Quantum Information Processing for SMD: A Vision

Lucas Braydwood (né Brady)
QuAIL Associate Lead
Eleanor Rieffel
QuAIL Lead
NASA Ames Research Center

July 2025

Prepared by NASA Quantum Artificial Intelligence Laboratory (QuAIL) team: M. Sohaib Alam, Namit Anand, Humberto Munoz Bauza, Lucas Brady (Associate Lead), Stephen Cotton (joint with NAS), Zoe Gonzalez Izquierdo, Shon Grabbe (Deputy Lead), Erik Gustafson, Stuart Hadfield, Aaron Lott, Filip Maciejewski, Gianni Mossi, Eleanor Rieffel (Lead), Jason Saied, Davide Venturelli, Zhihui Wang

Aim: **Leverage NASA QuAIL team's expertise and connections to expand quantum information processing (QIP) know-how across NASA centers**

- Guidance on Investment
- Increase awareness, generate SMD challenge problems, and foster collaborations across centers
- Cross-agency collaborations, identifying funding opportunities, collaboration with industry
- Workforce development

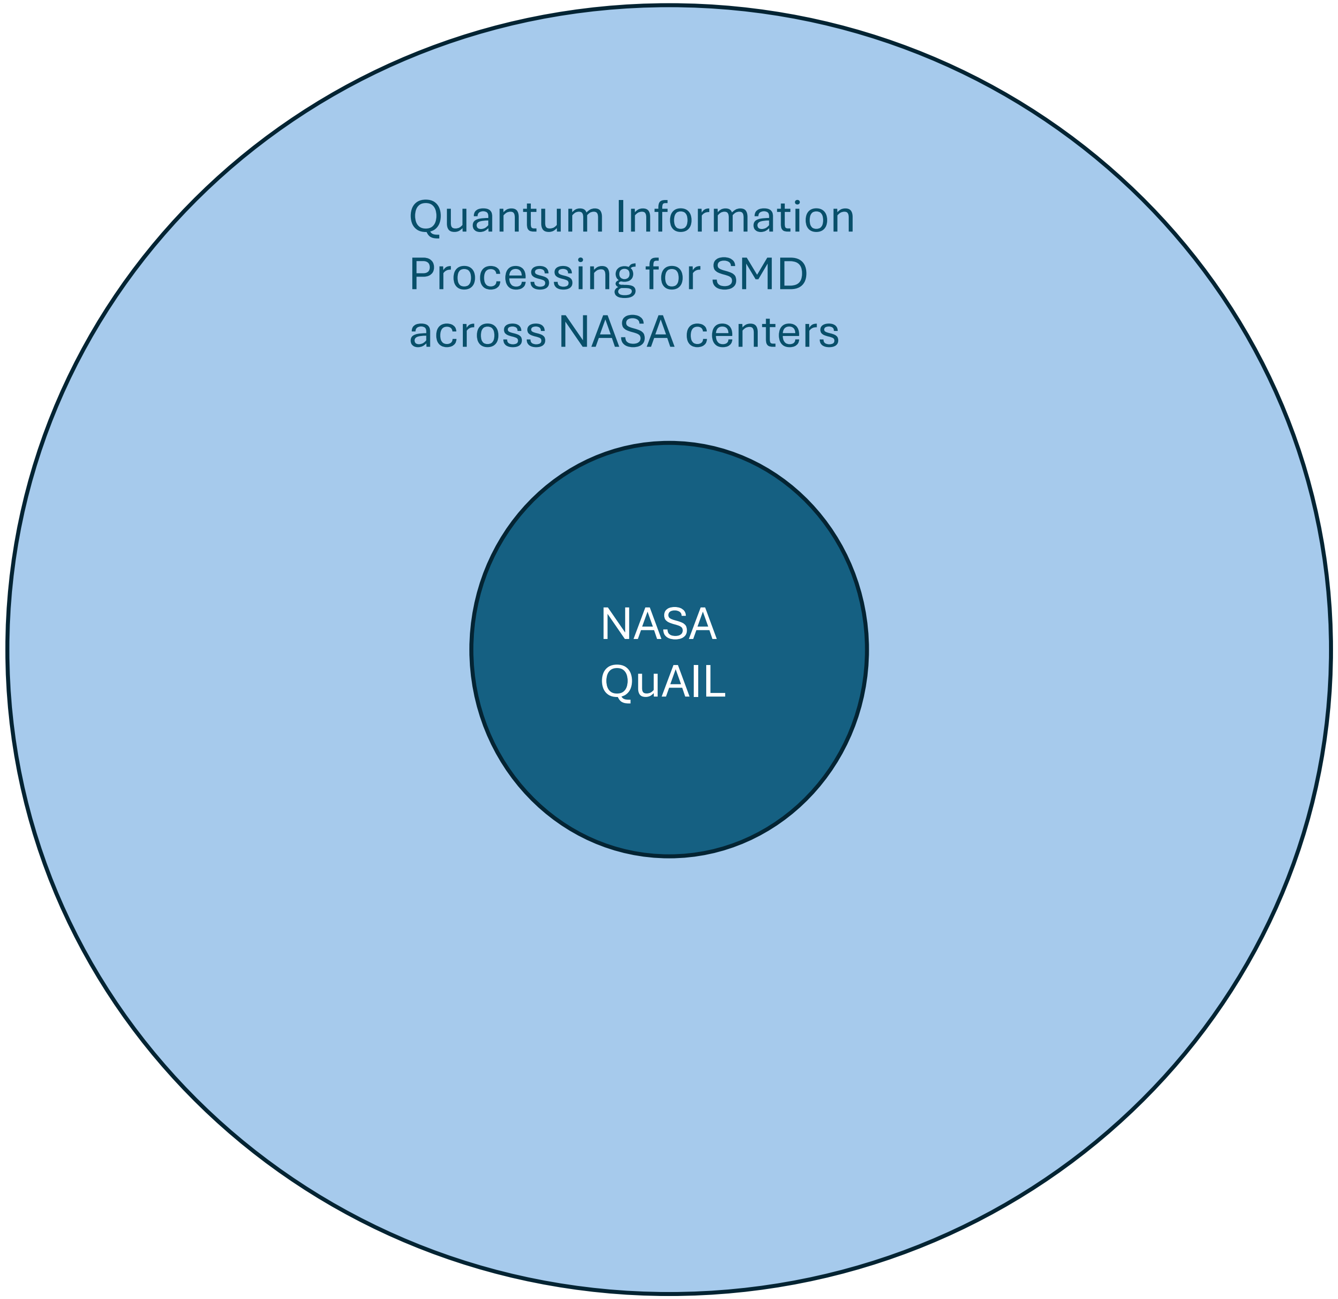

Quantum Information Processing

The power of Quantum Information Processing (QIP) comes from encoding information in a non-classical way, and processing it using uniquely quantum operations

- More than just using quantum effects for computing
- Includes quantum computing, but also much smaller and more targeted quantum information processing

Broad applications

- more time and energy efficient computation
- higher levels of security
- more precise sensing

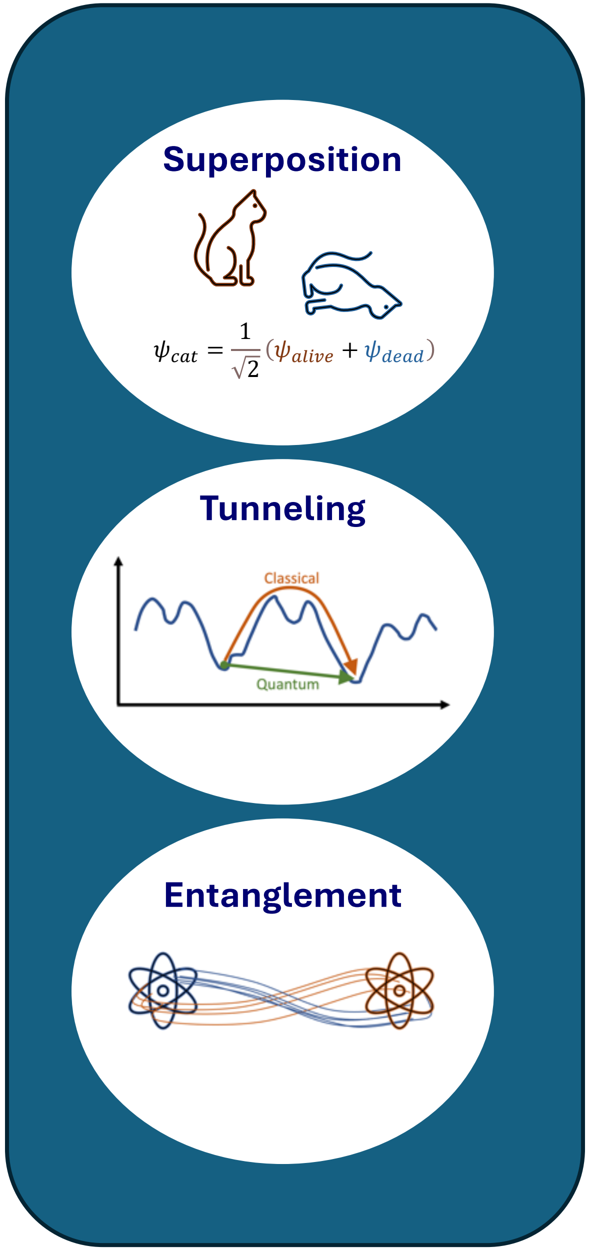

Quantum algorithms and subroutines

- E.g. Shor's factoring algorithm: days versus longer than the age of the universe. Breaks all deployed public key encryption

QIP for quantum sensor arrays

- Quantum sensors provide greater precision and require fewer samples
- Add QIP for even greater benefits
  - Pre-processing: generate and distribute entangled quantum states across array
  - Process quantum sensor data rather than immediately turning it into classical data

Quantum networks

- quantum repeaters extend range
- Utility-scale quantum computers will be quantum-networked QPUs

Examples of Potential SMD Quantum Applications

Past SMD Examples: ❖ **Machine Learning** ❖ **Combinatorial Optimization**

Quantum-Assisted Machine Learning for Wildfire **Image Processing**

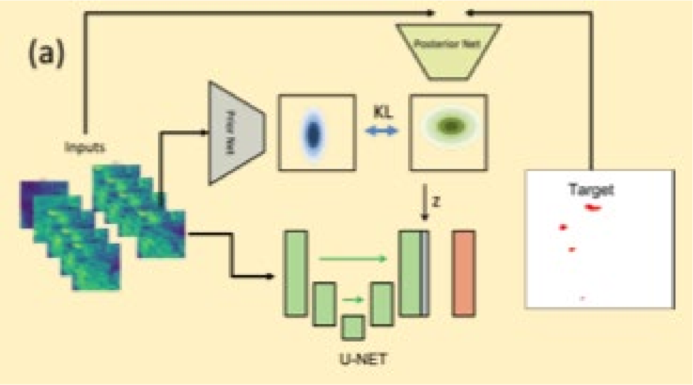

A. Asanjan *et al.* (2023), *Probabilistic Wildfire Segmentation Using Supervised Deep Generative Model from Satellite Imagery*, **Remote Sensing**

Scheduling of Observing Satellites

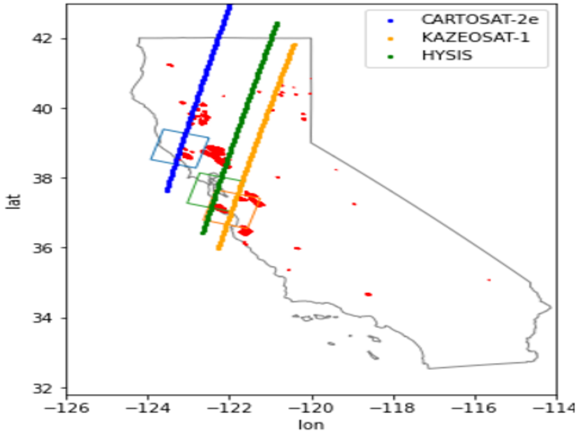

L. T. Brady *et al.* (2025), *Partitioned Iterative Quantum Scheduling of Satellites for Urgent Disaster Response*, accepted **IGARSS 2025**

MODIS Dataset **Similarity Search**

N. Gao et al. (2020), High-dimensional similarity search with quantum-assisted variational autoencoder, **KDD** '20, 956-964

Potential Future SMD Applications

- Fluid Dynamics
  - ❖ Ocean Dynamics
  - ❖ Atmospheric Dynamics
  - ❖ Heliophysics
- Chemistry and Materials
  - ❖ Biological Sciences
  - ❖ Stellar Evolution
  - ❖ Materials Design
- Processing for Sensing
  - ❖ Data Analysis and Anomaly Detection
  - ❖ Direct Integration into Quantum Sensors
  - ❖ Entangled State Generation and Distribution for Sensor Arrays
  - ❖ Quantum-Enhanced Very Long Baseline Interferometry
- Other Applications
  - ❖ Natively Quantum Tasks
  - ❖ And beyond!

New Era for Quantum Computing

Quantum supremacy achieved

- Perform computations not possible on even largest supercomputers in reasonable time
- Google – NASA – ORNL collaboration

F. Arute et al. (2019), Quantum supremacy using a programmable superconducting processor, Nature ***574****, 505-510*

Samples obtained in a few minutes from a quantum processor would take over 47 years to obtain on the largest supercomputer using today's best methods

2023 Update

A. Morvan, B. Villalonga, X. Mi, ***S. Mandrà****, et al., (2024) Phase transition in Random Circuit Sampling, Nature 634 (8033), 328-333*

Exp.	1 amp.	1 million noisy samples		
	FLOPs	FLOPs	XEB fid.	Time
SYC-53 [9]	$6.44 \cdot 10^{17}$	$2.60 \cdot 10^{17}$	$2.24 \cdot 10^{-3}$	6.18 s
ZCZ-56 [10]	$6.24 \cdot 10^{19}$	$6.40 \cdot 10^{19}$	$6.62 \cdot 10^{-4}$	25.3 min
ZCZ-60 [11]	$1.32 \cdot 10^{21}$	$1.41 \cdot 10^{23}$	$3.66 \cdot 10^{-4}$	38.7 days
This work	$4.74 \cdot 10^{23}$	$6.27 \cdot 10^{25}$	$1.68 \cdot 10^{-3}$	47.2 yr

… but so far only for an artificial problem

- Quantum hardware currently too small for solving practical problems intractable on classical supercomputers
- These devices need to scale up and become more reliable

So what to do in the interim?

- Unprecedented opportunity to invent, explore, and evaluate quantum algorithms *empirically*

NASA QuAIL Focus

- Algorithms and applications to enable more ambitious and time- and energy-efficient missions
- Codesign of hardware/algorithms
- Tools for advancing quantum computing, from quantum circuit simulation, noise characterization, error correction, compilation to realistic hardware

Quantum Info Processing Research Stack

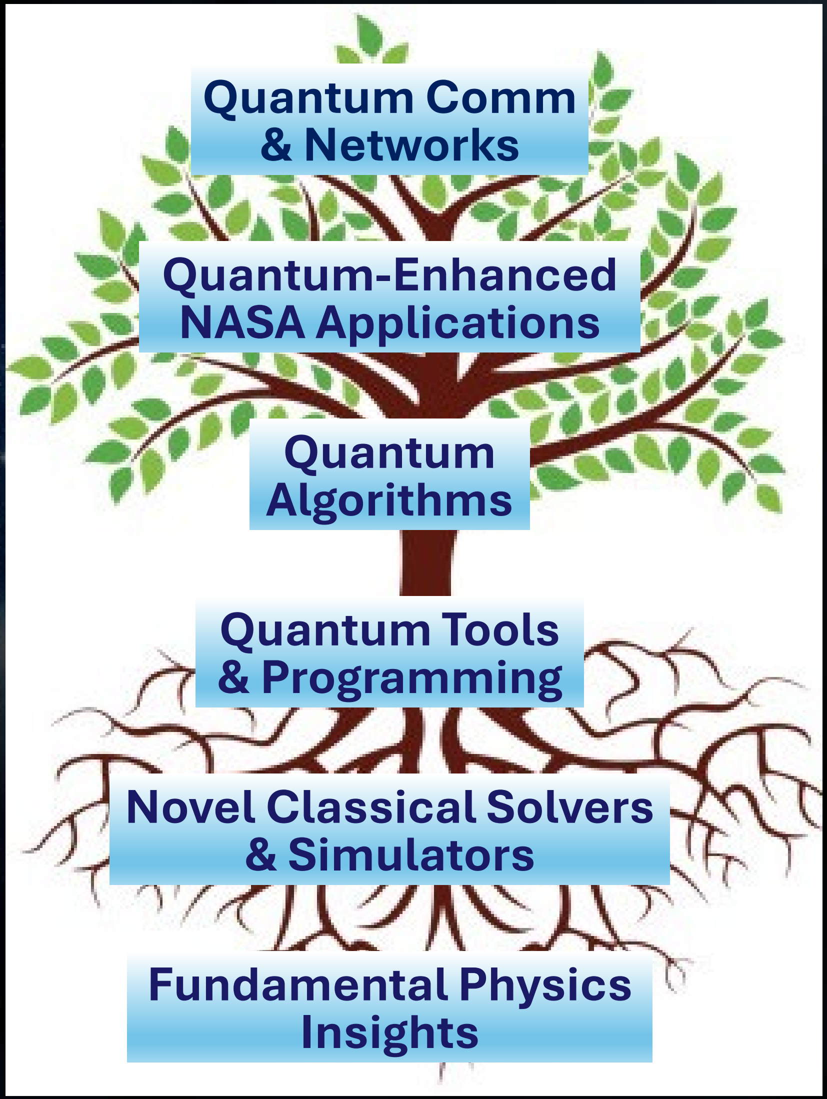

Communication & Networks

Quantum networking   Distributed QC

Application Focus Areas

Logistics Optimization   Planning & scheduling
Machine learning   Fault diagnosis   Material Discovery
Computational Fluid Dynamics   Differential Equation Solving

Software Tools & Algorithms

Quantum algorithm design   Compiling to hardware
Mapping, parameter setting, error mitigation
Hybrid quantum-classical approaches

Solvers & Simulators

Physics-inspired classical solvers
HPC quantum circuit simulators

Physics Insights

Co-design quantum hardware

E. Rieffel et al. (2024), Assessing and Advancing the Potential of Quantum Computing: A NASA Case Study, FGCS 160,598-618
E. Rieffel *et al.* (2019), From Ansätze to Z-gates: A NASA view of quantum computing, *Adv. in Parallel Computing* **34**, 133–160
R. Biswas *et al.* (2017), A NASA perspective on quantum computing: Opportunities and challenges, *Parallel Computing* **64**, 81–98

What challenges do you face where quantum information processing could provide benefits?

How best to get such conversations started?

Example of Quantum Integration:
Quantum and Classical Generative Machine Learning

A. A. Asanjan, L. T. Brady, Z. G. Izquierdo, N. Suri, P. A. Lott, E. Rieffel, R. Biswas

- Wildfires cause massive economic and health damage each year and California and other states
- Goal: create classical and hybrid-quantum machine learning models for image-to-image translation of satellite imagery

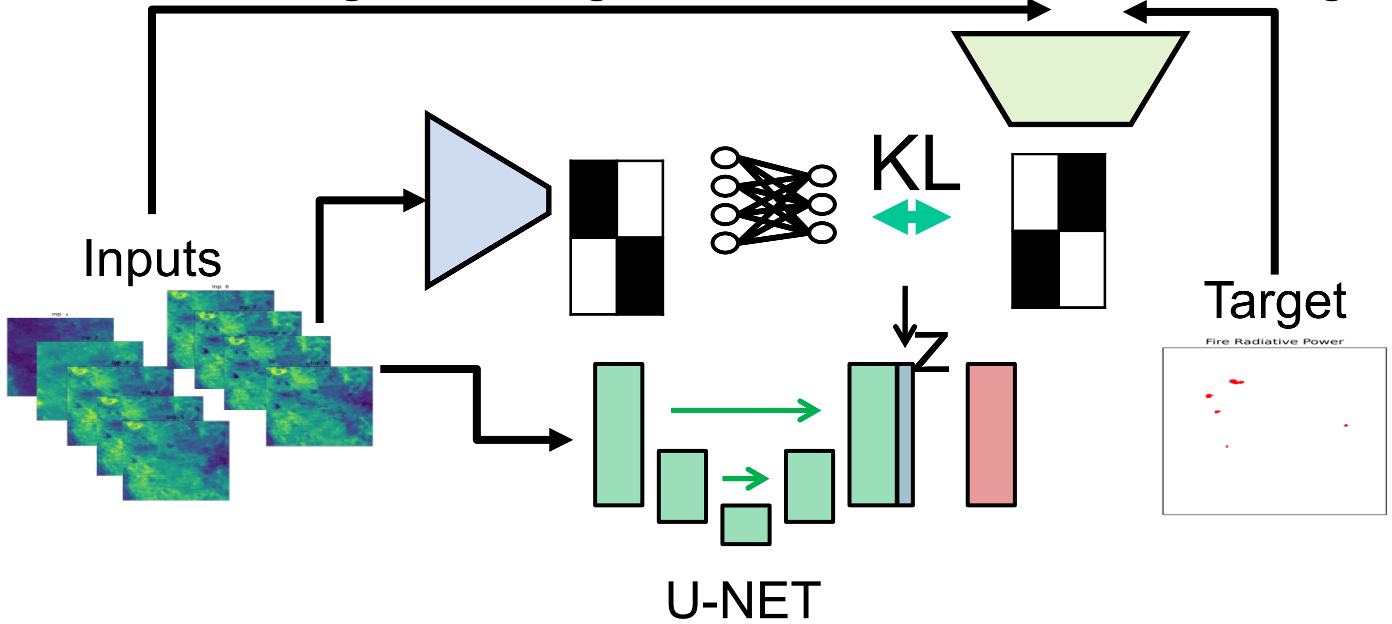

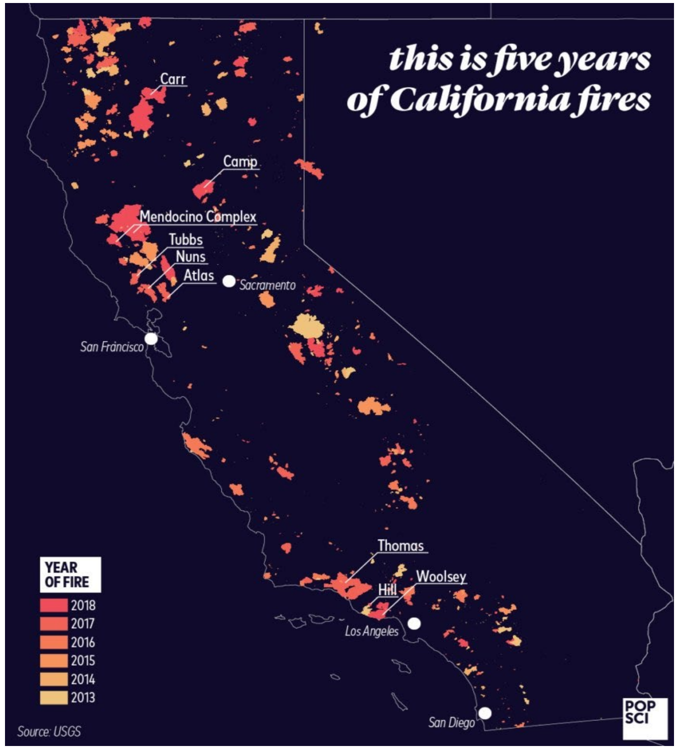

Credit: USGS

- We use a probabilistic U-Net architecture to get a quantification of uncertainty.
- Even in the classical model, we innovate by including a classical Boltzmann machine as a finite latent space rather than a typical continuous latent space

Classical Results:

A. A. Asanjan, et al., Probabilistic Wildfire Segmentation Using Supervised Deep Generative Model from Satellite Imagery, Remote Sensing,15,2718 (2023)

Quantum and Classical Generative Machine Learning: Quantum Additions and Results

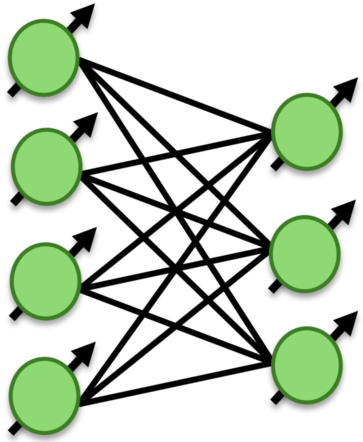

- We used a custom-developed Born Machine, that uses quantum rather than thermal effects to generate its distribution
- Our classical methods also sampled using an in-house QuAIL open-source package, PySA that works off parallel tempering
- We tested multiple models but were limited in training time due to our quantum simulator
- Even with a smaller latent space, the quantum model was competitive with classical methods

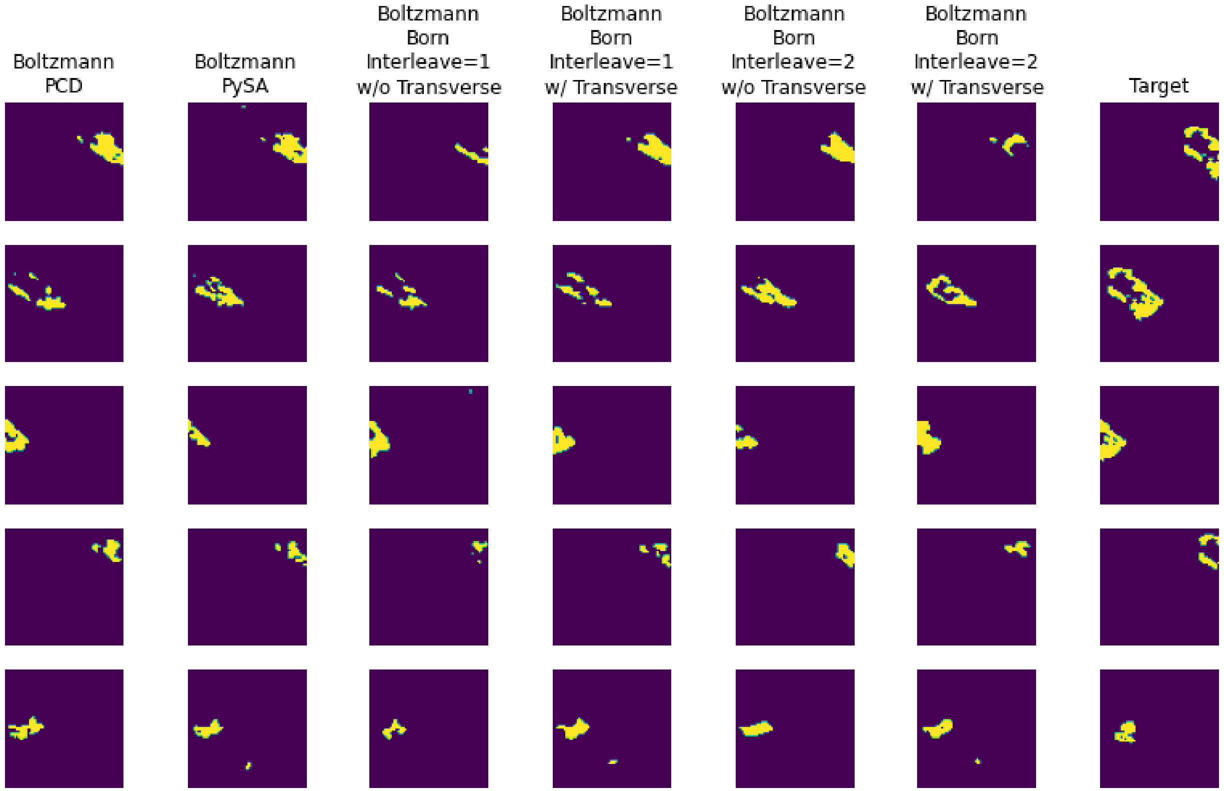

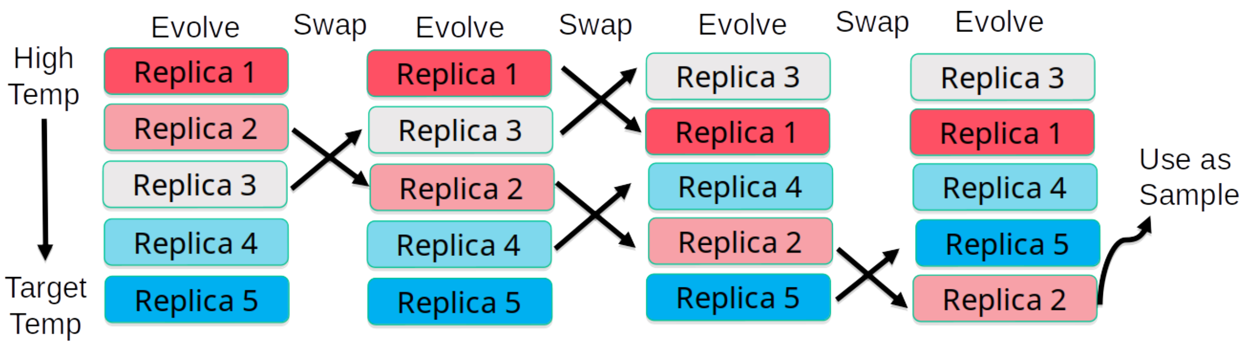

Why Integrate Quantum Thinking now?

Quantum-Ready	Quantum-Inspired	Quantum-Enabled
Classical algorithms with pieces that can be exchanged for quantum components when ready	Classical-algorithms whose design was inspired by quantum phenomena	Quantum-ready algorithms that have had quantum pieces integrated
Implementable in full today	Implementable in full today	Implementable only on small systems today
Can mimic today's state-of-the-art with certification of quantum augmentation	Novel methods that have sometimes beaten existing state-of-the-art	Expected to eventually have scaling advantage
Have been deployed on real-world applications	Have been deployed on real-world application	Currently too small and noisy to benefit real-world applications

Partitioned Iterative Quantum Scheduling of Satellites for Urgent Disaster Response: Case study of Wildfire – AIST Project

L. T. Brady, T. Park, H. Hashimoto, Z. G. Izquierdo, A. Michaelis, E. Rieffel, S. Grabbe

- We curated a set of GOES (Geostationary Operational Environmental Satellites) data to identify historical firetrack data
- We also incorporated information on the Wildland Urban Interface (WUI) to create a sequence of simulated imaging requests that would have been needed for fire tracking

GOES image at 2020-08-18_00:01

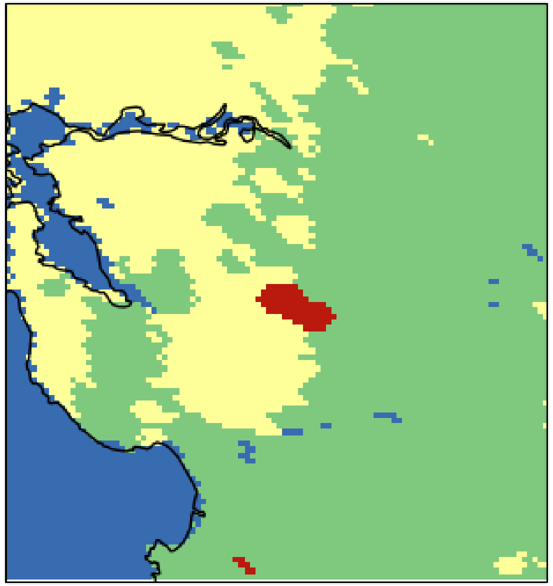

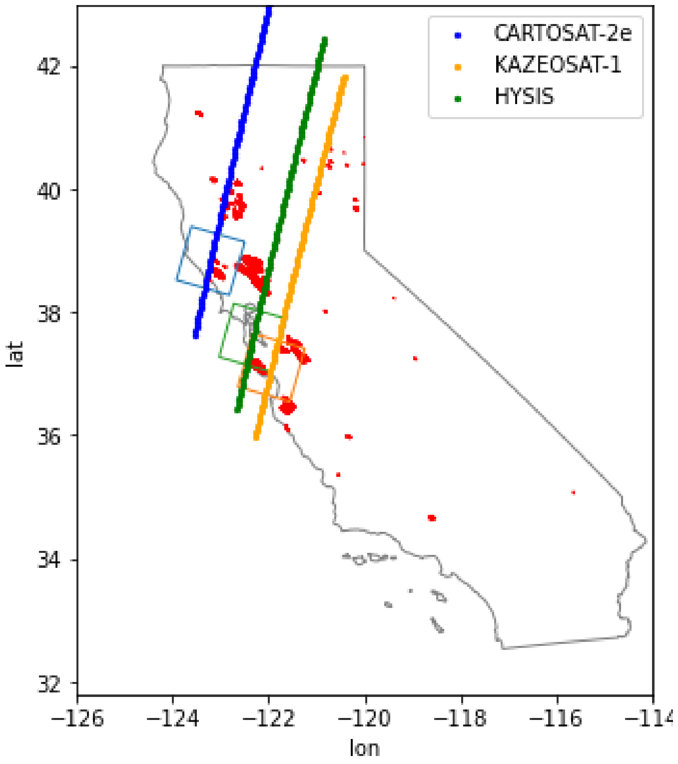

- Our goal was to create schedules for hypothetical satellites to most effectively satisfy these imaging requests
- We incorporated satellite orbit data from three real world satellites to specify the problem as well as hypothetical imaging constraints from postulated satellites
- This was formulated into a Maximum Independent Set problem to be solved by classical and quantum models

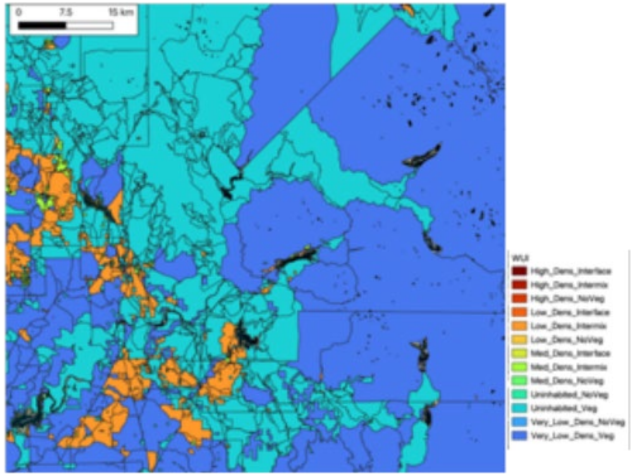

WUI classification over California

Partitioned Iterative Quantum Scheduling of Satellites for Urgent Disaster Response: Case study of Wildfire – Quantum Partitioning and Results

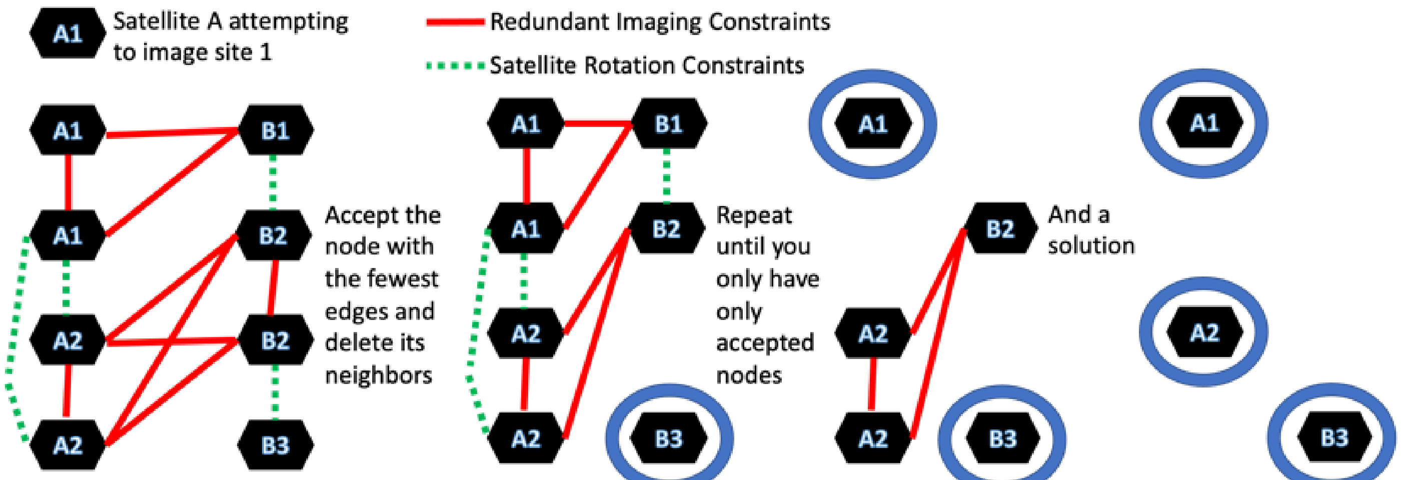

- Using the MIS formulation, we employed classical greedy algorithms
- A novel iterative quantum algorithm reduces to the classical greedy algorithm in the simplest case
- The quantum algorithm improves on classical greedy at higher depth

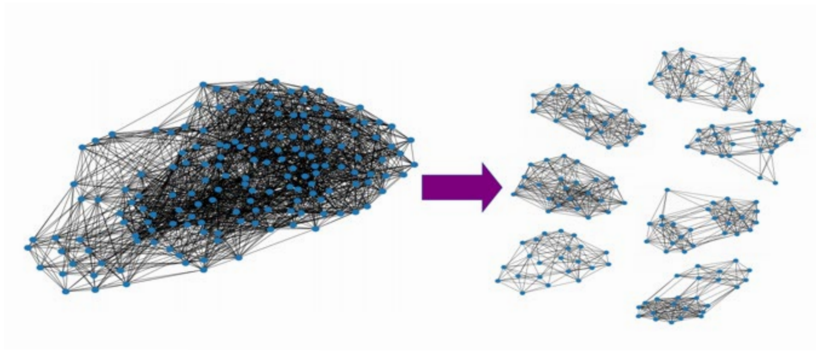

- Small quantum simulators necessitate graph partitioning, another novel quantum method

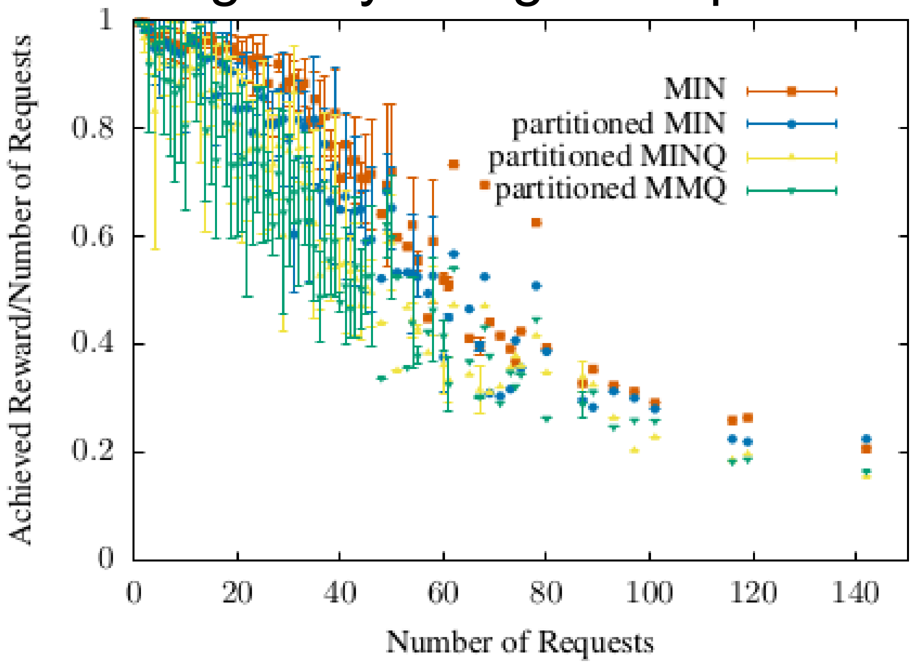

- MIN = greedy, MINQ/MMQ = new quantum algorithms
- The quantum algorithms perform worse than classical algorithms. This is predicted to be due to the partitioning methods, not the core algorithm

SMD Quantum Challenges and Collaborations

We seek to embed quantum experts in SME teams and embed SMEs in quantum teams

Formulate grand-challenge problems from science examples

Formulate stepping-stone problem instances to inspire and assess progress toward target challenges.

Bring together teams to work on challenge problems to develop novel solutions

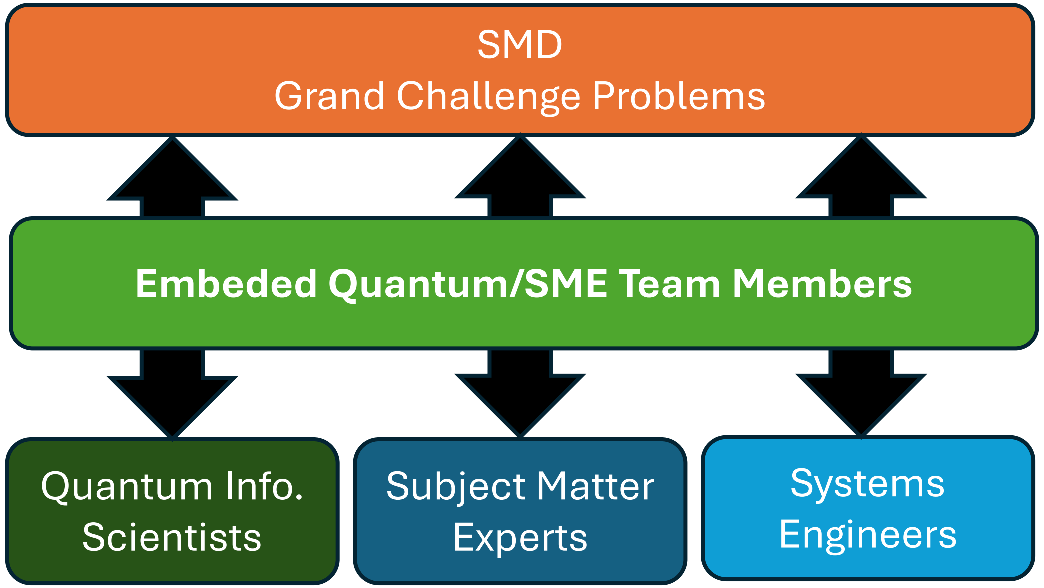

Problem Instance Specification (1 of 2)

Aim is precise specification of two types of instances

- Full application instances: if solved, would have significant impact
- Test instances: Smaller or simpler instances that are steppingstones to impactful instances

Precise Specification for Full Application Instances

- Is the requirements for solving a class of instances or a single instance?
- Is the request for a single solution to each instance, or a set of good instances?
  - What constitutes a sufficiently good solution?
  - For optimization, do we need provably optimal , or if not what approx. threshold do we need to pass?
- How quickly would the instance(s) need to be solved to meet application requirements?
- Is there a need to update instance solutions as new information comes in?
- If known, what are the underlying mathematical or computational bottlenecks that is limiting traditional computing solutions?

Problem Instance Specification (2 of 2)

Precise Specification for Full Application Instances (cont.)

- At what problem sizes are better computational approaches needed?
- What are the precision requirements?
- If data processing, what datasets are available?
- Can the solution algorithm be parallelized?
- How valuable is this to SMD?
  - How would calculating a better solution to the problem increase its value to SMD?
- Is there commercial value to solving this application instance? If so, what is commercial value to faster solution, better solution, etc.?
- What is the current cost to solution if solvable?

Precise Specification for Test Instances

- Relation between these instances and the full application instances
- Different classes for different aspects of full application instances
- Size, precision, accuracy for these problems
- Available datasets
- Classical reference solutions if known from literature, self-computed, or possibly derived from planted solutions
- Standardized schema for storing instances
- What are the problem inputs and outputs and associated characteristics of each (size/number, dimension, quality, accuracy, distribution, etc.)
- Can the solution be parallelized?

Provides targets and benchmarks for emerging computational technologies generally

Example Problem Instance Development Efforts

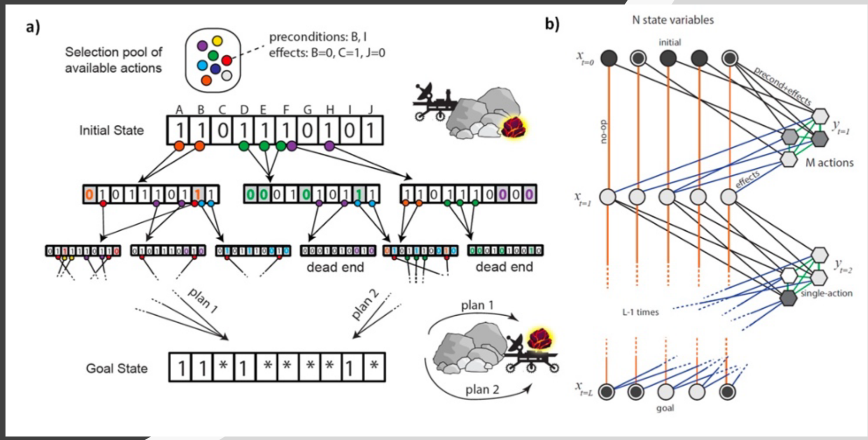

- DARPA Quantum Benchmarking (QB) explored potential utility scale benchmarks in three broad categories [QB24]: chemistry, materials science, and non-linear differential equations
  - QuAIL is T&E for this program and a contributor on the published benchmarks
- DARPA Quantum Inspired Classical Computing (QuICC) Test and Evaluation (T&E) Work
  - Novel technique for randomly generating hard instances of disordered Ising Hamiltonians with unique planted ground states using cryptographic protocols [MAN23]
- Quantum circuit layout synthesis benchmarks [SHA23] – references QuAIL's earlier work
- Two families of planning problems, navigation-types and scheduling-type, used to benchmark planning s/w [RIEF15]

[MAN23] S. Mandra,, G. Mossi, E. G. Rieffel, Generating Hard Ising Instances With Planted Solutions Using Post-Quantum Cryptographic Protocols, 2023. arXiv:2308.09704.

[RIEF15] E. G. Rieffel, D. Venturelli, B. O'Gorman, M. B. Do, E. M. Prystay, V. N. Smelyanskiy, A case study in programming a quantum annealer for hard operational planning problems, Quantum Information Processing 14 (1) (2015) 1–36

[QB24] https://www.darpa.mil/news/2024/quantum-computing-applications

[SHA23] I. Shaik and J. van de Pol, Optimal Layout Synthesis for Quantum Circuits as Classical Planning (full version), arXiv:2304.12014 (2023). https://github.com/ipc2023-classical/domain-quantum-layout

Poll of Activities to Increase SMD QIP Awareness

SMD QIP CoP	Community of Practice (CoP) to coordinate and advance QIP developments within SMD
SMD QIP "Gatherings of Experts"	Identify and/ or address specific SMD QIP problems of interest and foster collaboration with SMD practitioners
Industry and Academic Site Visits	Leverage the QuAIL team’s location and connections in Silicon Valley to organize small group tours for SMD people of local quantum hardware
SMD QIP Ask Me Anything (AMA)	Informal monthly opportunity for SMD management and staff to discuss QIP topics of interest with QuAIL staff
SMD QIP Hackathons	SMD and QuAIL researchers and scientists with QIS hardware/software developers work on specific SMD QIS challenge problems/project over a short period of time, usually 24-48 hours

External Collaboration

Vehicles to establish Industry partnerships

QuAIL has traded expertise and collaboration for access to quantum hardware and a seat during co-design, and can leverage this trade for the larger benefit of SMD researchers

Supported a wide array of quantum programs via Reimbursable & Non-Reimbursable Space Act Agreements, with further collaboration in process

These have also enabled collaboration exchange of QuAIL expertise for **no-cost access to quantum hardware**

This can be replicated to support other SMD tasks and with other QIP companies

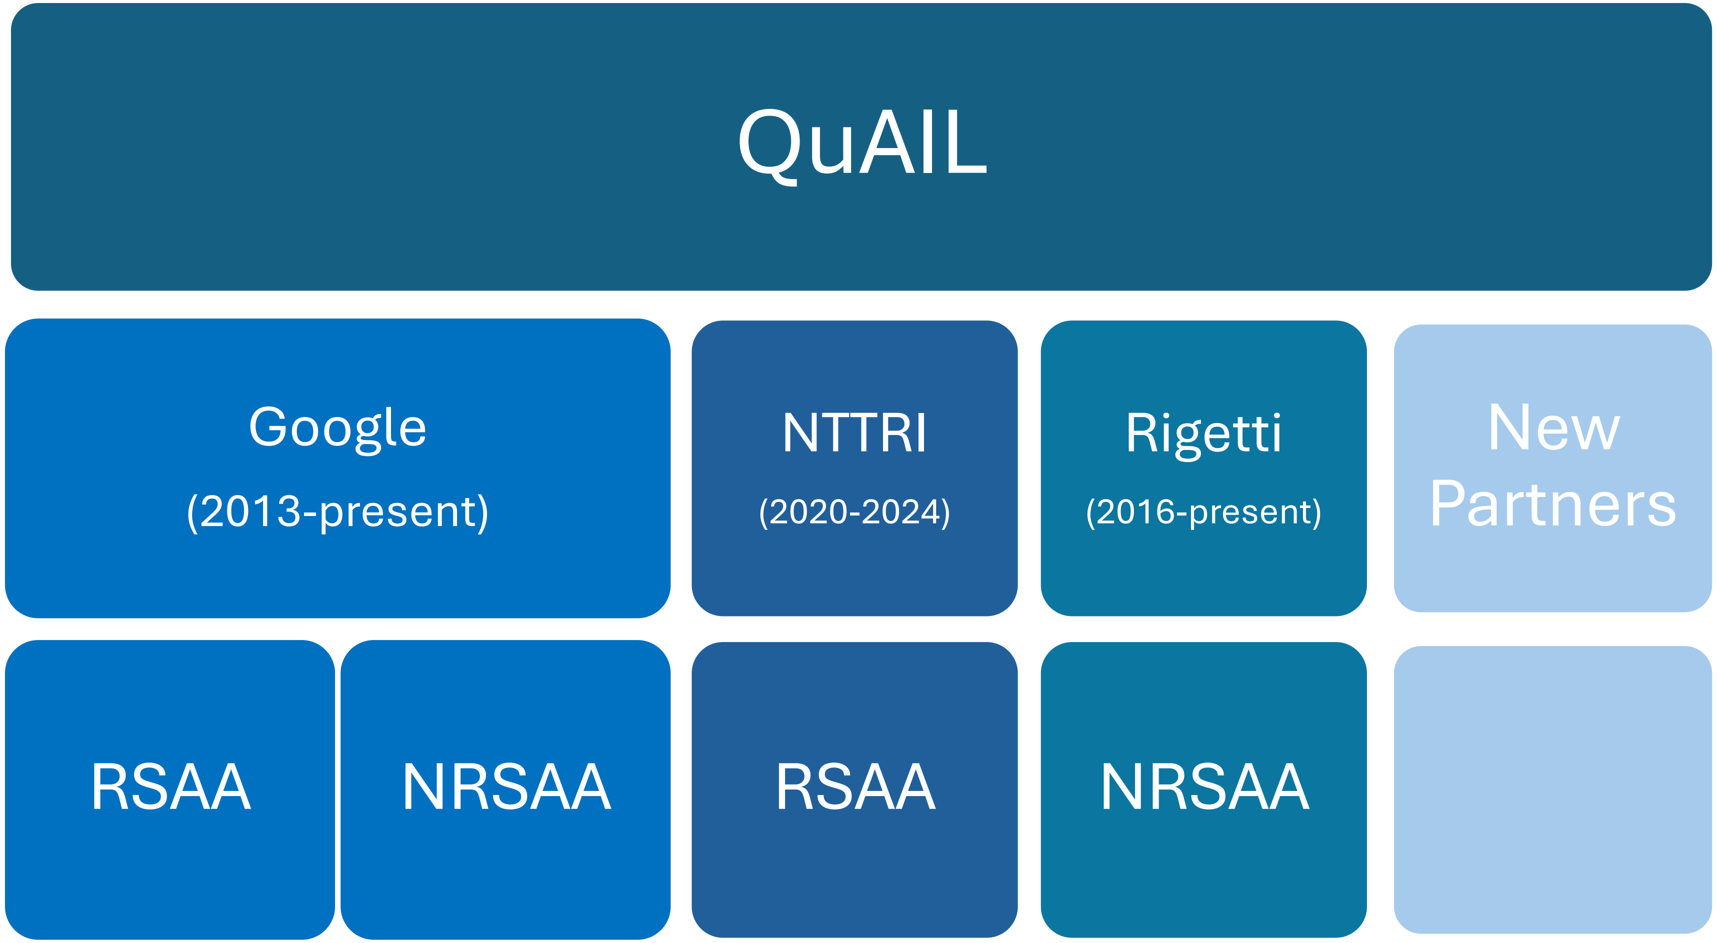

Vehicles to establish inter-agency partnerships

Similarly in the past, inter-agency partnerships have led to access to the broader community and funded work that we hope to replicate with future topics of mutual interest to the benefit of SMD SMEs

Communicate SMD challenge problems with other agencies to identify tech overlap/synergy to identify areas of synergistic/mutual interest.

Form IAA to collaborate, integrate, seek joint funding and/or sponsor shared funding opportunities based on problems of mutual interest.

QuAIL has helped many agencies quantum programs	Engage agencies in collaboration to jointly solve SMD problems

The National Quantum Initiative Reauthorization H.R.6213 is waiting for a full congressional vote. It includes a section on NASA.

How quantum experts would engage with subject matter experts (SMEs)

The Current Quantum Stack

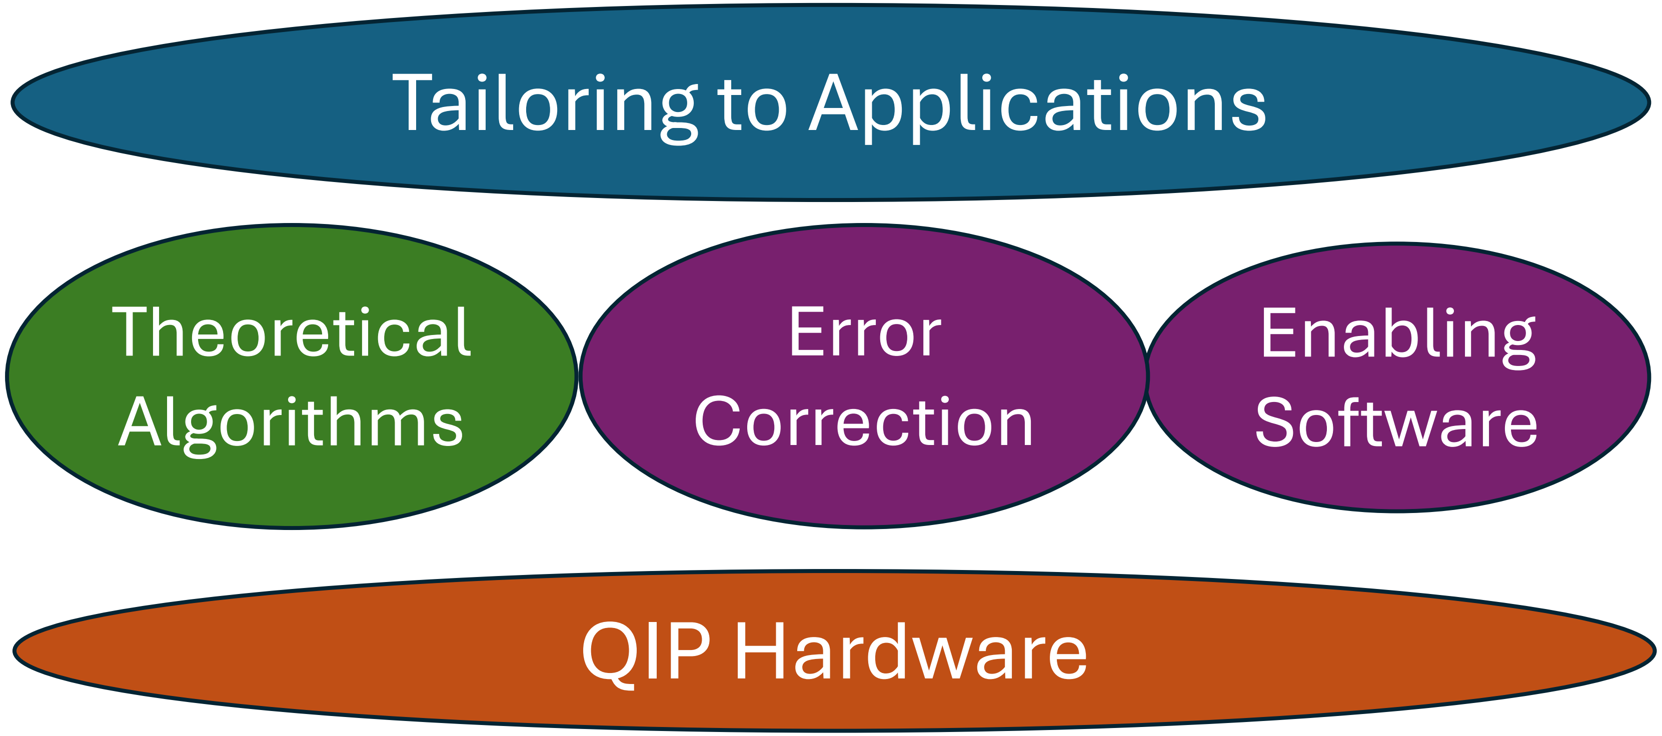

- Industry is focused on QIP Hardware Development
- Academia is focused on Theoretical Algorithms
- Error Correction and Enabling Software is being focused on across areas
- No one is focused on NASA specific applications

Improve NASA’s ability to respond to and benefit from advances in quantum technology in real-time

Quantum Engineers

Make NASA proposals more competitive in a landscape of increasing computational challenges and emerging technology

Engage with elements of the workforce that are excited about quantum to ensure high quality proposals and research in QIP

a) Identify quantum feasibility and timelines for SMD problems
b) Identify solutions (both quantum and classical) to their hard computational problem
c) Develop general quantum information processing code for core quantum capabilities that is deployable by end-users at NASA
d) Work with end-users to integrate quantum code and capabilities

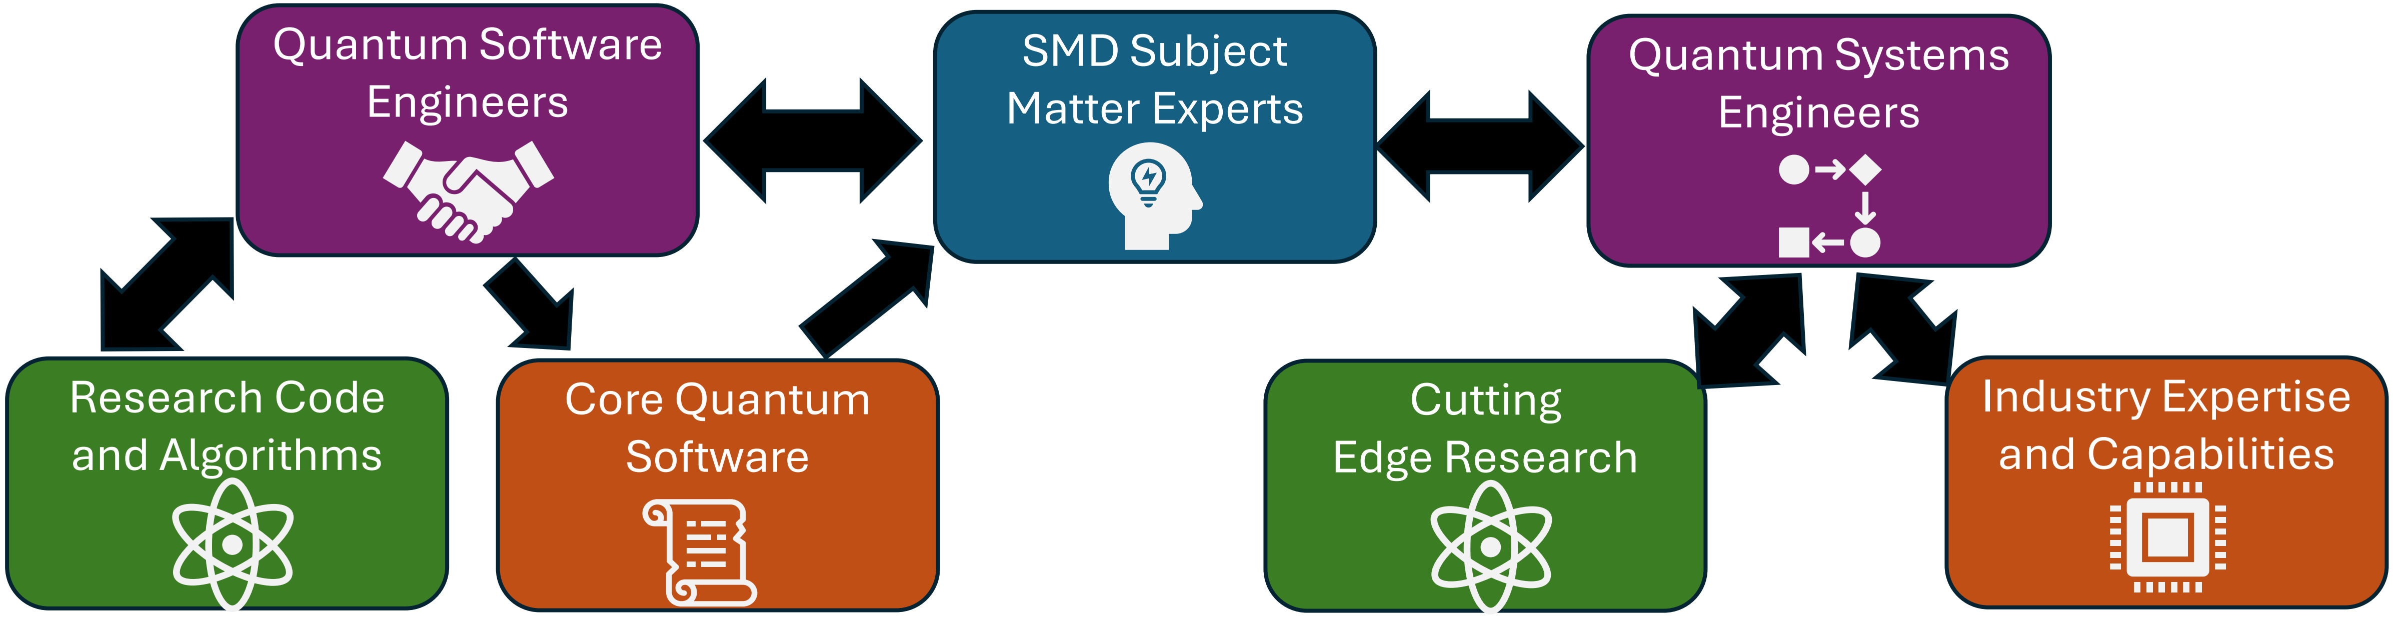

NASA Investment: Benchmarking and Lexicons

Exchange information on what tasks quantum is good (and bad) at and information of computational bottlenecks in applications

- Quantum experts understand the broad capabilities of this technology, and we will create written pipelines to communicate that to SMEs
- Benchmarking needs to occur with SMD needs in mind
  - Continual evaluation and interface with industry and other external partners to understand their real capabilities and how they can be leveraged.
- Lexicons and descriptions of quantum technologies that non-quantum experts can understand and use within the context of SMD/NASA.
- Ideally these lexicons will be targeted to specific science areas such as earth science, heliophysics, biological and physical sciences, planetary science or astrophysics with technologies and highlights specific to their interests

Where Quantum Computing fits in the AI ecosystem

Areas in Classical AI

- Artificial Intelligence is a wide field that includes far more than just machine learning or LLMs
- Current trends, especially in LLMs and foundation models, are focused on large general-purpose models with lots of ingested data and broad use cases
- Foundation models are complicated because of their generality

- AI also includes models for specific tasks, where good training data is hard to obtain and generating it is a core task
- AI includes models other than LMM that require less data through incorporation physical constraints and symmetries
- Lot of NASA application use cases for such models
  - E.g. physics-informed neural networks and operator learning

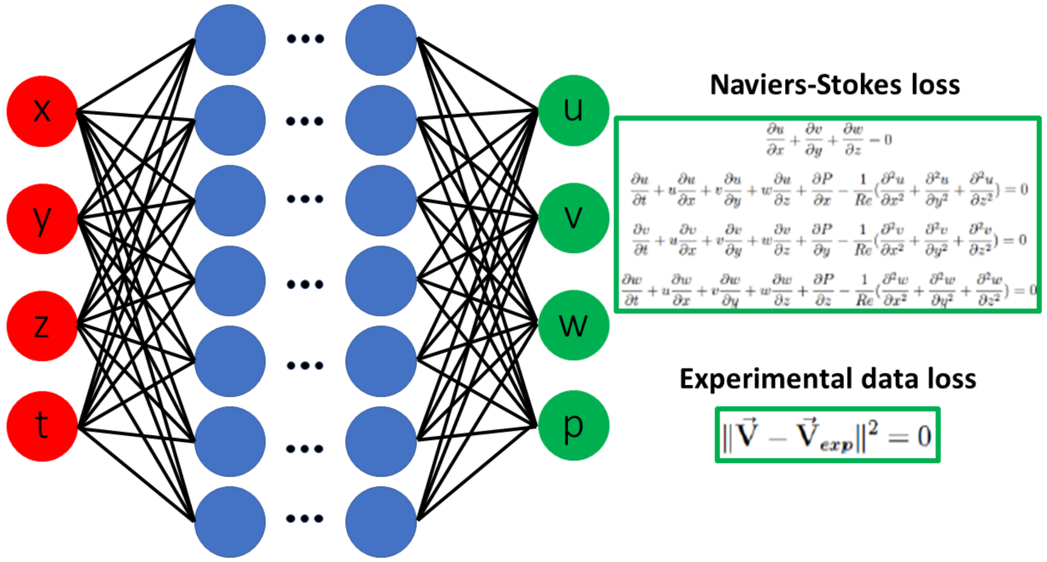

https://upload.wikimedia.org/wikipedia/commons/9/90/Physics-informed_nerural_networks.png

Quantum AI

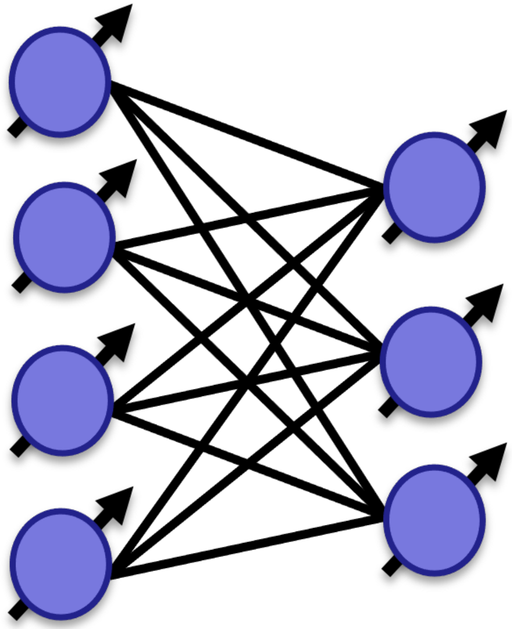

- Quantum AI includes
  - Quantum enhanced optimization
  - Quantum enhanced machine learning
  - Neural networks, and other non-ML forms of AI
- Quantum computers excel at large complicated calculations, rather than many small easy calculations
- Current work focuses on hybrid approaches where a small but difficult portion is handled by the quantum computer

- Quantum models are still relatively young with hurdles
- The training landscape for quantum models is a major difficulty that is an active area of research

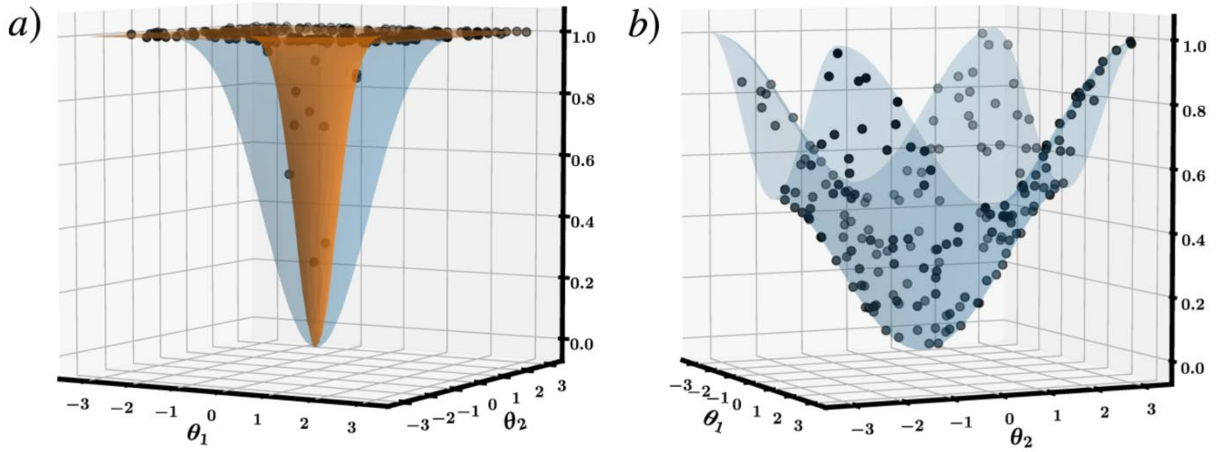

Cerezo, M., Sone, A., Volkoff, T. et al. Cost function dependent barren plateaus in shallow parametrized quantum circuits. Nat Commun 12, 1791 (2021).

Integration Points between Quantum and AI

- Near-term hybrid approaches are going to be needed
- Quantum information processing (QIP) is good in situations where the task itself is hard not in situations where generality is needed
- Therefore, we expect QIP to be useful for special purpose models, similar to how physics-informed or operator learning are used in classical machine learning
- Data ingestion is expected to be a problem for near- and intermediate-term quantum computers
- This means they are unlikely to be helpful with large foundation models and general methods in the near- and intermediate term

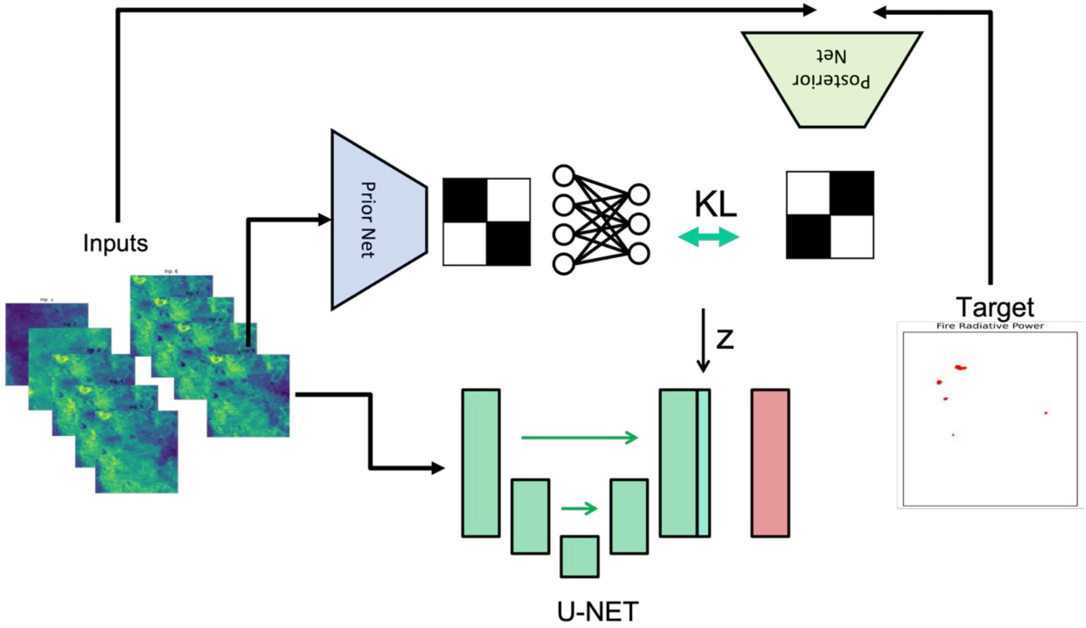

Very Recent External Survey Article on QC & AI

- Acampora et al., Quantum computing and artificial intelligence: status and perspectives, https://www.arxiv.org/abs/2505.23860 (May 29, 2025)
- Duncan et al., Taming quantum systems, https://arxiv.org/abs/2501.16436 (Jan. 27 2025)

The surveys are good but we believe it could have more emphasis on

- Hybrid quantum-classical approaches
- Quantum-ready classical algorithms
  - Advancing classical algorithms, with an eye to where quantum components can be added and quantum technologies mature
- Quantum-inspired classical approaches
  - E.g. Quantum Monte Carlo, dequantized quantum ML algorithms
- Physics-informed AI and ML approaches
  - Including applications to Differential Equations and Computational Fluid Dynamics

Example: QuAIL Innovation in Quantum Algorithms

Quantum Alternating Operator Ansatz (QAOA) (Hadfield, et al.)

- Generalization of Quantum Approximate Optimization Algorithm
- Supports hard constraints
- Introduced more general ”mixing operators”:
  - Bitflip mixers
  - XY mixers
  - Controlled Bitflip mixers
  - Controlled XY mixers

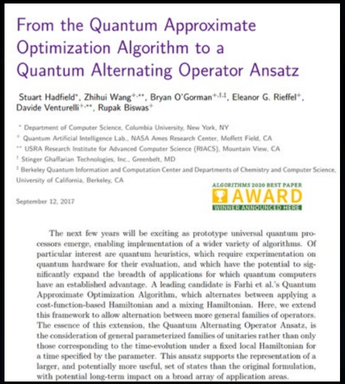

From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz

Stuart Hadfield*, Zhihui Wang+,**, Bryan O'Gorman+,†,‡, Eleanor G. Rieffel+, Davide Venturelli+,**, Rupak Biswas+

* Department of Computer Science, Columbia University, New York, NY
+ Quantum Artificial Intelligence Lab., NASA Ames Research Center, Moffett Field, CA
** USRA Research Institute for Advanced Computer Science (RIACS), Mountain View, CA
† Stinger Ghaffarian Technologies, Inc., Greenbelt, MD
‡ Berkeley Quantum Information and Computation Center and Departments of Chemistry and Computer Science, University of California, Berkeley, CA

September 12, 2017

ALGORITHMS 2020 BEST PAPER AWARD WINNER ANNOUNCED HERE

The next few years will be exciting as prototype universal quantum processors emerge, enabling implementation of a wider variety of algorithms. Of particular interest are quantum heuristics, which require experimentation on quantum hardware for their evaluation, and which have the potential to significantly expand the breadth of applications for which quantum computers have an established advantage. A leading candidate is Farhi et al.'s Quantum Approximate Optimization Algorithm, which alternates between applying a cost-function-based Hamiltonian and a mixing Hamiltonian. Here, we extend this framework to allow alternation between more general families of operators. The essence of this extension, the Quantum Alternating Operator Ansatz, is the consideration of general parameterized families of unitaries rather than only those corresponding to the time-evolution under a fixed local Hamiltonian for a time specified by the parameter. This ansatz supports the representation of a larger, and potentially more useful, set of states than the original formulation, with potential long-term impact on a broad array of application areas.

Iterative Quantum Algorithms (Brady & Hadfield)

- Combines quantum and classical, handles hard constraints

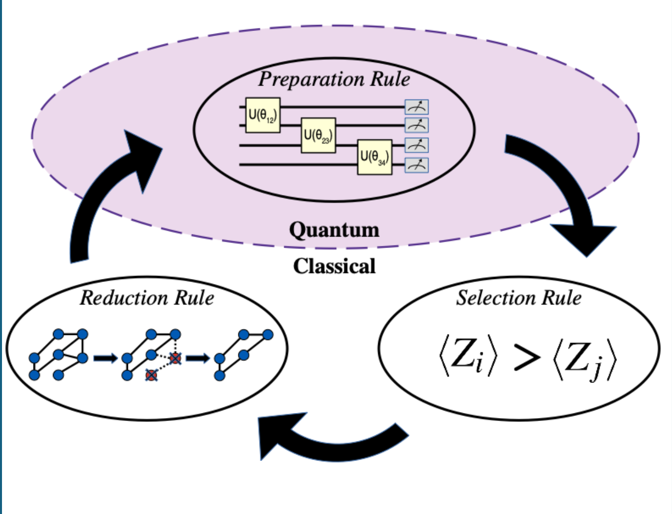

Quantum-accelerated constraint programming (Booth, et al.)

- Used to model and solve constraint satisfaction and optimization
- Shows how quantum algorithms, can accelerate CP, both inference and search

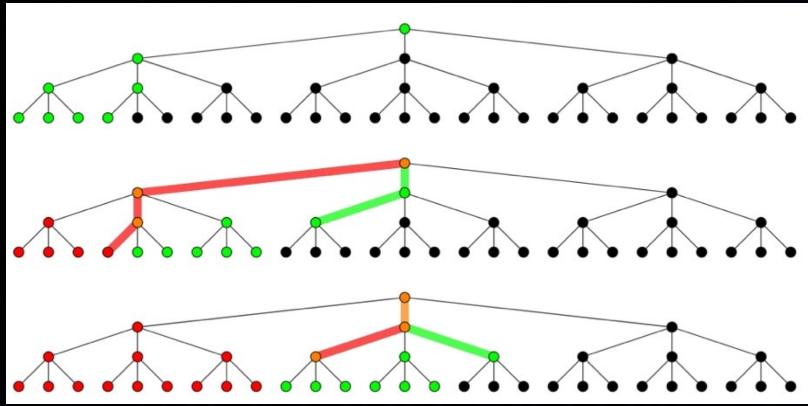

S. Hadfield *et al.* (2019), From the quantum approximate optimization algorithm to a quantum alternating operator Ansatz, *Algorithms* **12**, 34
L. T. Brady & S. Hadfield (2023) Iterative Quantum Algorithms for Maximum Independent Set: A Tale of Low-Depth Quantum Algorithms, arXiv:2309.13110
K.E.C. Booth et al. (2021), Quantum-accelerated constraint programming, Quantum, ***5****, pp. 550*

Very Long Baseline Interferometry (VLBI) telescopy

Goal: Higher resolution

Solution: Longer baseline

Except we are limited by loss and shot noise

Requires bringing info from both ends together, maintaining quantum properties, to enable quantum interference

Quantum Information Processing can help

- Quantum network enables distribution of entangled states
- Quantum teleportation enables robust transfer of quantum information, using entangled state and a classical (non-quantum) communication channel
- Combining N sources, beyond two, provides further advantages
- Additionally, using replacing Fourier transform with quantum Fourier transform enables reduction in number of samples needed

DARPA Quantum Benchmarking Program

- Feasibility of accelerating homogeneous catalyst discovery with fault-tolerant quantum computers
- Feasibility of accelerating incompressible computational fluid dynamics simulations with fault-tolerant quantum computers
- Prospects for NMR Spectral Prediction on Fault–Tolerant Quantum Computers
- Quantifying fault tolerant simulation of strongly correlated systems using the Fermi-Hubbard model
- Applications and resource estimates for open system simulation on a quantum computer
- Potential Applications of Quantum Computing at Los Alamos National Laboratory
- Fault-tolerant resource estimation using graph-state compilation on a modular superconducting architecture Quantum computing for corrosion-resistant materials and anti-corrosive coatings design
- Quantum Resources Required for Binding Affinity Calculations of Amyloid-Beta
- Fullerene-encapsulated Cyclic Ozone for the Next Generation of Nano-sized Propellants via Quantum Computation

These are just a few examples!

We look forward to exploring other computational challenges with you that will have an impact on NASA priorities going forward

Workforce development approaches

Cross-training of NASA staff

- Introductory high-level SMD-centric talks on quantum information processing
- SMD-focused tutorials?
- Full day quantum SMD bootcamp?
- Embedding quantum researchers in other SMD groups and/or embedding SMEs in QuAIL group?
- Quantum hackathaon?

SMD QIP intern program?

- Undergrad summer internships. Experiential learning in SMD-related QIP topics
- Capstone programs: undergrad full semester project
- Graduate summer internships
- Graduate 6+ months internships

Broad Audience

Focused impact/Scope

A Historical Perspective

NASA Ames director Hans Mark brought Illiac IV to NASA Ames in 1972

Illiac IV – first massively parallel computer

- 64 64-bit FPUs and a single CPU
- 50 MFLOP peak, fastest computer at the time

Finding good problems and algorithms was challenging

Questions at the time:

- How broad will the applications be of massively parallel computing?
- Will computers ever be able to compete with wind tunnels?

Back up slides