National Aeronautics and Space Administration

NASA Advanced Computing Environment for Science & Engineering

Dr. Rupak Biswas

Deputy Director, Exploration Technology Directorate
Manager, High End Computing Capability (HECC) Project

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

51st IDC HPC User Forum, Seoul, South Korea

1 October 2013

NASA Overview: Mission Directorates

- **Vision:** *To reach for new heights and reveal the unknown so that what we do and learn will benefit all humankind*
- **Mission:** *To pioneer the future in space exploration, scientific discovery, and aeronautics research*
- **Aeronautics Research (ARMD):** Pioneer and prove new flight technologies for safer, more secure, efficient, and environmentally friendly air transportation
- **Human Exploration and Operations (HEOMD):** Focus on ISS operations; and develop new spacecraft and other capabilities for affordable, sustainable exploration beyond low Earth orbit
- **Science (SCMD):** Explore the Earth, solar system, and universe beyond; chart best route for discovery; and reap the benefits of Earth and space exploration for society
- **Space Technology (STMD):** Rapidly develop, demonstrate, and infuse revolutionary, high-payoff technologies through transparent, collaborative partnerships, expanding the boundaries of the aerospace enterprise

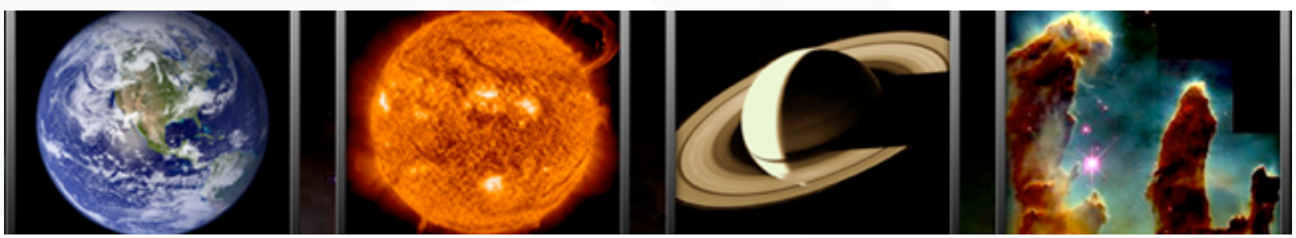

NASA Overview: Centers & Facilities

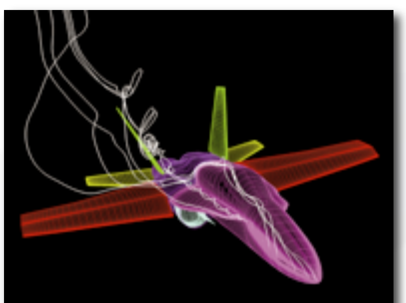

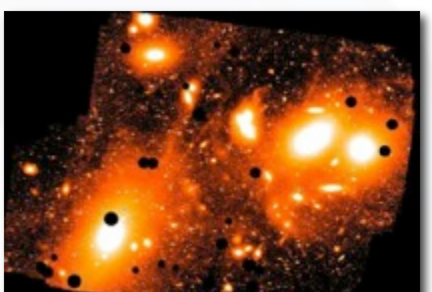

Need for Advanced Computing

Enables modeling, simulation, and analysis (MS&A)

– Digital experiments and physical experiments are tradable
– Physical systems and live tests are generally expensive & dangerous (e.g., extreme environments), require long wait times, and offer limited sensor data
– NASA collects and curates precious and limited observational (science) data; mission does not stop at collecting data, but requires understanding observed systems

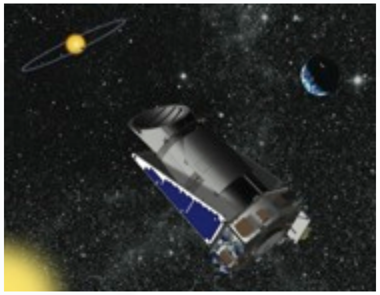

– Decades of exponentially advancing computing technology has enabled dramatic improvements in cost, speed, accuracy of MS&A – and provides predictive capability
– Numerous studies conclude that simulation is key to progress in science & engineering, and level of complexity is unattainable by strictly theoretical or experimental methods
– Aeronautics, Earth and space sciences, and human and robotic space exploration all require orders-of-magnitude increase in MS&A capability to enhance accuracy, reduce cost, mitigate risk, accelerate R&D, and heighten impact

MS&A essential to rapidly and cost-effectively advance NASA goals

Advanced Computing Environment

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

1) Engineering requires HPC resources that can process large ensembles of moderate-scale computations to efficiently explore design space **(high throughput / capacity)**
2) Research requires HPC resources that can handle high-fidelity long-running large-scale computations to advance theoretical understanding **(leadership / capability)**

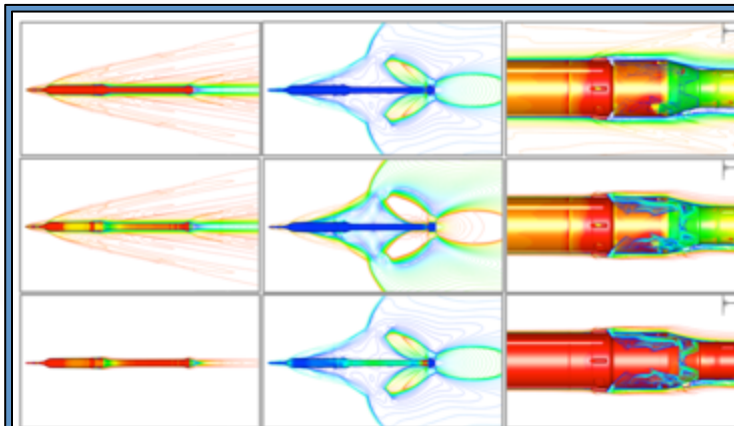

Aerodynamic database generation

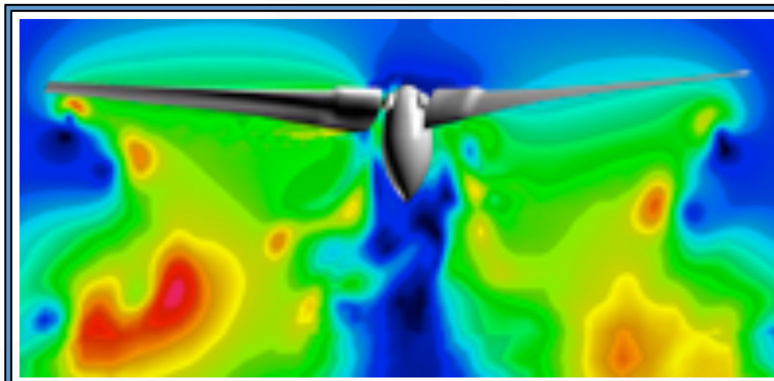

Rotary wing unsteady aerodynamics

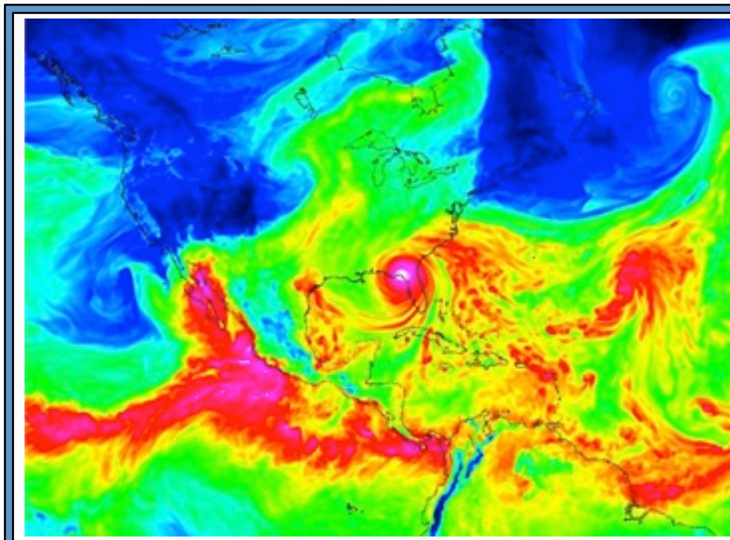

Real-time hurricane prediction

3) Time-sensitive mission-critical applications require HPC resources on demand **(high availability / maintain readiness)**

Integrated Spiral Support for MS&A

Develop and deliver the most productive integrated supercomputing environment in the world, enabling NASA to extend technology, expand knowledge, and explore the universe

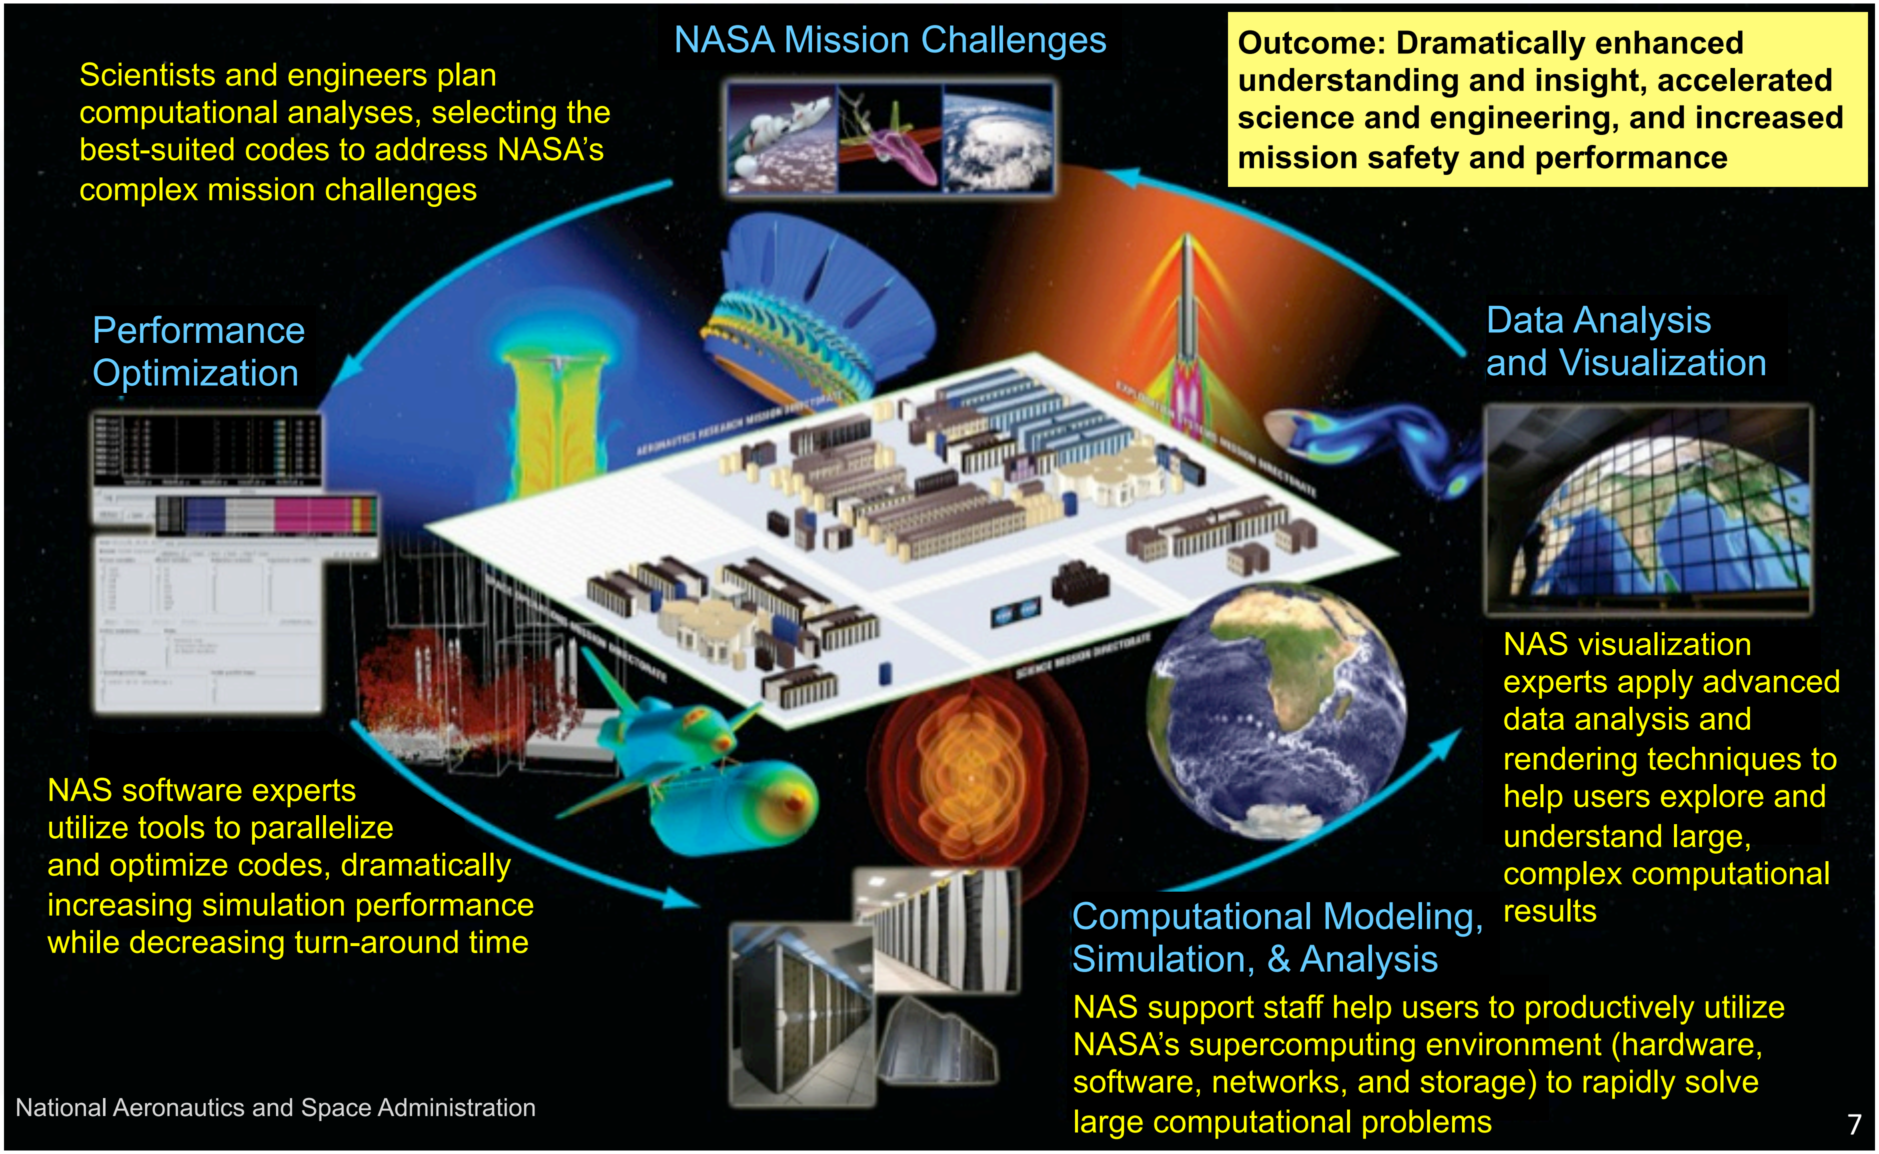

Supercomputing to Enable NASA Science and Engineering

Computing Systems

- **Pleiades** – 2.9 PF peak
  - 162,496-core SGI Altix ICE
  - 4 generations of Intel Xeon: Nehalem (4c), Westmere (6c), Sandy Bridge (8c), Ivy Bridge (10c)
  - QDR/FDR IB hypercube interconnect
  - 166 racks, 416TB of memory
  - Debuted as #3 on TOP500 in 11/08; now #19 (6/13)
  - 2 racks have 64 Nvidia M2090 GPU-enhanced nodes
- **Endeavour** – 32 TF peak
  - 2 SGI Ultra Violet 2 nodes with SB (one 1024-core with 4 TB shared memory; other 512-core and 2 TB)
  - Single System Image (SSI) via NUMALink-6
- **Maia** – 301 TF peak
  - 4 SGI Pyramid racks, 4 TB of memory
  - 128 nodes, each with two SB processors and two 60-core Phi accelerators
- **Mira** – 150 TF peak
  - 3 Cray racks, 2 TB of memory
  - 64 nodes, each with two SB processors and two 60-core Phi accelerators
- **hyperwall-2** – 182 TF peak
  - 136 node graphics cluster, 1,024 AMD Opteron processor cores, 136 Nvidia GeForce 480 GTX GPUs
  - Large-scale rendering, concurrent visualization

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
  - MDs select projects, determine allocations
  - Over 500 science & engineering projects
  - More than 1,200 user accounts

Advanced 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
  - Debuted as #1 resolution system in world
  - In-depth data analysis and software
- Two primary modes
  - Single large high-definition image
  - Sets of related images (e.g., a parameter space of simulation results)
- High-bandwidth to HEC resources
  - *Concurrent Visualization:* Runtime data streaming allows visualization of every simulation timestep – ultimate insight into simulation code and results with no disk i/o
  - *Traditional Post-processing:* Direct read/write access to Pleiades filesystems eliminates need for copying large datasets
- GPU-based computational acceleration R&D for appropriate NASA codes

Detailed HEC Utilization
(Normalized to 30-Day Months)

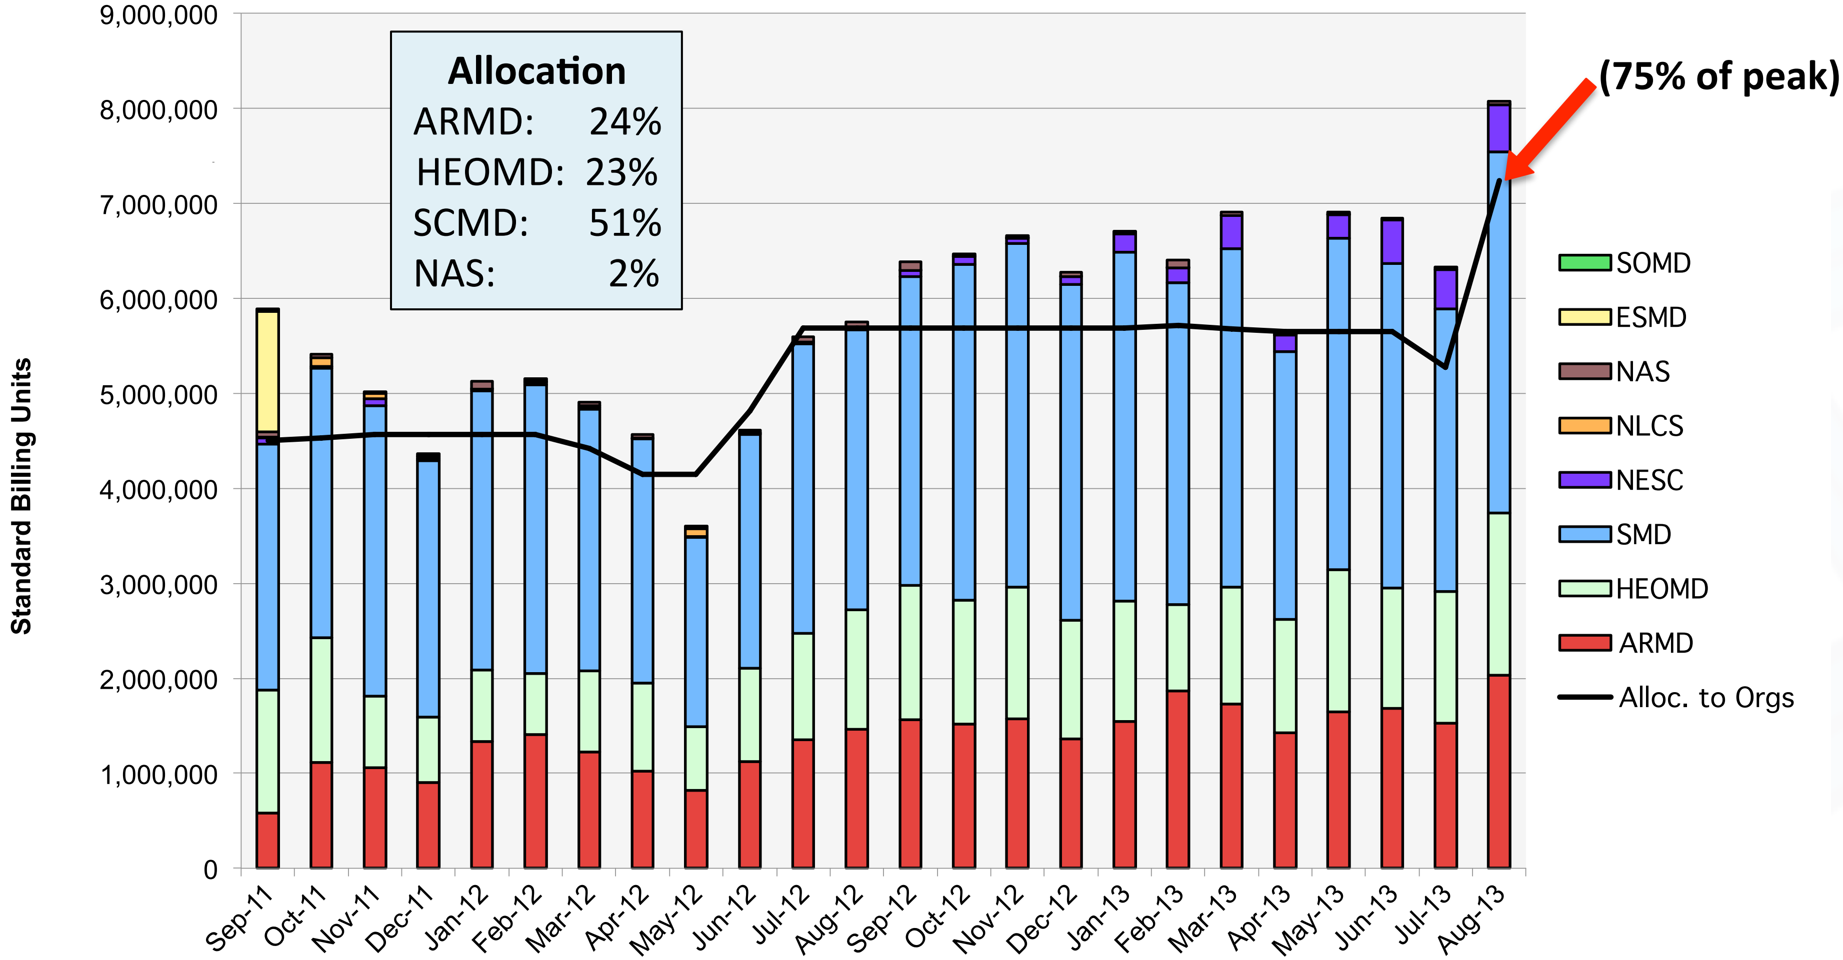

SBU: Used to normalize computational capability across systems
Equals one Westmere node-hour (24 hyperthreaded core-hours)

Strategic Support for NASA Programs

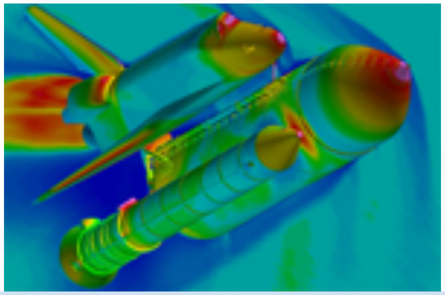

External tank redesign

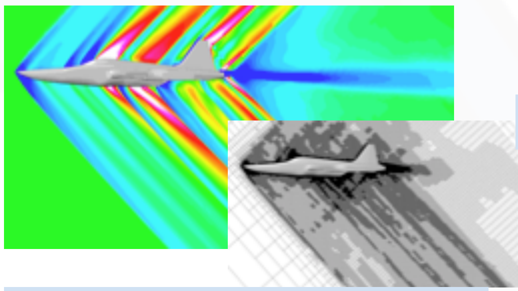

Sonic boom optimization

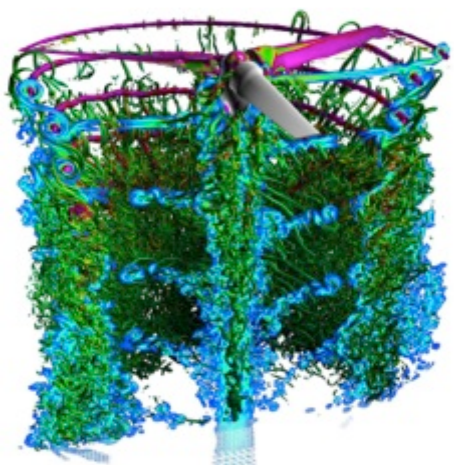

Rotary wing aerodynamics

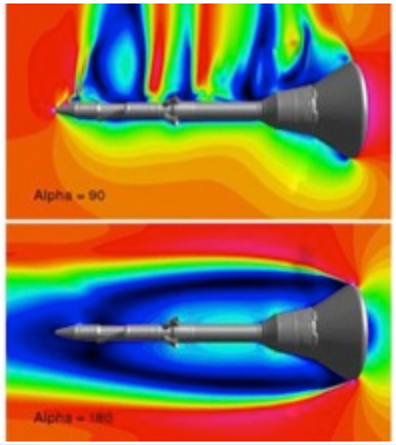

Launch abort system

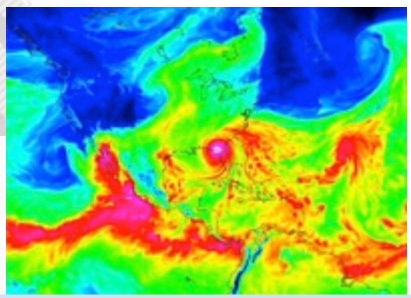

Hurricane prediction

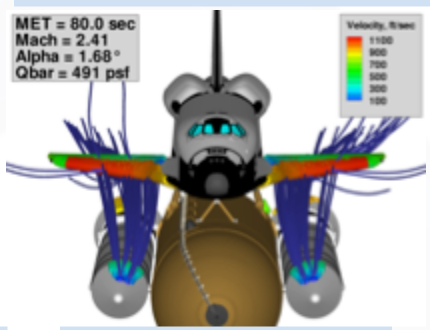

Debris transport

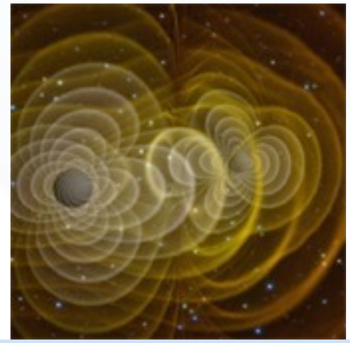

Merging black holes

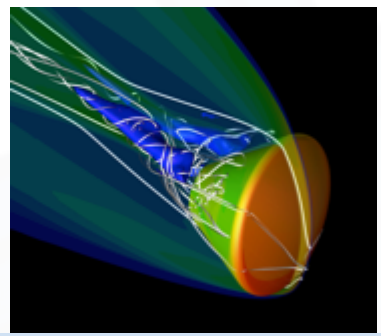

Orion/MPCV reentry

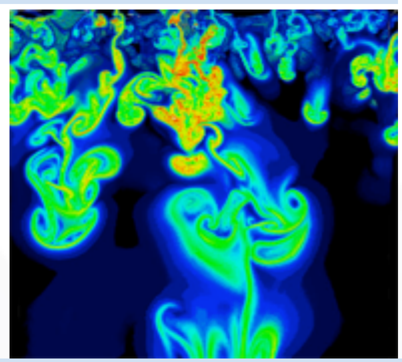

Solar magnetodynamics

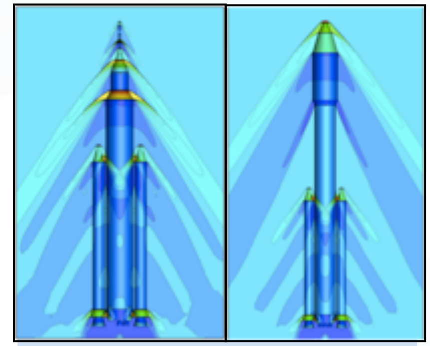

SLS vehicle designs

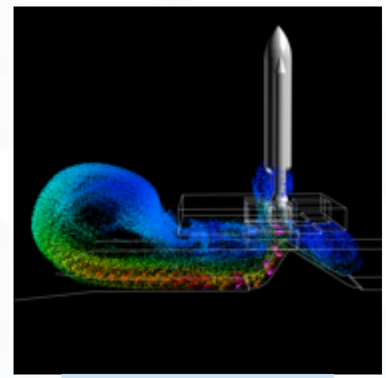

Flame trench

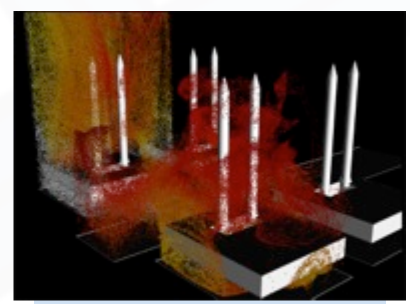

SRB burn in VAB

Accelerator Technologies

Significant performance potential for S&E applications

– Execute many threads simultaneously at relatively lower power

Two Primary Viable Options

– Nvidia GPGPU (general-purpose graphics processing unit)

– Intel MIC (many integrated core)

Issues

– *Performance:* To what extent can S&E apps exploit these technologies?

- Is this the road to exascale? Or other options emerging?

– *Programmability:* Are the programming models suitable to express the parallelism required to exploit these technologies?

- What is the nitial porting and optimization effort (including ramp-up time)?
- Is it possible to achieve code and performance portability across architectures (to future evolutions and to other existing platforms)?

– *Verification:* How much effort is necessary to ensure that the numerical results are identical?

Current Projects

Nvidia GPGPU

- 136 (GeForce 480GTX) in hyperwall and connected to Pleiades via InfiniBand
- 64 (Tesla M2090) integrated directly into Pleiades

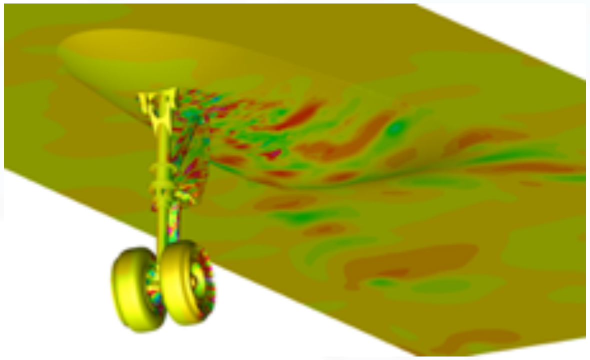

Airframe and Landing Gear Acoustic Modeling and Noise Reduction (EXA) – LaRC

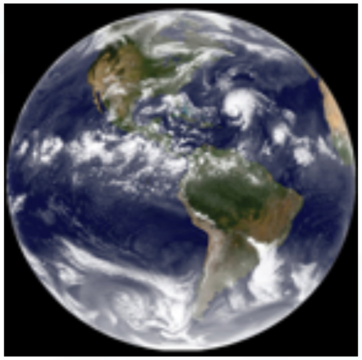

Atmospheric Assimilation System (GEOS-5) – GSFC

Intel MIC

- Performance and programmability study of cluster of heterogeneous nodes
  - Two systems: 64/128 nodes, each with 2 Sandy Bridges & 2 60-core Phis
  - Usage modes: Offload, Native, Symmetric
  - Hybrid programs (MPI + OpenMP)
- Experiences so far
  - System software and libraries (both host and Phi) are works in progress
  - Memory bandwidth may be limiting factor for most NASA applications
- Paper on single-node performance at SC'13
  - Saini et al., "An Early Performance Evaluation of MIC Architecture Based SGI Rackable Computing System"

Cloud Computing

Background

- Cloud computing becoming a commercial success with many companies providing public and private cloud environments
- These environments being extensively used for commercial and business tasks (e.g. email, data storage, human resources, etc.)
- Increasing use of clouds for certain S&E tasks suitable for such environments (e.g. bioinformatics, oil exploration, data analytics, etc.)
- NASA has extensive first-hand experience in cloud computing from developing and operating Nebula as well as co-founding and contributing significantly to OpenStack

Current Focus: Analyze scenarios of interest to NASA

- Cloud bursting to provide surge capacity for NASA HPC users
- Make cloud resources (including support) available to S&E users who currently do not have access to (or priority on) Pleiades
- Explore approaches to making NASA data more easily available

Current Cloud Work: Cloud Bursting

Focus

– Evaluate feasibility of integrating cloud services in the HPC environment to meet dynamic demand for HECC resources

Initial prototype completed

– Set up execution environment on cloud
– Stage input data sets and executables
– Start and monitor the execution
– Return output files to user
– Dismantle the execution environment

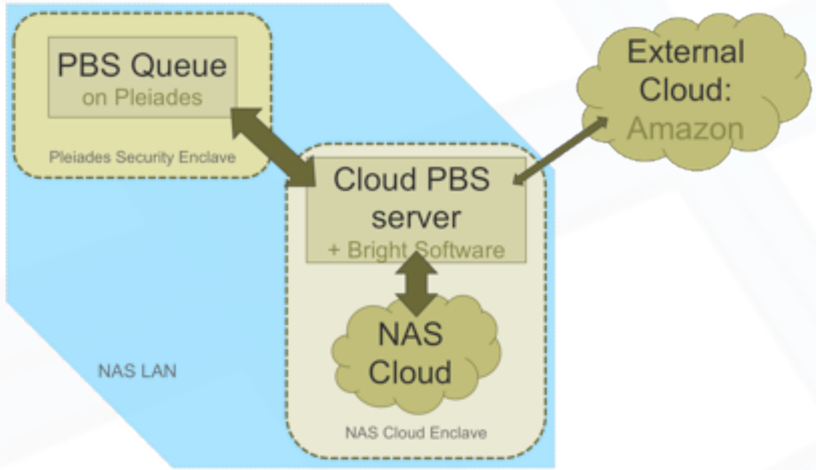

On-going work

– Build logic to decide if a job should be moved to the cloud based on certain parameters (e.g., appropriateness, cost, wait time, etc.)
– Develop fault tolerance if execution fails on the cloud for any reason
– Optimize data transfers by eliminating redundancy
– Implement multi-threaded support to simultaneously burst multiple jobs

NASA Earth Exchange (NEX)

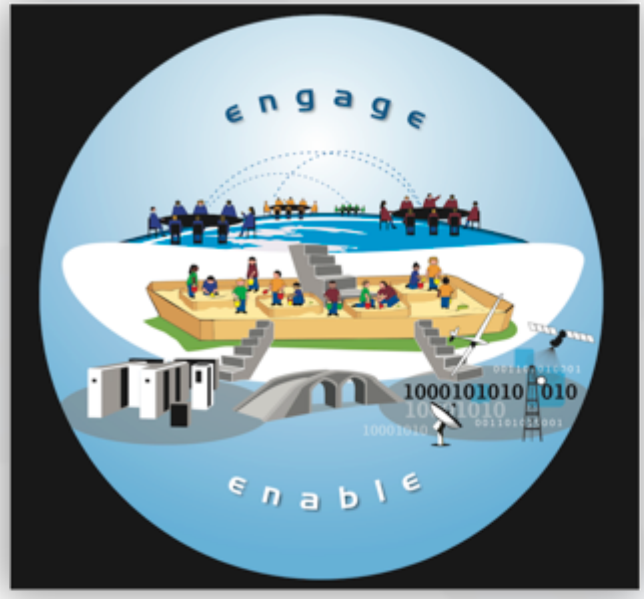

Collaborative Computing for Earth Sciences

- Provide "Science as a Service" by bringing together scientists and researchers in a knowledge-based social network to address global environmental challenges
- Encourage innovation, accelerate research, enhance transparency, and expand scope by centralizing all the necessary data, tools, and computing power

Engage
Network, share, collaborate and formulate new ideas

Enable
Access to data, computing, and knowledge

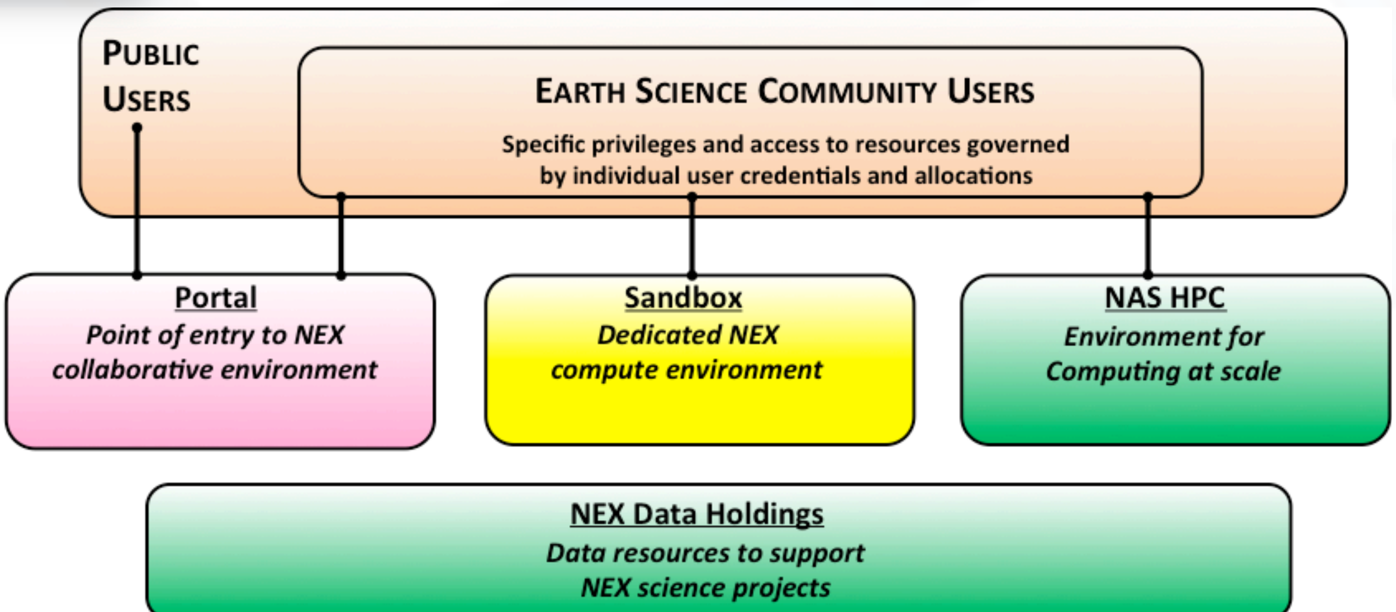

Big Data

- NASA has enormous collections of observational and model data
- Gaining knowledge and insight from these datasets involves Analysis and Analytics
  - Analysis: Process of converting data into information (apply known algorithms to well-understood data)
  - Analytics: Tool for discovering patterns from data and building predictive models (apply exploratory algorithms to data to uncover hidden information)
- Analysis and Analytics for large datasets face major challenges
  - Data Discovery: Finding the right data to answer questions in S&E hampered by lack of effective and efficient tools to search metadata, including domain-specific taxonomies / ontologies
  - Data Access and Movement: Data and processing need to be brought together to perform analysis and analytics; users need better tools to move and manage data
  - Data Manipulation: A new generation of algorithms and tools needs to be made available to users, including tools for data exploration
- Focus: Build NASA platforms & tools to support Analysis and Analytics
  - Engage with users of NASA Big Data and the Big Data research community
  - Conduct experiments and build prototypes to validate tools and architectures
  - Develop, implement infrastructure to support analysis & analytics on NASA Big Data

Big Data Analysis and Analytics

Engagement – Involve the community in planning

Interview current users, prospective users, technologists → What are they doing now? Where is their research going? What are the barriers? → Develop use cases to guide experimentation and implementation

Experimentation – Try things out to see what works

Provide Diverse Platforms and Tools

- Large Shared Memory
- Accelerators (MIC & GPU)
- In-house Cloud
- uRiKA Graph Appliance
- Sandbox Computing Environments
- Multiple Database Environments

Implementation – Build out the infrastructure

Balanced Infrastructure to support Simulation, Analysis, and Analytics

- Storage space for large datasets (multi PB)
- Bigger and faster filesystems
- Local hosting of data (Landsat, Modis, etc.)
- Tools for easing data transfer (SUP, SHIFT, etc.)
- Shared access to datasets from simulation on supercomputing and analysis/analytics resources
- Large memory nodes
- Software tools for analytics

All integrated via high speed networking (InfiniBand) within and across multiple facilities

Solving NASA's Optimization Problems

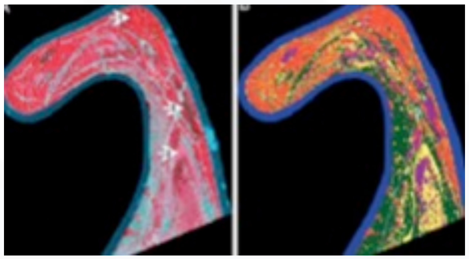

Data Fusion and Image Registration

Anomaly Detection and Decision Making

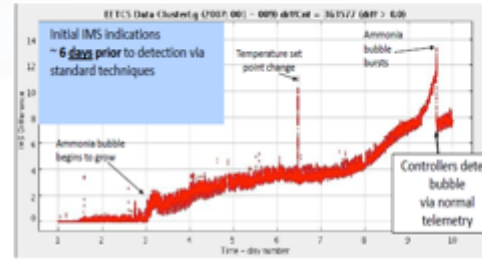

Air Traffic Management

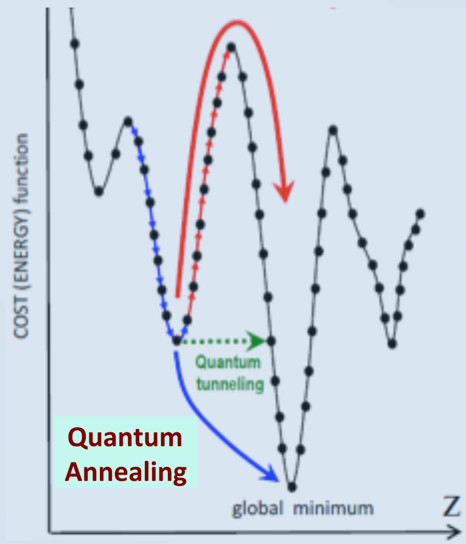

Distributed Coordination

Mission Planning and Scheduling

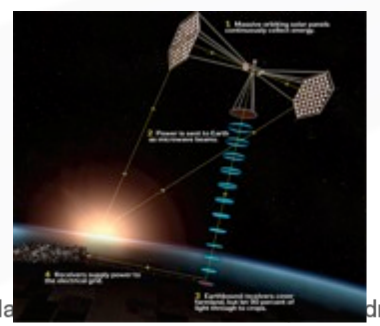

V&V and optimal sensor placement

Quantum Annealing using D-Wave Two

QA is the most natural quantum computing paradigm for solving hard optimization problems – it harnesses quantum tunneling as a computational resource

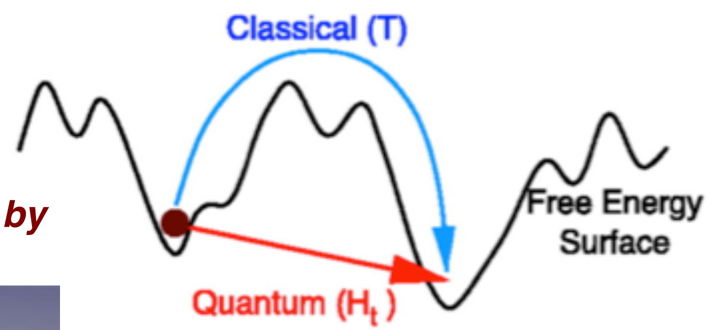

Quantum bit (qubit): Niobium superconducting loop encodes 2 states as tiny magnetic fields

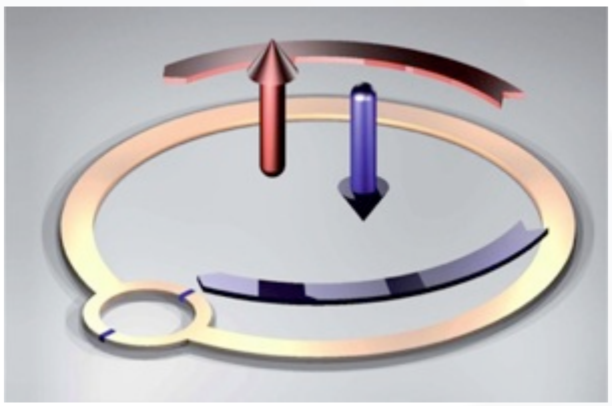

8 qubit loops connected by 16 coupling devices

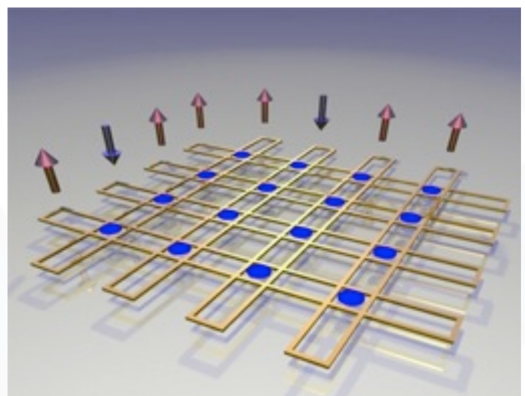

512 qubits & 1472 couplers

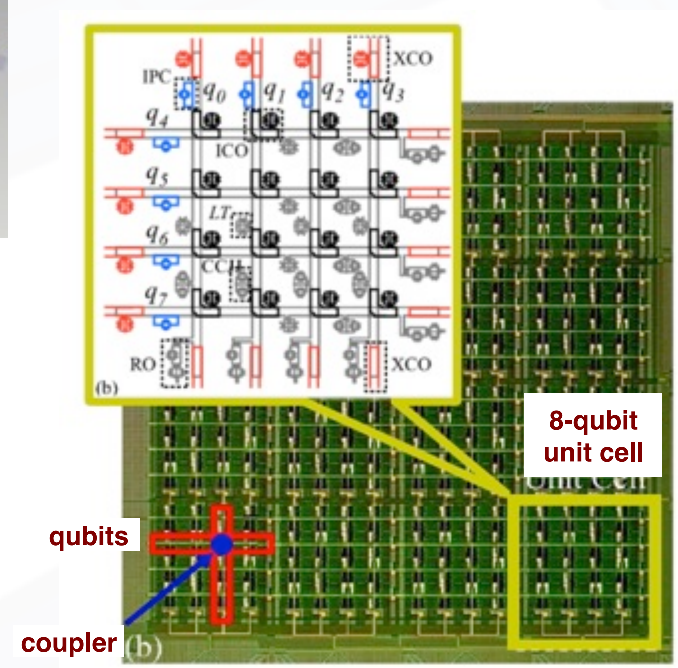

D-Wave processor solves a binary optimization problem

Given $\{ h_j , J_{ij} \}$, find $\{ s_k = \pm 1 \}$
that minimizes

$$\xi(s_1, \ldots, s_N) = \sum_{j=1}^{N} h_j s_j + \sum_{i,j \in E}^{N} J_{ij} s_i s_j$$

D-Wave Two

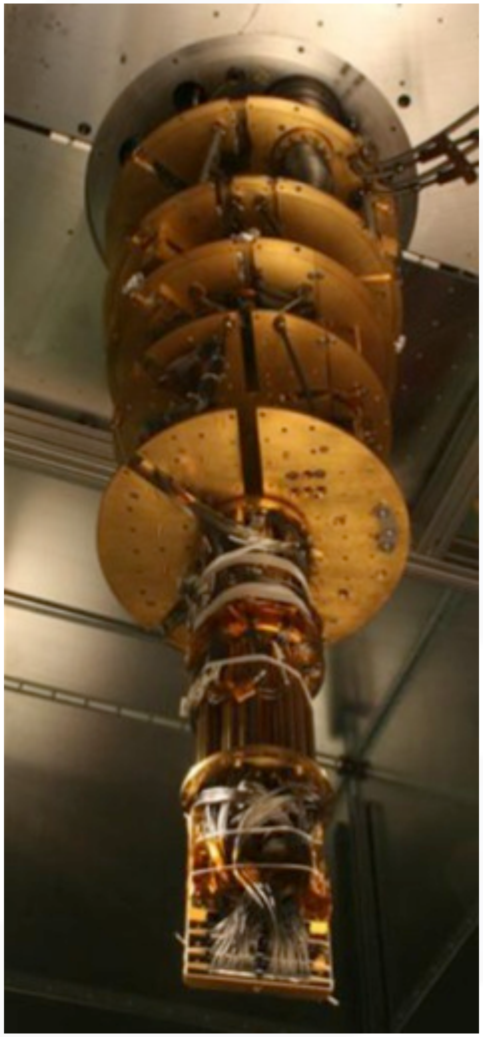

Cryogenics (10 kg of metal at 15 milliKelvin)

512-qubit Vesuvius processor

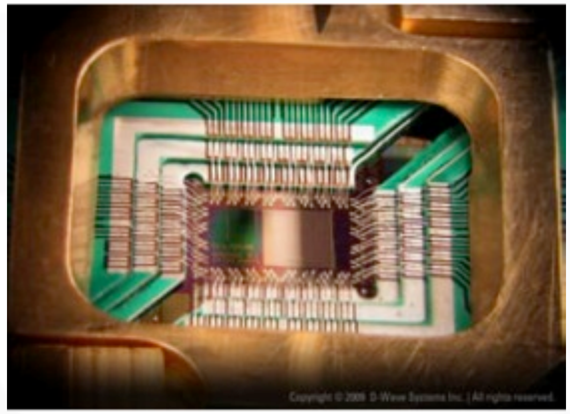

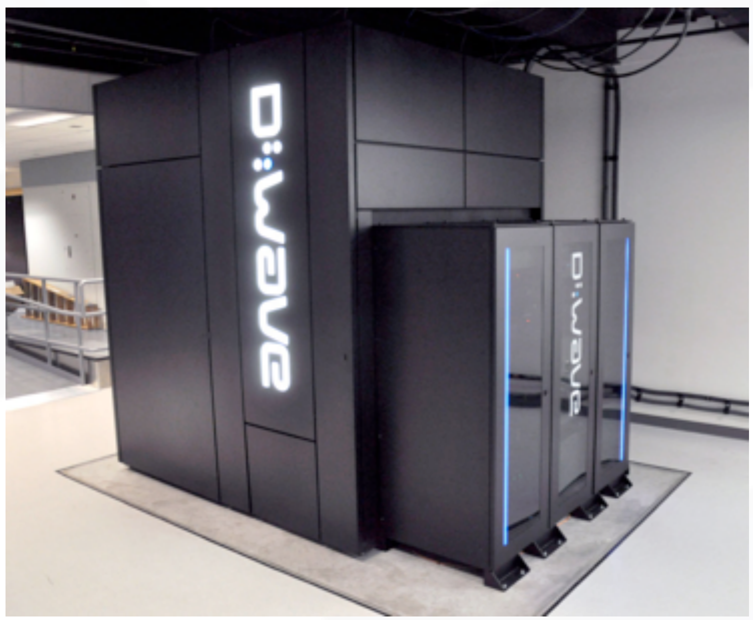

Motherboard packaging to refrigeration system

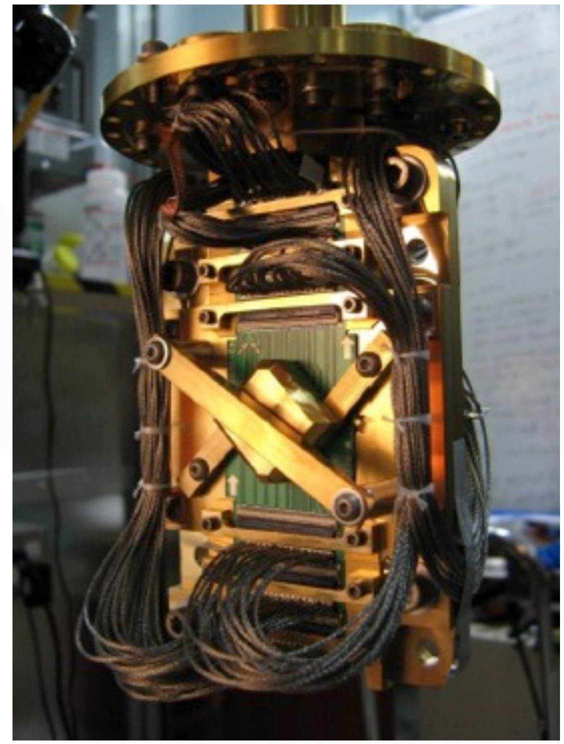

Magnetic Shielding (less than 1 nanoTesla)

Thank you!

Questions?