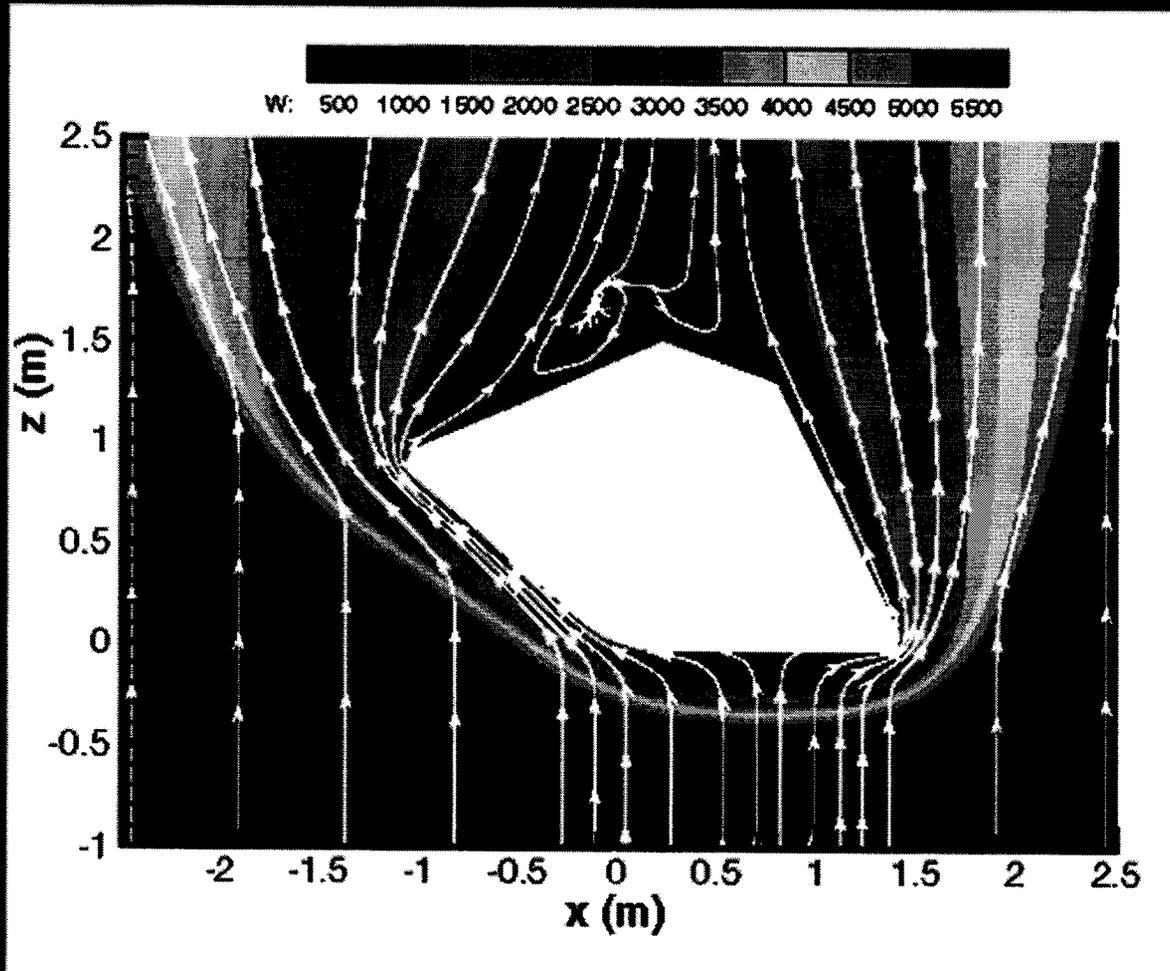
	$n_{CEX} \text{ cm}^{-3}$	$J_{CEX} \text{ A/cm}^2$
In-flight measurement	1.2E6 to 4.4E6	0.9E-7 to 3.3E-7
Simulation	1.43E6	1.5E-7

Charge-exchange plasma density and current density at the IDS instrument location for thruster level ML83.

- A fully 3-dimensional particle-in-cell simulation model was developed to predict the ion thruster plasma plume environment for the Deep Space 1 (DS1) spacecraft.
- The thruster exhaust plume is modeled as particle ions and fluid electrons. The particle trajectories, space charge and electric field are solved self-consistently. A typical simulation tracks about 5 million particles.
- The left panel shows the model setup, while the right panel shows contours of the electric potential surrounding the spacecraft. The table shows measurement and simulation results for the charge-exchange ion density and the current density for one typical case.
- Comparison of the simulation results with the in-flight measurements from the IDS instrument on DS1 shows excellent agreement. Thus, this model may be used to predict the thruster-induced environment for future spacecraft using ion propulsion.
- The simulations were performed on the JPL Cray SV1-1A, and a typical run took about 16 CPU-hours.

Mars Exploration Rover Atmospheric Entry Simulation

Mark Schoenenberger, NASA Langley Research Center

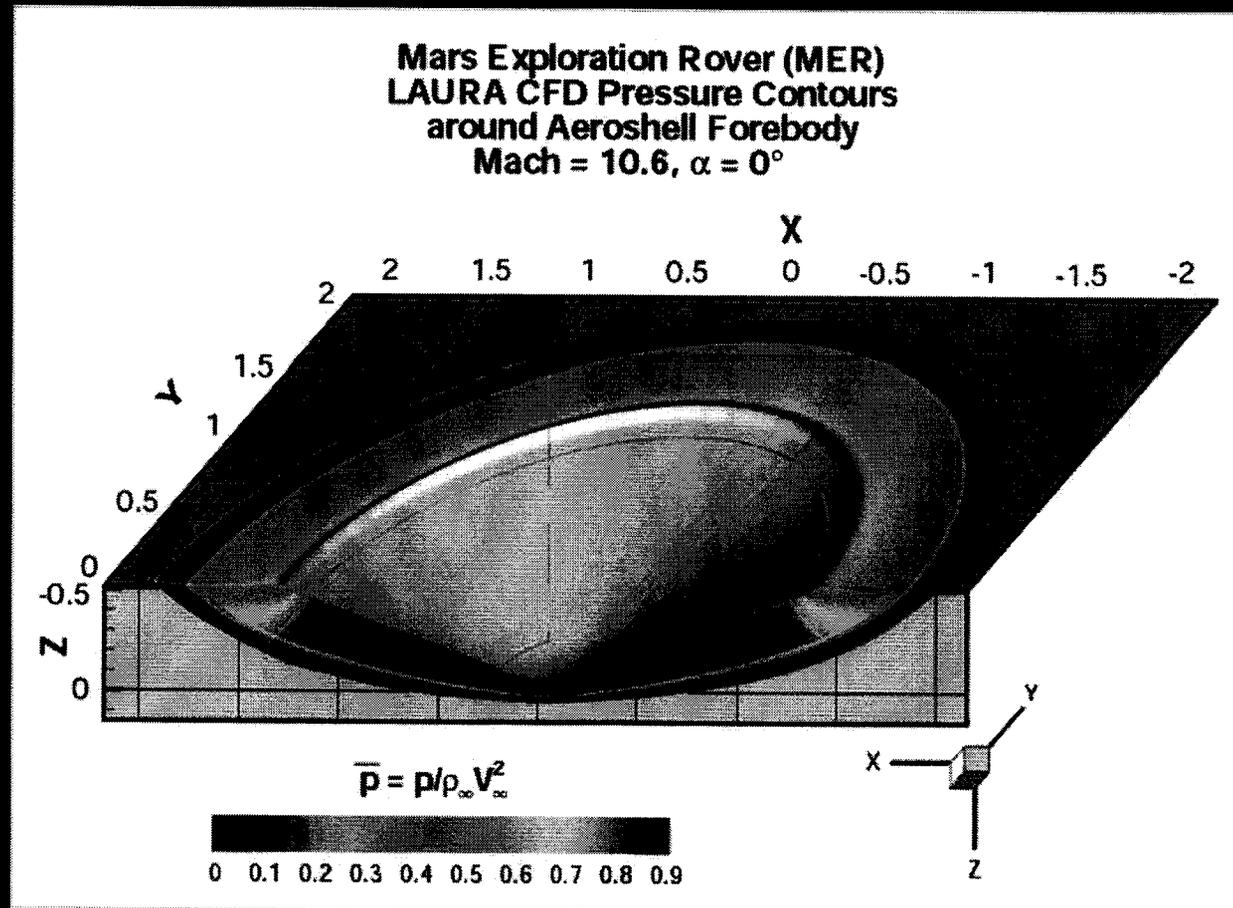


- The Mars Exploration Rover will enter the Martian atmosphere at hypersonic speeds, protected by an aeroshell/heat shield.
- This simulation uses the Direct Simulation Monte Carlo (DSMC) method, which employs statistical models of the individual molecules of a gas to calculate the flow field around the vehicle. This allows solutions in the low-density parts of the atmosphere, where the usual equations of computational fluid dynamics cannot be applied.
- The figure illustrates the aeroshell dropping through the atmosphere at a velocity of 5500 m/s, with an angle of attack of 20° and a Knudsen number of 0.01.
- The colors show the z-component of the gas velocity around the vehicle.
- This simulation involved 10 million molecules, and was run in 40 hours on 64 processors of Alhena, the JPL SGI Origin2000 supercomputer.

Mars Exploration Rover Atmospheric Descent Simulation

Mark Schoenenberger, NASA Langley Research Center

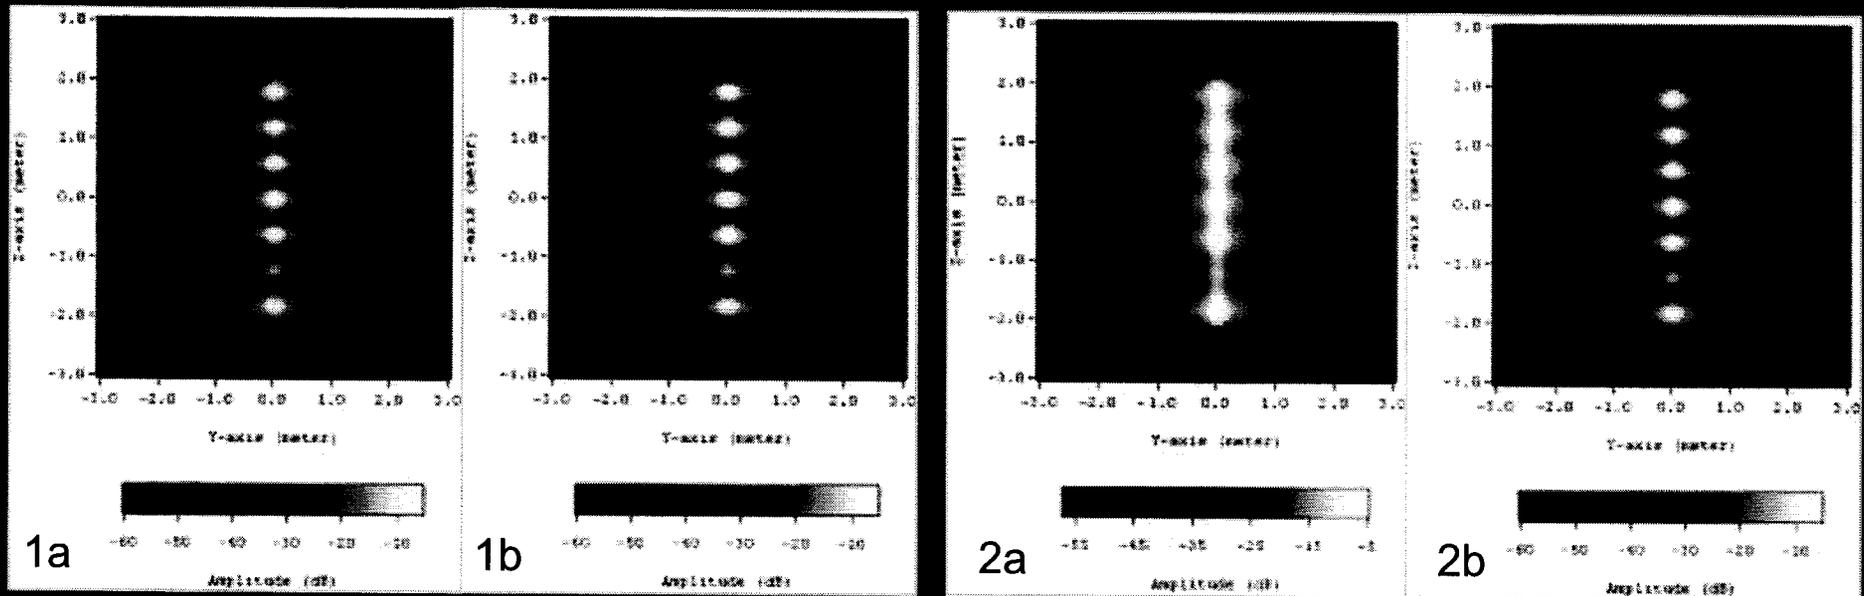
- The Mars Exploration Rover will descend through the Martian atmosphere at supersonic speeds, protected by an aeroshell/heatshield.
- This simulation uses the MPI version of NASA Langley's LAURA (Langley Aerothermodynamic Upwind Relaxation Algorithm) code, which is a finite-volume shock-capturing Navier-Stokes flow solver.



- The figure illustrates the aeroshell descending through the atmosphere at a velocity of Mach 10.6 and an angle of attack of 0 degrees. The colors show the non-dimensionalized gas pressures around the vehicle.
- This simulation involved 52,480 cells and was run in 20 hours on 7 processors of Alhena, the JPL SGI Origin2000 supercomputer.

Microwave Holography in Cylindrical Near-Field Scanning

Ziad A. Hussein, JPL



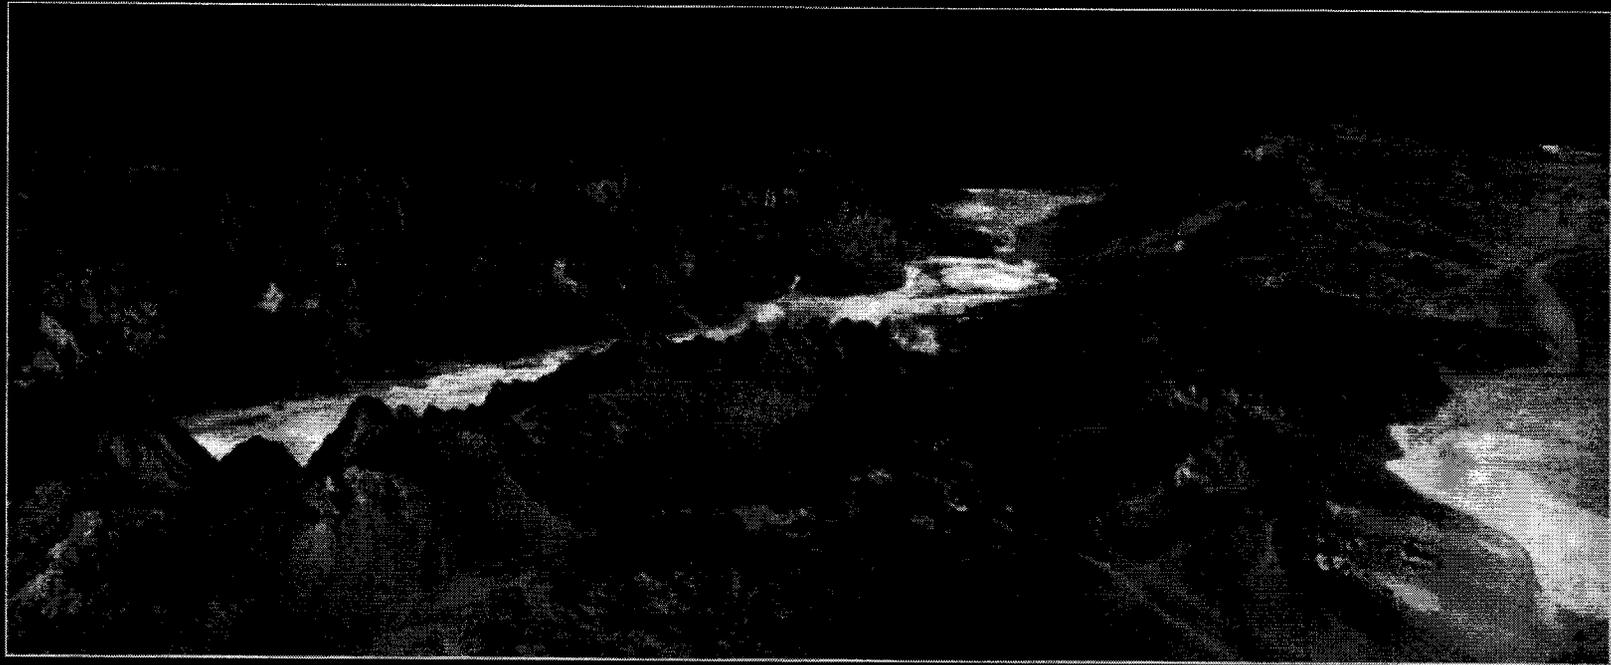
- This research develops a technique to diagnose defective elements in antenna arrays, and to apply probe compensation to improve measurement accuracy.
- The figures show holographic microwave images of a 7-element array constructed on a planar near-field surface (y-z plane), from the knowledge of the near-field on a cylindrical surface.
- Figures 1a and 2a are the non-compensated far-field patterns obtained from simulated near-field probe response on a cylindrical surface, while figures 1b and 2b are the compensated far-field patterns.
- The probe radius in figure 1 is $\frac{1}{2}$ wavelength, while in figure 2 the probe radius is 1 wavelength.
- Each element of the array was fed by unit amplitude (1 watt) and zero phase, with the exception of the second element which was fed 0.5 amplitude ($\frac{1}{2}$ watt) and zero phase.
- This work was carried out as part of the NASA Scatterometer/SeaWinds project.
- Both the Cray SV-1 and the SGI Origin 2000 supercomputers were used for this study.



RIVA – Remote Interactive Visualization and Analysis System

NASA

Peggy Li and Dave Curkendall, JPL

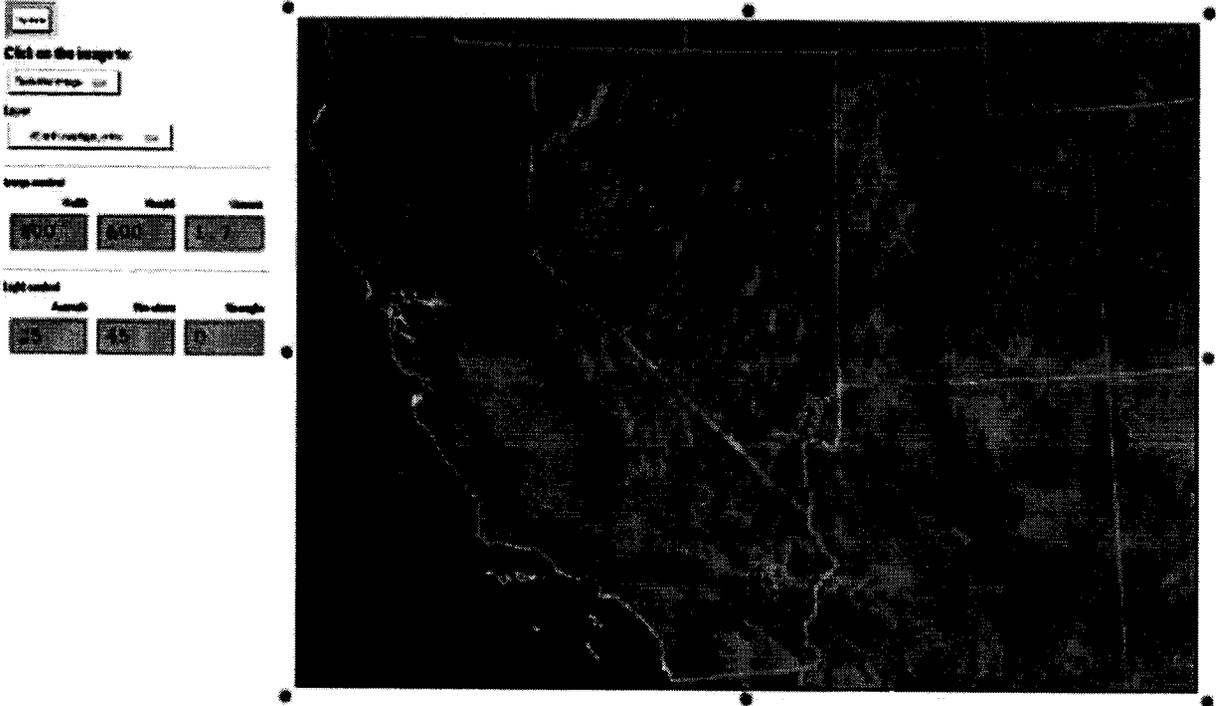


- RIVA is a parallel terrain rendering system for interactive visualization and exploration of large terrain datasets. It renders 3-D perspective views by wrapping an earth or planetary imagery around a digital elevation model. The above image is a 3-D view of Death valley, California, rendered by RIVA.
- It is a scalable and efficient rendering system using a ray-identification algorithm. The system is scalable to large machines, large datasets and large images. It has out-of-core rendering capability. Multiple viewport rendering allows it to render an image as large as 3000 x 4000 pixels or more. RIVA runs on the Origin2000.
- There is a built-in GUI tool. Flexible Flyer GUI program serves as a view finder for the parallel renderer. A key frame editor in the Flexible Flyer is for the design and preview of a flight path.
- RIVA is a great system to build fly-over animated movies for scientific exploration and outreach purposes.
- Many fly-over animated movies have been built using RIVA:
 - The Lewis & Clark Trail movie in HDTV format.
 - The San Diego movie for the Discovery Channel.
 - The History as Prolog movie, to celebrate the Mars Pathfinder landing.

LandSat Map of the U.S.

Lucian Plesea, JPL

US with overlays, ortho at 60 arcsecs/pixel, center (x,y) = (-115.768,37.433) <http://mapus.jpl.nasa.gov/>

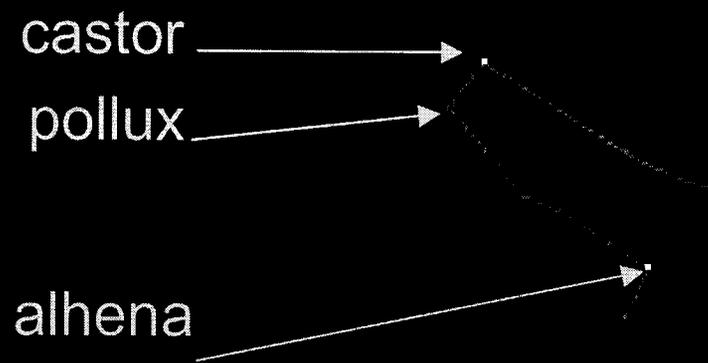


The screenshot shows a web interface for a satellite map of the United States. At the top, there is a URL and a title. Below the title, there are several control panels. The first panel has a 'Click on the image to:' label and a 'Submit map' button. The second panel is labeled 'Layer' and has a dropdown menu. The third panel is labeled 'Image control' and has three buttons: 'Full', 'Zoom', and 'Refresh'. The fourth panel is labeled 'Light control' and has three buttons: 'Auto', 'Manual', and 'Refresh'. The main part of the interface is a large satellite image of the United States, showing state boundaries and a grid. The image is centered on a specific location in the western US.

The viewport above measures width=13.334 and height=10 degrees on the ground. Each pixel you see measures 60x60 arc-seconds of longitude / latitude. The viewport is centered on X=-115.768, Y=37.433 degrees (Longitude / Latitude, WGS84).

- The MapUS web server offers a contiguous US LandSat image mosaic and DEM, compliant with the OpenGIS Web Mapping Server.
- The server is hosted on a four-processor SGI Onyx, and uses a 50 GB image database as its data source.
- It is an excellent example of custom server-side GIS image-processing, in which operations such as zooming, sub-setting, colorimetric operations, vector overlay drawing and 3D shading are all performed under user control.
- The server has been in continuous operation for public access since June 2000.
- Although it has never been advertised, in its first year of operation, the server generated more than 60,000 map images, representing more than 8 GB of compressed JPEG image data.

Conclusion





Need additional information?

For additional information,
please visit our web site:

<http://sc.jpl.nasa.gov/>