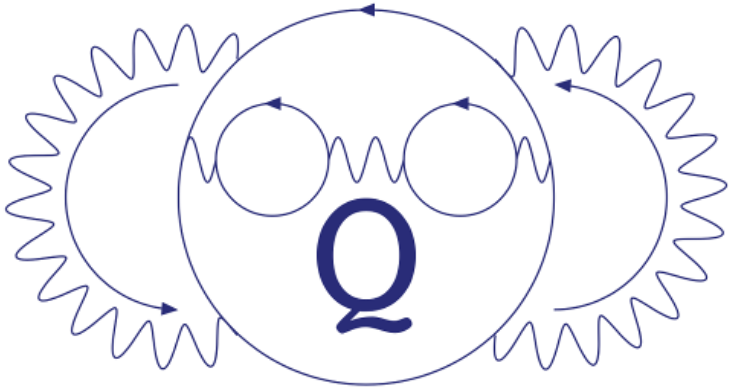

QOALAS

Quantum Optimization and Learning and Simulation
http://qoalas.sandia.gov

Lead PI: Ojas Parekh, Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Team and expertise

Sandia National Laboratories

- Ojas Parekh
- Andrew Baczewski
- Matthew Grace
- Kenneth Rudinger
- Mohan Sarovar

Quantum approaches to discrete optimization, Theoretical computer science, Quantum and classical simulation of many-body quantum systems, QCVV, Quantum dynamics and Control

University of Maryland, College Park

- Andrew Childs
- Stephen Jordan (also Microsoft Research)
- Yi-Kai Liu (also NIST)
- Brian Swingle
- Jacob Taylor (also NIST)
- Xiaodi Wu

Quantum computing, Quantum algorithms, Quantum complexity theory, Machine learning, Quantum many-body physics, Quantum gravity, Quantum machine learning, Optimization

Los Alamos National Laboratory

- Rolando Somma
- Yigit Subasi

Quantum computing, Condensed matter theory

California Institute of Technology

- John Preskill

Quantum computing, High-energy physics, Quantum error correction and fault tolerance

Affiliate:
Virginia Commonwealth University

- Sevag Gharibian (also Universität Paderborn, Germany)

Quantum algorithms and complexity theory

Project summary

Goal: New quantum techniques and algorithms from the interplay of quantum simulation, optimization, and machine learning

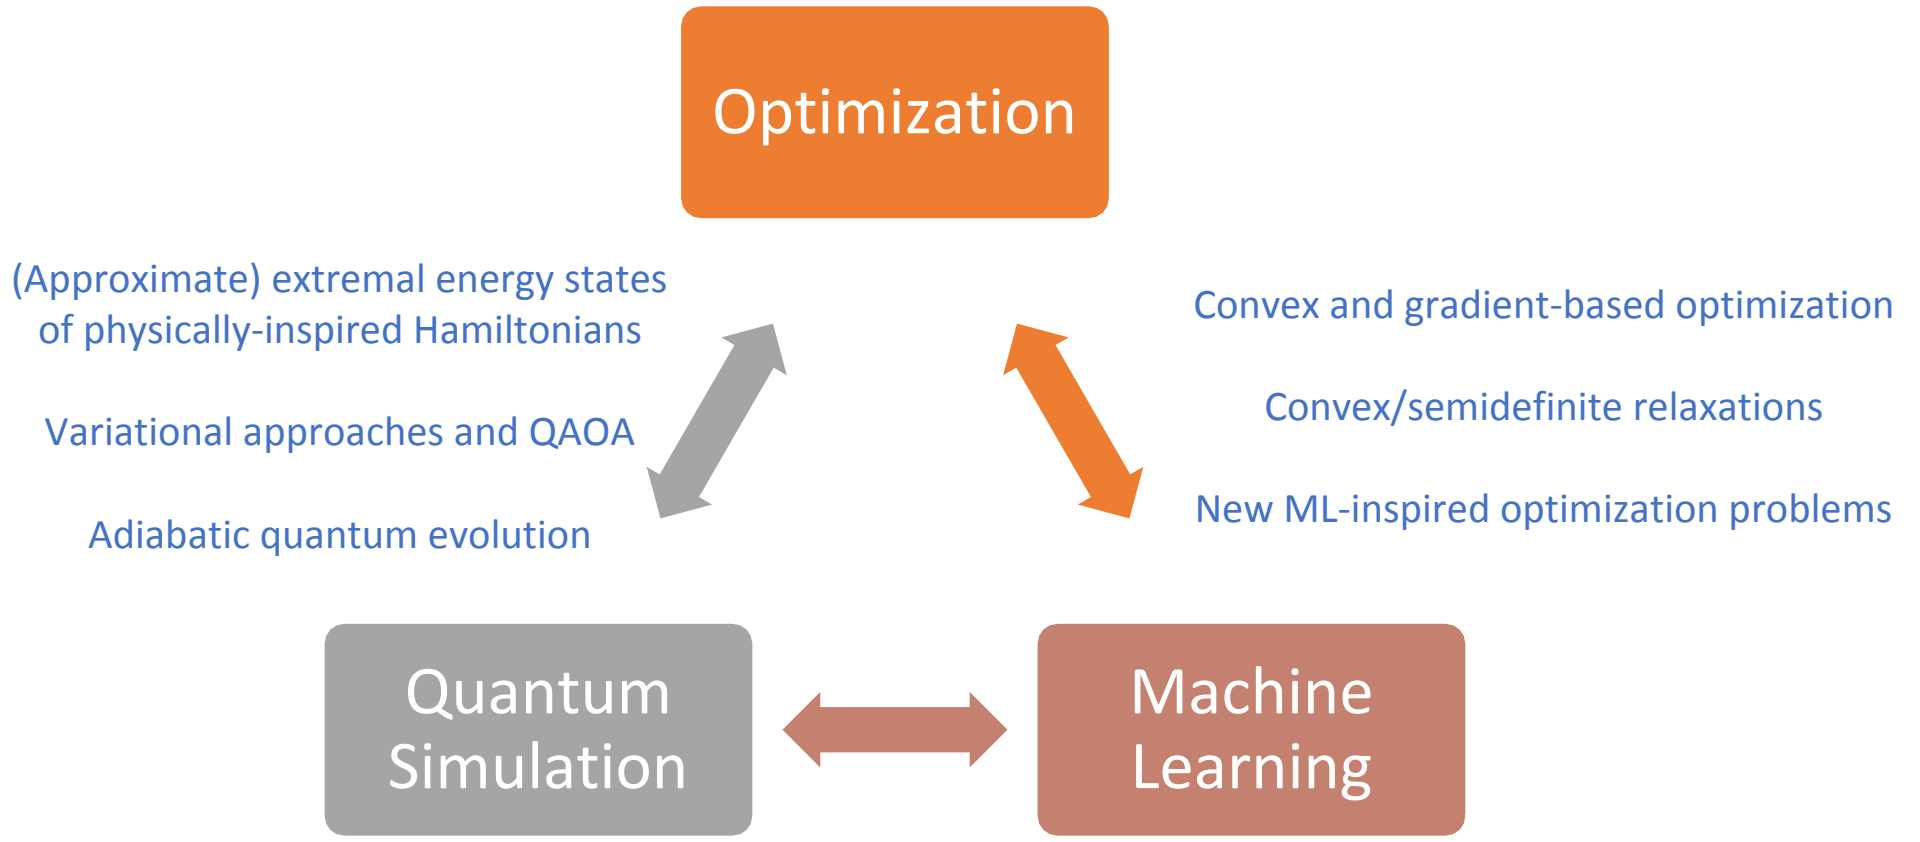

Research Strategy

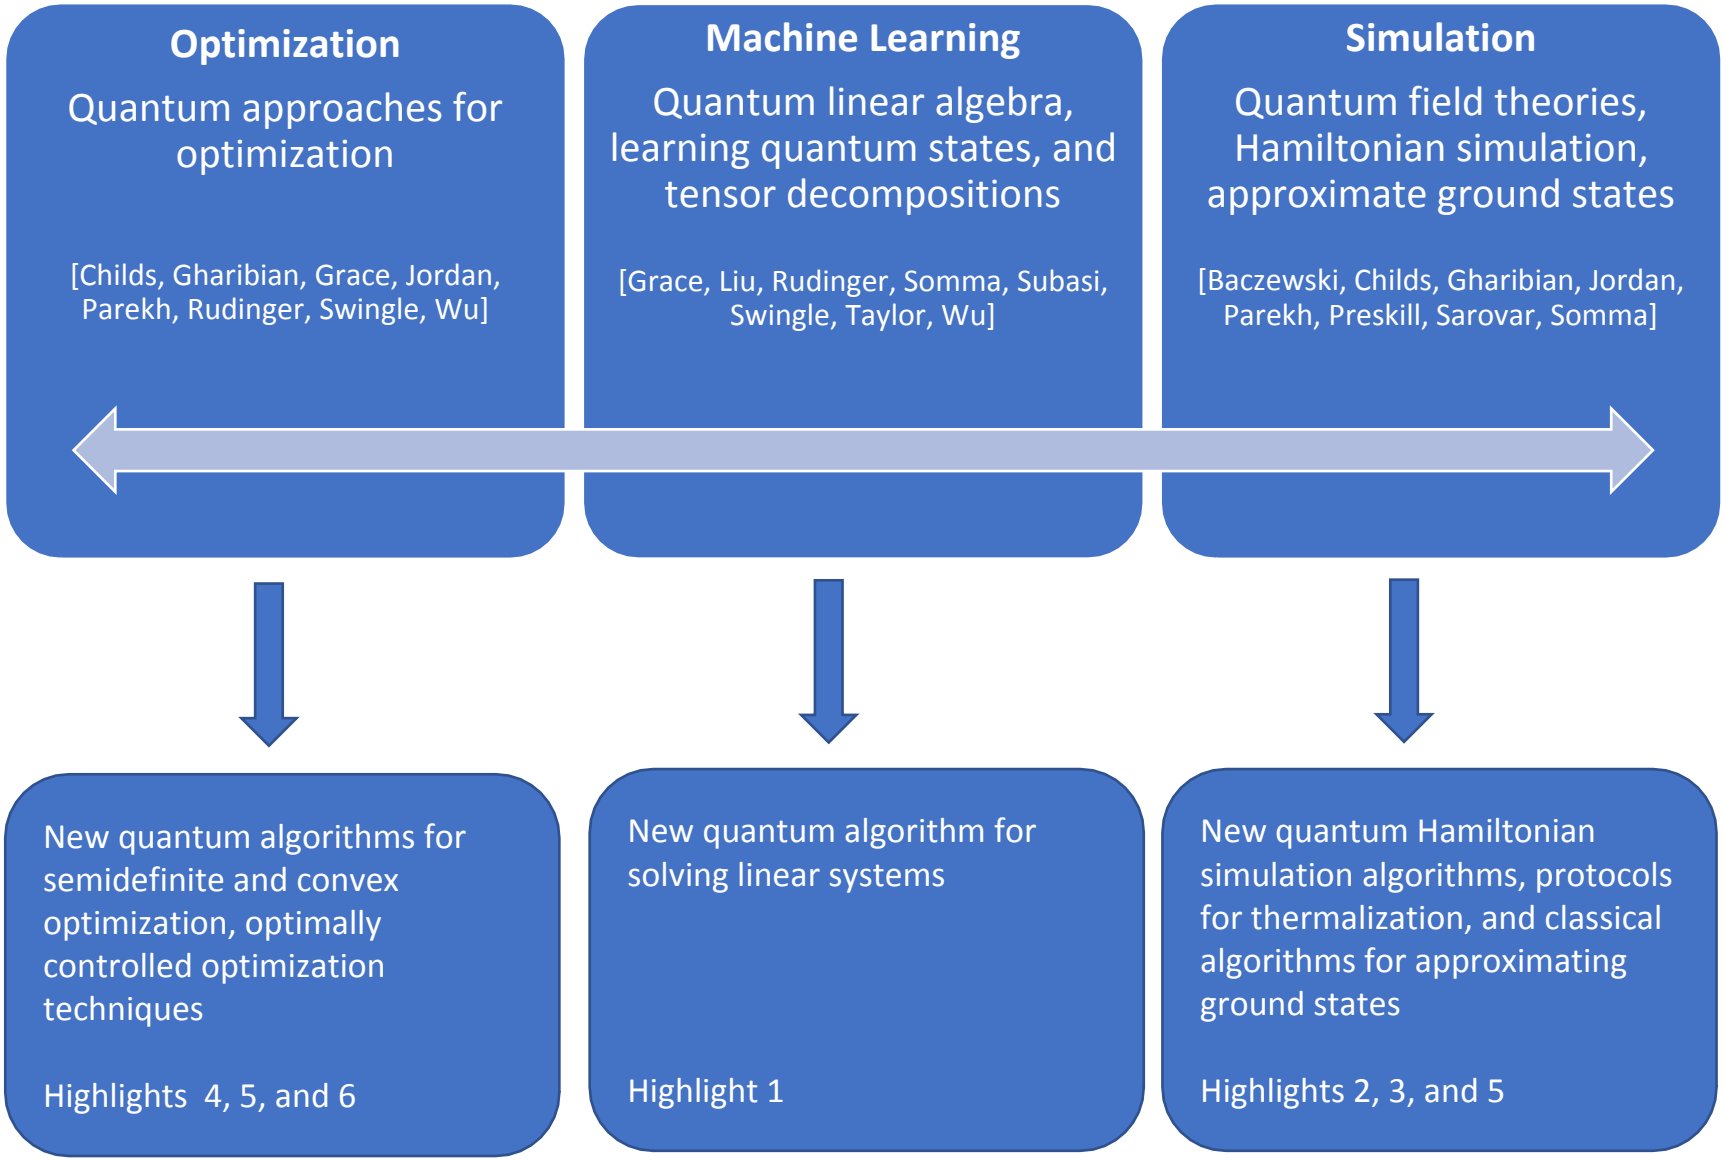

Highlights

Highlight 1 [Yigit Subasi, Rolando Somma, Davide Orsucci]

Quantum-Adiabatic like Algorithms for Solving Linear Systems of Equations

Scientific Achievement

A quantum algorithm to *solve* linear systems of equations much more efficiently than classical computers

Significance and Impact

Evolutions induced by simple Hamiltonians prepare a quantum state that encodes the solution to a linear system of equations. This new idea will find applications in Machine Learning and Optimization, where linear systems play an important role

Research Details

- A Hamiltonian path is built such that the ground states (or other eigenstates) encode the solution to an increasingly difficult linear system of equations
- A variant of adiabatic quantum computing allows one to prepare the ground states or eigenstates
- The complexity of our method is given by the total evolution time and is shown to be optimal

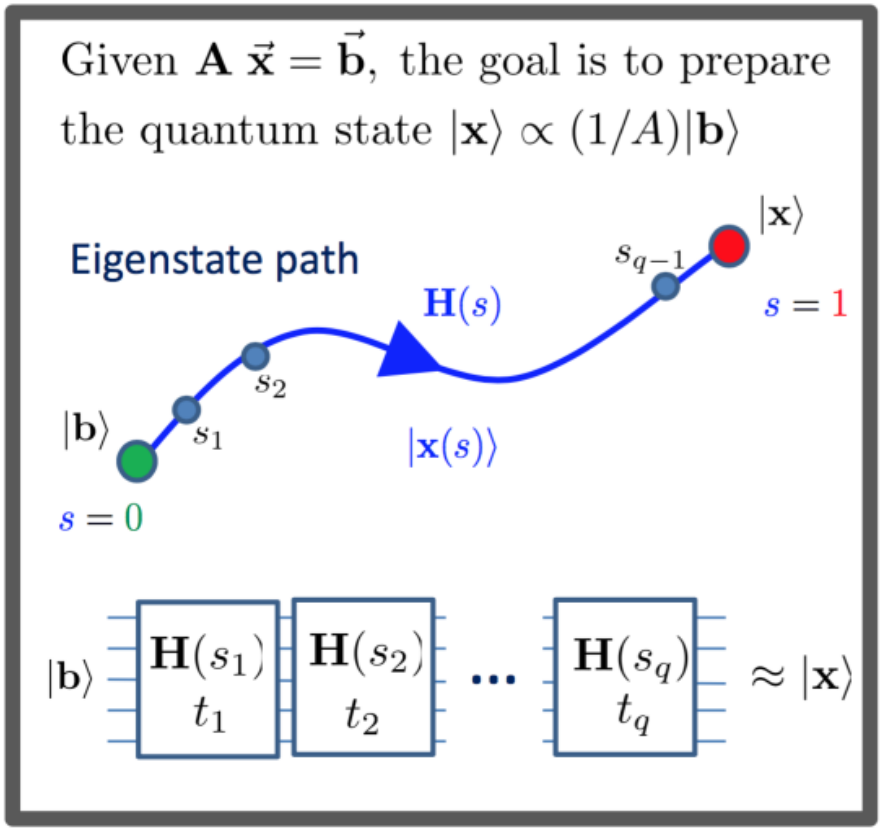

Schematic of the quantum algorithm for linear systems: To prepare a quantum state that is proportional to the solution of a linear system of equations, evolutions under certain Hamiltonians are required. These evolutions can be implemented efficiently on a quantum computer in many cases, achieving an exponential speedup.

Y. Subasi, R.D. Somma, and D. Orsucci, *arXiv:1805.10549* (2018).

Highlight 1 [Yigit Subasi, Rolando Somma, Davide Orsucci]

Quantum-Adiabatic like Algorithms for Solving Linear Systems of Equations

Hamiltonians and linear systems

We observe the property

$$\mathbf{A}\vec{\mathbf{x}} = \vec{\mathbf{b}} \Rightarrow \begin{cases} \mathbf{P}_{\mathbf{b}}^{\perp}.\mathbf{A}\ \vec{\mathbf{x}} = \mathbf{P}_{\mathbf{b}}^{\perp}\ \vec{\mathbf{b}} = 0 \\ \mathbf{B} \Rightarrow \left(\mathbf{B}^{\dagger}\mathbf{B}\right)\ \vec{\mathbf{x}} = 0 \end{cases}$$

The family of Hamiltonians is

$$\mathbf{H}(s) = \mathbf{B}^{\dagger}(s)\mathbf{B}(s)\ , \mathbf{B}(0) = \mathbf{P}_{\mathbf{b}}^{\perp}\ , \mathbf{B}(1) = \mathbf{B}$$

The randomization method (RM) S. Boixo, E. Knill, and R. Somma, Q. Inf. Comp. 9, 0833 (2009).

A variant of adiabatic quantum computing in which the parameters of the Hamiltonian are changed discretely and the evolution is for a random time. The time complexity is

$$T_{\mathbf{RM}} \propto \frac{L^2}{\epsilon.\Delta}$$

- L is the path-length of the state encoding the solution
- Δ is the smallest spectral gap of the Hamiltonians
- ϵ is the accuracy of the state preparation

Important variables for this problem

The path length satisfies $L = O(\log \kappa)$

The spectral gap satisfies $\Delta = O(1/\kappa^2)$

$$T_{\mathbf{RM}} \propto \frac{\kappa^2 \log^2(\kappa)}{\epsilon}$$

Same time complexity as HHL

κ is the condition number of A.

Polynomial dependence in condition number and no dependence in dimension imply quantum speedup

Y. Subasi, R.D. Somma, and D. Orsucci, *arXiv:1805.10549* (2018).

Highlight 1 [Yigit Subasi, Rolando Somma, Davide Orsucci]

Quantum-Adiabatic like Algorithms for Solving Linear Systems of Equations

Spectral gap amplification

R. Somma & S. Boixo, SIAM J. Comp. **42**, 593 (2013)

The gap-amplified family of Hamiltonians is:

$$\mathbf{H}(s) = \mathbf{B}(s) \otimes \sigma^{+} + \mathbf{B}^{\dagger}(s) \otimes \sigma^{-}$$

Important variables for this problem

The path length satisfies $L = O(\log \kappa)$

The spectral gap satisfies $\Delta = O(1/\kappa)$

$$T_{\mathbf{RM}} \propto \frac{\kappa \log^2(\kappa)}{\epsilon}$$

Linear dependence in condition number results in an optimal quantum algorithm for linear systems

Y. Subasi, R.D. Somma, and D. Orsucci, *arXiv:1805.10549* (2018).

Experimental verification of new algorithm:

Experimental realization of quantum algorithms for linear system inspired by adiabatic quantum computing

Jingwei Wen[1],* Xiangyu Kong[1],* Shijie Wei[4], Bixue Wang[1], Keren Li[1], Yuanye Zhu[1], Tao Xin[5],† and Guilu Long [1,2,3‡]

arXiv:1806.03295

NMR, 4 qubits

Highlight 2 [Mekena Metcalf, Jonathan Moussa, Mohan Sarovar]

Engineered thermalization of many-body quantum systems

Scientific Achievement

Designed and demonstrated protocol for generating thermal states of many-body quantum systems using coupled, driven ancilla spins.

Significance and Impact

Quantum simulation of thermal states is useful for extracting finite temperature properties of many-body quantum systems, and we have developed a practical protocol for doing this.

Research Details

Thermalizing a many-body quantum system requires coupling it to an engineered reservoir. The conditions required for thermalization are known if one has access to a macroscopic reservoir, but this is not typical.

We have developed a protocol that results in thermalization using driven, dissipated ancilla spins that are coupled to the system in a way that results in engineered energy exchange.

Numerical and analytical results demonstrate that the steady-state of the engineered dynamics is the desired thermal state of the many-body system.

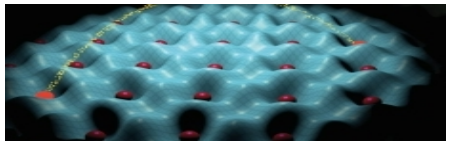

Example target system: cold atom lattice, Esslinger Lab, ETH Zurich

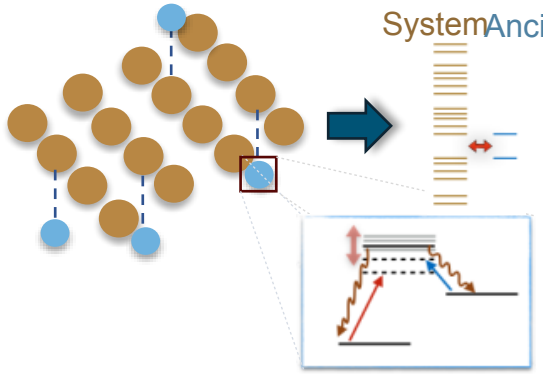

Schematic of protocol: Ancilla systems that are optically pumped to local thermal states are coupled to the system to thermalize. The resonant energy of the ancilla systems are swept across the spectrum of the primary system, and over a few sweeps this dynamics generates a thermal distribution in the principal system.

Highlight 2 [Mekena Metcalf, Jonathan Moussa, Mohan Sarovar]

Engineered thermalization of many-body quantum systems

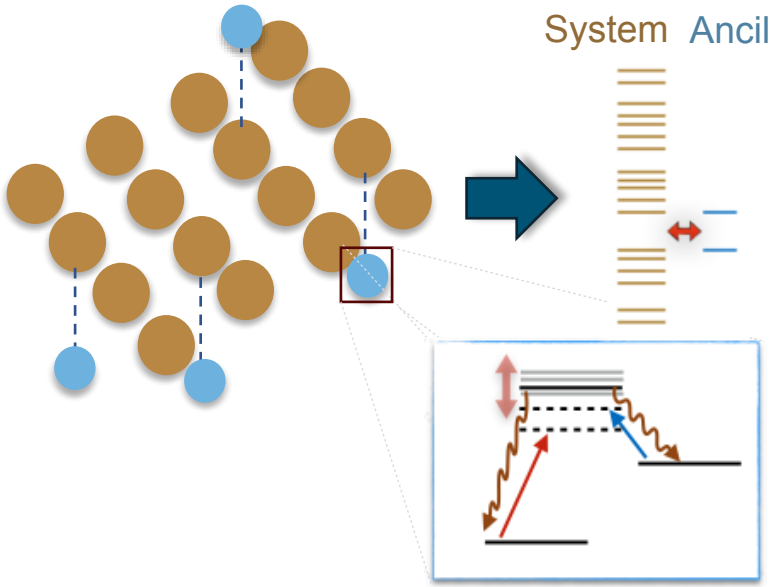

A time-dependent term is added to the ancilla energy to *sweep the system's full energy spectrum.*

$$\Omega_m(t) = \Delta_0 f(t)$$

$$H_I(t) = H_{sys} - \sum_{m=1}^{M} \frac{\Omega_m(t)}{2} \tau_z^m + \sum_{m=1}^{M} g_m \left(\sigma_x^{k_m} \tau_x^m\right)$$

Ancilla optically pumped to thermal state

Timescale Hierarchy

$$\left|\frac{df(t)}{dt}\right| << g_m \sim \Gamma^m << \|H_{sys}\|, \forall m, t$$

$$\Gamma^m = \gamma_+^m + \gamma_-^m$$

A reduced master equation describes system evolution

when the ancilla dynamics is averaged over

$$\dot{\rho} = g^2 \sum_{\omega} \gamma(t,\omega) \left(X(\omega) \rho X^\dagger(\omega) - \frac{1}{2} \{X^\dagger(\omega) X(\omega), \rho\} \right)$$

Operator on System

$$X(\omega) = \sum_{\epsilon' - \epsilon = \omega} |\epsilon\rangle \langle\epsilon| \sigma_x^k |\epsilon'\rangle \langle\epsilon'$$

Ancilla Correlation Functions

$$\gamma(t,\omega) = \frac{P_0(t)\Gamma}{\left(\frac{\Gamma}{2}\right)^2 - (\omega - \Omega(t))^2} + \frac{P_1(t)\Gamma}{\left(\frac{\Gamma}{2}\right)^2 - (\omega + \Omega(t))^2}$$

Highlight 2 [Mekena Metcalf, Jonathan Moussa, Mohan Sarovar]

Engineered thermalization of many-body quantum systems

Two Site System

$$H_{sys} = -\frac{\omega_1}{2}\sigma_z^1 - \frac{\omega_2}{2}\sigma_z^2 + J_x\sigma_x^1\sigma_x^2 + J_y\sigma_y^1\sigma_y^2 + J_z\sigma_z^1\sigma_z^2$$

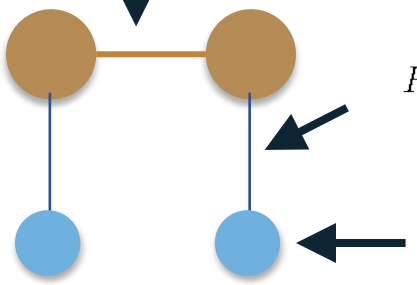

$$H_{s-a} = \sum_{i=1}^{2} g\left(\sigma_x^i\tau_x^i\right)$$

$$\Omega_i(t) = \Delta_{max}\sin^2(\omega t)$$

$$\Delta_{max} = E_{sys}^{max} - E_{sys}^{min}$$

How close is the time averaged state to thermal state?

$$\Delta = \|\rho_{th} - \bar{\rho}\|$$

$$\rho_{th}(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})}$$

$$\mathcal{T}\left[e^{\int_0^T dt\mathcal{L}(t)}\right]\rho_{th} = \bar{\rho}$$

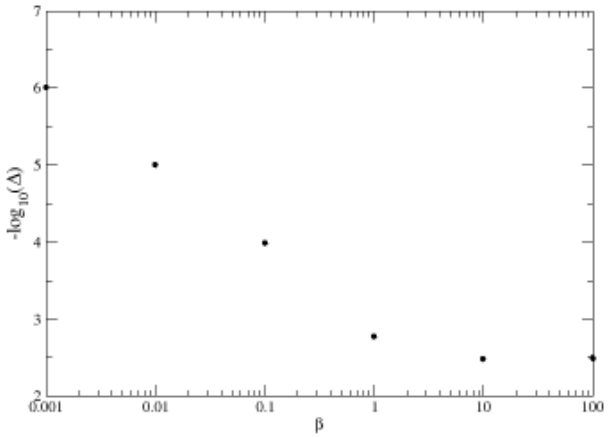

Markovian Evolution of Random Initial State

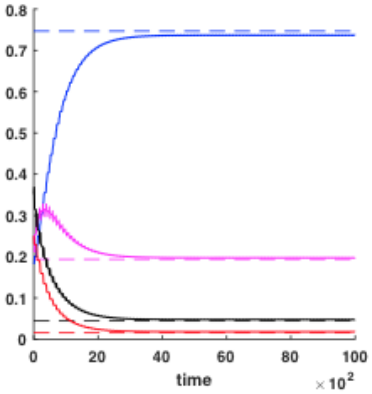

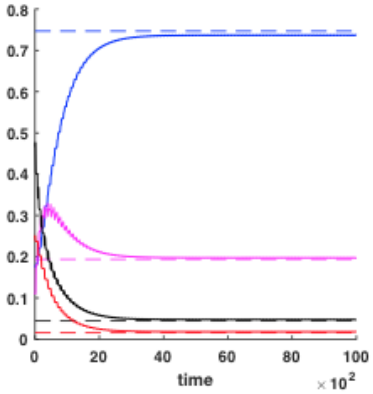

Ongoing Research:

- Analytical calculations to approximate the time-averaged, effective generator of evolution
- Prove fixed point of the dynamics and provide a bound on thermalization time.
- Demonstrate this scheme is generalizable to generic syste Hamiltonians.

Highlight 3 [Andrew Childs, Aaron Ostrander, Yuan Su]

Faster quantum simulation by randomization

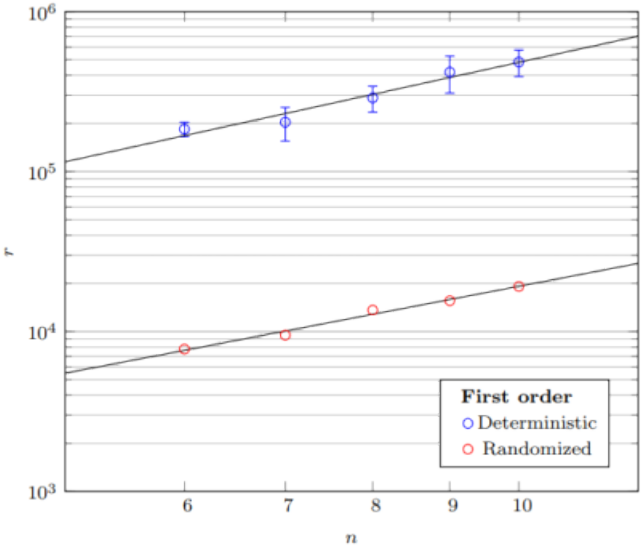

Comparison of the number of trotter steps between deterministic and randomized first-order product formulas. Error bars are omitted when they are negligibly small on the plot. Straight lines show powerlaw fits to the data.

Scientific Achievement

Stronger product-formula bounds were proved by randomization, and their advantage for near-term quantum simulation was evidenced by numerics.

Significance and Impact

Product formula algorithm is a straightforward yet surprisingly efficient approach to simulating quantum dynamics on a quantum computer. Recent results show that this approach can outperform more sophisticated algorithms, and it is important to understand why it has such outstanding performance.

Research Details

– Proved stronger error bounds for product formulas by randomizing how the terms in the Hamiltonian are ordered.

– Showed that randomized bounds can be asymptotically better than previous bounds that exploit commutation between the terms in the Hamiltonian.

– Numerically compared the deterministic and randomized product formulas, showing improvement even with respect to the empirical performance.

Andrew Childs, et al., arXiv:1805.08385 (2018)

Highlight 3 [Andrew Childs, Aaron Ostrander, Yuan Su]

Faster quantum simulation by randomization

Quantum (Hamiltonian) simulation

Given a description of a Hamiltonian *H* and an evolution time *t*, perform $U(t) = e^{-itH}$ up to some error ε.

Product formula algorithm

- Target system $H = \sum_{\ell=1}^{L} \alpha_\ell H_\ell$
  - $0 \leq \alpha_\ell \leq ;1$
  - H_ℓ is a tensor product of Paulis (up to a sign).
- Can use the first-order product formula

$$\left\| e^{-it\sum_{j=1}^{L}\alpha_j H_j} - \left[\prod_{j=1}^{L} e^{-i\frac{t}{r}\alpha_j H_j}\right]^r \right\| \leq O\left(\frac{(Lt)^2}{r}\right)$$

- Generalizations to *(2k)*th-order are known [Suzuki 92].
- Advantage: straightforward; can empirically outperform more sophisticated simulation algorithms.
- Problem: error bounds are loose in practice.

Andrew Childs, et al., arXiv:1805.08385 (2018)

Highlight 3 [Andrew Childs, Aaron Ostrander, Yuan Su]

Faster quantum simulation by randomization

New bounds by randomization

- We prove stronger bounds by randomizing how terms in the Hamiltonian are ordered, *e.g.,*

$$e^{\lambda H_1} e^{\lambda H_2} \to \frac{1}{2}\left(e^{\lambda H_1} e^{\lambda H_2} + e^{\lambda H_2} e^{\lambda H_1}\right)$$

- New error bound for first-order formula:

$$O\left(\frac{(Lt)^3}{r^2}\right)$$

improving over the old bound

$$O\left(\frac{(Lt)^2}{r}\right)$$

- We also derive new bounds for higher-order formulas.
- Numerical simulations show improvement even with respect to the empirical performance.

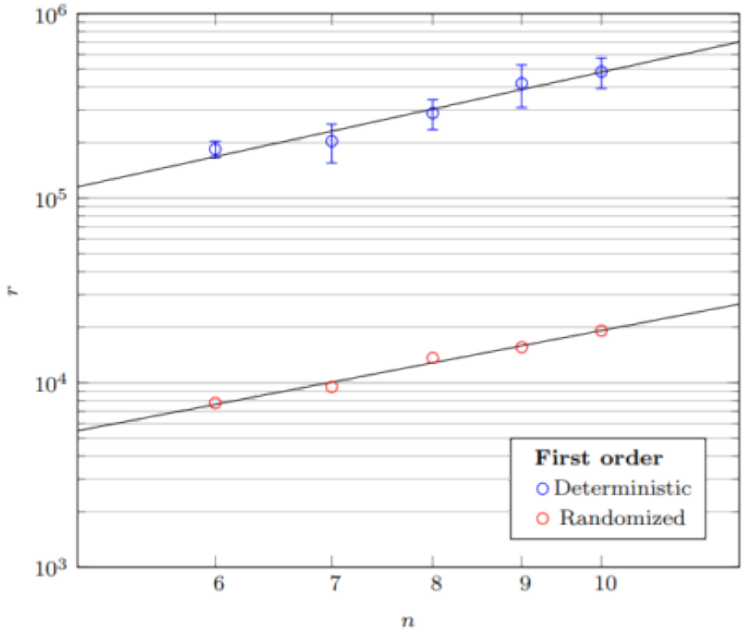

Comparison of the number of trotter steps between deterministic and randomized *first-order* product formulas.

Andrew Childs, et al., arXiv:1805.08385 (2018)

Highlight 4 [Shouvanik Chakrabarti, Andrew Childs, Tongyang Li, Xiaodi Wu]

Quantum algorithms and lower bounds for convex optimization

We give a quantum algorithm that can optimize a convex function over an n-dimensional convex body using $\tilde{O}(n)$ queries to oracles that evaluate the objective function and determine membership in the convex body. The best known classical algorithm [1] uses $\tilde{O}(n^2)$ queries.

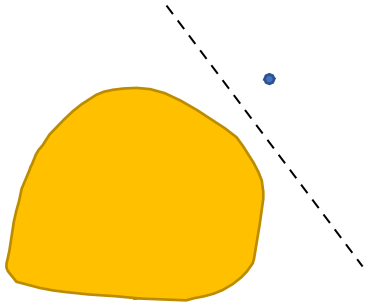

Problem: minimize $f(x)$ subject to $x \in K$, where f and K are convex.

Classical Oracle:
MEM(x): Decide whether $x \in K$
EVAL(x): Compute $f(x)$

Quantum Oracle:
MEM: $|x\rangle|0\rangle \rightarrow |x\rangle|\delta[x \in K]\rangle$
EVAL: $|x\rangle|0\rangle \rightarrow |x\rangle|f(x)\rangle$

- We use Jordan's gradient estimation algorithm[2] to construct a separation oracle for points outside the convex body, using only $\tilde{O}(1)$ queries to the oracles.
- We then use the ellipsoid method to solve the optimization problem by applying the separation oracle $\tilde{O}(n)$ times. The total query and time complexities are $\tilde{O}(n)$ and $\tilde{O}(n^3)$ [1], respectively.
- For both oracles, we also prove an $\widetilde{\Omega}(\sqrt{n})$ quantum lower bound.

Chakrabarti et al., arXiv:1809.01731

[1] Y.T. Lee, A. Sidford, and S. Vempala, Efficient convex optimization with membership oracles, Proceedings of the 31st Conference on Learning Theory (COLT 2018), 1292-1294.

[2] S.P. Jordan, Fast quantum algorithm for numerical gradient estimation, Phys. Rev. Lett. 95.5, 050501 (2005).

Highlight 5 [Sevag Gharibian, Ojas Parekh, Ciaran Ryan-Anderson]

Approximate Solutions for Quantum Heisenberg Models via Hyperplane Rounding

Discrete optimization techniques enable new rigorous approximations of low-energy states of quantum Heisenberg Hamiltonians, a central topic in condensed matter physics.

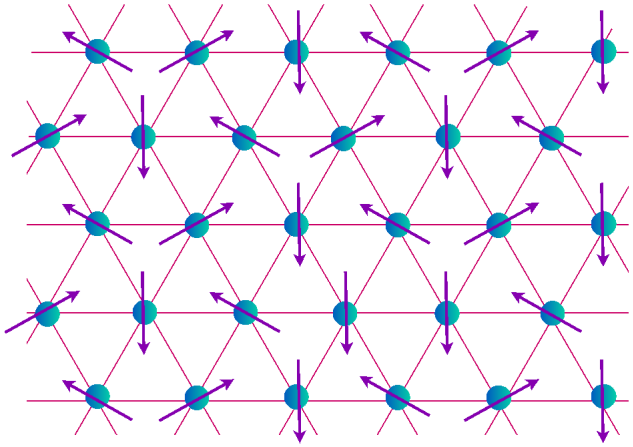

Anti-ferromagnetic Heisenberg model: roughly neighboring quantum particles aim to align in opposite directions. This kind of Hamiltonian appears, for example, as an effective Hamiltonian for so-called Mott insulators.
(Image: Sachdev, http://arxiv.org/abs/1203.4565)

Significance and Impact

The Heisenberg model is fundamental for describing quantum magnetism, superconductivity, and charge density waves. Beyond 1 dimension, the properties of the anti-ferromagnetic Heisenberg model are notoriously difficult to analyze. Exploiting analytical tools from discrete optimization, a team led by Sandia National Labs has developed new algorithms to rigorously approximate hard-to-compute properties of this model beyond 1-D.

Research Details

- The researchers introduce a new quantum Hamiltonian model that simultaneously generalizes the quantum Heisenberg anti-ferromagnet and hard classical graph partitioning problems.
- A new classical algorithm produces approximate solutions for the above model that are mathematically guaranteed to be relatively close in quality to optimal quantum solutions.

Highlight 5 [Sevag Gharibian, Ojas Parekh, Ciaran Ryan-Anderson]

Approximate Solutions for Quantum Heisenberg Models via Discrete Optimization

New maximum-energy version of the antiferromagnetic Heisenberg model, generalizing the fundamental Maximum Cut discrete optimization problem:

Max Cut Hamiltonian:

$$\sum(I - Z_iZ_j)$$

⇨

Quantum Heisenberg generalization:

$$\sum(I - X_iX_j - Y_iY_j - Z_iZ_j)$$

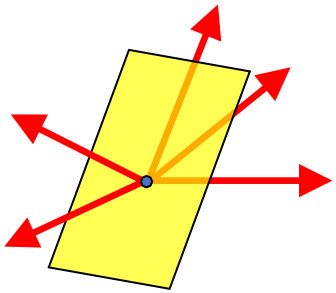

- The hard-to-solve Heisenberg maximum energy optimization problem is relaxed to an easier-to-solve semidefinite program
- A generalization of the celebrated Goemans-Williamson hyperplane rounding method produces an approximate max-energy product state

We obtain the first nontrivial rigorous approximations for these problems:
0.498-approx via a product state, where 1/2 is best possible for product states
(also 0.649-approx for XY model, where 2/3 is best possible for product states)

Our results extend to approximating max-energy of any "symmetric" 2-local Hamiltonian:

$$I - \sum_{\{k=1\}}^{3}(\alpha_{k,i}X_i + \beta_{k,i}Y_i + \gamma_{k,i}Z_i)(\alpha_{k,j}X_j + \beta_{k,j}Y_j + \gamma_{k,j}Z_j)$$

Highlight 6 [Aniruddha Bapat, Stephen Jordan]

Optimally Controlled Quantum Optimization

General Framework

$$\frac{\mathrm{d}}{\mathrm{d}t}|\psi\rangle = -iH(u)|\psi\rangle$$

$$H(u) = (1-u)H_0 + uH_1$$

H_0 : "Easy" Hamiltonian (e.g. TFIM)
Ground state = Initial state

H_1 : Cost Hamiltonian
Ground state = Global cost minimum

Adiabatic: QAO, Quantum annealing, ...

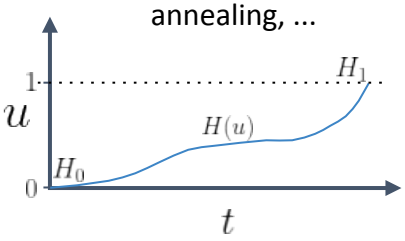

Bang-bang: QAOA, VQE, Spin pulse echo, ...

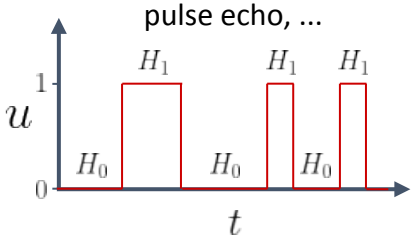

Pontryagin's Minimum Principle
⇓
Optimal controls are bang-bang [1]

Example: Ramp with Spike

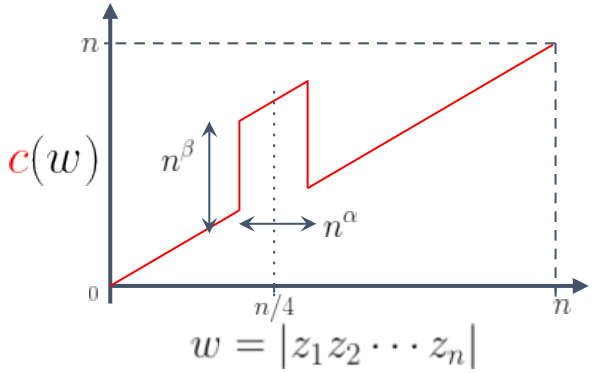

QAO [2]

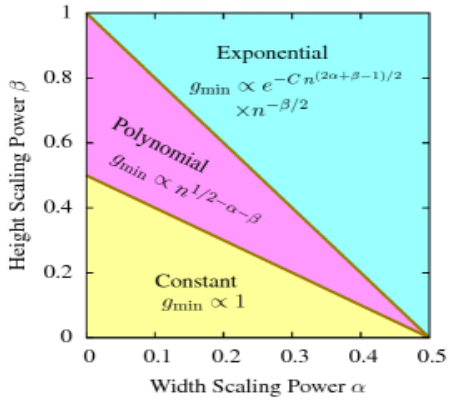

QAOA-1

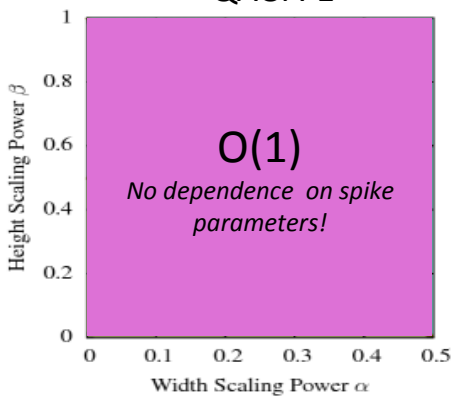

Key intuition: QAOA-1 samples local gradient on hypercube, and imparts a momentum kick to wavepacket, allowing barrier jumps.

?

Quantum-inspired Classical Heuristic

[1] Z.C. Yang *et al.*, Optimizing Variational Quantum Algorithms using Pontryagin's Minimum Principle, Phys. Rev. X 7, 021027 (2017).

[2] L.T. Brady, W. van Dam, Spectral Gap Analysis for Efficient Tunneling in Quantum Adiabatic Optimization, *arXiv:1601.01720v2* (2016).

Highlight 7 [Nai-Hui Chia, Tongyang Li, Han-Hsuan Lin, Chunhao Wang]

Quantum-inspired classical sublinear algorithm for solving SDPs with low-rank constraints

Scientific Achievement

A quantum-inspired classical algorithm that solves low-rank semidefinite programs (SDPs) as efficient as best-known quantum algorithms.

Significance and Impact

Since SDP is ubiquitous in optimization, our work provides general understanding about how quantum algorithms can motivate classical optimization algorithms.

Semidefinite programming

$$\begin{aligned} \max \quad & \mathrm{Tr}[CX] \\ \text{s.t.} \quad & \mathrm{Tr}[A_i X] \le a_i \quad \forall i \in [m] \\ & X \ge 0 \end{aligned}$$

Quantum-inspired classical input

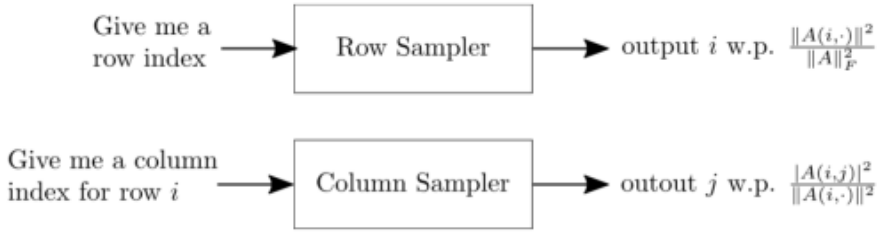

Research Details

We give a quantum-inspired classical algorithm that given an SDP with m constraint matrices, each of dimension n and rank r, solves the SDP in time $O\big(m \cdot \mathrm{poly}(\log n, r)\big)$.

Technically, our approach aligns with the recent studies of quantum-inspired machine learning algorithms and proposes new ingredients for matrix arithmetics.

Chia et al., arXiv:1901.03254

Highlight 7 [Nai-Hui Chia, Tongyang Li, Han-Hsuan Lin, Chunhao Wang]

Quantum-inspired classical sublinear algorithm for solving SDPs with low-rank constraints

$$\begin{aligned} \max \quad & \mathrm{Tr}[CX] \\ \text{s.t.} \quad & \mathrm{Tr}[A_i X] \le a_i \quad \forall i \in [m] \\ & X \ge 0 \end{aligned}$$

Input: Sample-based data structure

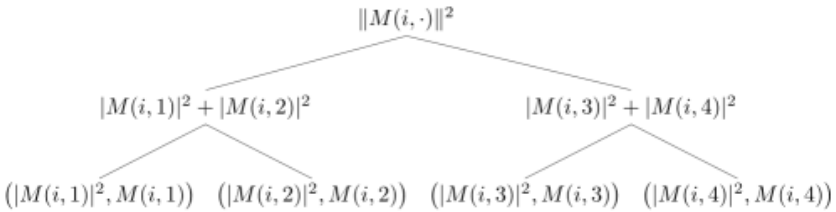

Algorithm:

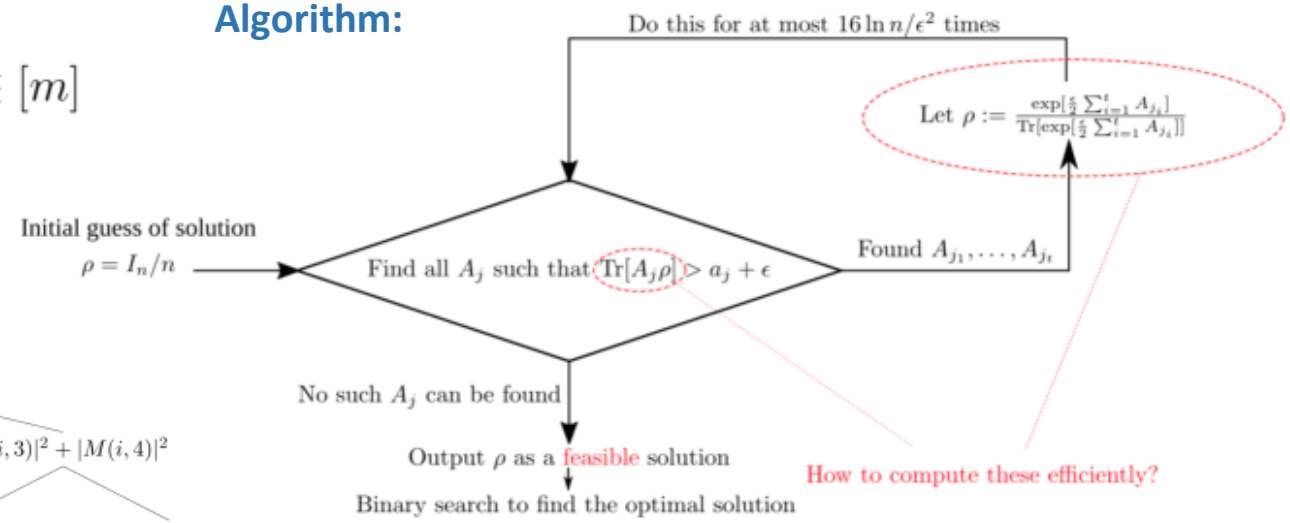

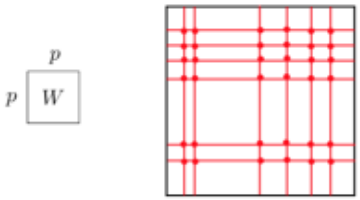

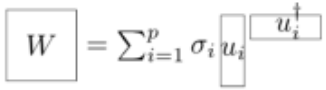

- Sample a small matrix, solve SVD; correctness promised by [FVK04][1]
- Two technical contributions for solving SDP:
  - Weighted sampling: sample from $A_1 + \cdots + A_m$;
  - Matrix Exponentiation: $\exp[A_1 + \cdots + A_m]$ by inner products of both the (truncated) U and V of the singular value decomposition (SVD)

[1] A. Frieze, R. Kannan, and S. Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations, Journal of the ACM 51 (2004), no. 6, 1025-1041.

[2] E. Tang. A quantum-inspired classical algorithm for recommendation systems, Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC 2019), 217-228, 2019, ACM.

Highlight 8 [Tongyang Li, Shouvanik Chakrabarti, Xiaodi Wu]

Sublinear quantum algorithms for training linear and kernel-based classifiers

Scientific Achievement: A quantum algorithm for training classifiers over n d-dimensional data running in time $\tilde{O}(\sqrt{n}+\sqrt{d})$, a quadratic speed-up in both n and d compared to the state-of-the-art. A tight lower bound is also established.

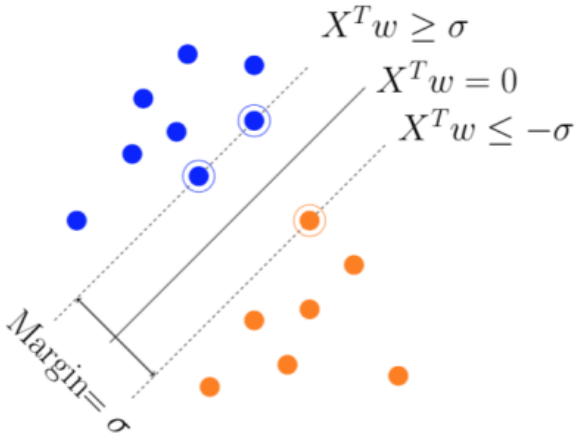

Problem: Given n data $X_1, \ldots, X_n \in R^d$ and a label vector $y \in R^d$, classification finds a separating hyperplane w such that $y_i X_i^T w \geq 0$ for any $i \in [n]$.

Input: Coherent oracle access:

$$O|i\rangle|j\rangle|0\rangle \rightarrow |i\rangle|j\rangle|X_{ij}\rangle$$

Generality of results: The classifier can be kerneralized: $X \rightarrow \Psi(X)$, where Ψ can be Gaussian or polynomial kernel.

Research Details:

- We follow the classical state-of-the-art[1], a primal-dual algorithm; we give quadratic quantum speed-ups for both the primal (multiplicative weight) and the dual (online gradient descent)

- Our algorithms use the standard quantization of the classical input and generate the same classical output, suggesting minimal overheads for end-to-end applications

[1] K.L. Clarkson, E. Hazan, and D.P. Woodruff, Sublinear optimization for machine learning, Journal of the ACM 59 (2012), no. 5, 23.

Li et al., arXiv:1904.02276, ICML 2019