Quantum Linear Algebra with Near-Optimal Complexities

Lin Lin

Department of Mathematics, UC Berkeley
Lawrence Berkeley National Laboratory

Joint work with Dong An and Yu Tong (Berkeley)

Quantum Seminar,
Simons Institute, May 2020

arXiv: 1909.05500; 1910.14596; 2002.12508

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

A ritual

- There is perhaps a widespread belief that a quantum talk should start with a picture of Feynman

Figure: A superposition of Feynmans

Quantum linear algebra

- Solving linear systems, eigenvalue problems, matrix exponentials, least square problems, singular value decompositions etc on a quantum computer.

- Many interesting, exciting progresses in the past few years.

- Reasonable way towards “quantum advantage”. “Quantum machine learning”.

- Solving linear equations (MATH 54 at Berkeley, first class)

$$Ax = b$$

- Quantum linear system problem (QLSP)

$$A|x\rangle = |b\rangle$$

Voila!

Quantum linear system problem (QLSP)

- All vectors must be normalized. $A \in \mathbb{C}^{N\times N}, |b\rangle \in \mathbb{C}^N, N = 2^n$. $\||b\rangle\|_2^2 := \langle b|b\rangle = 1$. WLOG $\|A\|_2 = 1$.
- Solution vector
$$|x\rangle \propto A^{-1} |b\rangle .$$
- How to put the information in $A, |b\rangle$ into a quantum computer? read-in problem. Oracular assumption.
- Query complexity: the number of oracles used. Gate complexity. Rely on implementation of query models.

Quantum speedup for QLSP

- κ: condition number of A. ϵ: target accuracy. Proper assumptions on A (e.g. d-sparse) so that oracles cost $\mathrm{poly}(n)$.
- (Harrow-Hassadim-Lloyd, 2009): $\widetilde{\mathcal{O}}(\kappa^2/\epsilon)$.
- Exponential speedup with respect to n? Answer could depend on read-in / read-out models (Tang, 2019)
- (Childs-Kothari-Somma, 2017): Linear combination of unitary (LCU). $\widetilde{\mathcal{O}}(\kappa^2 \operatorname{poly}\log(1/\epsilon)))$
- (Low-Chuang, 2017) (Gilyén-Su-Low-Wiebe, 2019): Quantum signal processing (QSP). $\widetilde{\mathcal{O}}(\kappa^2 \log(1/\epsilon)))$

Comparison with classical iterative solvers

- Positive definite matrix. Error in A-norm
- Steepest descent: $\mathcal{O}(N\kappa \log(1/\epsilon))$; Conjugate gradient: $\mathcal{O}(N\sqrt{\kappa} \log(1/\epsilon))$
- Quantum algorithms can scale better in N but worse in κ.
- Lower bound: Quantum solver cannot generally achieve $\mathcal{O}(\kappa^{1-\delta})$ complexity for any $\delta > 0$ (Harrow-Hassadim-Lloyd, 2009)
- Goal of near-optimal quantum linear solver: $\widetilde{\mathcal{O}}(\kappa \operatorname{poly}\log(1/\epsilon))$ complexity.

LCU for QLSP: Basic idea

- $A \in \mathbb{C}^{N\times N}$, Hermitian. $\|A\|_2 = 1$. Condition number κ.
- $\mathrm{spec}(A) \subset D_\kappa = [-1, -\kappa^{-1}] \cup [\kappa^{-1}, 1]$.
- A^{-1} is non-unitary. Matrix function expansion

$$A^{-1} \approx \sum_{k=0}^{M-1} c_k e^{-\mathrm{i}At_k}$$

- Hamiltonian simulation problem. Linear combination of unitaries (LCU). Efficient: M terms with $\log M$ ancilla qubits.
  (Berry-Childs-Cleve-Kothari-Somma, 2014)
  (Childs-Kothari-Somma, 2017)

LCU for QLSP: cost

- Cost of $e^{-\mathrm{i}At}|\psi\rangle$ (for longest t)

$$\mathcal{O}(t\log(t/\epsilon)) \sim \widetilde{\mathcal{O}}(\kappa\,\mathrm{poly}\log(1/\epsilon))$$

- Overall cost (suitable implementation of the select oracle)

$$\underbrace{\widetilde{\mathcal{O}}(\kappa\,\mathrm{poly}\log(1/\epsilon))}_{\text{Cost of each simulation}} \times \underbrace{\widetilde{\mathcal{O}}(\kappa^2\,\mathrm{poly}\log(1/\epsilon))}_{\substack{\text{\# Repetition} \\ \text{(due to success prob.)}}} = \widetilde{\mathcal{O}}(\kappa^3\,\mathrm{poly}\log(1/\epsilon)))$$

- Using amplitude amplification, can be improved to

$$\underbrace{\widetilde{\mathcal{O}}(\kappa\,\mathrm{poly}\log(1/\epsilon))}_{\text{Cost of each simulation}} \times \underbrace{\widetilde{\mathcal{O}}(\kappa\,\mathrm{poly}\log(1/\epsilon))}_{\substack{\text{\# Repetition} \\ \text{(due to success prob.)}}} = \widetilde{\mathcal{O}}(\kappa^2\,\mathrm{poly}\log(1/\epsilon)))$$

Compare the complexities of QLSP solvers

Algorithm	Query complexity	Remark
HHL (Harrow et al 2009)	$\widetilde{\mathcal{O}}(\kappa^2/\epsilon)$	w. VTAA, complexity becomes $\widetilde{\mathcal{O}}(\kappa/\epsilon^3)$ (Ambainis 2010)
Linear combination of unitaries (LCU) (Childs et al 2017)	$\widetilde{\mathcal{O}}(\kappa^2 \operatorname{poly}\log(1/\epsilon))$	w. VTAA, complexity becomes $\widetilde{\mathcal{O}}(\kappa \operatorname{poly}\log(1/\epsilon))$
Quantum signal processing (QSP) (Gilyén et al 2019)	$\widetilde{\mathcal{O}}(\kappa^2 \log(1/\epsilon))$	Queries the RHS only $\widetilde{\mathcal{O}}(\kappa)$ times
Randomization method (RM) (Subaşi et al 2019)	$\widetilde{\mathcal{O}}(\kappa/\epsilon)$	Prepares a mixed state; w. repeated phase estimation, complexity becomes $\widetilde{\mathcal{O}}(\kappa \operatorname{poly}\log(1/\epsilon))$
Time-optimal adiabatic quantum computing (AQC(exp)) (An-Lin, 2019)	$\widetilde{\mathcal{O}}(\kappa \operatorname{poly}\log(1/\epsilon))$	No need for any amplitude amplification. Use time-dependent Hamiltonian simulation.
Eigenstate filtering (Lin-Tong, 2019)	$\widetilde{\mathcal{O}}(\kappa \log(1/\epsilon))$	No need for any amplitude amplification. Does not rely on any complex subroutines.

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

Reformulating QLSP into an eigenvalue problem

- Weave together linear system, eigenvalue problem, differential equation (Subasi-Somma-Orsucci, 2019)
- $Q_b = I_N - |b\rangle \langle b|$. If $A|x\rangle = |b\rangle \quad \Rightarrow \quad Q_b A|x\rangle = Q_b|b\rangle = 0$
- Then

$$H_1 = \begin{pmatrix} 0 & AQ_b \\ Q_b A & 0 \end{pmatrix}, \quad |\widetilde{x}\rangle = |0\rangle|x\rangle = \begin{pmatrix} x \\ 0 \end{pmatrix}$$

$$\mathrm{Null}(H_1) = \mathrm{span}\{|\widetilde{x}\rangle, |\bar{b}\rangle\}, \quad |\bar{b}\rangle = |1\rangle|b\rangle = \begin{pmatrix} 0 \\ b \end{pmatrix}$$

- QLSP $\quad \Rightarrow \quad$ Find an eigenvector of H_1 with eigenvalue 0.

Adiabatic computation

- Known eigenstate $H_0 |\psi_0\rangle = \lambda_0 |\psi_0\rangle$ for some H_0.
- Interested in some eigenstate $H_1 |\psi_1\rangle = \lambda_1 |\psi_1\rangle$
- $H(s) = (1-s)H_0 + sH_1$,

$$\frac{1}{T}\mathrm{i}\partial_s |\psi_T(s)\rangle = H(s) |\psi_T(s)\rangle , \quad |\psi_T(0)\rangle = |\psi_0\rangle$$

- $|\psi_T(1)\rangle \approx \psi(1)$ (up to a phase factor), T sufficiently large?
- Gate-based implementation: time-dependent Trotter, for near-optimal complexity (Low-Wiebe, 2019)

Adiabatic computation

- (Born-Fock, 1928)

  A physical system remains in its instantaneous eigenstate if a given perturbation is acting on it slowly enough and if there is a gap between the eigenvalue and the rest of the Hamiltonian's spectrum.

- Albash, Avron, Babcock, Cirac, Cerf, Elgart, Hagedorn, Jansen, Lidar, Nenciu, Roland, Ruskai, Seiler, Wiebe...

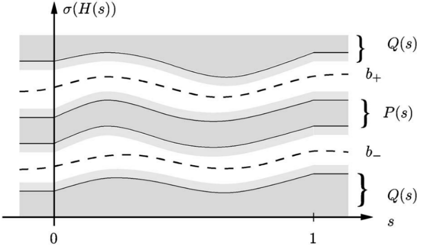

Adiabatic quantum computation (AQC) for QLSP

- Introduce

$$H_0 = \begin{pmatrix} 0 & Q_b \\ Q_b & 0 \end{pmatrix}, \quad \mathrm{Null}(H_0) = \mathrm{span}\{|\widetilde{b}\rangle, |\bar{b}\rangle\}$$

$$|\widetilde{b}\rangle = |0\rangle |b\rangle = \begin{pmatrix} b \\ 0 \end{pmatrix}, \quad |\bar{b}\rangle = |1\rangle |b\rangle = \begin{pmatrix} 0 \\ b \end{pmatrix}$$

- Adiabatically connecting $|\widetilde{b}\rangle$ (zero eigenvector of H_0) to $|\widetilde{x}\rangle$ (zero eigenvector of H_1) (Subasi-Somma-Orsucci, 2019)

- Only one eigenvector in the null space is of interest: transition to $|\bar{b}\rangle$ is prohibited during dynamics

Eigenvalue gap and fidelity

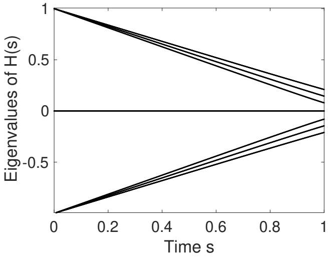

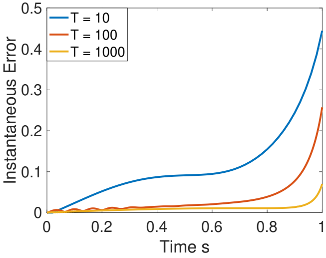

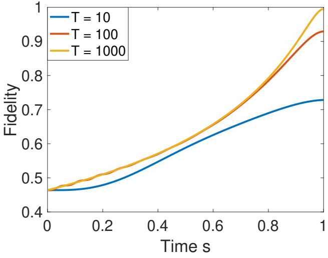

Fidelity:

$$F(|\varphi\rangle, |\psi\rangle) = |\langle\varphi|\psi\rangle|^2 = \mathrm{Tr}[P_\varphi P_\psi].$$

Adiabatic quantum computation

Theorem (Jansen-Ruskai-Seiler, 2007)

Hamiltonian $H(s)$, $P(s)$ projector to eigenspace of $H(s)$ separated by a gap $\Delta(s)$ from the rest of the spectrum of $H(s)$

$$|1 - \langle \psi_T(s)|P(s)|\psi_T(s)\rangle\,| \leq \eta^2(s), \quad 0 \leq s \leq 1$$

where

$$\eta(s) = \frac{C}{T}\Big\{ \frac{\|H^{(1)}(0)\|_2}{\Delta^2(0)} + \frac{\|H^{(1)}(s)\|_2}{\Delta^2(s)} + \int_0^s \left(\frac{\|H^{(2)}(s')\|_2}{\Delta^2(s')} + \frac{\|H^{(1)}(s')\|_2^2}{\Delta^3(s')} \right) ds' \Big\}.$$

T: time complexity; $1/T$ convergence.
$\Delta(s) \geq \Delta_*$, $T \sim \mathcal{O}((\Delta_*)^{-3}/\epsilon)$ (worst case)

Implication in QLSP

- Lower bound of gap (Assume $A \succ 0$ for now, can be relaxed)

$$\Delta(s) \geq \Delta_*(s) = 1 - s + s/\kappa \geq \kappa^{-1}$$

- Worst-case time complexity $T \sim \mathcal{O}(\kappa^3/\epsilon)$
- AQC inspired algorithm: randomization method (Subasi-Somma-Orsucci, 2019),

$$T \sim \mathcal{O}(\kappa \log(\kappa)/\epsilon)$$

ϵ : 2-norm error of the density matrix.

- Rescheduled dynamics.

Accelerate AQC for QLSP: Scheduling

- Goal: improve the scaling AQC w.r.t. κ.
- Adiabatic evolution with $H(f(s)) = (1 - f(s))H_0 + f(s)H_1$

$$\frac{1}{T}\mathrm{i}\partial_s \left|\psi_T(s)\right\rangle = H(f(s))\left|\psi_T(s)\right\rangle, \quad \left|\psi_T(0)\right\rangle = |\widetilde{b}\rangle$$

- $f(s)$: scheduling function. $0 \le f(s) \le 1, f(0) = 0, f(1) = 1$.
- allow $H(f(s))$ to slow down when the gap is close to 0, to cancel with the vanishing gap.
- (Roland-Cerf, 2002) for time-optimal AQC of Grover search.

Choice of scheduling function: AQC(p)

- Schedule (Jansen-Ruskai-Seiler, 2007; Albash-Lidar, 2018)

$$\dot{f}(s) = c_p \Delta_*^p(f(s)), \quad f(0) = 0, \quad 1 \le p \le 2.$$

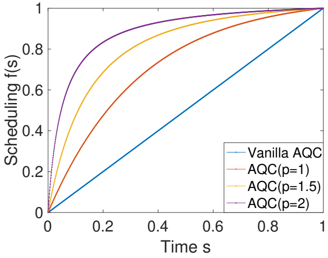

AQC for QLSP

Theorem (An-L., 1909.05500)

$A \succ 0$, condition number κ. For any $1 < p < 2$, the error of the AQC(p) scheme is

$$\| P_T(1) - |\widetilde{x}\rangle \langle \widetilde{x}| \|_2 \le C\kappa/T.$$

Therefore in order to prepare an $\epsilon-$approximation of the solution of QLSP it suffices to choose the runtime $T = \mathcal{O}(\kappa/\epsilon)$.
Furthermore, when $p = 1, 2$, the bound for the runtime becomes $T = \mathcal{O}(\kappa \log(\kappa)/\epsilon)$.

Similar results for Hermitian indefinite and non-Hermitian matrices.

Improve the dependence on ϵ

- AQC(exp): modified schedule (slow at beginning and end)

$$f(s) = c_e^{-1} \int_0^s \exp\left(-\frac{1}{s'(1-s')}\right) \mathrm{d}s'$$

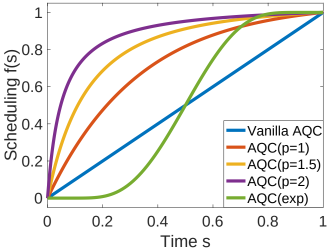

- Intuition: error bound of (Jansen-Ruskai-Seiler, 2007) and integration by parts (Wiebe-Babcock, 2012)

- Rigorous proof of exponential convergence: follow the idea of (Nenciu, 1993), asymptotic expansion of $P(s)$

Improve the dependence on ϵ

Theorem (An-L., 1909.05500)

$A \succ 0$, condition number κ. Then for large enough $T > 0$, the error of the AQC(exp) scheme is

$$\| P_T(1) - |\widetilde{x}\rangle \langle \widetilde{x}| \|_2 \leq C \log(\kappa) \exp \left(-C \left(\frac{\kappa \log^2 \kappa}{T} \right)^{-\frac{1}{4}} \right).$$

Therefore the runtime $T = \mathcal{O}\left(\kappa \log^2(\kappa) \log^4 \left(\frac{\log \kappa}{\epsilon}\right)\right)$.

Near-optimal complexity (up to $\operatorname{poly}\log$ factors).
Similar results for Hermitian indefinite and non-Hermitian matrices.

Implications on QAOA

- Quantum approximate op timization algorithm (QAOA) (Farhi-Goldstone-Gutmann, 2014)

$$|\psi_\theta\rangle := e^{-i\gamma_P H_1} e^{-i\beta_P H_0} \cdots e^{-i\gamma_1 H_1} e^{-i\beta_1 H_0} |\psi_i\rangle$$

- Trotterize AQC $\Rightarrow$: one implementation of QAOA
- Hybrid quantum-classical optimization.
- The optimal protocol of QAOA yields near-optimal complexity
- QAOA is expected to follow a non-adiabatic shortcut (Brady et al, 2020)

Numerical results: positive definite matrix

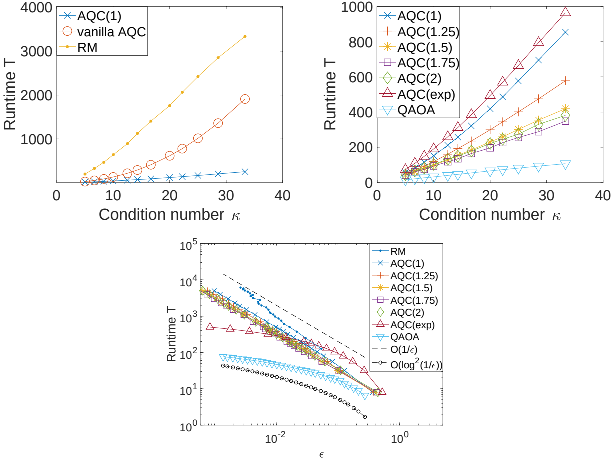

Figure: Top: the runtime to reach desired fidelity (left: 0.99, right: 0.999) as a function of the condition number. Bottom: a log-log plot of the runtime as a function of the accuracy with $\kappa = 10$.

Numerical results: positive definite matrix

methods	scaling w.r.t. κ	scaling w.r.t. $1/\epsilon$
vanilla AQC	2.2022	/
RM	1.4912	/
AQC(1)	1.4619	1.1205
AQC(1.25)	1.3289	1.0530
AQC(1.5)	1.2262	1.0010
AQC(1.75)	1.1197	0.9724
AQC(2)	1.1319	0.9821
AQC(exp)	1.3718	0.5377
AQC(exp)	/	1.7326 (w.r.t. $\log(1/\epsilon)$)
QAOA	1.0635	0.6555
QAOA	/	1.5889 (w.r.t. $\log(1/\epsilon)$)

Table: Numerical scaling of the runtime as a function of the condition number and the accuracy, respectively, for the Hermitian positive definite example.

Numerical results: non-Hermitian matrix

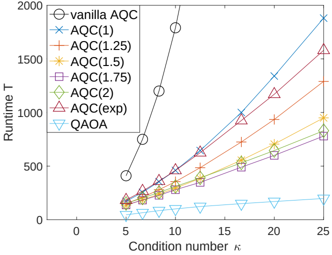

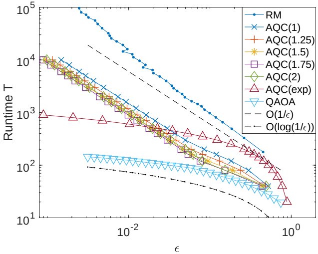

Figure: Left: the runtime to reach 0.999 fidelity as a function of the condition number. Right: a log-log plot of the runtime as a function of the accuracy with $\kappa = 10$.

Numerical results: non-Hermitian matrix

methods	scaling w.r.t. κ	scaling w.r.t. $1/\epsilon$
vanilla AQC	2.1980	/
RM	/	/
AQC(1)	1.4937	0.9611
AQC(1.25)	1.3485	0.9249
AQC(1.5)	1.2135	0.8971
AQC(1.75)	1.0790	0.8849
AQC(2)	1.0541	0.8966
AQC(exp)	1.3438	0.4415
AQC(exp)		0.9316 (w.r.t. $\log(1/\epsilon)$)
QAOA	0.8907	0.5626
QAOA	/	0.8843 (w.r.t. $\log(1/\epsilon)$)

Table: Numerical scaling of the runtime as a function of the condition number and the accuracy, respectively, for the non-Hermitian example.

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

Block-encoding

- A “grey box” for the read-in problem.
- Example: $A \in \mathbb{C}^{N \times N}$. Unitary matrix $U \in \mathbb{C}^{2N \times 2N}$.

$$U_A = \begin{pmatrix} A & \cdot \\ \cdot & \cdot \end{pmatrix}$$

  U_A block-encodes A, which can be non-unitary.

- Given $A \in \mathbb{C}^{N \times N}$, can we find U_A? Block-encoding problem.
- Clearly not possible if $\|A\|_2 > 1$.

Block-encoding

Definition

Given an n-qubit matrix A, if we can find $\alpha, \epsilon \in \mathbb{R}_+$*, and an* $(m+n)$*-qubit matrix* U_A *so that that*

$$\|A - \alpha\left(\langle 0^m| \otimes I_n\right) U_A \left(|0^m\rangle \otimes I_n\right)\| \leq \epsilon,$$

then U_A *is called an* (α, m, ϵ)*-block-encoding of A.*

- Example: $m = 1$,

$$U_A = \begin{pmatrix} \widetilde{A} & \cdot \\ \cdot & \cdot \end{pmatrix}, \quad \left\|A - \alpha\widetilde{A}\right\| \leq \epsilon.$$

- Many examples of block-encoding: density operators, POVM operators, d-sparse matrices, addition and multiplication of block-encoded matrices (Gilyén-Su-Low-Wiebe, 2019)

Quantum signal processing

- A is Hermitian with eigenvalue decomposition $A = VDV^\dagger$. Compute matrix function $f(A) = Vf(D)V^\dagger$.

- Quantum signal processing: powerful, general, low-cost tool for block-encoding $f(A)$, where $f \in \mathbb{C}[x]$ is a polynomial satisfying certain parity constraints. (Low-Yoder-Chuang,2016) (Low-Chuang, 2017) (Gilyén-Su-Low-Wiebe, 2019)

- Generalizable to quantum singular value transformation.

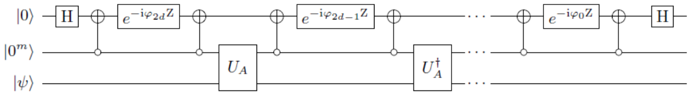

Eigenstate filtering problem

- H is Hermitian. λ is an eigenvalue of H, separated from the rest of the spectrum by a gap Δ.
- P_λ: projection operator into the λ-eigenspace of H. How to find a polynomial P to approximate P_λ?
- Requirement: $P(\lambda) = 1$ and $|P(\lambda')|$ is small for $\lambda' \in \sigma(H)\backslash\{\lambda\}$.

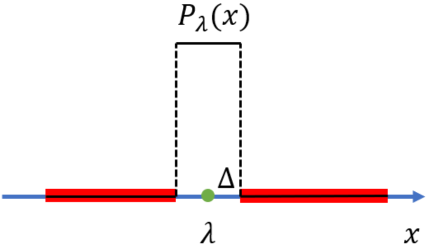

Eigenstate filtering

Theorem (**L.**-Tong, 1910.14596)

H is Hermitian, U_H is an $(\alpha, m, 0)$-block-encoding of H. λ is an eigenvalue of H separated from the rest of the spectrum by a gap Δ. Then we can construct a $(1, m+2, \epsilon)$-block-encoding of P_λ, by $\mathcal{O}((\alpha/\Delta)\log(1/\epsilon))$ applications of (controlled-) U_H and $U_H^\dagger$, and $\mathcal{O}((m\alpha/\Delta)\log(1/\epsilon))$ other primitive quantum gates.

Best polynomial approximation.

Eigenstate filtering

- Minimax polynomial

$$R_\ell(x;\Delta) = \frac{T_\ell\left(-1+2\frac{x^2-\Delta^2}{1-\Delta^2}\right)}{T_\ell\left(-1+2\frac{-\Delta^2}{1-\Delta^2}\right)}.$$

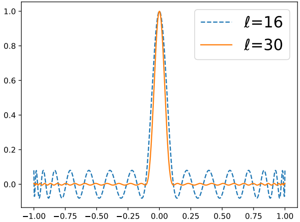

- Quantum algorithm based on quantum signal processing (Low-Chuang, 2017) (Gilyén-Su-Low-Wiebe, 2019)

Application of eigenstate filtering: Accelerating AQC(p) for QLSP

Theorem (**L.**-Tong, 1910.14596)

A is a d-sparse Hermitian matrix with condition number κ, $\|A\|_2 \leq 1$. *The solution* $|x\rangle \propto A^{-1}|b\rangle$ *can be obtained with fidelity* $1-\epsilon$ *using*

1. $\mathcal{O}\left(d\kappa\left(\frac{\log(d\kappa)}{\log\log(d\kappa)}+\log(\frac{1}{\epsilon})\right)\right)$ *oracle queries to* $A, |b\rangle$,
2. $\mathcal{O}\left(d\kappa\left(n\log(\frac{1}{\epsilon})+(n+\log(d\kappa))\frac{\log(d\kappa)}{\log\log(d\kappa)}\right)\right)$ *other primitive gates,*
3. $\mathcal{O}(n+\log(d\kappa))$ *qubits.*

- Complexity of AQC(p) is $T=\mathcal{O}(\kappa\log(\kappa)/\epsilon)$. Obtain solution $|x_0\rangle$ with $\epsilon\sim\mathcal{O}(1)$ accuracy using time $\mathcal{O}(\kappa\log(\kappa))$.

- Perform eigenstate filtering $|x\rangle \approx P_{\lambda=0}(H_1)|x_0\rangle$.

- Near-optimal complexity!

Numerical results

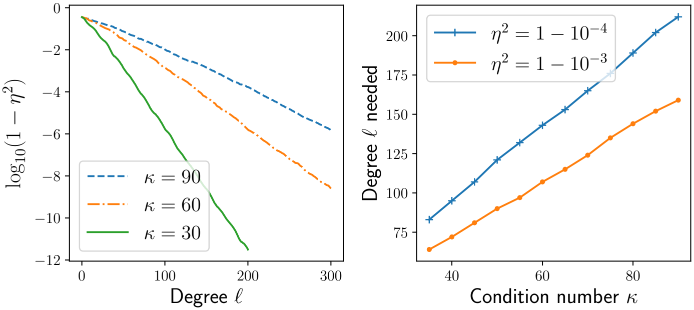

Figure: Left: fidelity η^2 converges to 1 exponentially as ℓ in the eigenvalues filtering algorithm increases, for different κ. Right: the smallest ℓ needed to achieve fixed fidelity η^2 grows linearly with respect to condition number κ. The initial state in eigenstate filtering is prepared by running AQC(p) for $T = 0.2\kappa$, with $p = 1.5$, which achieves an initial fidelity of about 0.6.

Application of eigenstate filtering: Quantum Zeno effect for QLSP

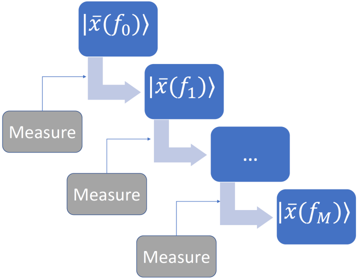

- Start with $|\bar{x}(0)\rangle = |0\rangle |b\rangle$ and end with $|\bar{x}(1)\rangle = |1\rangle |x\rangle$.
- At each step measure the state $|\bar{x}(f_{j-1})\rangle$ in the eigenbasis of $H(f_j)$.
- Fidelity approaches 1 as step size decreases.

- Quantum Zeno effect (QZE): (Childs et al, 2002) (Aharonov, Ta-Shma, 2003) (Boixo-Knill-Somma, 2009)

Application of eigenstate filtering: Quantum Zeno effect for QLSP

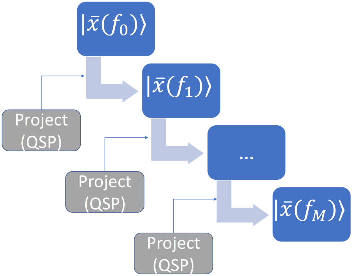

- Start with $|\bar{x}(0)\rangle = |0\rangle |b\rangle$ and end with $|\bar{x}(1)\rangle = |1\rangle |x\rangle$.
- At each step measure the state $|\bar{x}(f_{j-1})\rangle$ in the eigenbasis of $H(f_j)$.
- Fidelity approaches 1 as step size decreases.
- Replace measurement with eigenstate filtering (projection).

- Quantum Zeno effect (QZE): (Childs et al, 2002) (Aharonov, Ta-Shma, 2003) (Boixo-Knill-Somma, 2009)

Application of eigenstate filtering: Solving QLSP via quantum Zeno effect (QZE)

Theorem (**L.**-Tong, 1910.14596)

A is a d-sparse Hermitian matrix with condition number κ, $\|A\|_2 \leq 1$. Then $|x\rangle \propto A^{-1}|b\rangle$ can be obtained with fidelity $1-\epsilon$ using

1. $\mathcal{O}\left(d\kappa\left(\log(\kappa)\log\log(\kappa)+\log(\frac{1}{\epsilon})\right)\right)$ *queries to* $A, |b\rangle$,
2. $\mathcal{O}\left(nd\kappa\left(\log(\kappa)\log\log(\kappa)+\log(\frac{1}{\epsilon})\right)\right)$ *other primitive gates,*
3. $\mathcal{O}(n)$ *qubits.*

- Fully-gate based implementation (does not rely on adiabatic computing for the initial guess.
- Successive projection along the carefully scheduled adiabatic path.
- Near-optimal complexity!

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

Finding ground energy

- Hamiltonian H and its $(\alpha, m, 0)$-block-encoding U_H.
- Initial state $|\phi_0\rangle$ prepared by unitary U_I.
- Find λ_0 and the corresponding eigenstate $|\psi_0\rangle$.
- Assumptions

  (P1) Lower bound for the overlap: $|\langle\phi_0|\psi_0\rangle| \geq \gamma$,

  (P2) Bounds for the ground energy and spectral gap:
  $\lambda_0 \leq \mu - \Delta/2 < \mu + \Delta/2 \leq \lambda_1$.

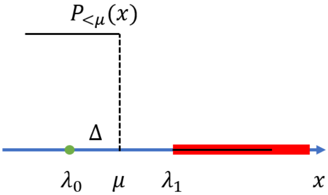

Binary search for ground energy

Polynomial $p(x)$ satisfies ($\deg p(x) = \mathcal{O}(\frac{1}{\delta}\log(\frac{1}{\epsilon}))$)

$$1-\epsilon \le p(x) \le 1,\ x \in [\delta, 1],$$
$$0 \le p(x) \le \epsilon,\ x \in [-1, -\delta].$$

$p(x)$ can be constructed by approximating erf (Low-Chuang, 2017).

- H is given in its $(\alpha, m, 0)$-block-encoding.
- Apply $p(\frac{H-x}{2\alpha})$ to an initial state with large overlap with the ground state.
- Can tell from the amplitude whether $E_0 \le x - h$ or $E_0 \ge x + h$ with high confidence, provided $E_0 \notin (x-h, x+h)$.

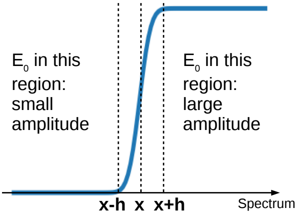

What if E_0 is here?

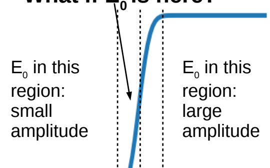

Binary search for ground energy

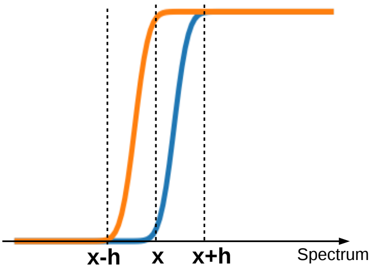

- Solution: apply two shifted polynomials.
- We can now return one of the two (not mutually exclusive) results with high confidence: $E_0 \geq x - h$ or $E_0 \leq x + h$.
- Perform binary search for E_0.

Near-optimal algorithm for finding the ground energy

- Well-known result: phase estimation (Kitaev, 1995)
- Previous best results: (Ge-Tura-Cirac, 2019)
- Our work: (**L.**-Tong, 2002.12508)

		Preparation (bound known)	Ground energy	Preparation (bound unknown)
U_H	This work	$\mathcal{O}\left(\frac{\alpha}{\gamma\Delta}\log(\frac{1}{\epsilon})\right)$	$\widetilde{\mathcal{O}}\left(\frac{\alpha}{\gamma h}\log(\frac{1}{\vartheta})\right)$	$\widetilde{\mathcal{O}}\left(\frac{\alpha}{\gamma\Delta}\log(\frac{1}{\vartheta\epsilon})\right)$
	GTC19	$\widetilde{\mathcal{O}}\left(\frac{\alpha}{\gamma\Delta}\right)$	$\widetilde{\mathcal{O}}\left(\frac{\alpha^{3/2}}{\gamma h^{3/2}}\right)$	$\widetilde{\mathcal{O}}\left(\frac{\alpha^{3/2}}{\gamma\Delta^{3/2}}\right)$
U_I	This work	$\mathcal{O}\left(\frac{1}{\gamma}\right)$	$\widetilde{\mathcal{O}}\left(\frac{1}{\gamma}\log(\frac{\alpha}{h})\log(\frac{1}{\vartheta})\right)$	$\widetilde{\mathcal{O}}\left(\frac{1}{\gamma}\log(\frac{\alpha}{\Delta})\log(\frac{1}{\vartheta})\right)$
	GTC19	$\widetilde{\mathcal{O}}\left(\frac{1}{\gamma}\right)$	$\widetilde{\mathcal{O}}\left(\frac{1}{\gamma}\sqrt{\frac{\alpha}{h}}\right)$	$\widetilde{\mathcal{O}}\left(\frac{1}{\gamma}\sqrt{\frac{\alpha}{\Delta}}\right)$
Extra qubits	This work	$\mathcal{O}(1)$	$\mathcal{O}(\log(\frac{1}{\gamma}))$	$\mathcal{O}(\log(\frac{1}{\gamma}))$
	GTC19	$\mathcal{O}(\log(\frac{1}{\Delta}\log(\frac{1}{\epsilon})))$	$\mathcal{O}(\log(\frac{1}{h}))$	$\mathcal{O}(\log(\frac{1}{\Delta}\log(\frac{1}{\epsilon})))$

h: precision of the ground energy estimate; $1-\vartheta$: success probability

Optimality of the algorithm (lower bound)

Theorem (**L.**-Tong, 2002.12508)

Given a generic Hamiltonian H and its $(\alpha, m, 0)$-block-encoding U_H, and $\alpha = \Theta(1)$. Initial state $|\phi_0\rangle$ is prepared by U_I with known lower bound of the initial overlap γ and the energy gap Δ. Then to prepare the ground state

1. *When $\Delta = \Omega(1)$, and $\gamma \to 0^+$, the number of queries to U_H is $\Omega(1/\gamma)$,*
2. *When $\gamma = \Omega(1)$, and $\Delta \to 0^+$, the number of queries to U_H is $\Omega(1/\Delta)$,*
3. *When $\Delta = \Omega(1)$, and $\gamma \to 0^+$, the number of queries to U_I cannot be $\mathcal{O}(1/\gamma^{1-\theta})$ while the number of queries to U_H is $\mathcal{O}(\mathrm{poly}(1/\gamma))$ for any $\theta > 0$.*

Outline

Introduction

Near-optimal quantum linear solver: adiabatic quantum computing

Near-optimal quantum linear solver: eigenstate filtering

Near-optimal algorithm for ground energy

Future works

Challenges

- Large-scale fully error-corrected quantum computer remains at least really, really, really hard in the near future. Think about both near-term and long-term for quantum linear algebra.

- Efficient gate-based implementation of adiabatic quantum computing (AQC).
  1. Time-dependent Hamiltonian simulation problem.
  2. Commutator-based error bounds (Childs et al, 2019)

- Quantum signal processing: approximation theory in SU(2).
  1. How to obtain the phase factors: optimization based approach (Dong-Meng-Whaley-**L.**, 2002.11649)
  2. Polynomial approximation with nontrivial constraints.
  3. Decay of phase factors and regularity of the function.

Challenges

- Fast-forwarding of certain Hamiltonians, and preconditioning. Simulation in the interaction picture.

- Quantum speedup in terms of solving ODEs / PDEs / open quantum systems.

- Explore the power of the block-encoding model:
  1. Block-encoding based Hamiltonian simulation can be much tricker than Trotter based approaches in practice.
  2. Connection with supremacy type circuits.

- Beyond the oracular assumption and demonstrate the advantage of QLSP solvers for real applications.

- What is the proper counterpart of dense matrices in the quantum setting? What should be the proper “quantum LINPACK benchmarks” in the post-supremacy era?

References

- L. Lin and Y. Tong, Near-optimal ground state preparation [arXiv:2002.12508]
- Y. Dong, X. Meng, K. B. Whaley, L. Lin, Efficient Phase Factor Evaluation in Quantum Signal Processing [arXiv:2002.11649]
- L. Lin and Y. Tong, Optimal quantum eigenstate filtering with application to solving quantum linear systems [arXiv:1910.14596]
- D. An and L. Lin, Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm [arXiv:1909.05500]

Acknowledgement

Thank you for your attention!

Lin Lin
https://math.berkeley.edu/~linlin/

U.S. DEPARTMENT OF ENERGY | Office of Science   QAT4Chem   Google AI