Supporting Science @ NAS: HPC, Data Analysis and Collaboration

Piyush Mehrotra

Chief, NASA Advanced Supercomputing (NAS) Division

NASA Ames Research Center, Moffett Field, Calif., USA

March 22, 2014

Advanced Computing @ NAS

In addition to production supercomputing for NASA S&E applications, also researching, evaluating, and developing candidate advanced computing technologies for maturation

NASA's Diverse HPC Requirements

High throughput / capacity
Engineering requires HEC resources that can handle large ensembles of moderate-scale computations to efficiently explore design space

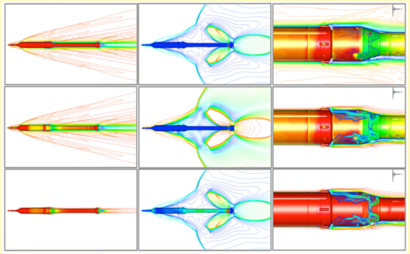

Aerodynamic database generation

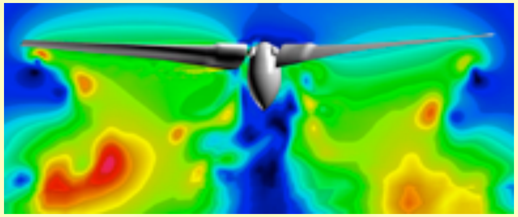

Rotary wing aerodynamics

Leadership / capability Research requires HEC resources that can handle high-fidelity long-running large-scale computations to advance theoretical understanding

High availability
Time-sensitive mission-critical applications require HEC resources on demand. readily available

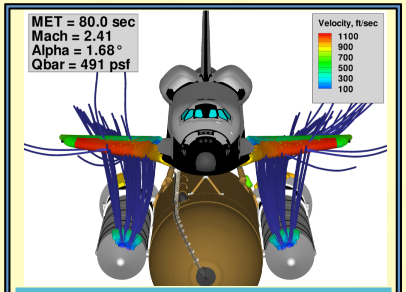

Debris transport analysis

NAS provides the primary supercomputing resources to meet all of NASA's HPC requirements

Strategic Support for NASA Programs

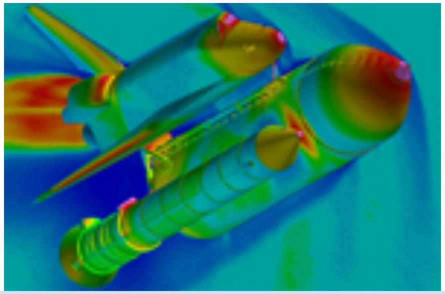

External tank redesign

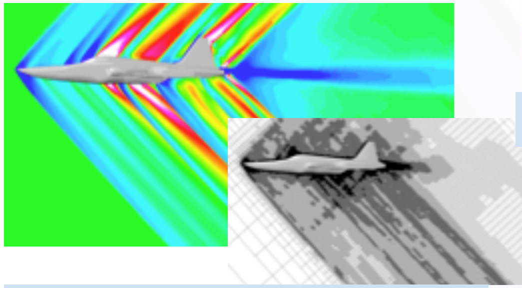

Sonic boom optimization

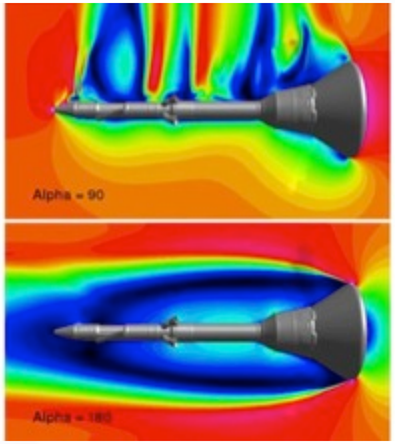

Launch abort system

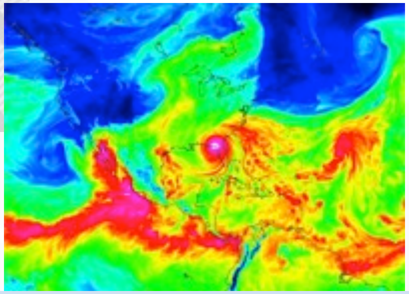

Hurricane prediction

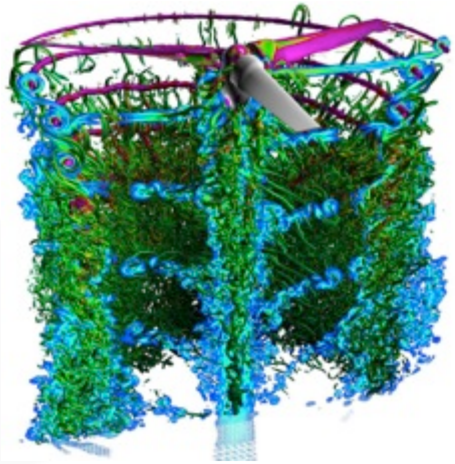

Rotary wing aerodynamics

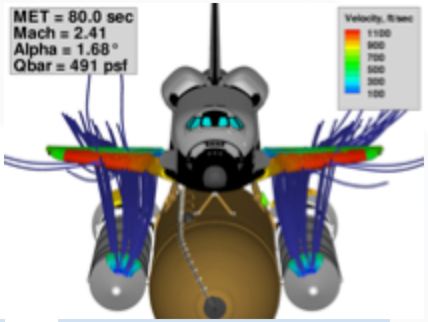

Debris transport

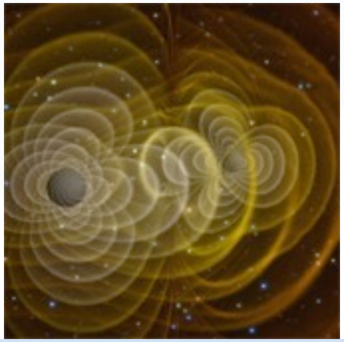

Merging black holes

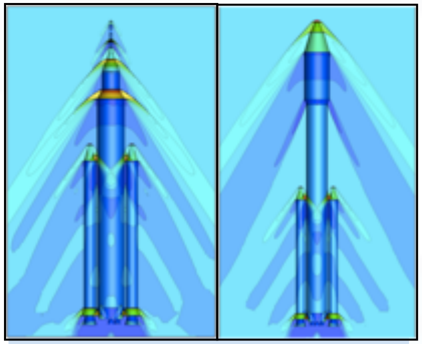

SLS vehicle designs

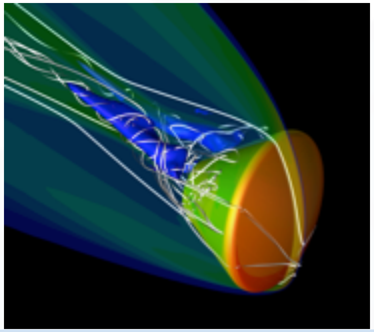

Orion/MPCV reentry

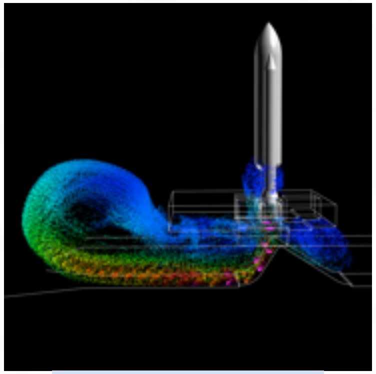

Flame trench

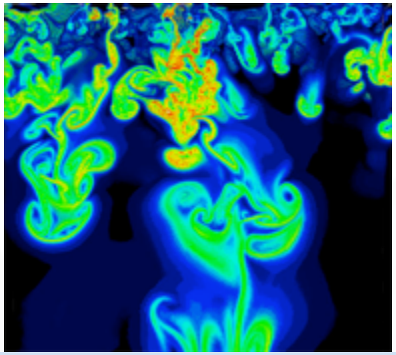

Solar magnetodynamics

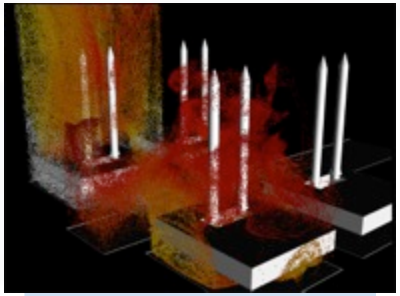

SRB burn in VAB

High-End Computing Capability (HECC)

Supercomputing to support NASA Science & Engineering

Computing Systems

- **Pleiades** – 3.6 PF peak
  - 184,800-core 11,176 node SGI Altix ICE
  - 3 generations of Intel Xeon: Westmere (6c), Sandy Bridge (8c), Ivy Bridge (10c)
  - 163 racks, 502TB of memory
  - 2 racks have 64 Nvidia M2090 GPU-enhanced nodes
  - #16 on TOP500 (11/13)
- **Endeavour** – 32 TF peak
  - 2 SGI UV 2000 Sandy Bridge based nodes
    - 1024-core with 4 TB shared memory
    - 512-core with 2 TB shared memory
  - Single System Image (SSI) via NUMALink-6
- **Maia** – 300 TF peak
  - SGI Rackable 128 nodes, each with two Sandy Bridge processors and two 60-core Phi accelerators
- **Mira** – 150 TF peak
  - Cray 64 nodes, each with two two Sandy Bridg processors and two 60-core Phi accelerators

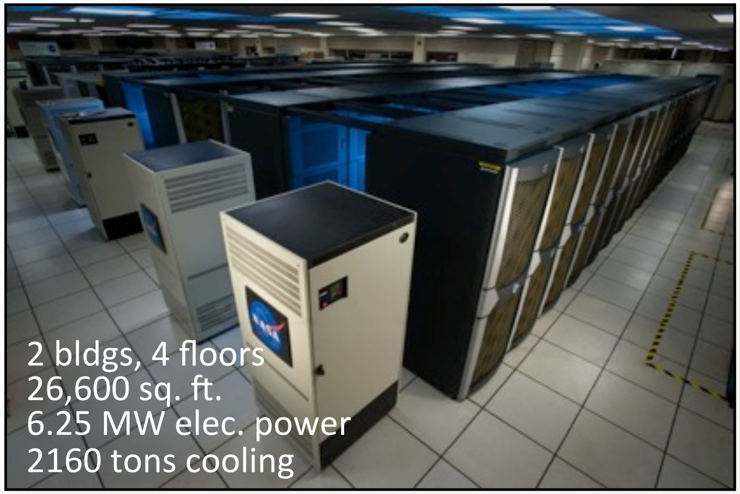

Balanced Environment

- Storage: 20.1 PB disk; 115 PB tape
  - Archiving ~1 PB/month
- WAN: 10 Gb/s external peering
  - Transferring 200 TB/month to users
- Resources enable broad mission impact
  - Mission Directorates select projects, determine allocations
  - Over 500 S & Eprojects
  - More than 1,200 user accounts

Advanced Visualization: hyperwall-2 and Concurrent Visualization

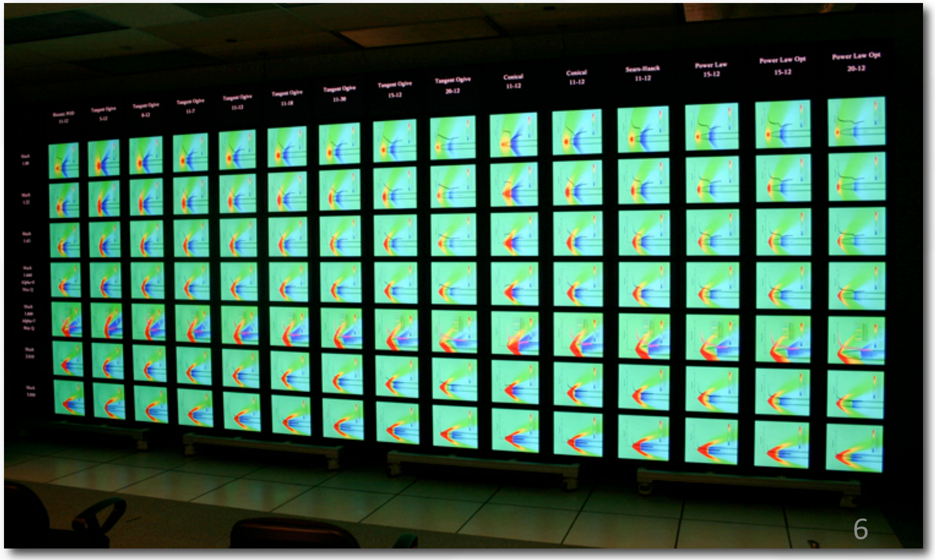

- Supercomputer-scale visualization system to handle massive size of simulation results and increasing complexity of data analysis needs
  - 8x16 LCD tiled panel display (23 ft x 10 ft)
  - 245 million pixels
  - IB interconnected to NAS supercomputer
- Two primary modes
  - Single/multiple large high-definition image(s)
  - Sets of related images (e.g., a parameter space of simulation results)
- *Traditional Post-processing:* Direct read/write access to Pleiades file systems eliminates need for copying large datasets
- *Concurrent Visualization:* Runtime data streaming increases temporal fidelity at much lower storage costs:
  - ECCO: images every integration time step as opposed to every 860+ time steps originally

Simulation: Padoan (UCSD)
Viz: Henze/Sandstrom (NAS)

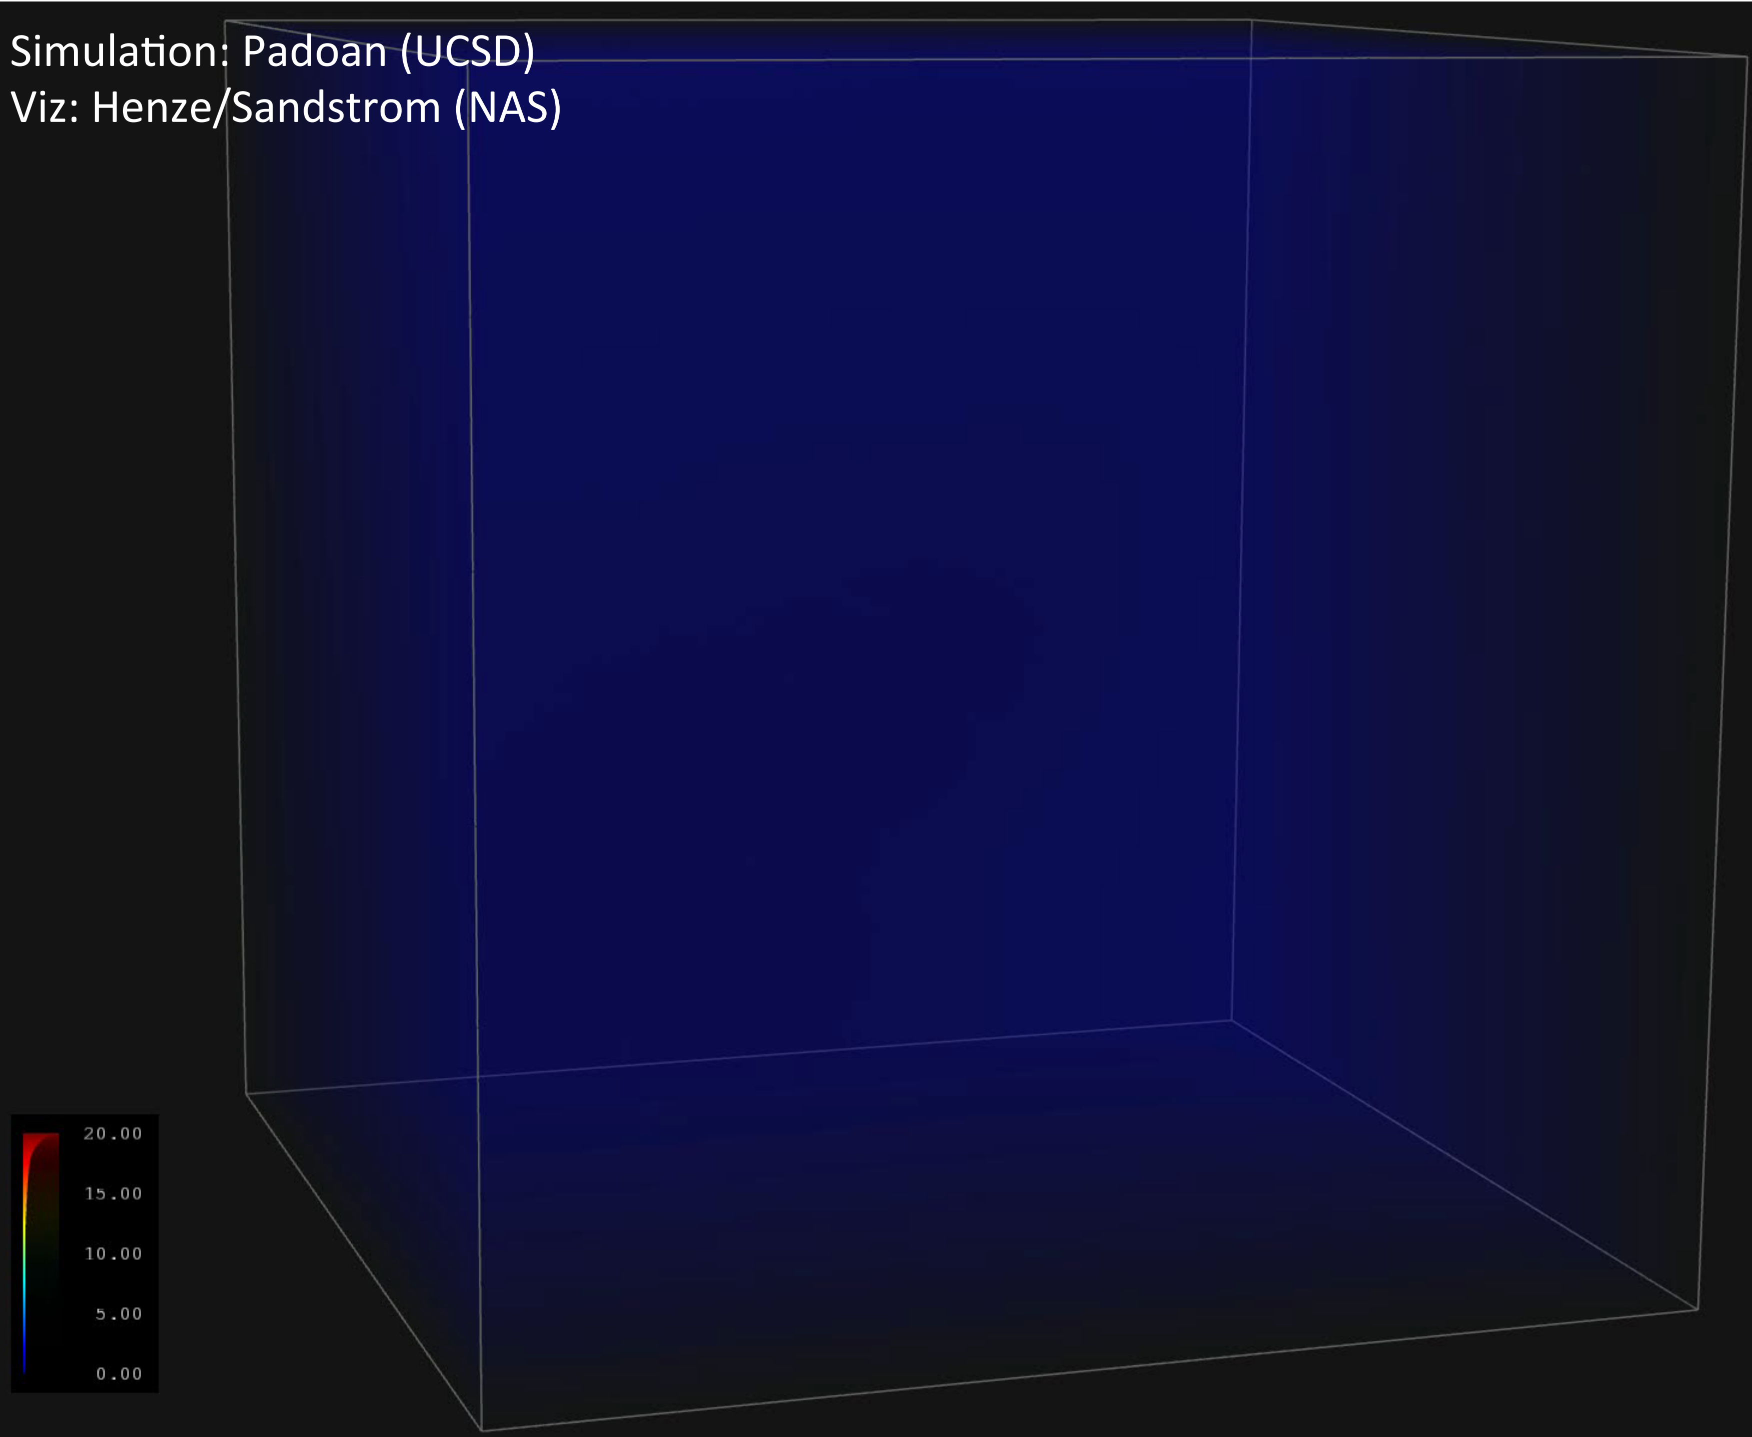

Simulation: Klein (Berkely)
Viz: Henze/Sandstrom (NAS)

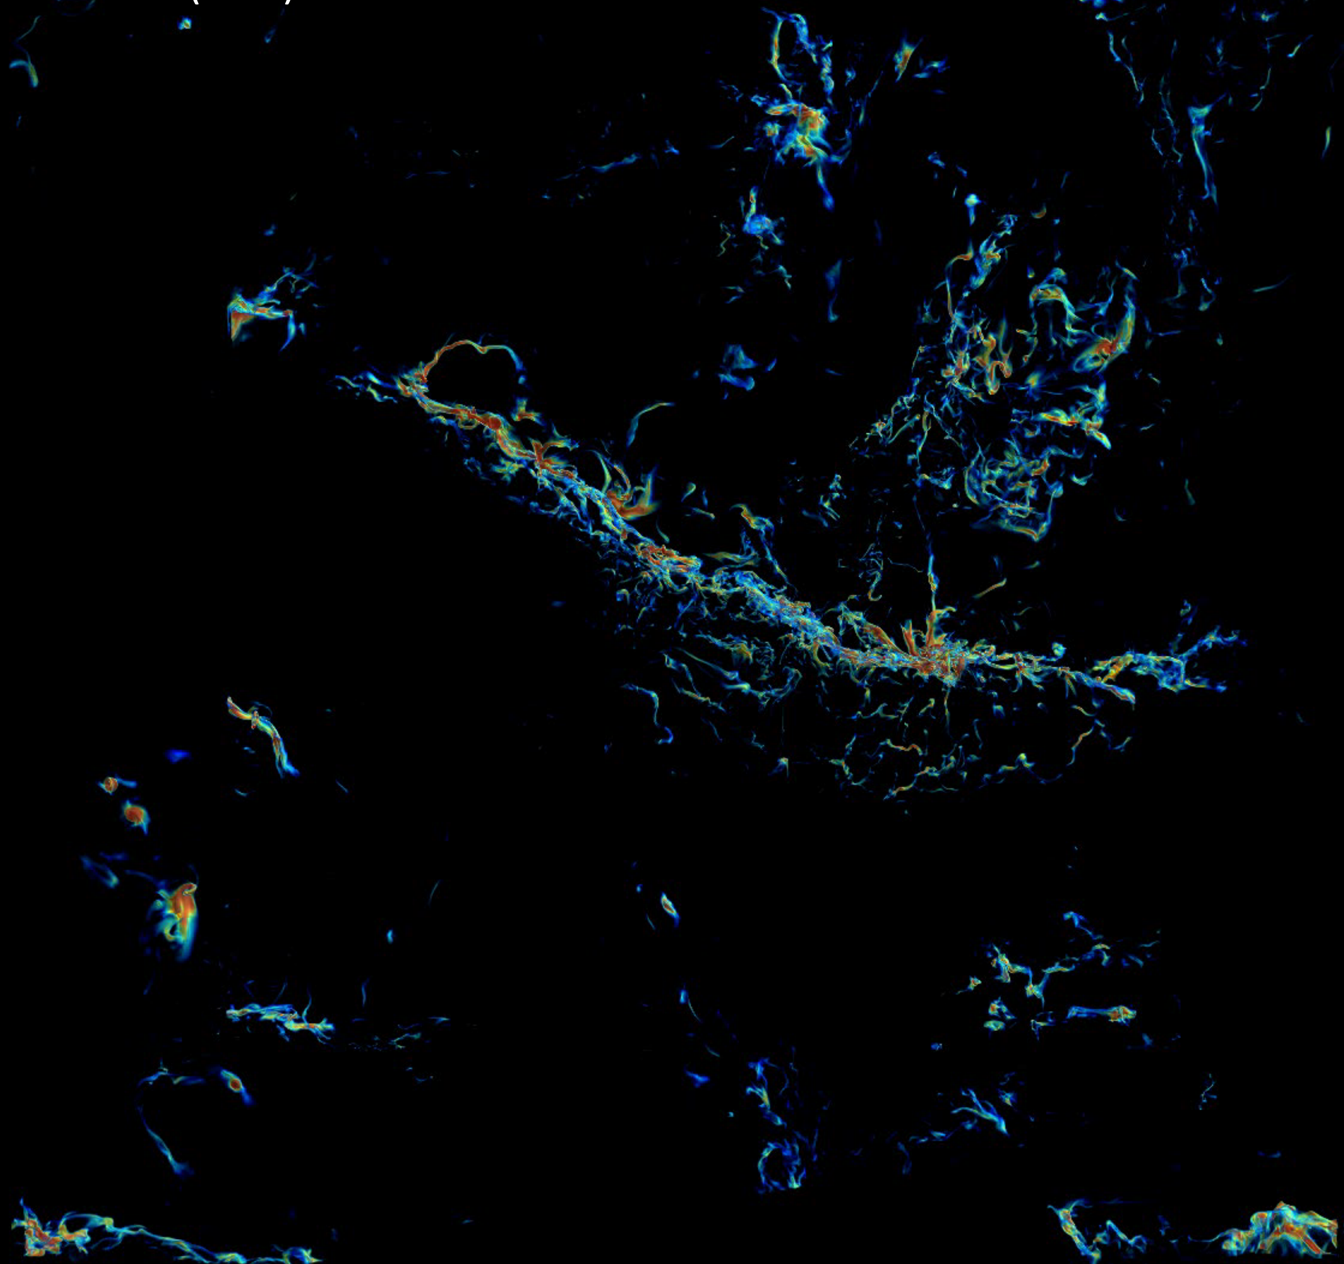

Big Data @ NASA

- NASA has enormous collections of observational and model data

- Model/Simulation Data:
  - NAS has 20+ PB storage; 115 PBs archive storage & archiving 1+ PB per month
    - MITgcm 35K core run produced 1.44 PB in its 5 day run; full run will produce 9-18 PB

- Observational Data:
  - Estimate 100+ active satellites producing 50PBs per year
    - Solar Dynamics Observation (SDO) satellite produces 1 GB per minute => > 1/2 PB/ year ; ~ 3PB in its 5 year life cycle
  - NASA Earth Science operates 12 DAACs (archive centers); National Space Science Data Center

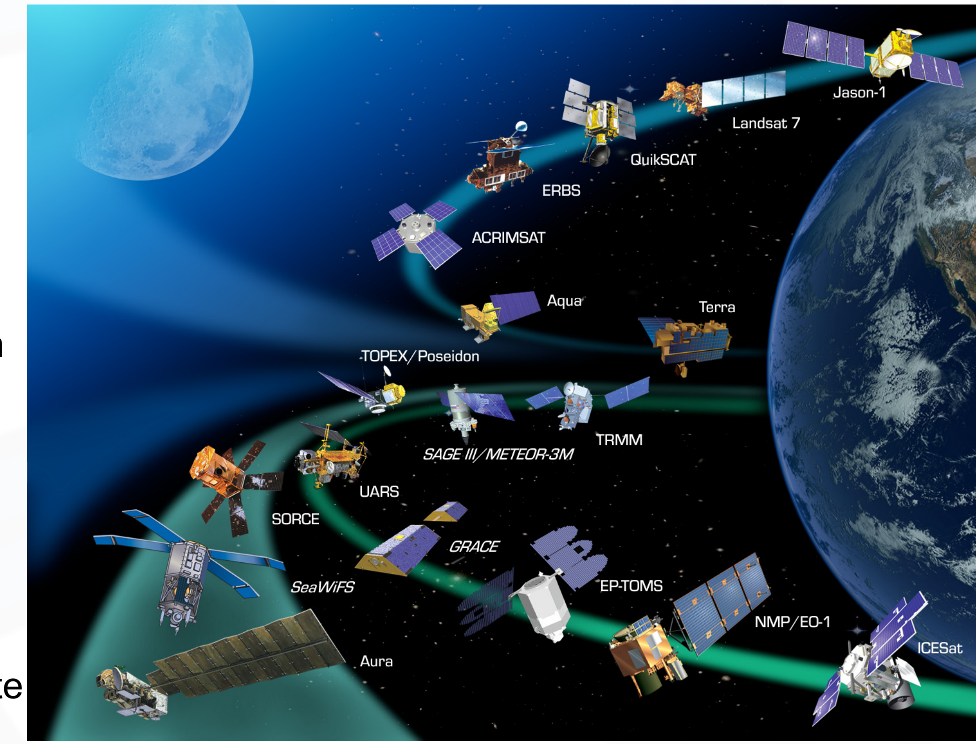

Big Data Effort @ NAS

Focus: *Develop and implement a roadmap for an infrastructure to support analysis & analytics*

- Conducted survey of projects dealing with big data (Available under Publications on http:www.nas.nasa.gov)
- Big Data Challenges:
  - *Data management* – storage/access/transport
  - *Data discovery* - Indexing/archiving, metadata – requires semantic reasoning
  - *Tools/models/algorithms* - development & discovery
  - *Data Analysis/Analytics infrastructure*
    - Most NASA data is structured, gridded, geospatial
    - Shared memory systems with large I/O pipes; data preferably co-located with compute
    - Visualization support
  - *Workflow description and management* - to tie all components together
  - *Collaboration environments*
    - Dissemination and sharing of results/tools/models/algorithms
- Currently conducting deep dives and prototype experiments
- Develop requirements and an environment to support the entire workflow of data scientists

NASA Earth Exchange (NEX)

NEX A collaborative environment that brings scientists and researchers together in a knowledge-based social network along with tools, computing power and data to accelerate research, innovation and provide transparency.

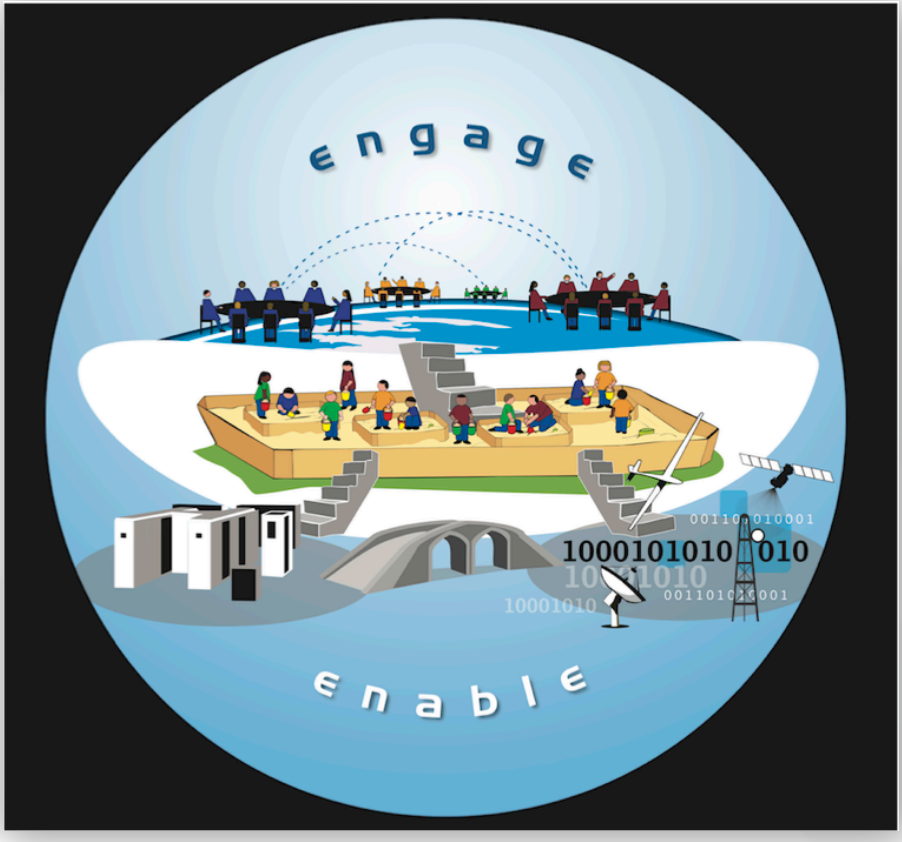

VISION

To provide "*science as a service*" to the Earth Science community addressing global environmental challenges.

GOAL

To improve efficiency and expand the scope of NASA Earth science technology, research and applications programs.

Project Manager: Piyush Mehrotra
Principal Scientist: Ramakrishna Nemani

NEX Environment

Collaboration Portal
(Over 400 members)

Tools
Models
Workflows
Papers
Data

Compute Resources

Sandbox for experimentation

HPC

Centralized Data Repository
(over 1.3 PBs of Data)

Working copies of observational data
Project data

- Collaboration portal open to all Earth Scientists
- Sandbox currently available to a subset with NASA credentials
- HPC resources available only to approved projects with allocation
- OpenNEX a collaboration with Amazon to provide access to some NEX data
- In the process of setting up an OpenSandbox to allow access to all NEX datasets along with some compute resources to a wider Earth Science community

QuAIL
(Quantum Artificial Intelligence Lab)

- Collaboration between NASA / Google/ USRA
- D-Wave 2 system installed at
  – Vesuvius processor – 512 qubits (quantum bits – niobium superconducting loops encoding 2 magnetic states)
  – Physical characteristics
    - 10 kg of metal in vacuum at 15 mK
    - Uses 12 kW electrical power
    - Magnetic shielding to 1 nanoTesla (50,000x less than Earth's magnetic field)
  – Focused on solving discrete optimization problems via quantum annealing based on the quantum physical principles of superposition, entanglements, tunneling

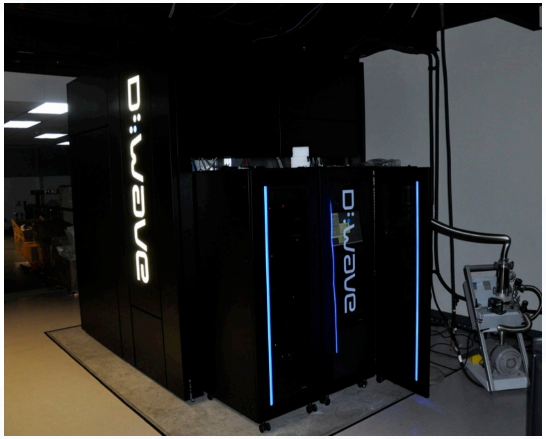

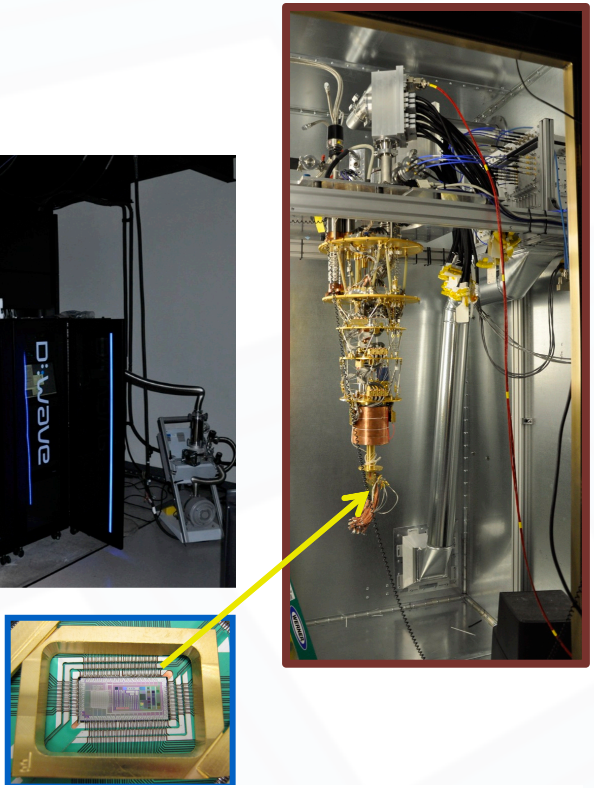

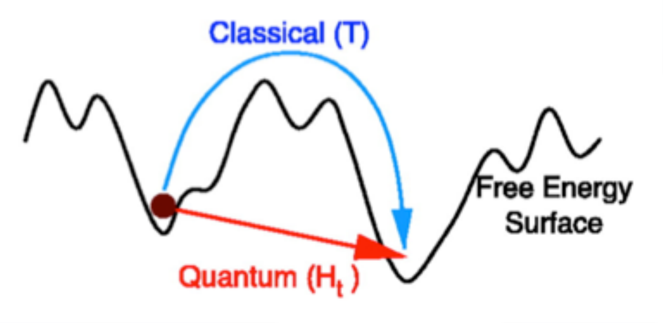

How to Program D-Wave Two

1 **Map the target combinatorial optimization problem into Quadratic Unconstrained Binary Optimization (QUBO) form**

No general algorithms, smart mathematical tricks (penalty functions, locality reduction..)

$$\alpha_{ijk} z_i z_j z_k$$

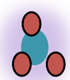

$$\alpha_{ijk} y_{ij} z_k +$$

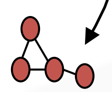

$$\beta_{ijk} \left(3y_{ij} - 2z_i y_{ij} - 2z_j y_{ij} + z_i z_j\right)$$

$$\sum_{ij} Q_{ij} z_i z_j$$

2 **Embed the QUBO coupling matrix in the hardware graph of interacting qubits**

The D-Wave hardware qubit connectivity is a "Chimera Graph", so embedding methods mostly based on heuristics

Q_{ij} =

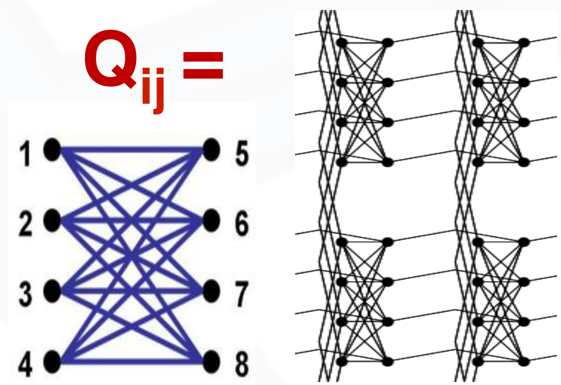

Note: D-Wave provides a heuristic blackbox compiler that bypasses embedding

3 **Run the problem many times and collect statistics**

Use symmetries, permutations, and error correction to eliminate the systemic hardware errors and check the solutions

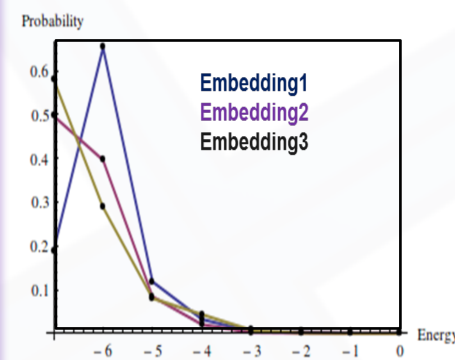

NASA and Quantum Computing

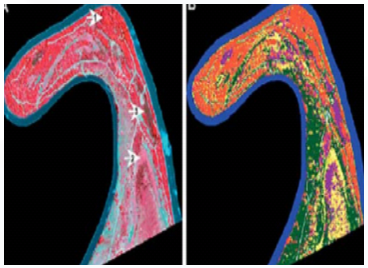

Data Analysis and Data Fusion

Anomaly Detection and Decision Making

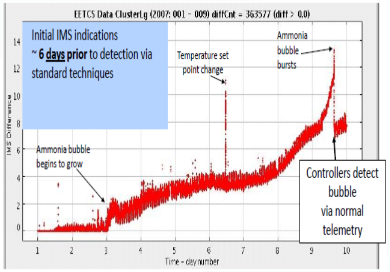

Air Traffic Management

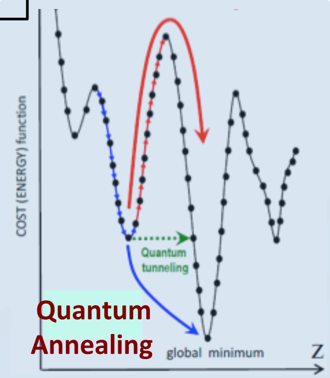

V&V and optimal sensor placement

Mission Planning and Scheduling, and Coordination

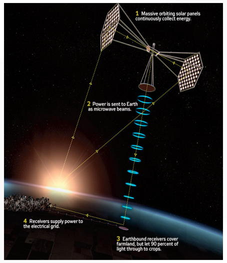

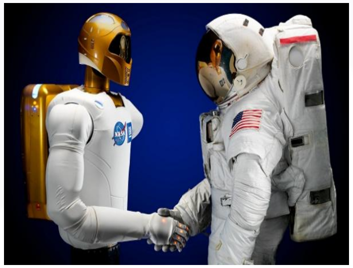

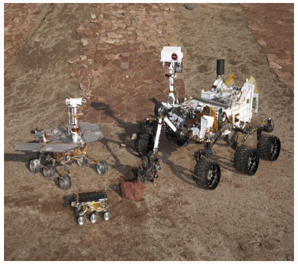

Topologically aware Parallel Computing

Thank you!

Questions?

piyush.mehrotra@nasa.gov