National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

High Performance Computing for Flight Projects at JPL

Chris Catherasoo
High Performance Computing
Office of the CIO

8 October 2010

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Jet Propulsion Laboratory overview

- JPL is part of NASA and Caltech
    - Owned by NASA – a “Federally-Funded Research and Development Center” (FFRDC)
    - Operated by Caltech, under contract to NASA
- $1.7 billion business base
- 5,000 employees
- Site area: 0.75 km^2

JPL's mission for NASA is robotic space exploration

- Mars
- Solar system
- Exoplanets
- Space science
- Earth science
- Interplanetary network

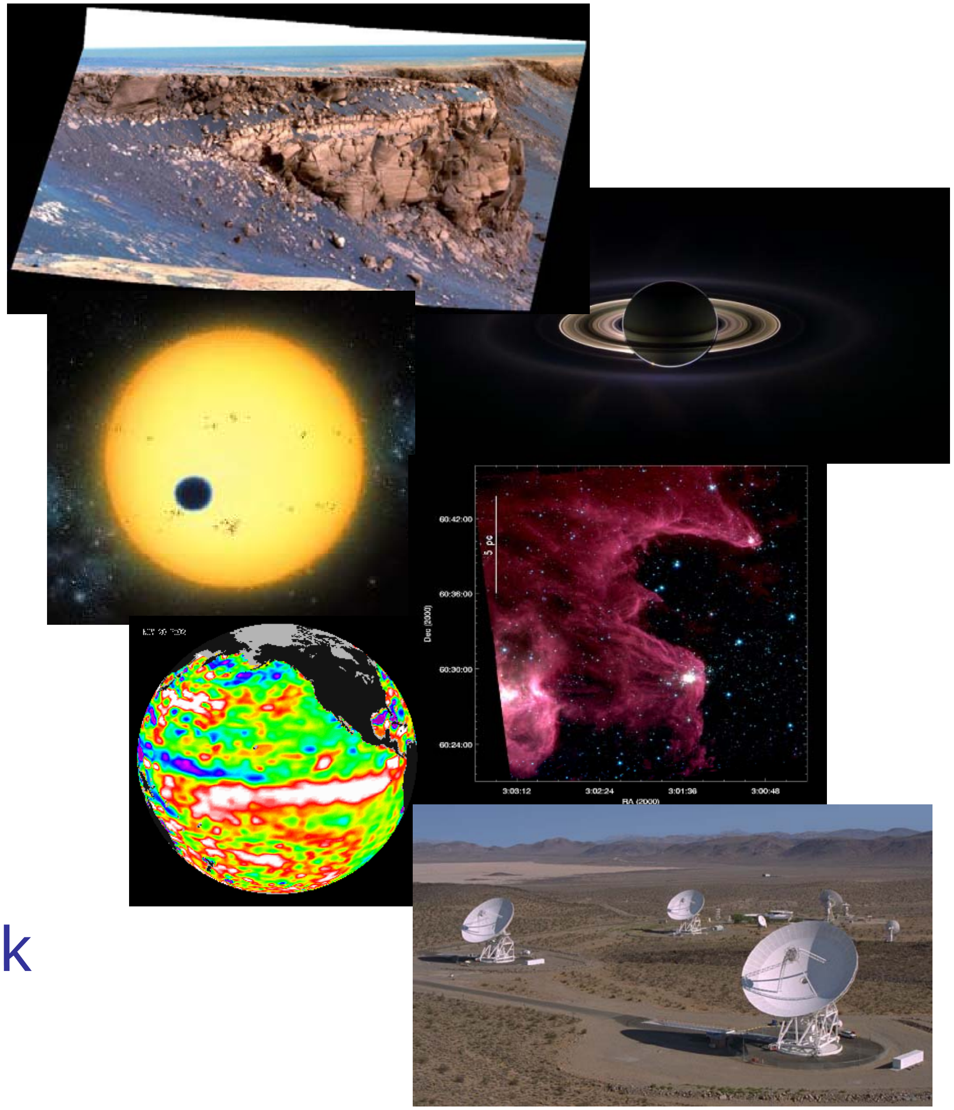

Robotic explorers in space

Voyagers 1 and 2 in interstellar space

Cassini at Saturn

Mars Reconnaissance Orbiter

Two Mars Exploration Rovers

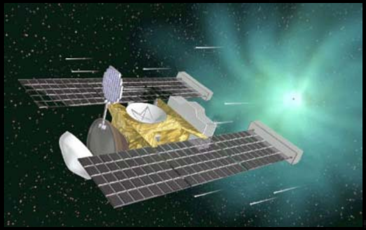

Stardust-NExT to comet Tempel 1

EPOXI to comet Hartley 2

Mars Science Laboratory

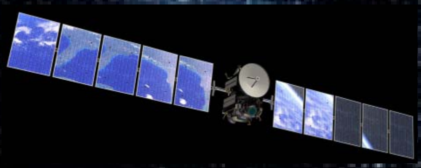

Dawn to asteroids Ceres and Vesta

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Robotic remote sensing on earth

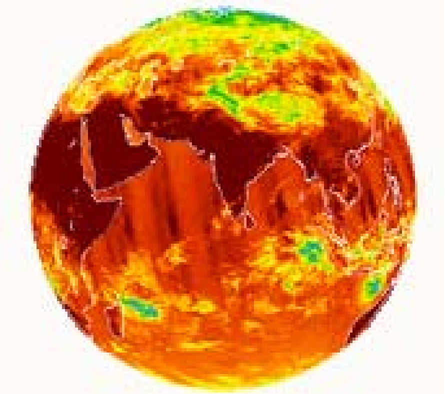

Atmospheric Infrared Sounder (AIRS) provides monthly global temperature maps

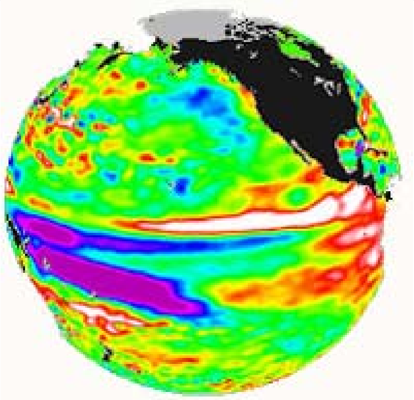

Jason provides global sea surface height maps every ten days

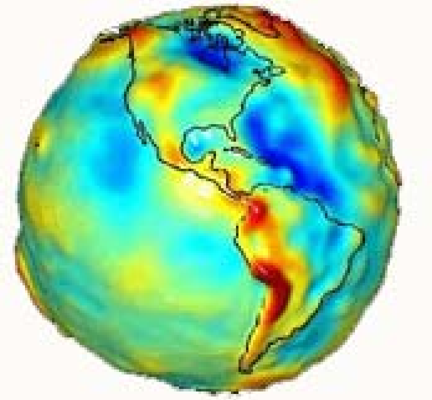

Gravity Recovery and Climate Experiment (GRACE) provides monthly maps of Earth's gravity

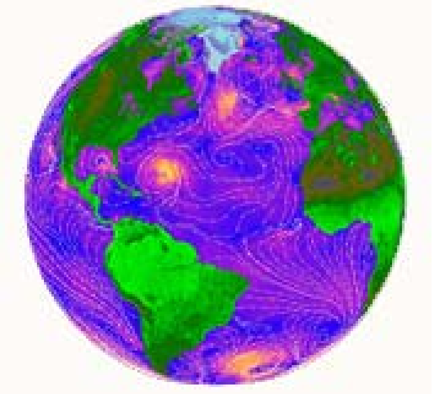

QuikSCAT provides near global (90%) ocean surface wind maps every 24 hours

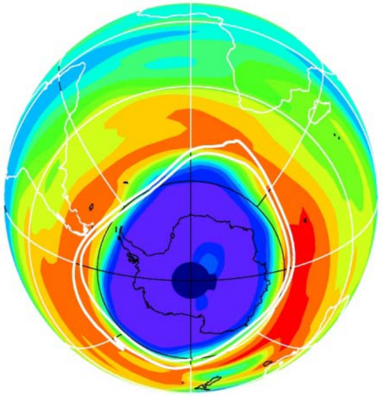

Microwave Limb Sounder (MLS) provides daily maps of stratospheric chemistry

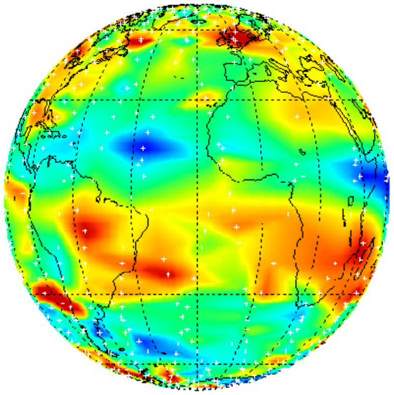

Tropospheric Emission Spectrometer (TES) provides monthly global ozone maps

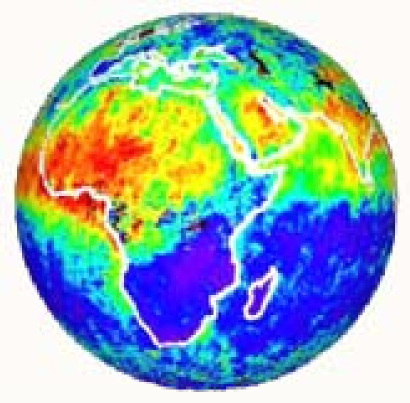

Multi-angle Imaging Spectro Radiometer (MISR) provides monthly global aerosol maps

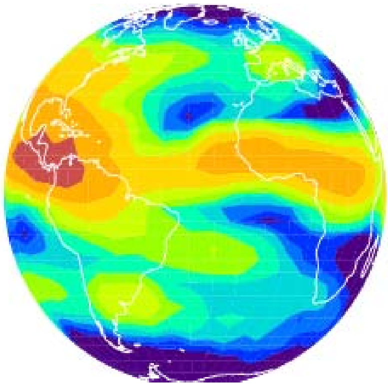

CloudSat provides monthly maps of cloud ice water content

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Institutional HPC resources

Dell Xeon Cluster

- 1024 3.2 GHz processors
- 2 GB per CPU distributed memory
- Myrinet interconnect

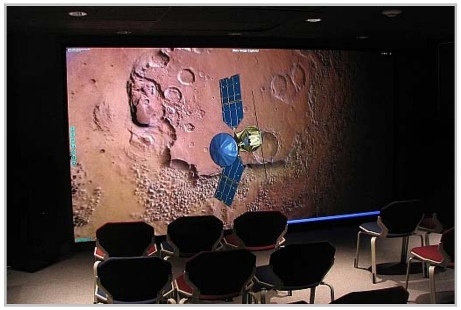

Visualization Center

- Sony SRX-R110 projector
- 12’ x 7’ display
- Resolution: 4096 x 2160 (8 MPixels)

HP SFS File System

- 2 MDS/Admin servers
- 16 OSS servers
- Read / write > 2 GB/s
- 32 TBytes

Dell Xeon Clusters

- 2 x 512 3.06 GHz processors
- 2 GB per CPU distributed memory
- Gigabit ethernet interconnect

Online Storage

- RAID6 system
- 1 PByte

SGI Altix 3700s

- 2 x 256 and 1 x 64 1.5 GHz processors
- 2 GB per CPU shared memory
- ccNUMA interconnect

HPC resources at NASA Ames

- SGI ICE cluster
  - Total cores: 84,992
    - 2,304 Westmere (Xeon X5670) nodes
      - 2 six-core processors per node
    - 1,280 Nehalem (Xeon X5570) nodes
      - 2 quad-core processors per node
    - 5,888 Harpertown (Xeon E5472) nodes
      - 2 quad-core processors per node
  - Total memory: 133 TB
  - Infiniband DDR, QDR interconnect
    - 11-D hypercube topology
- SGI Altix 4700 system
  - Total cores: 4,608 (originally 10,240)
    - Four compute nodes
  - Total memory: 9 TB
  - NUMALink interconnect
    - Single-system image on each compute node

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Spacecraft components in cruise configuration

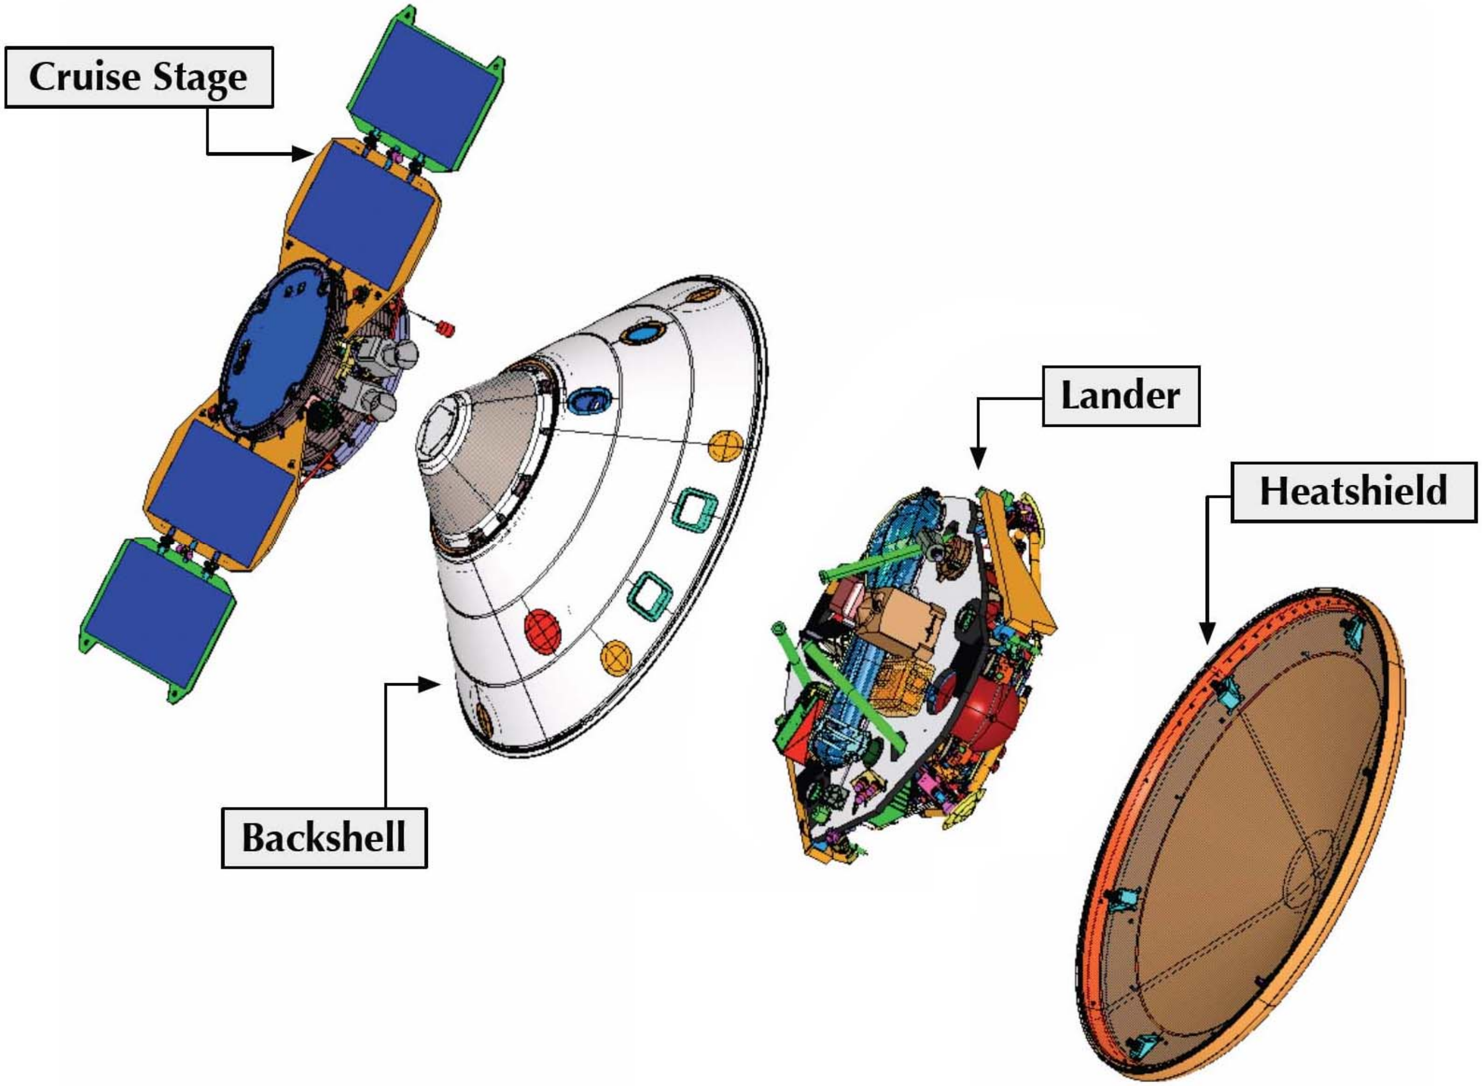

Entry, descent and landing

- Entry, Descent and Landing (EDL) is the sequence of events that brings a spacecraft safely to the surface of a planet
- It consists of several phases
  - Cruise stage separates before entering the atmosphere
  - Entry phase
    - Aerobraking – friction with the planetary atmosphere is used to slow the spacecraft from over 5,500 m/s to 500 m/s in about 220 s
  - Descent phase
    - Parachute braking – slows the spacecraft down to 100 m/s in 20 s
  - Landing phase
    - The parachute separates and spacecraft lands
      - Retro rockets
      - Airbags
      - Sky crane
- For Mars, the EDL sequence takes about 7 minutes
  - Signal time from Mars to Earth is about 10 minutes

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Typical entry, descent and landing sequence

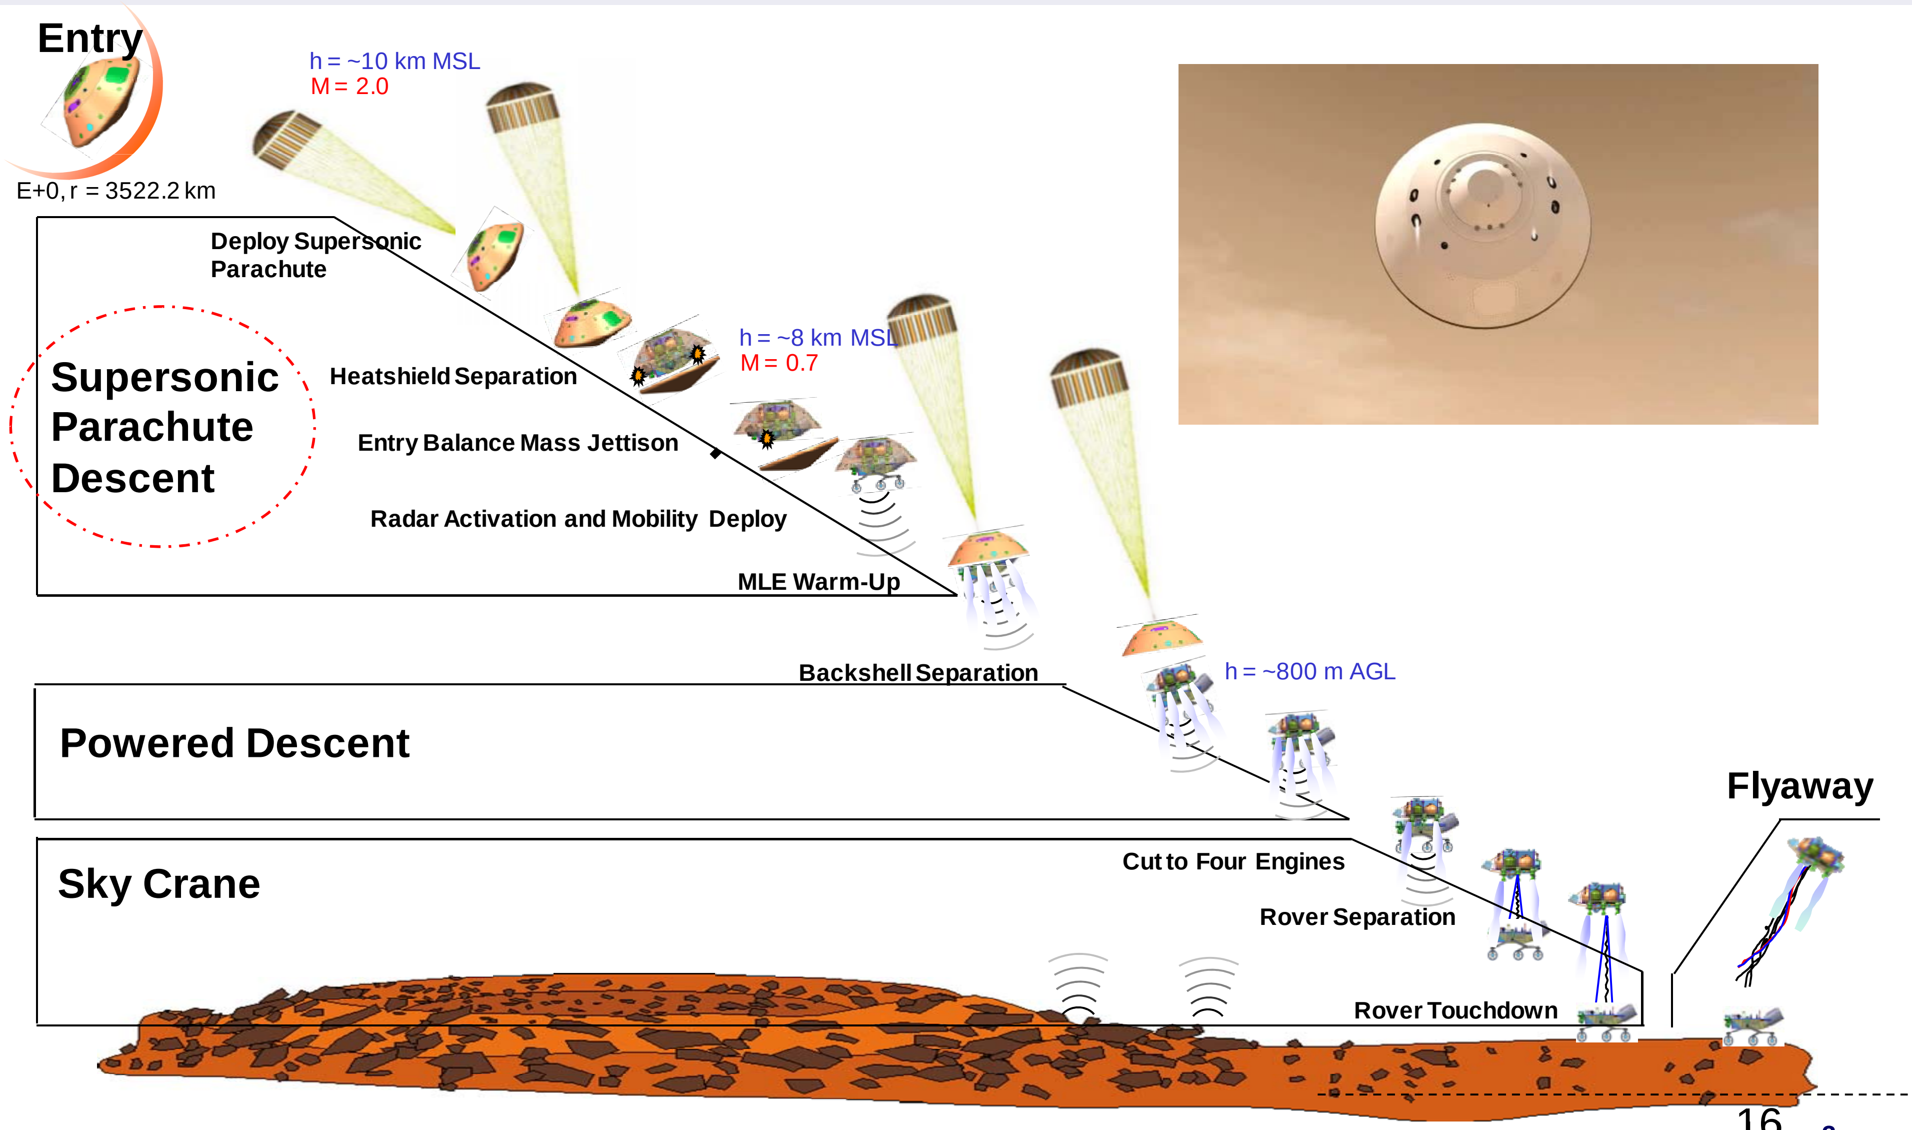

Entry

E+0, r = 3522.2 km

h = ~10 km MSL
M = 2.0

Deploy Supersonic Parachute

Supersonic Parachute Descent

Heatshield Separation

Entry Balance Mass Jettison

Radar Activation and Mobility Deploy

MLE Warm-Up

h = ~8 km MSL
M = 0.7

Backshell Separation

h = ~800 m AGL

Powered Descent

Sky Crane

Cut to Four Engines

Rover Separation

Rover Touchdown

Flyaway

2000 m above MOLA areoid

Entry, descent and landing simulations

- EDL simulations are one of the most mission-critical HPC applications run at JPL
  - The simulations involve multi-body dynamics of the parachute, backshell and lander system
  - The EDL application is the “Program to Optimize Simulated Trajectories” (POST)
    - Was developed at NASA Langley
    - Uses a 6 degrees of freedom modelling scheme
    - Inputs include ambient atmospheric conditions and wind speeds
  - Monte Carlo simulations are performed to determine spacecraft entry, descent and landing characteristics to evaluate safety metrics for landing
  - EDL simulations are used to
    - Down-select landing sites, and to choose the final landing site
    - Apply final trajectory maneuver corrections to the spacecraft prior to cruise stage separation

EDL simulations used for JPL missions to Mars

- EDL simulations used successfully for
    - Mars Pathfinder
        - Landed: 4 July 1997
        - Length: 0.65 m; weight: 10.6 kg
    - Mars Exploration Rovers – twin rovers
        - Landed: 3 January 2004 and 24 January 2004 respectively
        - Length: 1.6 m; weight: 180 kg
    - Mars Phoenix Lander
        - Landed: 25 May 2008
        - Length: 1.5 m; weight: 350 kg
- Upcoming
    - Mars Science Laboratory
        - Launch: November 2011
        - Length: 2.7 m; weight: 950 kg

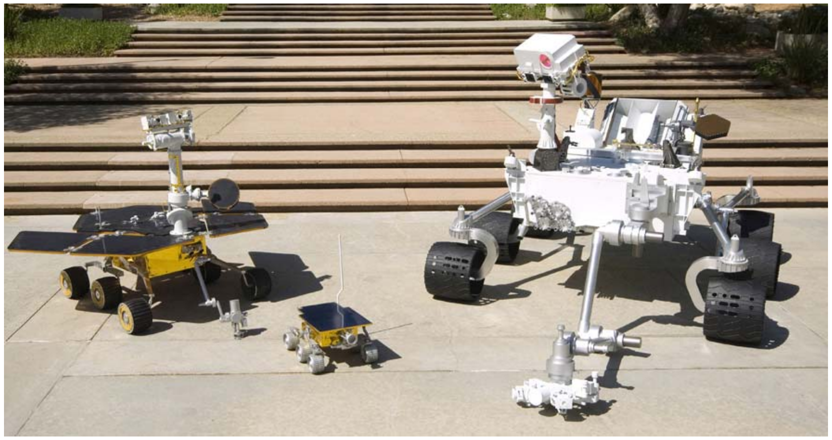

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

The Phoenix Mars Lander radar ambiguity

- The Phoenix Mars Lander was launched in August 2007
  - Mission was to explore the Martian polar region for evidence of water
- The lander used a radar to obtain ground-relative altitude and velocity during terminal descent
- Both helicopter field tests and simulations were used to validate the radar performance

- Analysis of the radar simulation data showed that the presence of the jettisoned heat shield could cause radar to lock on a range ambiguity
  - The radar was not locking on to the heat shield
  - Did not occur in the absence of the heat shield

The Phoenix Mars Lander radar ambiguity

- The radar erroneously reported an altitude that was much lower than the true lander altitude
  - Could not be distinguished from the expected altitude behavior
  - Would have caused the premature separation of the lander from the backshell
  - Result would have been catastrophic loss of the mission
- The problem was impossible to characterize analytically
  - Too many contributing parameters – lander altitude, heatshield range, heatshield radar cross-section, heatshield attitude, attitude rate
- With only eight months to go before launch, resolving this problem became a critical activity
- Hundreds of thousands of radar simulation runs were made to characterize the design space
  - Phoenix was given highest priority on all the laboratory’s clusters

The Phoenix Mars Lander radar ambiguity

- Results were plotted on radar ambiguity maps
  - Each dot is the result of a single simulation that took about 3.5 core hours to run
    - Gray: No target acquisition
    - Green: Radar correctly locked on the ground
    - Red: Radar incorrectly (ambiguously) locked on the ground
    - Cyan: Points at which radar begins making measurements

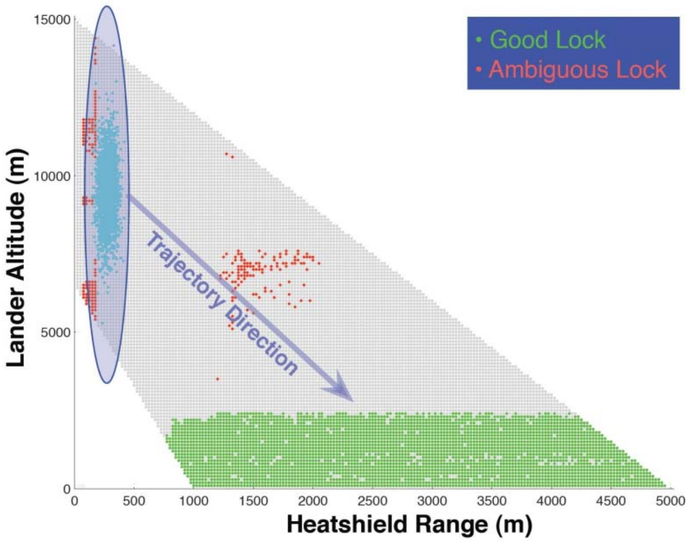

The Phoenix Mars Lander radar ambiguity

- The problem was resolved by
  - Delaying the start of the radar search
  - Modifying the radar Pulse Repetition Interval (PRI)
- The modified radar was field tested, and the updated radar-model simulation results were used to verify that the problem had been eliminated

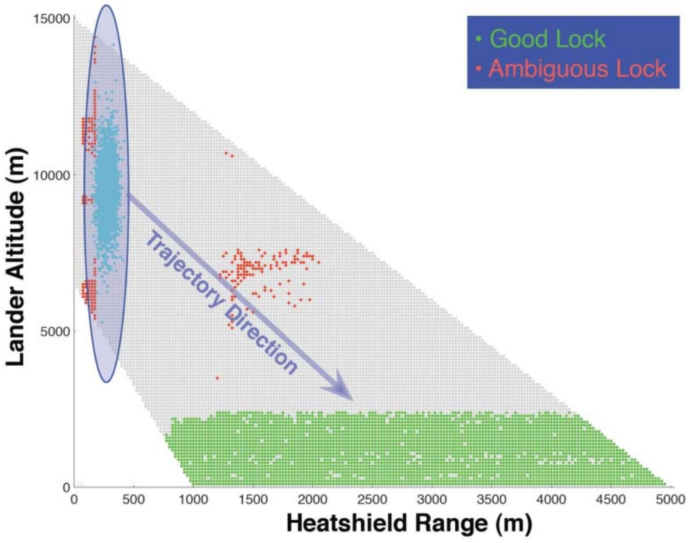

Problem

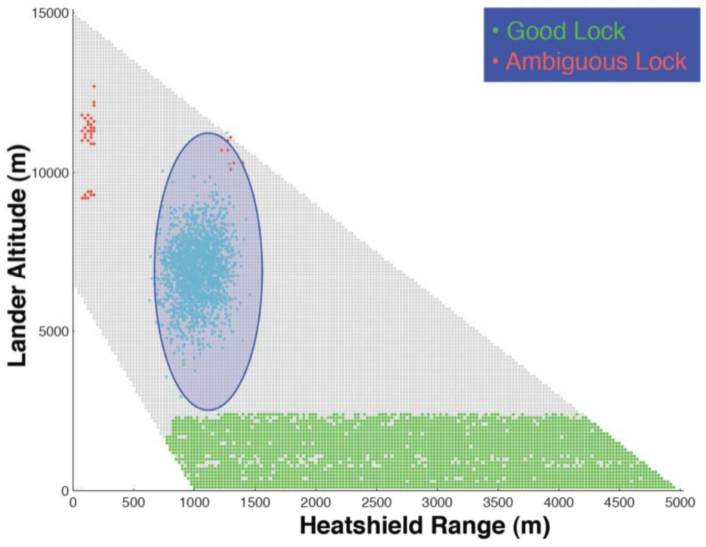

Resolution

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

The Mars Science Laboratory supersonic parachute design

- The Mars Science Laboratory (MSL) will be launched in November 2011
  - Mission is to detect and study organic molecules on the surface of Mars
- Will employ advanced entry, descent and landing techniques
  - 21.5 m diameter supersonic parachute
  - Powered descent vehicle with 8 Mars Landing Engines (MLEs)
  - Sky-crane tethered landing of rover

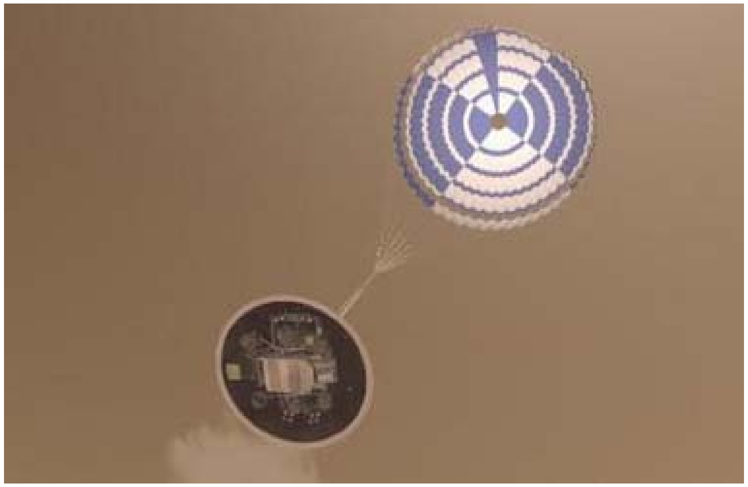

The Mars Science Laboratory supersonic parachute design

- A Parachute Decelerator System (PDS) provides the most efficient means of slowing an entry vehicle from supersonic to subsonic speeds
  - Prepares the vehicle for safe heatshield separation
  - Prepares the vehicle for powered descent
    - Attitude and velocity

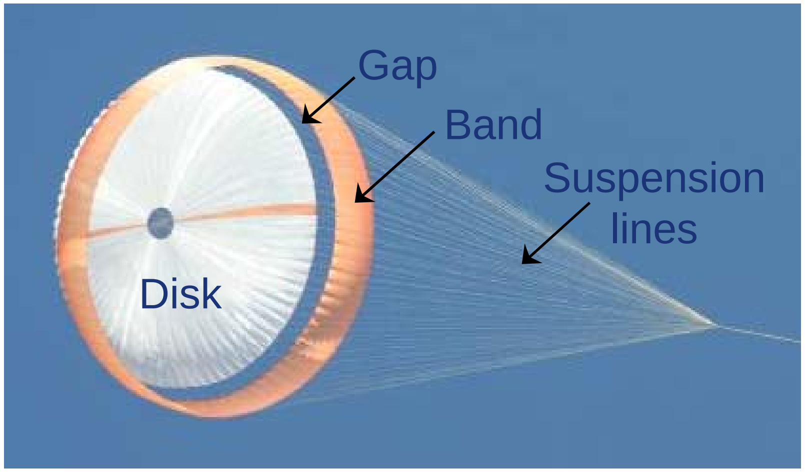

- MSL PDS characteristics
  - 21.5 m Viking-type Disk-Gap-Band type parachute
    - Viking parachute was 16.1 m
  - Similar capsule to parachute diameter scaling as Vikings
  - Deployed at Mach 2.3
  - Limits time above Mach 1.5 (~10s)
  - Modern materials – nylon, Kevlar and Technora

The Mars Science Laboratory supersonic parachute design

- Supersonic parachute instability
  - In 1960s, high altitude (~50 km) supersonic parachute tests were performed
    - Showed canopy instabilities at Mach Numbers above 1.5
      - Partial inflations and collapses of the parachute, termed “Area Oscillations”
      - Resulting in projected area and drag fluctuations of the canopy
    - Resulted in load spikes after the first full open

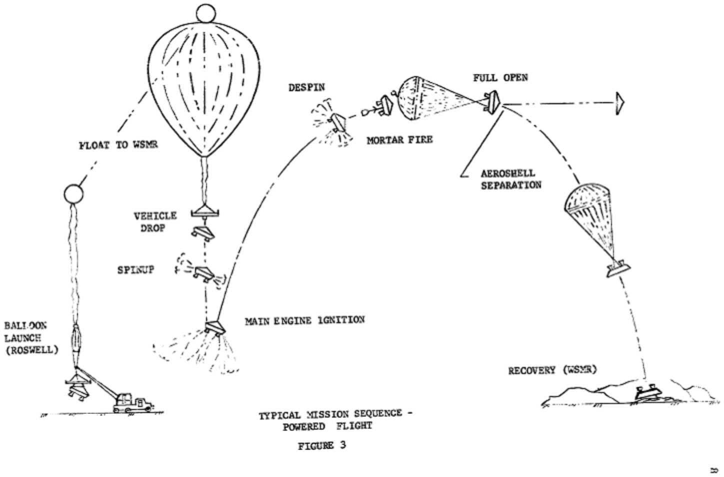

TYPICAL MISSION SEQUENCE - POWERED FLIGHT

FIGURE 3

The Mars Science Laboratory supersonic parachute design

- Scaling up from Viking to MSL
  - A simulation capability was developed to extrapolate the Viking data to the larger scale and with different materials
    - Alternative would have been expensive high-altitude tests
  - Aerodynamic and dynamic performance of the MSL parachute in the supersonic regime is determined from
    - Subscale wind tunnel testing
    - Computational simulations

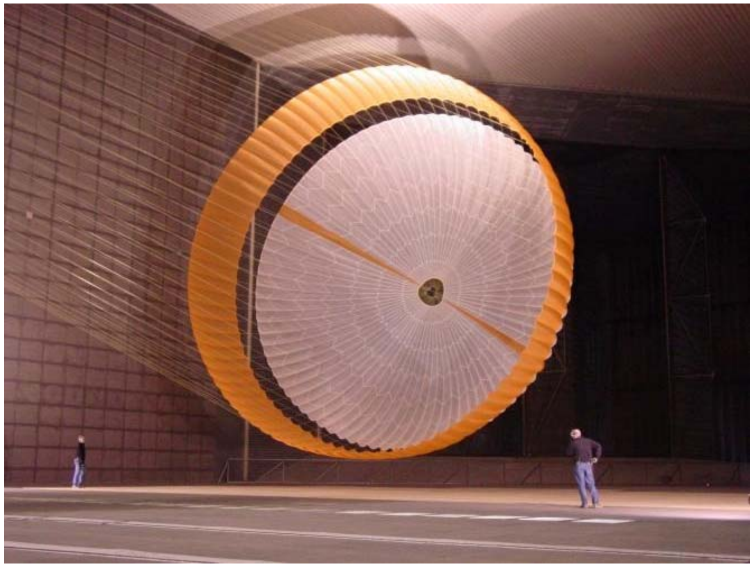

The Mars Science Laboratory supersonic parachute design

- Computational qualification approach
  - Developed CFD and FSI tools to model the physics of interest
  - Validated the simulations using wind tunnel data from scaled models
    - Piecewise validation approach
      - Capsule only
      - Rigid parachute only
      - Capsule and rigid parachute
      - Capsule and flexible parachute
  - Used simulations to explore the parachute behavior at different sizes and materials
    - Validate the Viking parachute behavior over a range of sizes, materials and flight conditions
    - Extrapolate to the MSL parachute size, materials and flight conditions

The Mars Science Laboratory supersonic parachute design

- The parachute simulation application
  - Based on the Virtual Test Facility (VTF), a CFD/FEM toolkit
    - Originally developed at the California Institute of Technology for the Department of Energy
    - Further developed by University of Illinois and Cambridge University
  - Uses a 3-D Large Eddy Simulation solver coupled to FEM solver
    - Fluid is simulated using unsteady, compressible, large-eddy simulations in an Eulerian-Cartesian mesh ~ 50 million cells
    - Canopy is simulated using large-deformation thin-shell Kirchhoff-Love finite elements on a Lagrangian mesh ~ 10,000 elements
    - Four levels of adaptive mesh refinement are used for finer resolution as needed
  - Validated by comparison to the (4%) scaled wind-tunnel experiments
  - Following validation, the code is being used to simulate the full-scale parachute

The Mars Science Laboratory supersonic parachute design

- Simulations
  - Simulation domain is [-3,5] x [-1,1] x [-1,1] m
    - Initially run on a Dell-Myrinet Xeon cluster (64 to 206 processors)
    - Subsequently run on an SGI Altix 3700 system with 96 processors allocated to the fluid and 4 to the solid
  - Simulation results showed that the parachute supersonic behaviour and performance were as expected

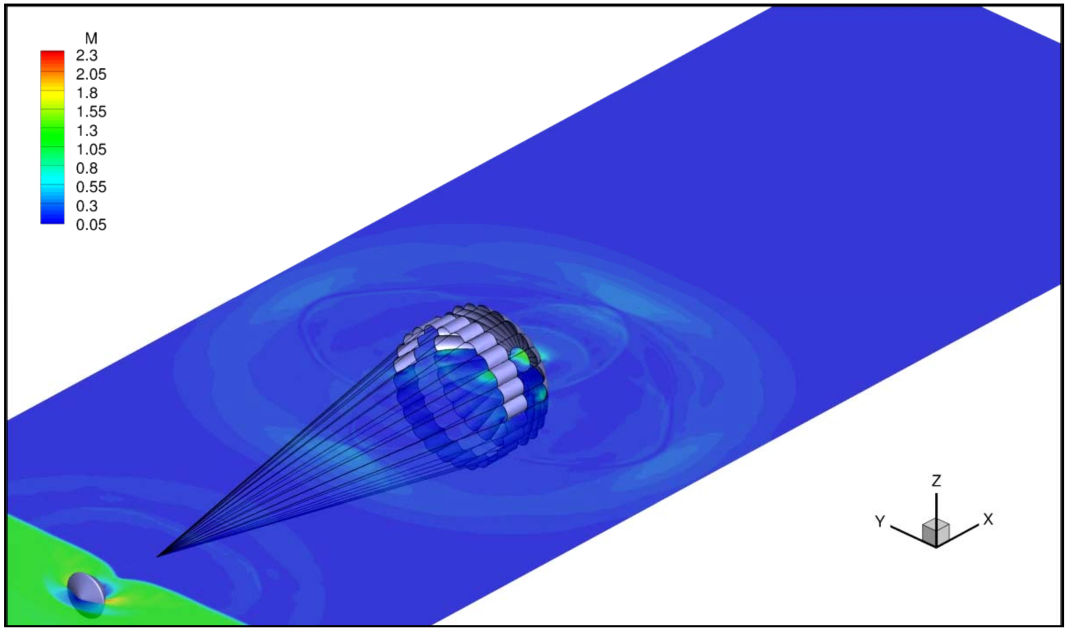

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Juno planetary protection trajectory analysis

- The Juno spacecraft will be launched in August 2011
    - Mission is to orbit Jupiter to study its origin and evolution
    - Will measure Jupiter’s gravity field, and explore the Jovian atmosphere and magnetosphere
- Juno’s highly eccentric orbit could lead to potential impact with the Galilean satellites (Io, Europa, Ganymede and Callisto)
    - These large icy moons are of particular interest for future exobiology and astrobiology exploration
        - Potentially contain biological and/or organic materials

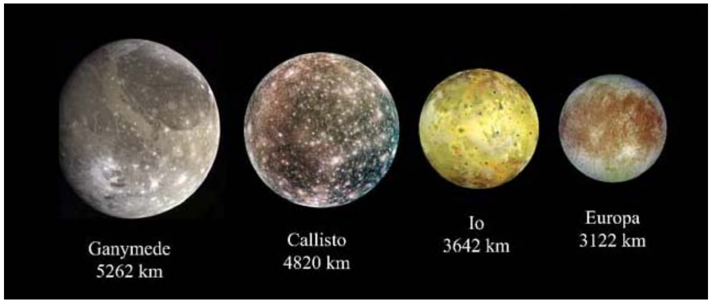

Juno planetary protection trajectory analysis

- Planetary protection requirements dictate that during its prime mission Juno must not collide with any of the Galilean satellites
    - Any collision would cause contamination that would jeopardize future explorations
- Juno’s planned mission is for one year, after which it will be de-orbited into Jupiter’s atmosphere
    - In case de-orbiting is unsuccessful, planetary protection requirements must be met for a further period of 150 years

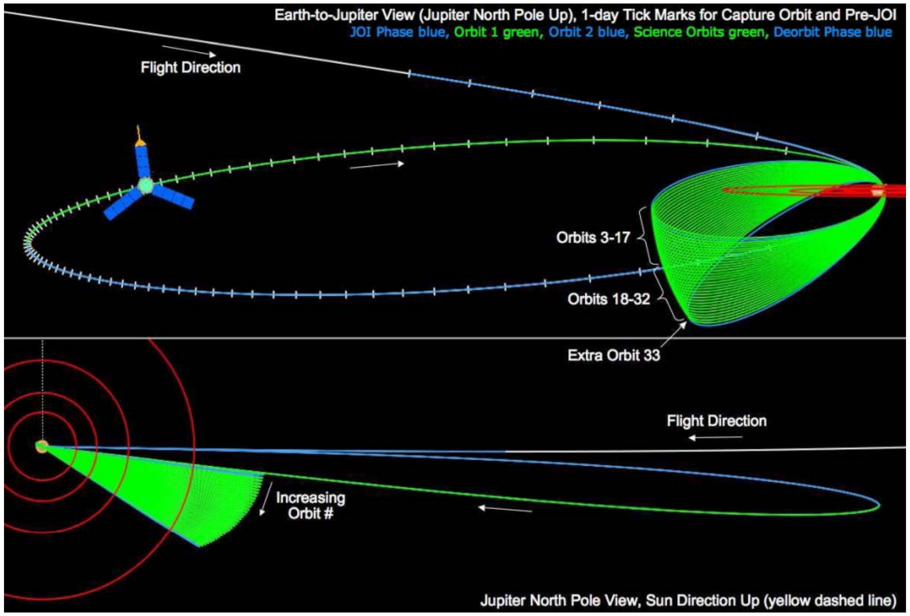

Juno planetary protection trajectory analysis

- Monte-Carlo techniques were employed to determine the collision probabilities
  - Required the analysis of thousands of trajectory states for hundreds of years for each case
  - The wall-clock time for a single trajectory propagation was about 10 hours
  - On one CPU, a single case would have taken over a year to complete
  - Were able to complete each Monte-Carlo run in less than 12 hours, instead of the estimated year

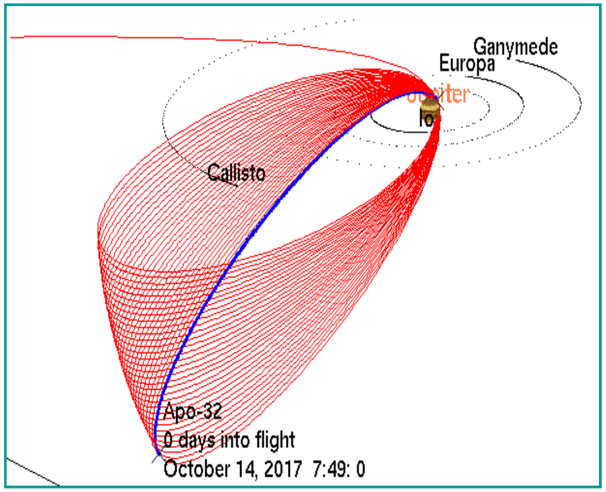

Example of a Juno trajectory

- HPC enabled the investigation of many failure scenarios and potential baseline trajectories

Outline

- Introduction
  - JPL and its mission
  - Current flight projects
- HPC resources at JPL
  - Institutional HPC resources
  - HPC resources at NASA Ames
- Examples of HPC usage by flight projects
  - Entry, descent and landing simulations
  - The Phoenix Mars Lander radar ambiguity
  - Mars Science Laboratory supersonic parachute design
  - Juno planetary protection trajectory analysis
- Future work
  - Evolutionary computing
    - Low-thrust orbit optimization

Evolutionary computing

- Evolutionary computing seeks to mimic processes used in nature to optimize multi-parameter engineering designs
  - Uses sophisticated biological operators
    - Selection
    - Mutation
    - Recombination
- Advantages
  - Enables larger design spaces to be explored than could be examined manually or by computational brute force
  - Results have shown competitive advantages over human-created designs

Low-thrust orbit optimization

- Low-thrust propulsion is more efficient than chemical propulsion
  - Uses less propellant
  - Demonstrated on Deep Space 1 and currently on Dawn
- Requires different trajectory optimization techniques
  - Involves many revolutions and continuous thrust arcs
- Goal is to optimize the trade-off between propellant mass and flight time for orbit transfers
  - Thrust angles and thrust arcs are optimized with Q-law
  - Q-law has 12 independent parameters
- In this work, the Q-law parameters are optimized using evolutionary algorithms
- Evolutionary algorithms are amenable to parallel computing implementation

Low-thrust orbit optimization

Circle-to-circle orbit transfers

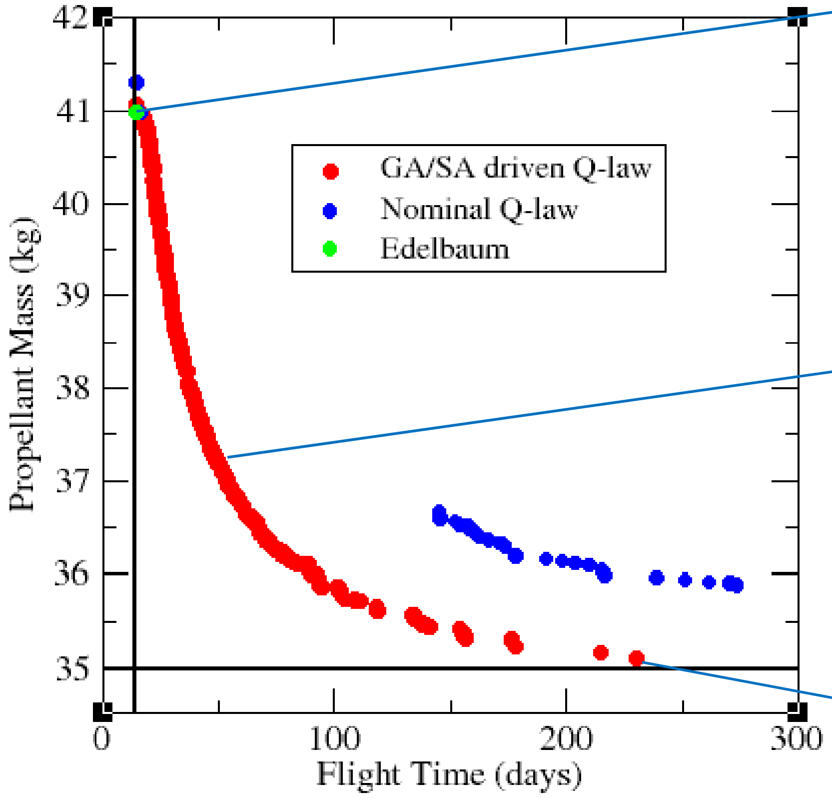

Flight-time optimal
Edelbaum transfer

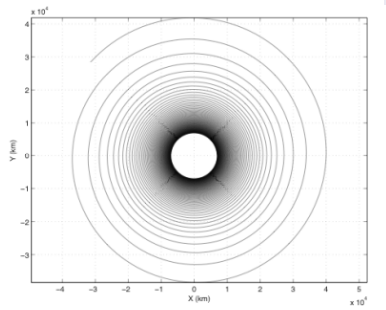

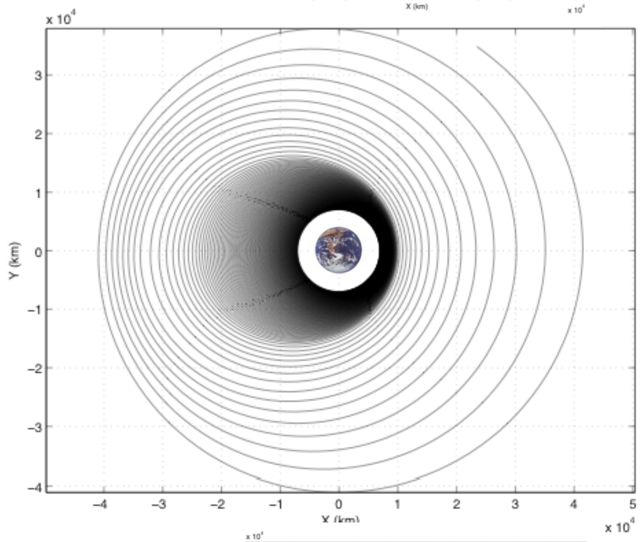

Propellant optimal
Hohmann transfer

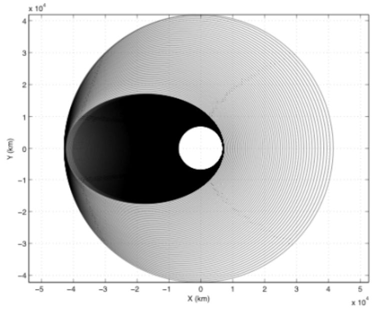

References

- Desai, P.N., et al., “Mars Exploration Rovers Entry, Descent and Landing Trajectory Analysis”, AIAA paper 2004-5092, 16 - 19 August 2004.
- Skulsky, E.D., et al., “Under the Radar: Trials and Tribulations of Landing on Mars”, AAS paper #08-043, 1 - 6 February 2008.
- Sengupta, A., et al., “Results from the Mars Science Laboratory Parachute Decelerator System Supersonic Qualification Program”, IEEEAC paper #1435, Version 3, Updated December 20, 2007.
- Lam, T., et al., “Planetary Protection Trajectory Analysis for the Juno Mission”, AIAA paper 2008-7368, 18 - 21 August 2008.
- Terrile, R.J., et al., “Evolutionary Computation Technologies for Space Systems”, IEEEAC paper #1257, Version 6, Updated December 29, 2004.
- Lee, S,. et al., “Design and Optimization of Low-Thrust Orbit Transfers”, IEEEAC paper #2.1204, 5 - 12 March 2005.

Thank you!

Questions and/or comments?

Contact information:
Christopher.J.Catherasoo@ jpl.nasa.gov
+1 818 354 7180

Backup Slides

End-to-end mission development

Project formulation

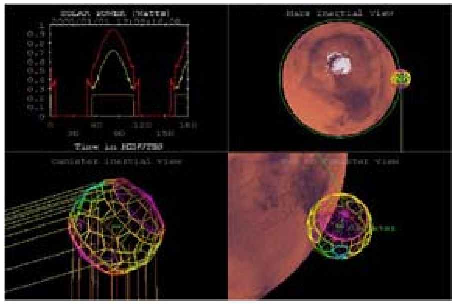

Mission design

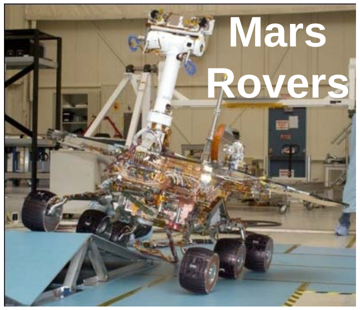

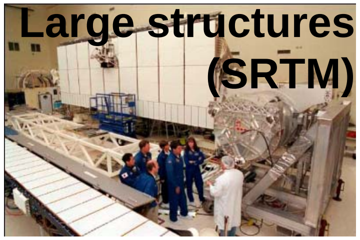

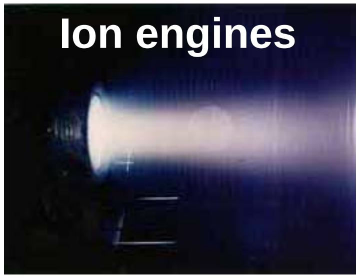

Spacecraft development

Integration and test

Environmental test

Real time operations

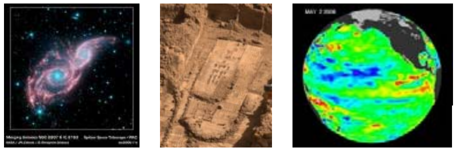

Scientific research

The Mars Science Laboratory supersonic parachute design

- Geometry and mesh
  - Fluid flow region had approximately 50 million cells
  - Canopy had 92,016 finite elements
  - Four levels of mesh refinement were used
  - The grid was iteratively designed with knowledge of the wake and shock structures
  - Well-designed grids are essential for correct representation of flow fields

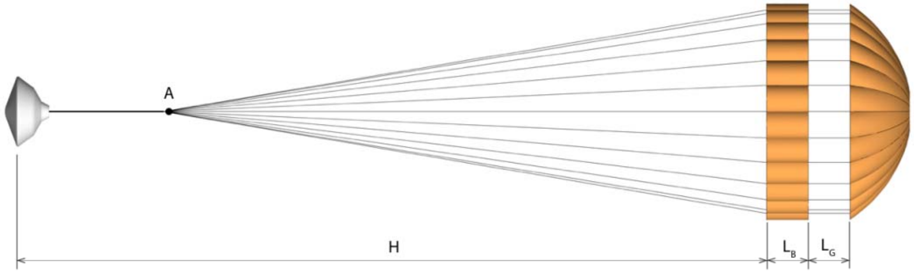

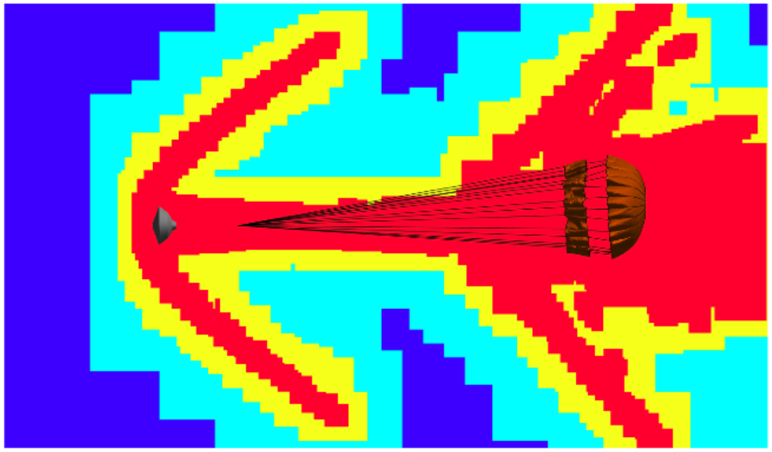

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Critical engineering achievements using HPC

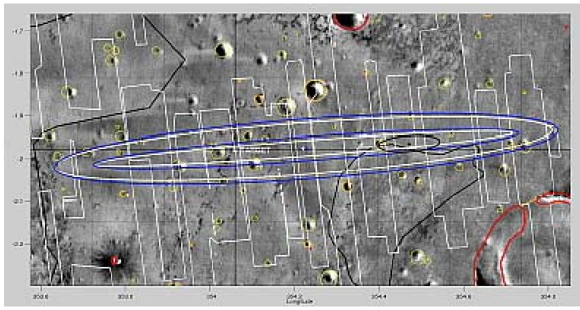

Landing-site analyses

CIELO high-fidelity integrated thermal/ structural/optical aberration analyses

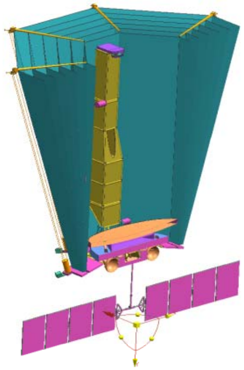

Entry, descent and landing simulations

Supersonic parachute design coupled CFD-FEM simulations

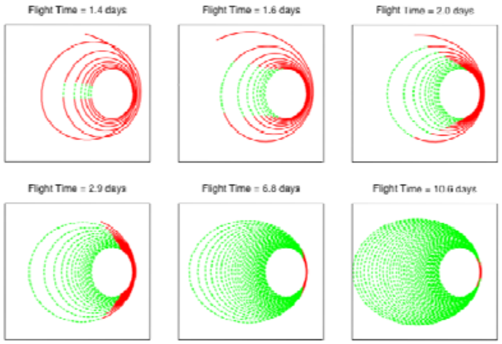

Optimization of low-thrust orbit transfers using evolutionary techniques

Juno spacecraft impact probabilities with Jupiter’s Galilean satellites