Digital Technologies at NASA for Science & Engineering

Dr. Rupak Biswas

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

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

Digital Thailand Big Bang 2018

Bangkok, Thailand, 22 September 2018

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 collaborative partnerships, expanding the boundaries of aerospace enterprise

NASA Overview: Centers & Facilities

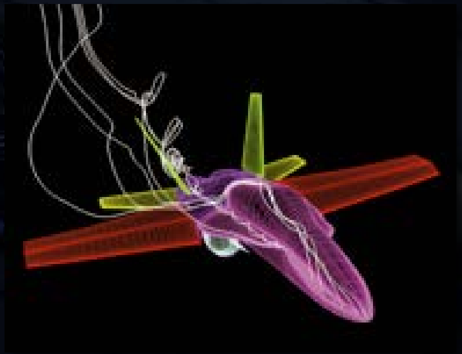

Current NASA Focus Areas

NASA Ames Research Center

- **<u>Occupants</u>: ~1,130 civil servants; ~2,100 contractors; ~1,650 tenants; 855 summer students (in 2016)**
- **<u>FY2016 Budget</u>: ~$915M (including reimbursable and Enhanced Use Lease (EUL) revenue)**
- **<u>Real Estate</u>: ~1,900 acres (400 acres security perimeter); 5M building ft^{2}; Airfield: ~9,000 and 8,000 ft runways**

Ames Core Competencies Today

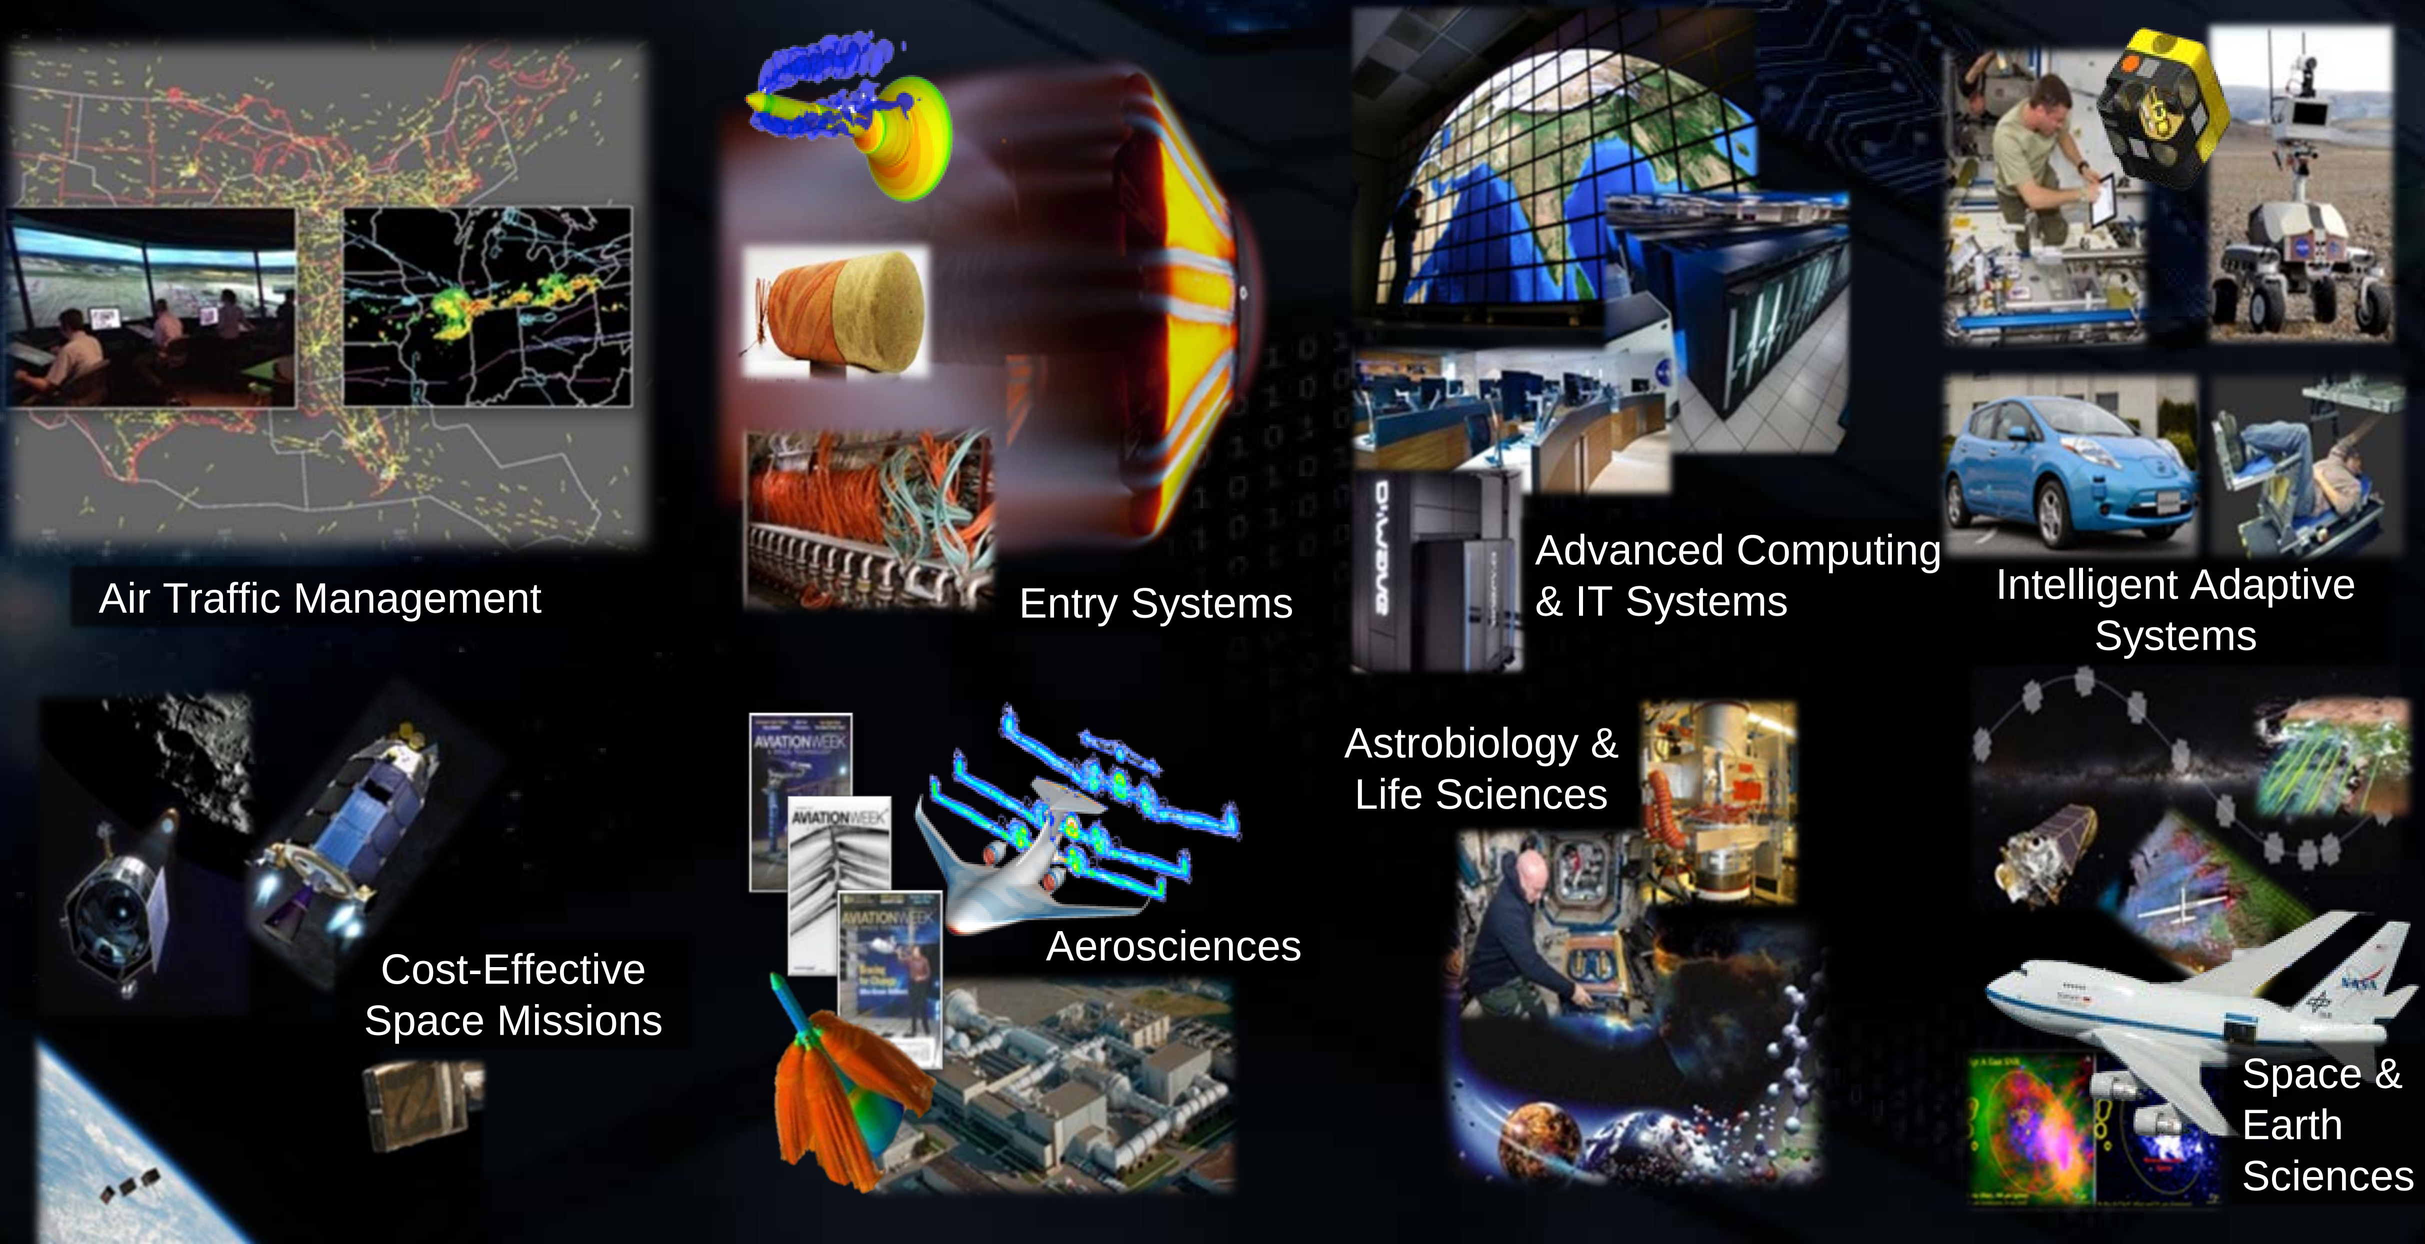

Need for Advanced Computing

Enables modeling, simulation, analysis, and decision-making

- Digital experiments and physical experiments are tradable
- Physical systems and live tests generally expensive & dangerous (e.g., extreme environments), require long wait times, and offer limited sensor data
- NASA collects and curates vast amounts of observational science data that require extensive analysis and innovative analytics to advance our understanding

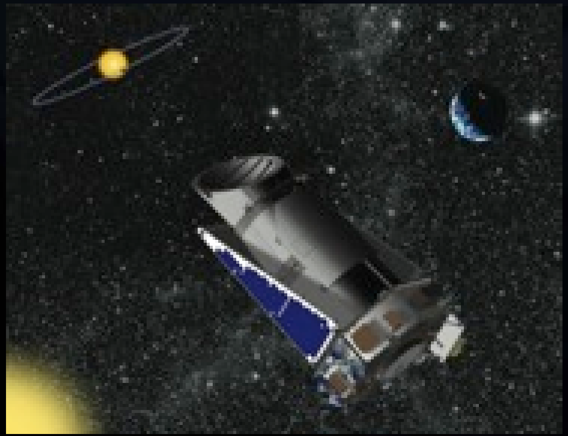

- Decades of exponentially advancing computing technology has enabled dramatic improvements in cost, speed, and accuracy – in addition to providing a predictive capability
- Many problems pose extremely difficult combinatorial optimization challenges that can only be solved accurately using advanced technologies such as quantum computing
- NASA's goals in aeronautics, Earth & space sciences, and human & robotic exploration require orders-of-magnitude increase in computing capability to enhance accuracy, reduce cost, mitigate risk, accelerate R&D, and heighten societal impact

Advanced Computing Environment

NASA's Diverse HPC Requirements

- Engineering requires HPC resources that can process large ensembles of moderate-scale computations to efficiently explore design space ***(high throughput / capacity)***
- Research requires HPC resources that can handle high-fidelity long-running large-scale computations to advance theoretical understanding ***(leadership / capability)***
- Time-sensitive mission-critical applications require HPC resources on demand ***(high availability / maintain readiness)***

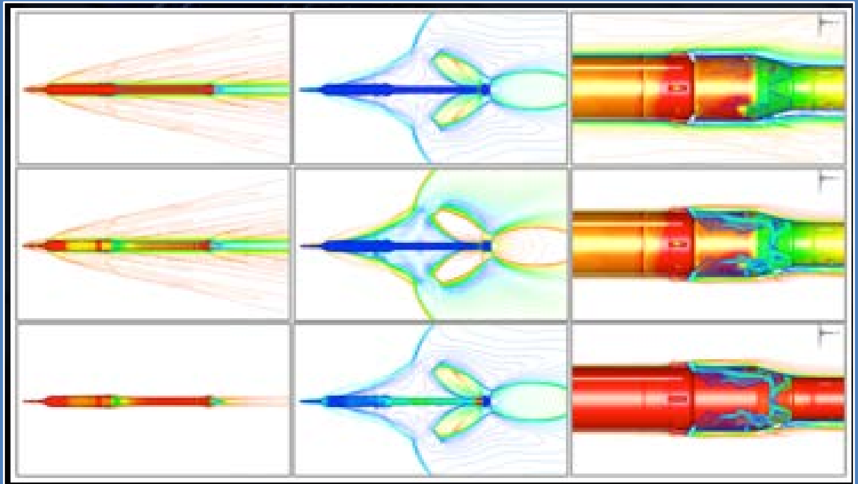

Aerodynamic database generation

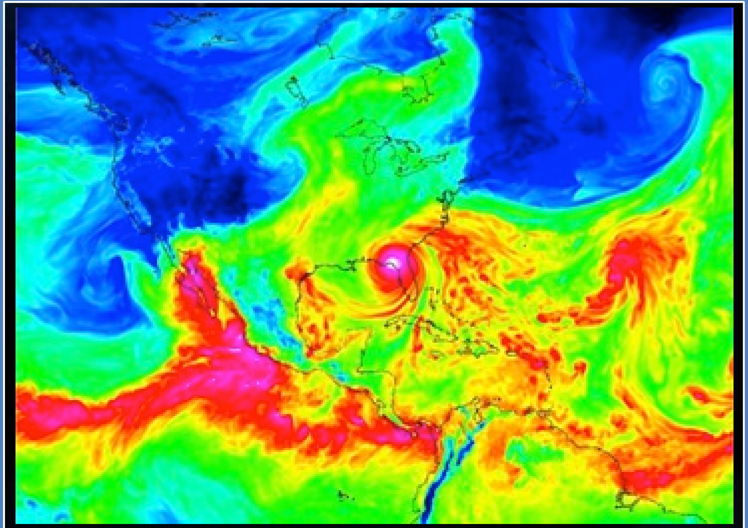

Real-time hurricane prediction

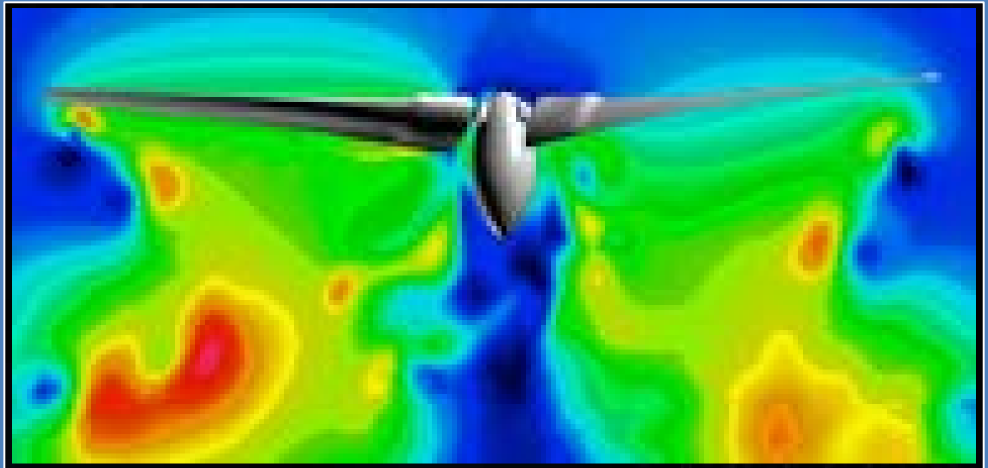

Rotary wing unsteady aerodynamics

Balanced HPC Environment

Computing Systems

- ***Pleiades:*** 246K-core SGI Altix ICE (now HPE) with 4 generations of Intel Xeon (64 nodes GPU-enhanced: Nvidia M2090, K40; 32 nodes have Phi 5110P); 938 TB RAM; 7.25 PF peak (#15 on TOP500, #10 on HPCG)
- ***Electra:*** 32K-core Altix ICE with Intel Broadwell; modular container; 147 TB RAM; 1.24 PF peak
- ***Merope:*** 22K-core Altix ICE with Intel Westmere; 86 TB RAM; 252 TF peak
- ***Endeavour:*** Two SGI UV2000 nodes with 2 and 4 TB shared memory SSI via NUMALink-6; 32 TF peak
- ***hyperwall:*** 2560-core Intel Ivy Bridge, 128-node Nvidia GeForce GTX78 cluster for large-scale rendering & concurrent visualization (240M pixels)

Data Storage

- 49 PB of RAID over 7 Lustre filesystems
- 490 PB of tape archive

Networks

- InfiniBand interconnect for Pleiades in partial hypercube topology; connects all other HPC components as well
- 10 Gb/s external peering

Modular Supercomputing Facility (MSF)

Current HEC Facility

- Limited to 6 MW electrical power of which 25% used for cooling
- Open-air cooling tower with four 450 T chillers

Prototype MSF (FY17)

- Modular container currently holds Electra (16 Broadwell-based racks)
- External air fan cooling; switch to adiabatic evaporative cooling when needed
- PUE of 1.03 resulting in 93% power savings and 99.4% water use reduction over our traditional computer floor
- Pad has 2.5 MW of electrical power and can accommodate 2 modules
- In production use since Jan '17
- Second module being added with 4 E-Cells, bringing Electra to 4.78 PF peak

Full MSF (FY18 – FY22)

- Larger second pad with 30 MW electrical power and associated switchgear
- Ability to hold up to 16 modular units (and 1 M cores)
- Flexibility to rapidly modify and react to changes in NASA requirements, computing technology, and facility innovations

Prototype MSF hosting Electra

Artist's rendering of future MSF

Integrated Spiral Support Services

NASA Mission Challenges

<u>Outcome</u>: Dramatically enhanced understanding and insight, accelerated science and engineering, and increased mission safety and performance

Scientists and engineers plan computational analyses, selecting the best-suited codes to address NASA's complex mission challenges

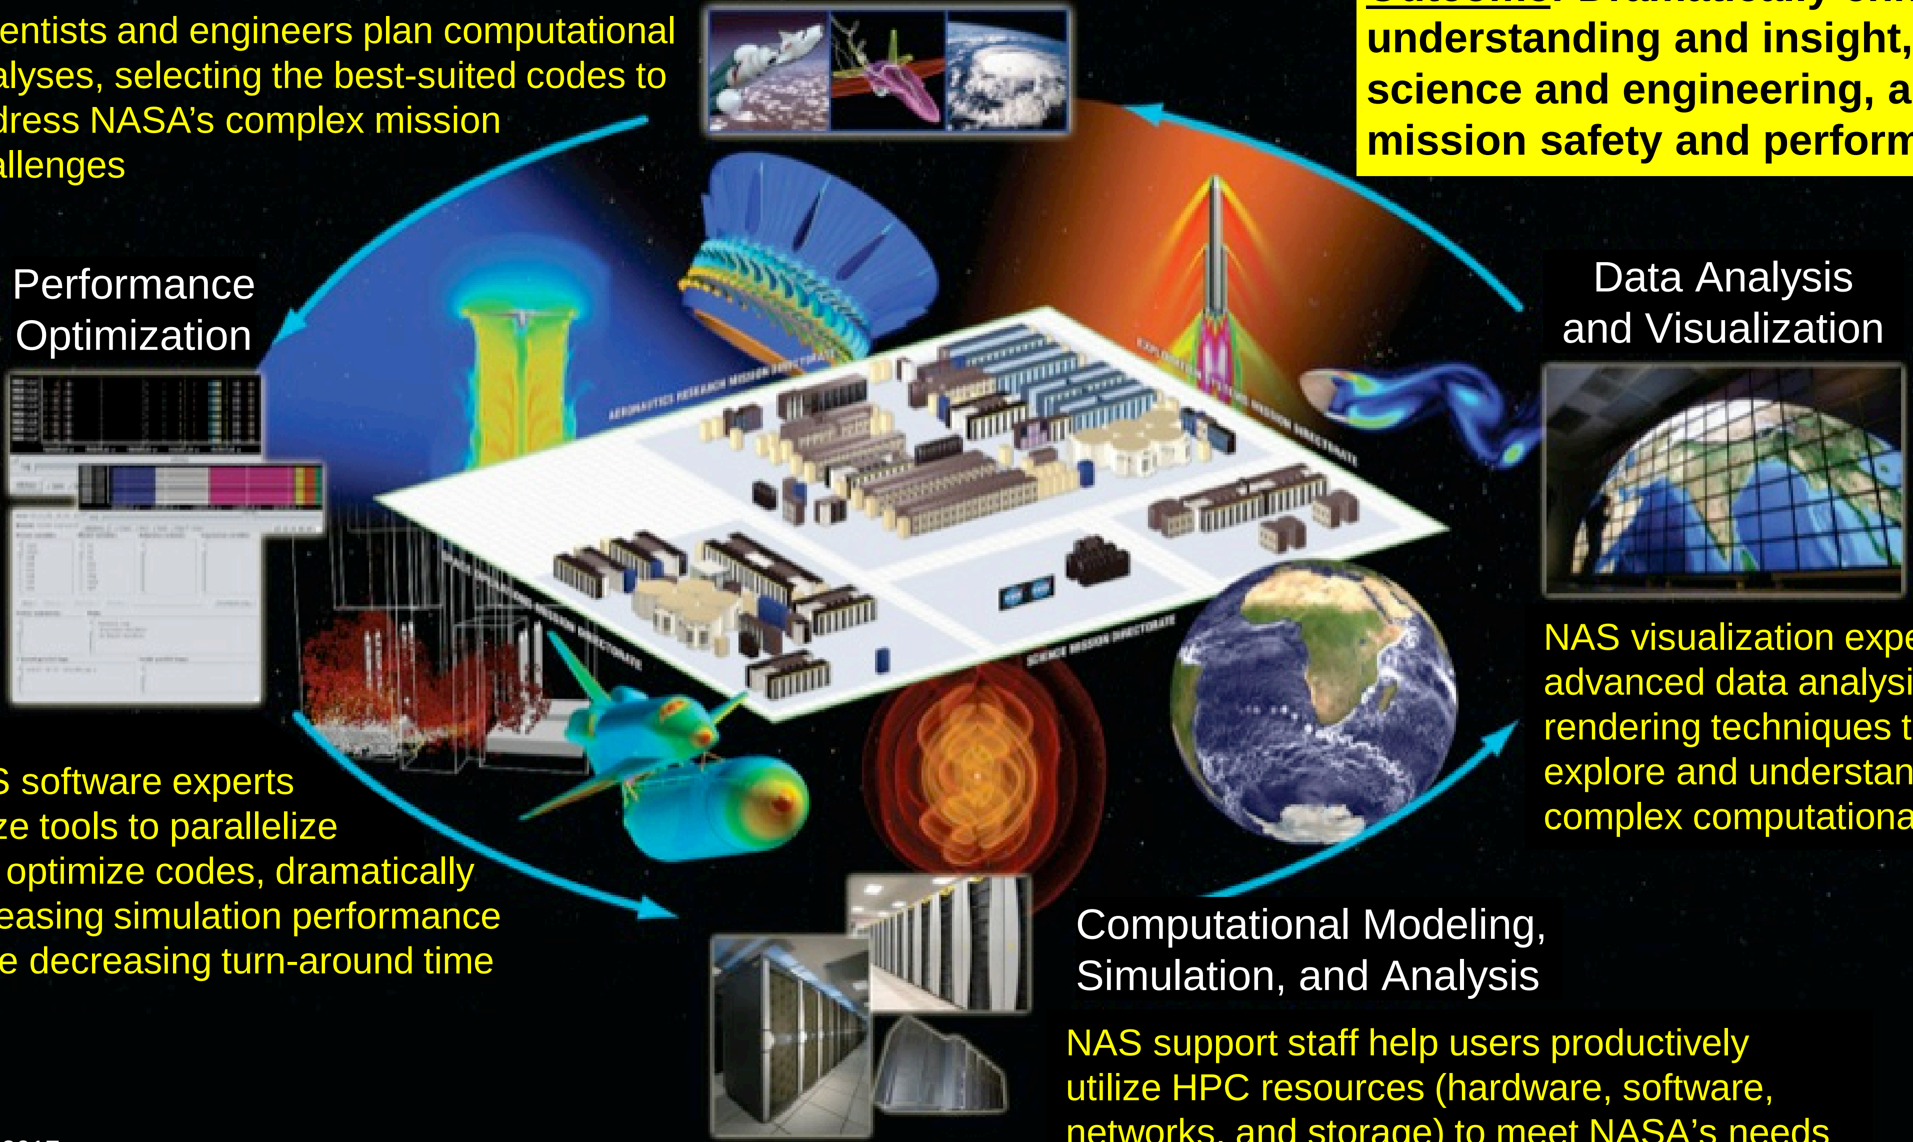

Performance Optimization

Data Analysis and Visualization

NAS visualization experts apply advanced data analysis and rendering techniques to help users explore and understand large, complex computational results

NAS software experts utilize tools to parallelize and optimize codes, dramatically increasing simulation performance while decreasing turn-around time

Computational Modeling, Simulation, and Analysis

NAS support staff help users productively utilize HPC resources (hardware, software, networks, and storage) to meet NASA's needs

NAS Supercomputing Facility at ARC

- Current computer floor (Bldg. N258) limited by power and cooling
- Electrical System
  - Facility limited to 6 MW
  - About 25% used for cooling
  - Hence, approx. 4.5 MW for computing
- Cooling System
  - Open-air cooling tower with 4 50 HP pumps
  - 4 450-Ton Chillers
  - 7 pumps for outbound chilled water
  - 4 pumps for inbound warm water

Phase 1: Prototype MSF

- Engineered and deployed a single module adjacent to Bldg N258
- Houses Electra supercomputer system (debuted at #96 on Nov ‘16 TOP500)
  - 16 ICE-X racks of Intel Broadwell (1152 nodes; 1.24 PF peak; 147 TB memory)

Phase 2: Expand Current MSF

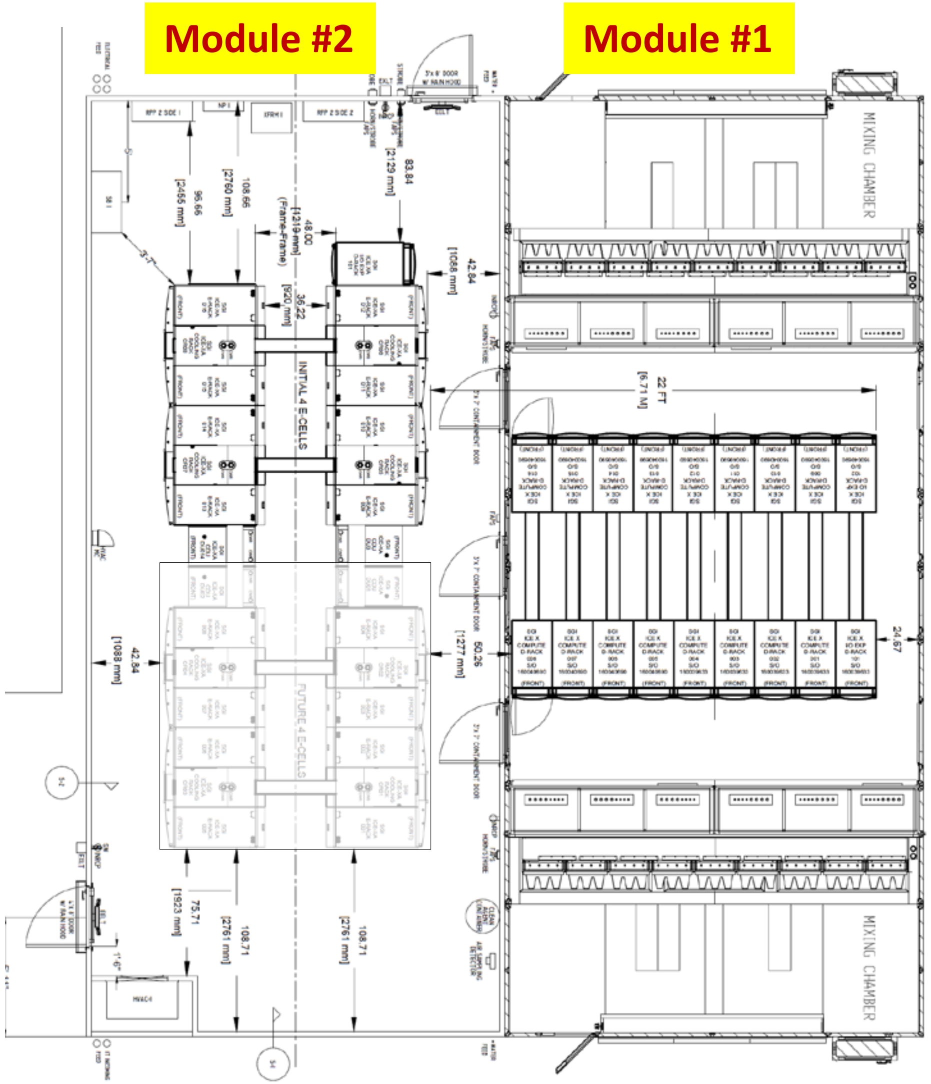

- Facility advances allow more power to 2nd module (from 500 KW to 1.2 MW)
- Doubled compute rack density
  - Electra added 4 E-Cells of ICE-XA with Intel Skylake (total 4.78 PF peak; 369 TB memory)

NASA Earth Exchange (NEX)

A virtual collaborative environment that brings scientists and researchers together in a knowledge-based social network along with observational data, necessary tools, and computing power to provide transparency and accelerate innovation: **Science-as-a-Service**

VIRTUAL COLLABORATION

Over 650 members

CENTRALIZED DATA REPOSITORY

Over 3.5 PB of observational data

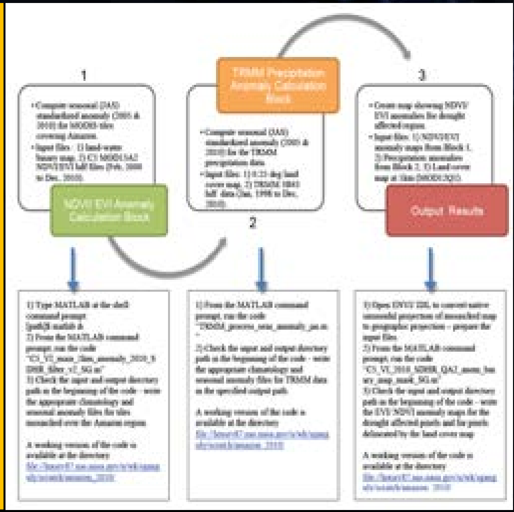

SCALABLE COMPUTING

Heterogeneous and remote, secure access

KNOWLEDGE

Workflows, virtual machine images, model codes, and reusable software

Science via NEX

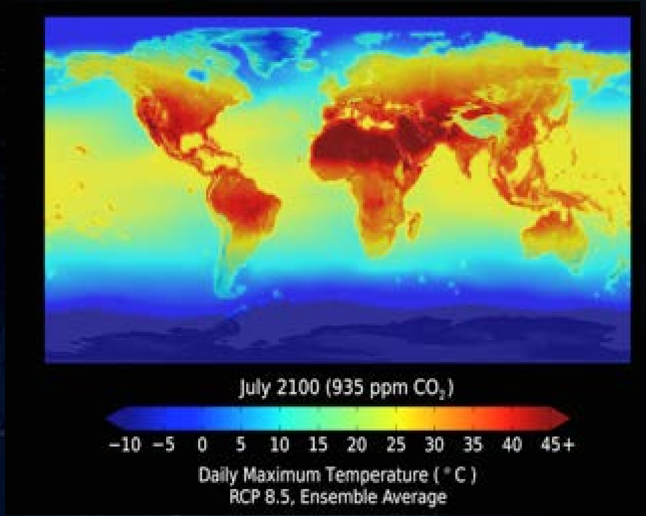

High-resolution projections for climate impact studies

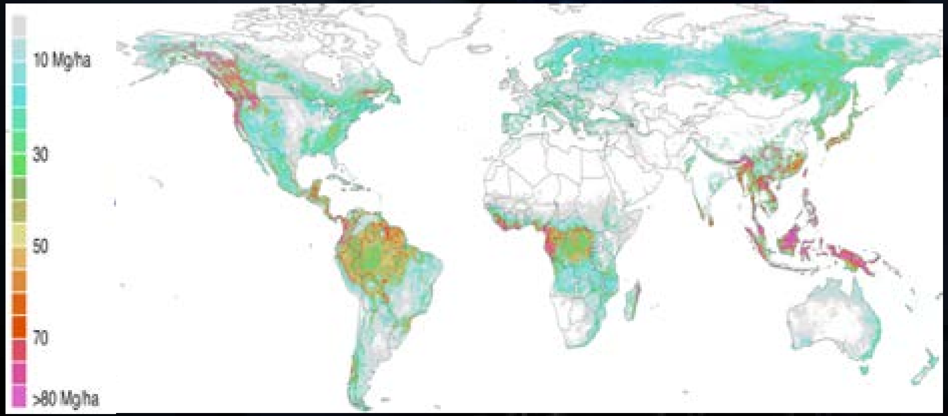

Global vegetation biomass at 100m resolution by blending data from 4 different satellites

High-resolution monthly global data for monitoring crops, forests, and water resources

Sample publication using NEX environment:
Nature 532.7599 (2016): 357

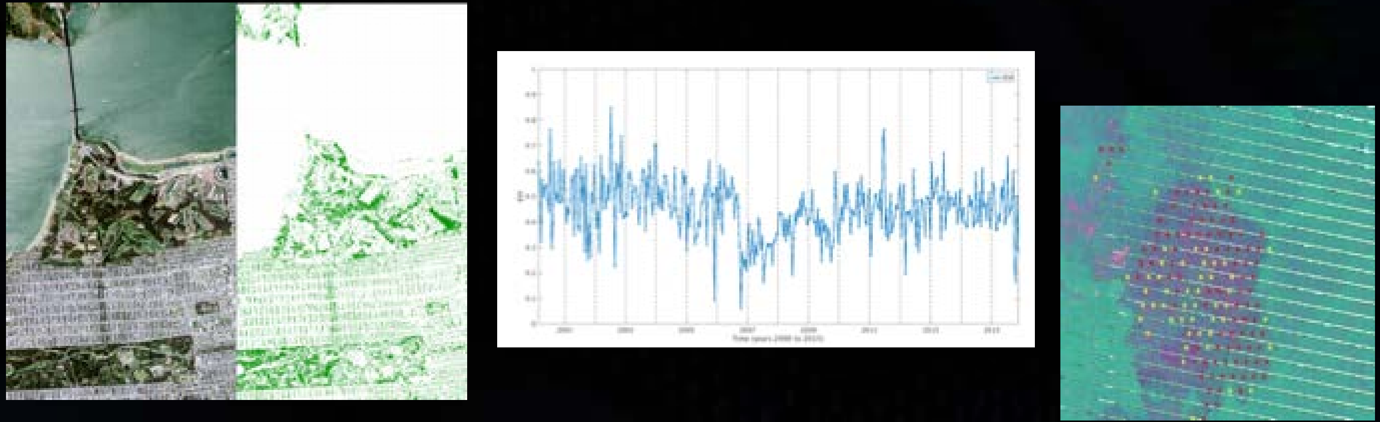

Machine learning and data mining – moving toward data-driven approaches

Quantum Computing 101

- Quantum mechanics deals with physical phenomena at very small scales (~100nm) and at very low temperatures (few K) where actions are quantized
- The outcome of a quantum experiment is probabilistically associated both with what was done before the measurement and how the measurement was conducted
- Qubits (quantum bits) can exist in a superposition of states, allowing *n* qubits to represent 2^n states simultaneously
- At the end of a computation, on measurement, the system collapses to a classical state and returns only one bit string as a possible solution

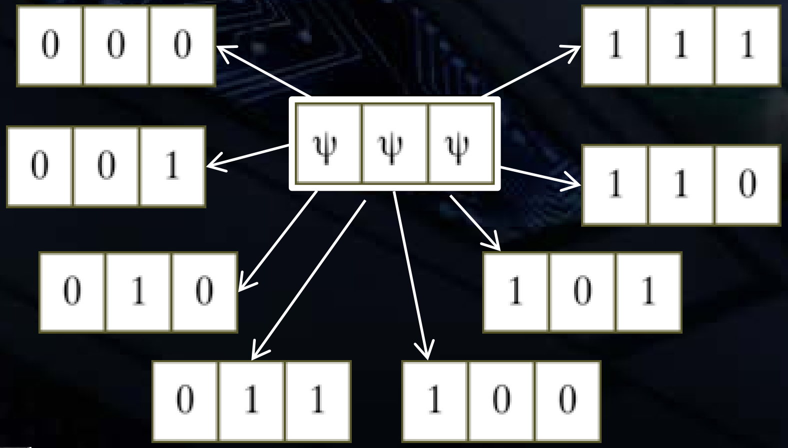

Numerous Implementations

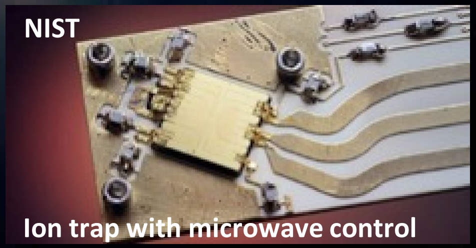

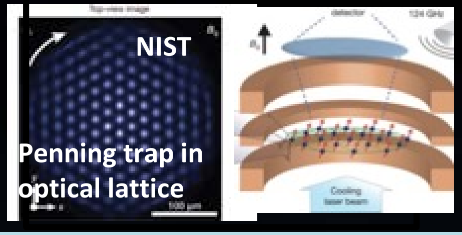

Trapped Ions and Neutral Atoms

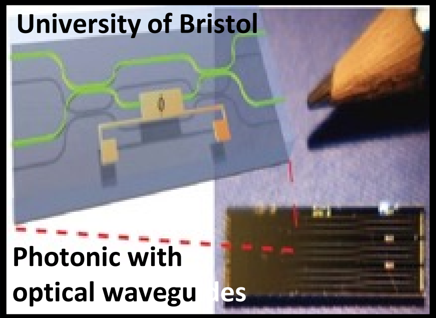

Photonic Quantum Chips

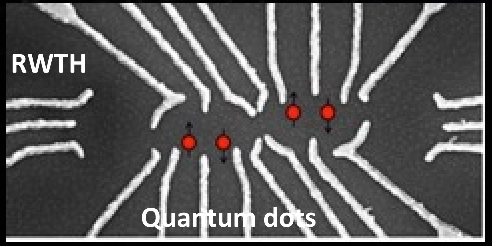

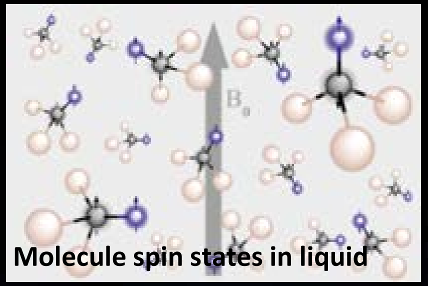

Nanoelectronics, NMR, Diamond Chips, etc.

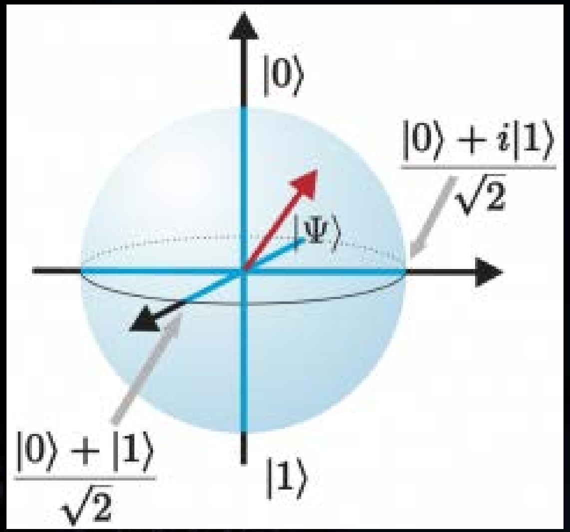

Quantum Computing for NASA Applications

Objective: Find “better” solution
- Faster
- More precise
- Not found by classical algorithm

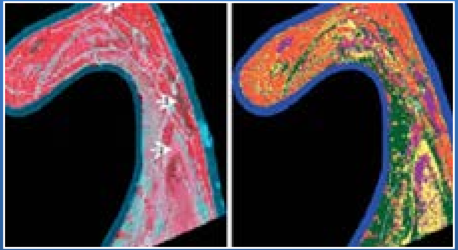

Data Analysis and Data Fusion

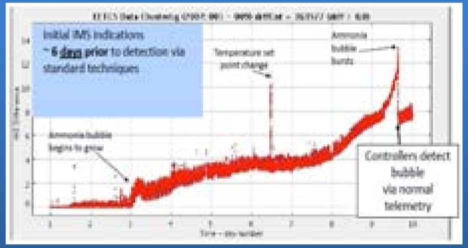

Anomaly Detection and Decision Making

Air Traffic Management

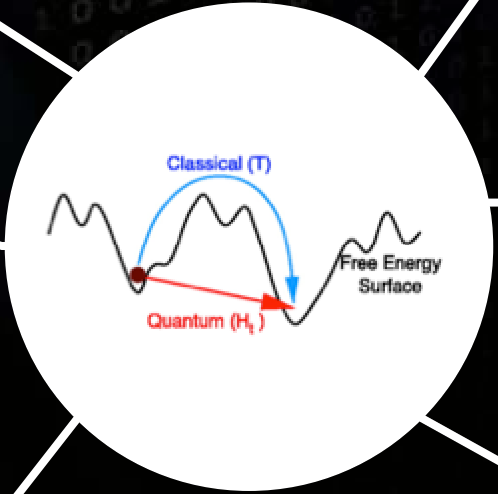

V&V and Optimal Sensor Placement

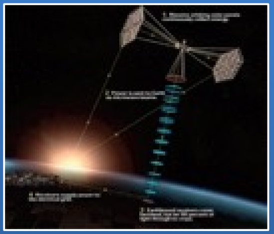

Mission Planning, Scheduling, and Coordination

Topologically-aware Parallel Computing

Common Feature: Intractable (NP-hard / NP-complete) problems!

Quantum Annealing

A physical technique to solve combinatorial optimization problems

$$E(z_1, z_2, \ldots z_n) = \underbrace{\left(1 - \frac{t}{T}\right) H_0(\{z\})}_{A(t)} + \underbrace{\frac{t}{T} H_P(\{z\})}_{B(t)}$$

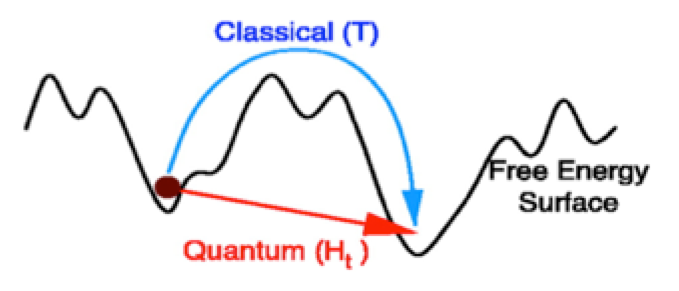

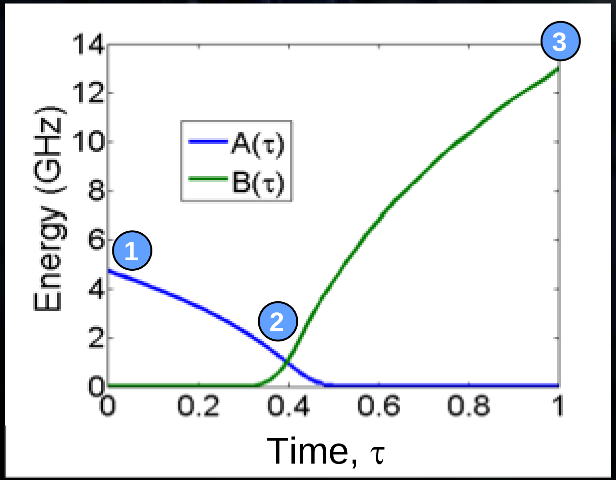

- N-bit string of unknown variables $\{z\}$
- H_0: Hamiltonian with known ground state
- H_P: Hamiltonian whose ground state represents solution to the problem
- Large *A(t)* responsible for quantum fluctuations slowly (adiabatically) lowered to zero while maintaining minimum energy of the system at all times
- In conjunction, cost function of interest *B(t)* gradually turned on
- Transitions between states occur via tunneling through barriers due to quantum fluctuations
- Solution is configuration $\{z\}$ that produces minimum E with non-zero probability
- Method similar to simulated annealing where transitions between states occur via jumping over barriers due to thermal fluctuations

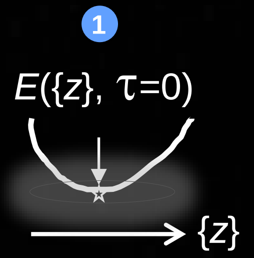

Initialize in an easy-to-prepare full quantum superposition

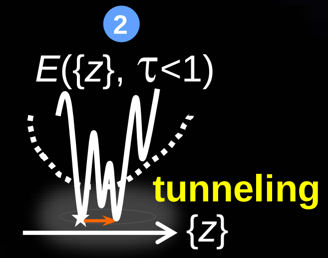

Quantum states explored by quantum tunneling

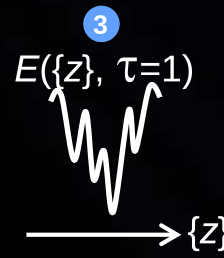

Final state a bit-string encoding solution with probability

D-Wave System Hardware

- Collaboration with Google and USRA via Space Act Agreement led to installation of system at NASA Ames in early 2013
- Started with 512-qubit Vesuvius processor – currently 2031-bit Whistler
- 10 kg of metal in vacuum at ~15 mK
- Magnetic shielding to 1 nanoTesla
- Protected from transient vibrations
- Single annealing typically 20 μs
- Typical run of 10K anneals (incl. reset & readout takes ~4 sec)
- Uses 15 kW of electrical power

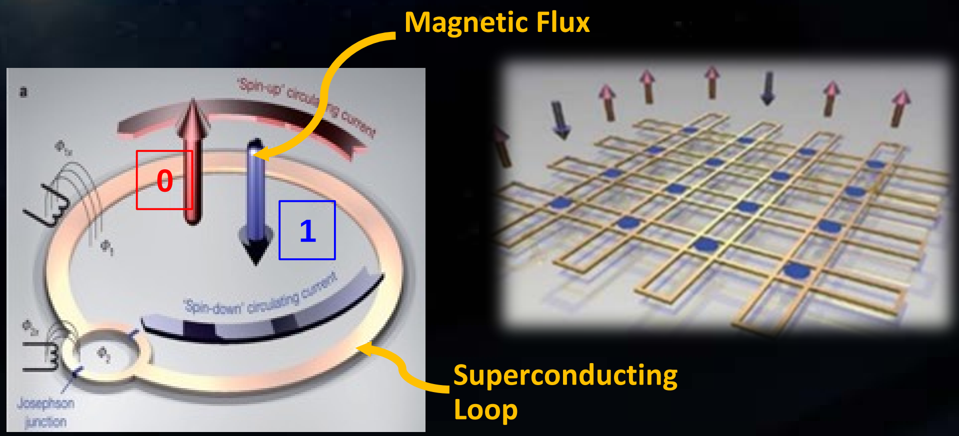

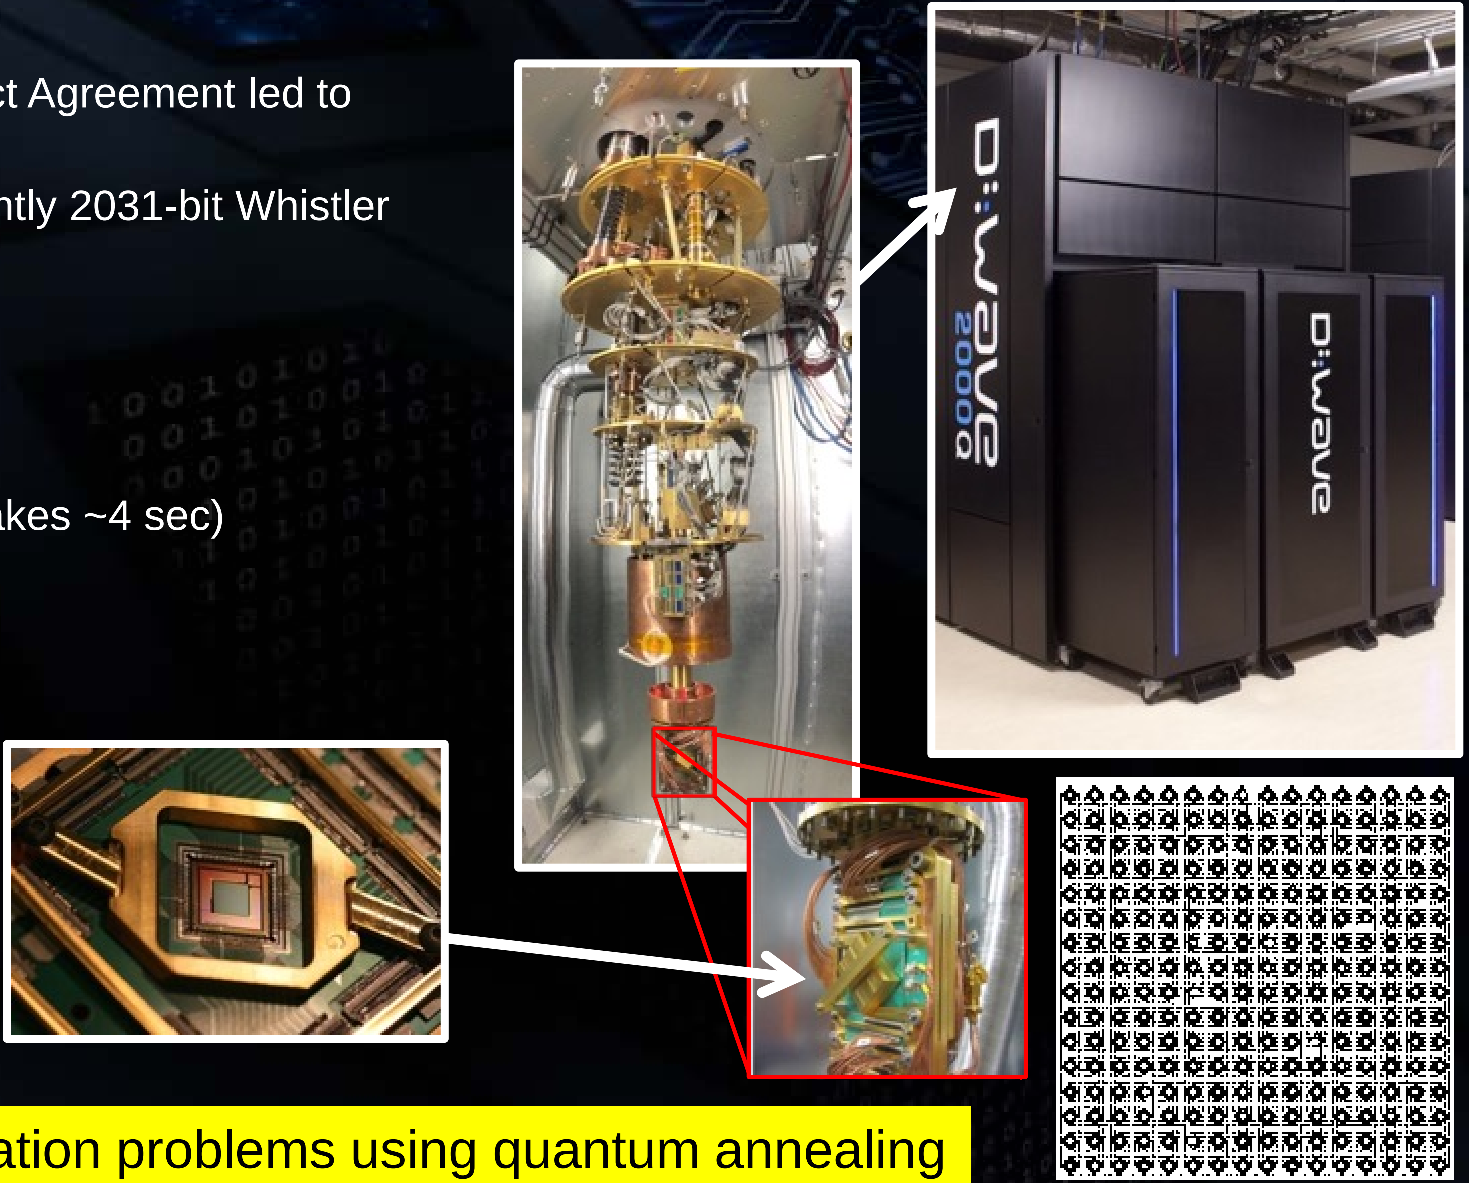

Focus on solving discrete optimization problems using quantum annealing

Programming the D-Wave System

1 Map the target combinatorial optimization problem into QUBO

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

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

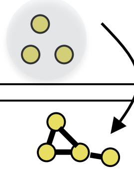

$$\alpha_{ijk}\, y_{ij} z_k + \beta_{ijk}(3y_{ij} - 2z_i y_{ij} - 2z_j y_{ij} + z_i z_j)$$

$$\sum_{ij} Q_{ij} z_i z_j \rightarrow \sum_i h_i s_i + \sum_{i,j} J_{ij} s_i s_j$$

Mapping not needed for random spin-glass models

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

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

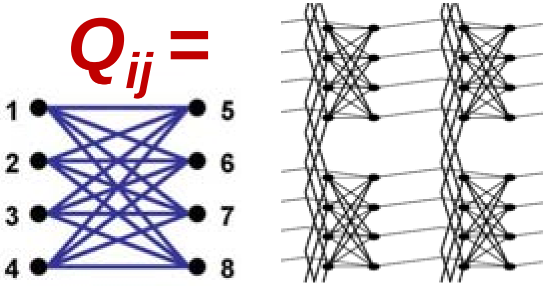

Embedding not needed for native Chimera problems

3 Run the problem several times and collect statistics

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

Probability

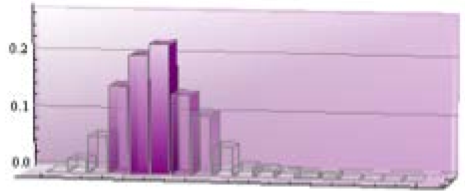

Solution's energy/cost

Performance can be improved dramatically with smart pre-/post-processing

Mapping to QUBO: Graph Coloring Example

Graph Coloring Problem:
Assign one of *k* colors to each vertex so that no two vertices sharing an edge have the same color

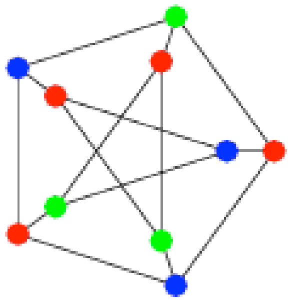

Costing cases

(1) No color or Multi-colored

(2) Same color for connected vertices

Binary variable:

$$x_{v,c} = \begin{cases} 1 & \text{vertex } v \text{ with color } c \\ 0 & \text{vertex } v \text{ not with color } c \end{cases}$$

Violation of requirements encoded as cost:

- (1) unique assignment: Each vertex v must be assigned exactly one color:

$$H_v^{(unique)} = \left(\sum_{c \in C} x_{v,c} - 1\right)^2 \Leftarrow \sum_{c \in C} x_{v,c} = 1$$

- (2) Connected vertices cannot use the same color

$$H_{v,v',c}^{(exclude)} = x_{v,c} x_{v',c} \text{ if } vv' \in E$$

Final QUBO form:

$$H = \sum_v H_v^{(unique)} + \sum_{v,v' \in E} \sum_c H_{v,v',c}^{(exclude)}$$

H = 0 corresponds to a valid coloring

Embedding the QUBO

Embed a triangle onto a bipartite graph

original QUBO

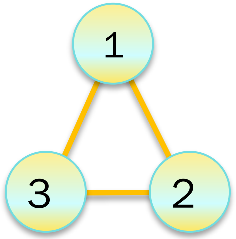

hardware connectivity

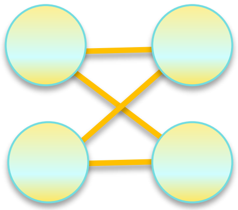

QUBO embedded

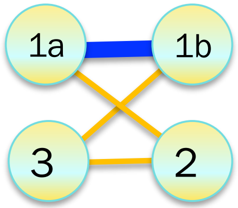

$$H_0 = J_{12}x_1x_2 + J_{23}x_2x_3 + J_{13}x_1x_3$$

$$H_1 = J_{12}x_{1a}x_2 + J_{23}x_2x_3 + J_{13}x_{1b}x_3 + \underline{J_{\texttt{Ferro}}}x_{1a}x_{1b}$$

Strong, but not too strong, ferromagnetic coupling between physical qubits x_{1a} and x_{1b} encourages them to take the same value, thus acting as a single logical qubit x_1

Embedding a realistic problem instance:

Physical qubits on each colored path represent one logical qubit

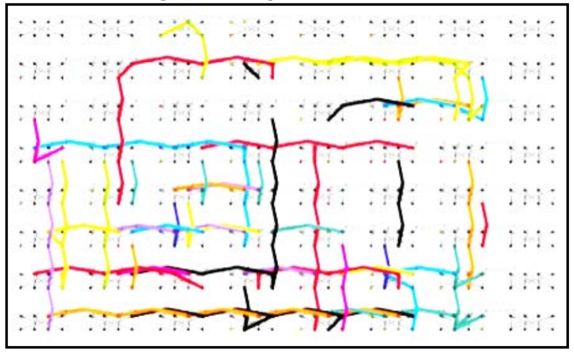

H_0 and H_1 have the same ground state but the energy landscape of the search space differs

Current research investigation: How best to set the magnitude of these “strong” couplings to maximize probability of success

Other NASA Research in Applications

Complex Planning and Scheduling

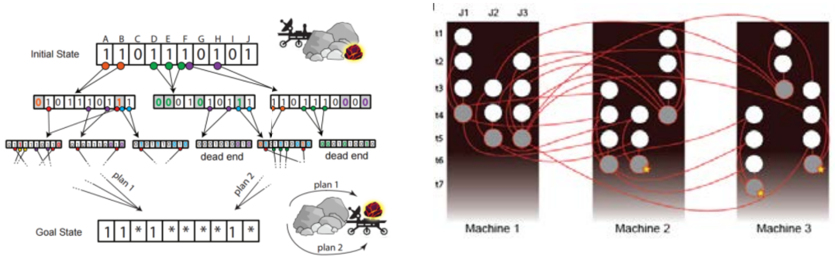

Graph-based Fault Detection

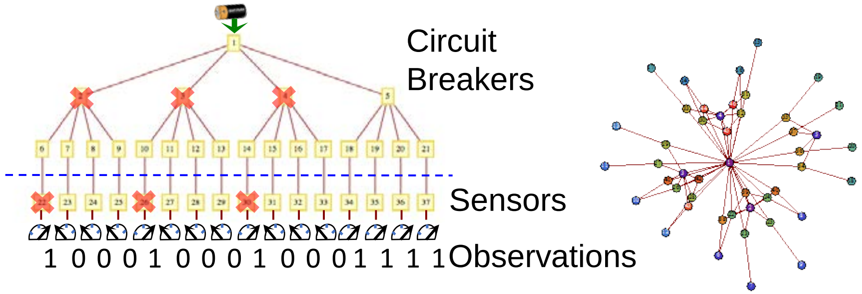

Graph Isomorphism

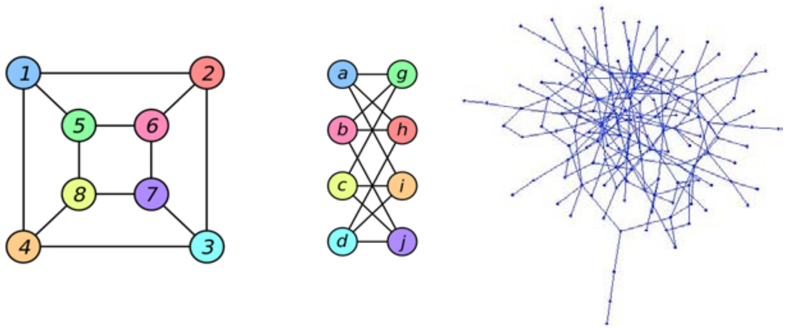

- General **Planning Problems** (e.g., navigation, scheduling, asset allocation) can be solved on a quantum annealer
- Developed a quantum solver for **Job Shop Scheduling** that pre-characterizes instance ensembles to design optimal embedding and run strategy – tested at small scale (6x6) but potentially could solve intractable problems (15x15) with 10x more qubits
- Analyzed simple graphs of **Electrical Power Networks** to find the most probable cause of multiple faults – easy and scalable QUBO mapping, but good parameter setting (e.g., gauge selection) key to finding optimal solution – now exploring digital circuit **Fault Diagnostics and V&V**
- **Subgraph Matching Problems** are common in applications of interest to the intelligence community – similarly, finding **Longest Matching Sequences** important in genomics and bioinformatics

Current NASA Research in Quantum Physics

Calibration of Quantum Annealers

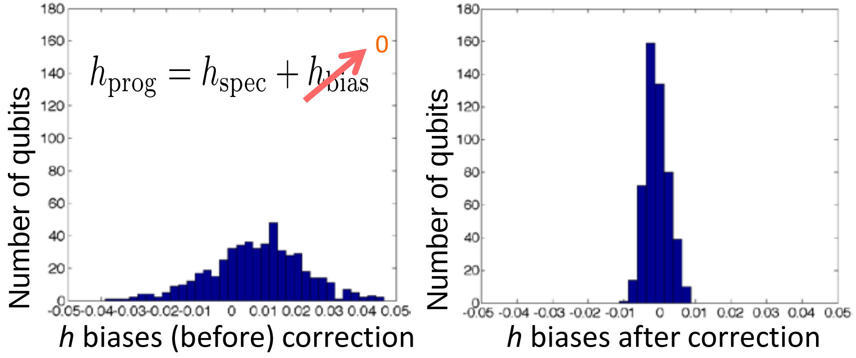

Effect of Noise on Quantum Annealing

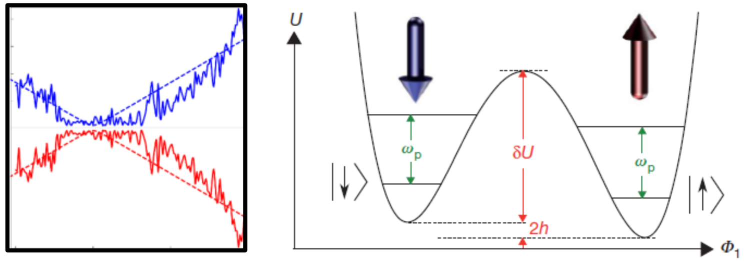

Optimal Embedding & Parameter Setting

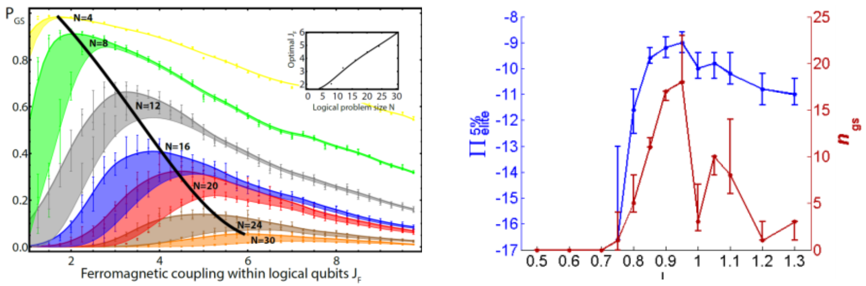

- Developed technique to determine and correct **residual persistent biases** in the programmable parameters of quantum annealers (*h* and *J*) – correction significantly improves performance and reliability (reduction in variability)
- First **realistic noise analyses** show how low-frequency noise dramatically affects the performance of quantum annealers – results being used to design hardware improvements
- Limited hardware connectivity makes embedding challenging – good runtime parameters determined by considering the **nature and dynamics of chains** – quick scans can be used to predict performance of extensive scans
- Small instances of **hard problems at phase transitions** in combinatorial optimization are intractable – they can be designed by looking at solvability phase transitions
- Predict **tractability of application problems** by studying the scaling of energy gaps and density of bottlenecks in spin glass phase

Emerging Quantum Hardware

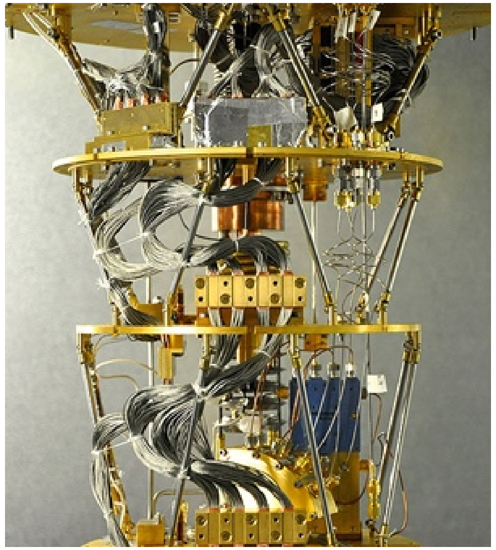

D-Wave Systems

- What should we do with the emerging and very exciting, but quite limited, quantum computational devices?
  - Still too small for solving practical problems
- Couple of possibilities:
  - Quantum supremacy
  - ***Develop intuitions for quantum heuristics***

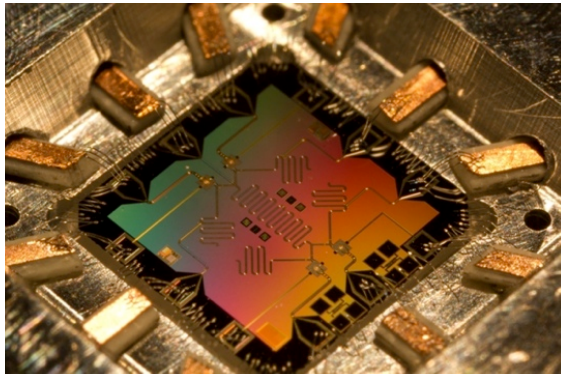

Google

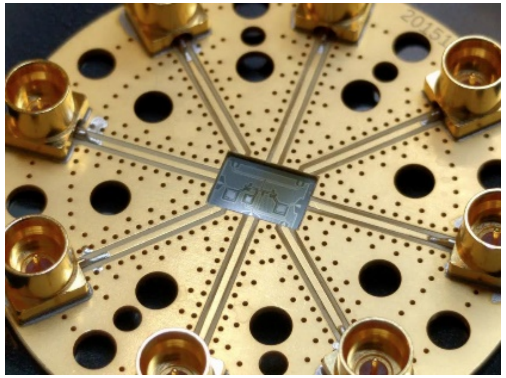

Rigetti

IARPA

Current NASA Research in Quantum

Complex Planning and Scheduling

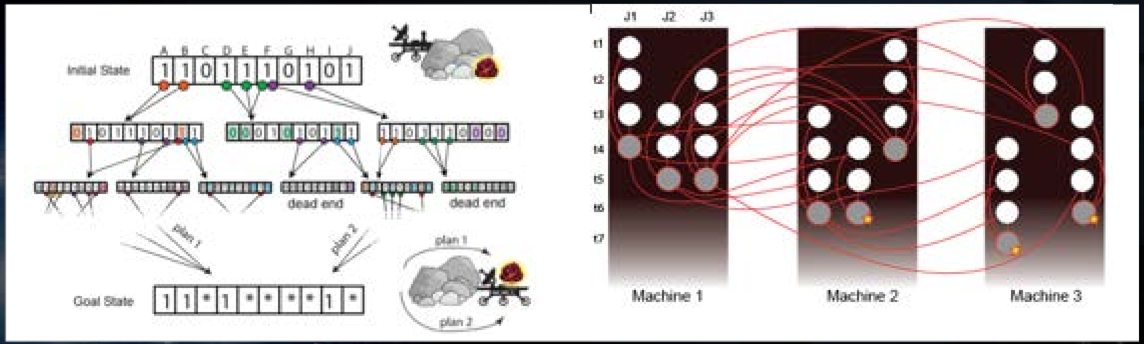

Calibration of Quantum Annealers

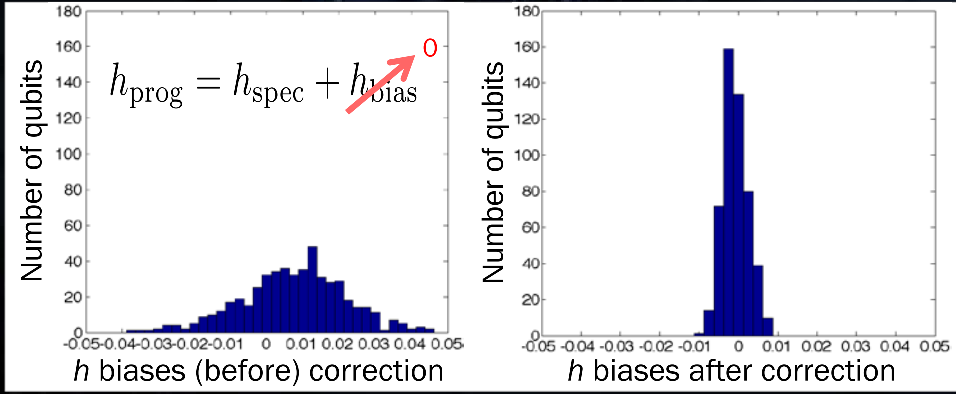

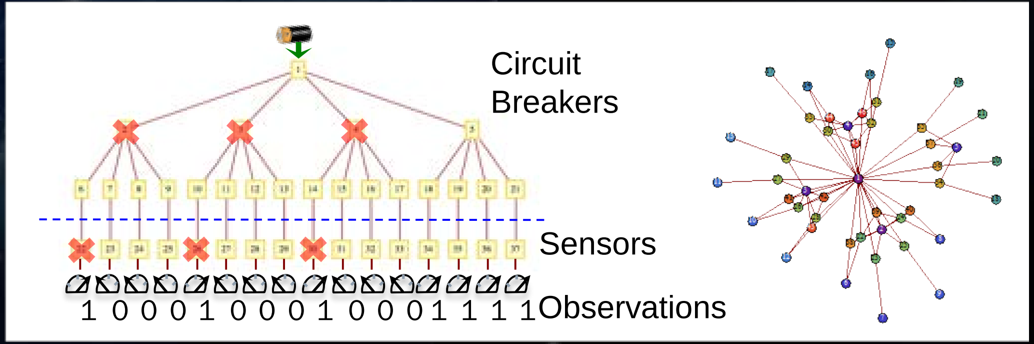

Graph-based Fault Detection

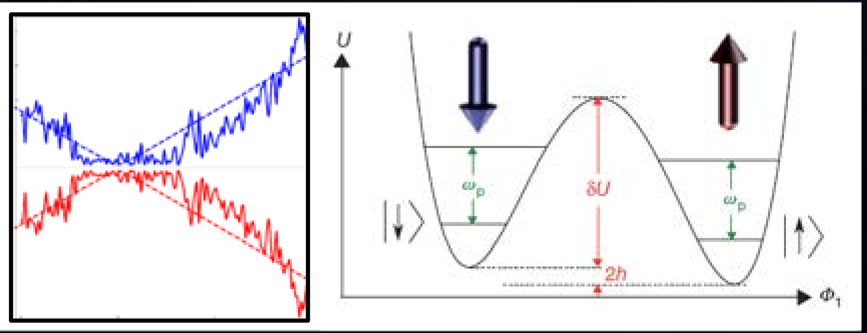

Effect of Noise on Quantum Annealing

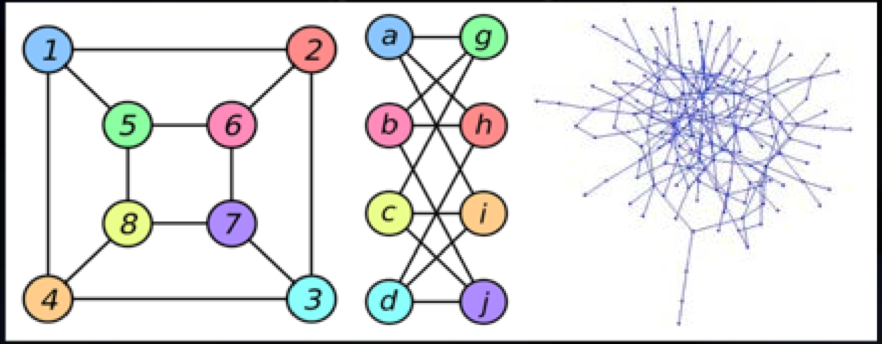

Graph Isomorphism

Optimal Embedding and Parameter Setting

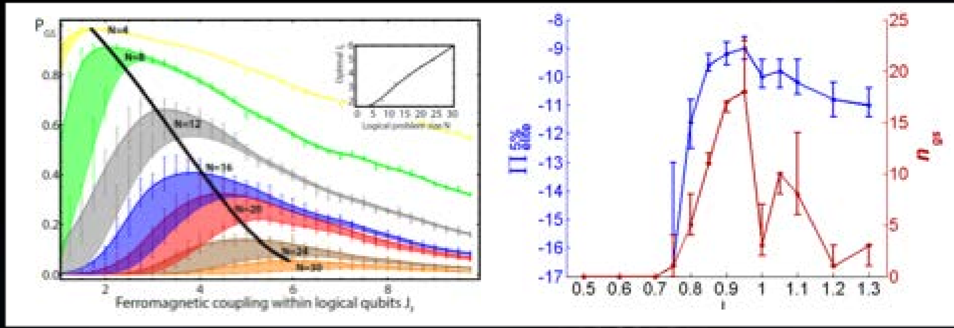

Assured Availability of UTM Network

UTM: UAS Traffic Management

UAS: Unmanned Aerial System (includes UAV, ground control, and comm.)

Future

- Higher vehicle density
- Heterogeneous air vehicles
- Mixed equipage
- Increased autonomy
- Greater vulnerability to communication disruptions

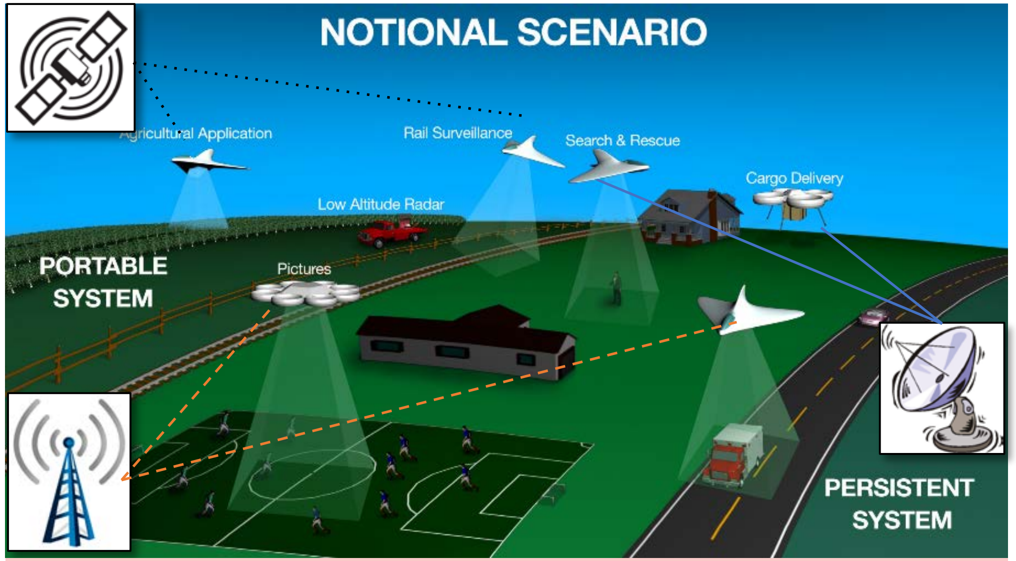

Explore quantum approaches for:

- Robust network design
- Track and locate moving jammers
- Secure communication of codes supporting anti-jamming protocols

- Harness power of quantum technologies to address cybersecurity challenge of assured availability
- Leverage work on QKD for spread spectrum codes

P. Kopardekar, et al., ***Unmanned Aircraft System Traffic Management (UTM) concept of operations***, AIAA 2016-3292

Quantum Communication

- Quantum Cryptography
  - Uses single or entangled photons to distribute a secure random key → QKD
- Teleportation
  - Uses entangled photons to send quantum information → necessary for quantum repeaters and networks
  - However, requires pre-shared entanglement, 3 photons, and hard (impossible?) to implement 2-photon gate
- Superdense Teleportation (SDT)
  - Uses entangled photons and restricted space to send quantum information
  - Remotely prepare quantum information with reduced resources
  - Requires only 2 photons of which only one has to be sent
- NASA funding development of SDT demo from ISS for last 5 years
  - Lab work performed at UIUC under auspices of Paul Kwiat
  - Independent assessment and flight system engineering conducted by MIT-LL
    - 2023: SDT from ISS to ground
    - ~2026: Demo on multi-spacecraft (i.e., swarm of small sats)
    - Late 2020’s: QKD from Mars orbiter

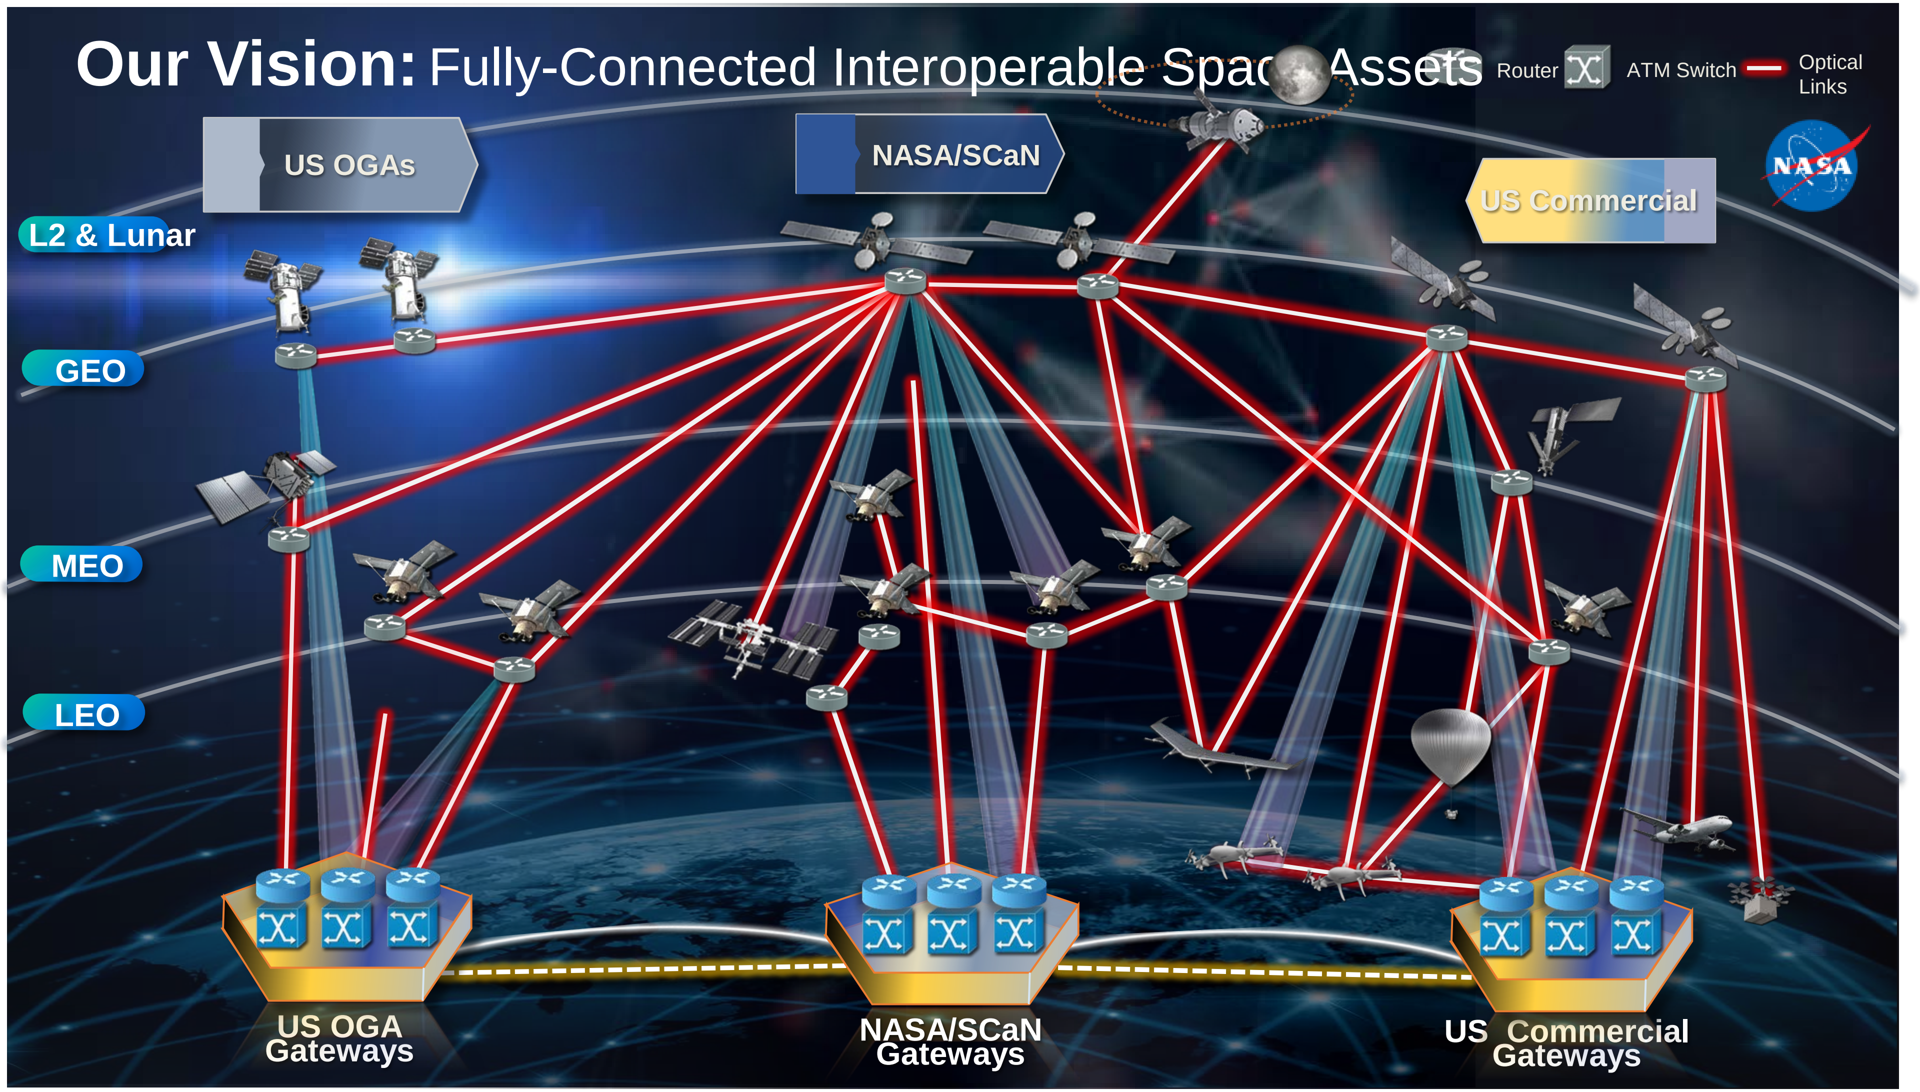
Our Vision: Fully-Connected Interoperable Space Assets
Router
ATM Switch
Optical Links
US OGAs
NASA/SCaN
US Commercial
NASA
L2 & Lunar
GEO
MEO
LEO
US OGA Gateways
NASA/SCaN Gateways
US Commercial Gateways

Advanced Computing Mission

Enable the science & engineering required to meet NASA's missions and goals

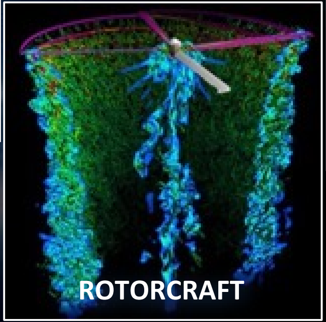

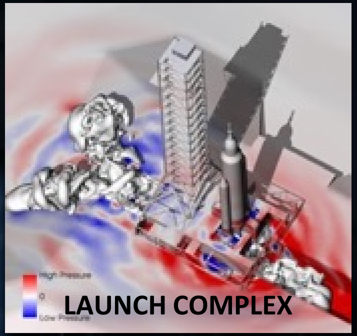

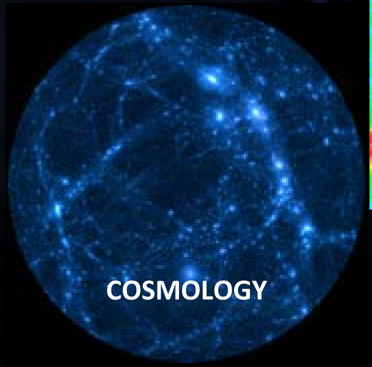

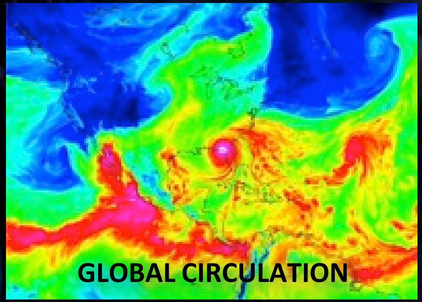

Effective, stable, production-level HPC environment

Advanced technologies to meet future goals