Quantum Linear System Algorithm for Dense Matrices

Leonard Wossnig,[1,2] Zhikuan Zhao,[3,4,*] and Anupam Prakash[4]
[1]Theoretische Physik, ETH Zürich, 8093 Zurich, Switzerland
[2]Department of Materials, University of Oxford, Oxford OX1 3PH, United Kingdom
[3]Singapore University of Technology and Design, 487372 Singapore
[4]Centre for Quantum Technologies, National University of Singapore, 117543 Singapore

(Received 25 May 2017; revised manuscript received 10 November 2017; published 31 January 2018)

Solving linear systems of equations is a frequently encountered problem in machine learning and optimization. Given a matrix A and a vector $\mathbf{b}$ the task is to find the vector $\mathbf{x}$ such that $A\mathbf{x} = \mathbf{b}$. We describe a quantum algorithm that achieves a sparsity-independent runtime scaling of $\mathcal{O}(\kappa^2\sqrt{n}\mathrm{polylog}(n)/\epsilon)$ for an $n \times n$ dimensional A with bounded spectral norm, where κ denotes the condition number of A, and ϵ is the desired precision parameter. This amounts to a polynomial improvement over known quantum linear system algorithms when applied to dense matrices, and poses a new state of the art for solving dense linear systems on a quantum computer. Furthermore, an exponential improvement is achievable if the rank of A is polylogarithmic in the matrix dimension. Our algorithm is built upon a singular value estimation subroutine, which makes use of a memory architecture that allows for efficient preparation of quantum states that correspond to the rows of A and the vector of Euclidean norms of the rows of A.

DOI: 10.1103/PhysRevLett.120.050502

A common bottleneck in statistical learning and machine learning algorithms is the inversion of high-dimensional matrices in order to solve linear systems of equations. Examples include covariance matrix inversions in Gaussian processes and support vector machines, as well as data matrix inversions in large scale regression problems [1,2].

Recent advances in the field of quantum information processing have provided promising prospects for the efficient solution of high-dimensional linear systems. The breakthrough work of Harrow, Hassidim, and Lloyd [3] introduced the quantum linear system algorithm (QLSA) that computes the quantum state $|\mathbf{x}\rangle = |A^{-1}\mathbf{b}\rangle$ corresponding to the solution of a linear system $A\mathbf{x} = \mathbf{b}$, where $A \in \mathbb{R}^{n\times n}$ and $\mathbf{x}, \mathbf{b} \in \mathbb{R}^n$, in time $\mathcal{O}(\mathrm{polylog}(n))$ for a sparse and well-conditioned A. Unlike the output $A^{-1}\mathbf{b} \in \mathbb{R}^n$ of a classical linear system solver, a copy of $|A^{-1}\mathbf{b}\rangle$ does not provide access to the coordinates of $A^{-1}\mathbf{b}$. Nevertheless, it allows us to perform useful computations such as sampling from the solution vector. The QLSA has inspired several works [4–13] in the emerging research area of quantum machine learning.

In the classical setting, the best known algorithm for the sampling task performed by the QLSA requires solving the linear system. The running time for a classical linear system solver scales as $\mathcal{O}(n^\omega)$, where the matrix multiplication exponent $\omega \leq 2.373$ [14,15]. However, as the subcubic scaling is difficult to achieve in practice, linear system solvers typically use the Cholesky decomposition and require time $\mathcal{O}(n^3)$ for dense matrices. The QLSA [3] has running time $\tilde{O}(\kappa^2 s(A)^2/\epsilon)$, where κ is the condition number, $s(A)$ is the sparsity or the maximum number of nonzero entries in a row of A, and ϵ is the precision to which the solution is approximated. There have been several improvements to the QLSA since the original proposal that have improved the running time to linear in κ and $s(A)$ and to polylogarithmic in the precision parameter ϵ [16,17]. The work in Ref. [18] introduced preconditioning for the QLSA and extended its applicability.

Quantum machine learning is an emerging research area that attempts to harness the power of quantum information processing to obtain speedups for classical machine learning tasks. A number of quantum machine learning algorithms have been proposed [4–13,19,20]. Most of these algorithms use a quantum linear system solver as a subroutine. However, as mentioned in Ref. [3], and later also pointed out in Refs. [21,22], the QLSA potentially has a few caveats. In particular, the QLSA achieves an exponential speedup over classical algorithms when the matrix A is sparse and well conditioned, due to the sparsity-dependent Hamiltonian simulation subroutine.

It was shown in Ref. [23] that given black-box access to the elements of matrix A, that is, given an oracle that supports queries of the form $O_A|j,k\rangle|z\rangle \to |j,k\rangle|z \oplus A_{jk}\rangle$ Hamiltonian simulation with error δ_h can be performed in time $\mathcal{O}(n^{2/3}\mathrm{polylog}(n)/\delta_h^{1/3})$. This yields a quantum linear system algorithm for dense matrices with running time $\mathcal{O}(\kappa^2 n^{2/3}\mathrm{polylog}(n)/\epsilon)$. It was also observed empirically in Ref. [23] that this algorithm scales as $\mathcal{O}(\sqrt{n}\mathrm{polylog}(n)/\delta_h^{1/2})$ for several classes of matrices. However, the $\tilde{\mathcal{O}}(\sqrt{n})$ runtime for black-box Hamiltonian simulation was not formally proven and is known not to

hold in the worst case. The oracle O_A can be implemented efficiently if A is well structured, that is, given j, k there is a small sized quantum circuit to compute the entry A_{jk}. However, for the dense linear systems that arise in a large class of interesting machine learning applications, we do not have access to such a circuit. Examples include kernel methods [24] and artificial neural networks where convolutional neural network architectures in particular rely on subroutines for manipulating large, dense matrices [25,26]. A possible solution in these cases is to use the quantum random access memory (QRAM) [27,28] as a general purpose device to implement the desired oracle.

In this Letter we present an alternative quantum algorithm for solving linear systems of equations. Instead of assuming black-box access to the matrix elements, we work in a memory model where the entries of A are stored as a suitable data structure in a QRAM. We show that a $\tilde{\mathcal{O}}(\sqrt{n})$ linear system solver is provably achievable even in the worst case with this augmented memory model. Moreover, an exponential advantage can be obtained with our proposed approach when rank of A is polylogarithmic in n.

Our linear system solver uses the quantum singular value estimation (QSVE) algorithm introduced in Ref. [29]. The algorithm achieves a runtime $\mathcal{O}(\kappa^2\|A\|_F\mathrm{polylog}(n)/\epsilon)$, where κ denotes the condition number of A, $\|A\|_F$ is the Frobenius norm, and ϵ is the precision parameter. When the spectral norm $\|A\|_*$ is bounded by a constant, the scaling becomes $\mathcal{O}(\kappa^2\sqrt{n}\mathrm{polylog}(n)/\epsilon)$, which amounts to a polynomial speed-up over the scaling achieved by Ref. [3] when applied to dense matrices. We assume without loss of generality that the matrix A is Hermitian as a general rectangular matrix M can be embedded into a block antidiagonal Hermitian matrix with the elements of $M^\dagger$ and M in the lower and upper half, respectively [3]. Furthermore, if A is not invertible, our algorithm applies the Moore-Penrose pseudoinverse and the running time is bounded by $1/|\lambda_{\min}|$, where $\lambda_{\min}$ is the nonzero eigenvalue for A with the smallest absolute value instead of κ.

We start by introducing some preliminaries. For a symmetric matrix $A \in \mathbb{R}^{n\times n}$ with spectral decomposition $A = \sum_{i\in[n]}\lambda_i \mathbf{s}_i\mathbf{s}_i^\dagger$, the singular value decomposition is given by $A = \sum_{i\in[n]}|\lambda_i|\mathbf{u}_i\mathbf{v}_i^\dagger$, where the left and right singular vectors $\mathbf{u}_i$ and $\mathbf{v}_i$ are equal to the eigenvectors $\mathbf{s}_i$ up to a sign, such that $\mathbf{s}_i = \mathbf{u}_i = \pm\mathbf{v}_i$. We also need the well-known quantum phase estimation algorithm:

Theorem 1.—(Phase estimation [30]) Let unitary $U|\mathbf{v}_j\rangle = \exp(i\theta_j)|\mathbf{v}_j\rangle$ with $\theta_j \in [-\pi,\pi]$ for $j \in [n]$. There is a quantum algorithm that transforms $\sum_{j\in[n]}\alpha_j|\mathbf{v}_j\rangle \to \sum_{j\in[n]}\alpha_j|\mathbf{v}_j\rangle|\bar{\theta}_j\rangle$ such that $|\bar{\theta}_j - \theta_j| \le \delta$ for all $j \in [n]$ with probability $1 - 1/\mathrm{poly}(n)$ in time $\mathcal{O}(T_U \log(n)/\delta)$, where T_U defines the time to implement U.

Quantum singular value estimation can be viewed as an extension of phase estimation to nonunitary matrices. It is the main algorithmic primitive required for our linear system solver. Given the matrix A with singular value decomposition $A = \sum_i \sigma_i \mathbf{u_i}\mathbf{v_i}^\dagger$ and an arbitrary input state $\sum_i \alpha_i|\mathbf{v}_i\rangle$, QSVE amounts to performing the map $\sum_i\alpha_i|\mathbf{v}_i\rangle \to \sum_i\alpha_i|\mathbf{v}_i\rangle|\bar{\sigma}_i\rangle$, where $|\overline{\sigma_i}\rangle$ encodes the precision δ estimates for the singular values of A with high probability, such that $|\overline{\sigma_i} - \sigma_i| \le \delta$ for all i. A QSVE algorithm with running time of $\tilde{\mathcal{O}}(\|A\|_F/\delta)$ was presented in Ref. [29], where it was used for quantum recommendation systems. An SVE algorithm applied to a symmetric matrix estimates $|\lambda_i|$ but does not recover $\mathrm{sgn}(\lambda_i)$. However, in order to solve linear systems we also need to recover $\mathrm{sgn}(\lambda_i)$. In this Letter, we provide a simple procedure for recovering the sign given an SVE algorithm. Our procedure thus provides a way to construct a quantum linear system solver from a QSVE algorithm in a black-box manner.

The main result of this Letter is a quantum linear system solver based on the QSVE algorithm [29] that achieves a running time of $\mathcal{O}(\kappa^2\|A\|_F\mathrm{polylog}(n)/\epsilon)$. We briefly describe the QSVE algorithm in the next section. We then present the quantum linear system solver and provide a complete analysis for the linear system solver as well as a comparison with other approaches to the QLSA in the discussion.

The QSVE algorithm.—The QSVE algorithm requires the ability to efficiently prepare the quantum states corresponding to the rows and columns of matrix A. The matrix entries are stored in the following data structure, such that a quantum algorithm with access to this data structure has this ability.

Lemma 1.—(Data structure [29]) Let $A \in \mathbb{R}^{m\times n}$ be a matrix with entries A_{ij} which arrive in an arbitrary order. There exists a data structure with the following properties: (i) A quantum computer with access to the data structure can perform the following mappings in $\mathcal{O}(\mathrm{polylog}(mn))$ time.

$$U_{\mathcal{M}}\colon |i\rangle|0\rangle \to |i,\mathbf{A_i}\rangle = \frac{1}{\|\mathbf{A_i}\|}\sum_{j=1}^{n} A_{ij}|i,j\rangle,$$
$$U_{\mathcal{N}}\colon |0\rangle|j\rangle \to |\mathbf{A}_F,j\rangle = \frac{1}{\|A\|_F}\sum_{i=1}^{m}\|\mathbf{A_i}\||i,j\rangle, \quad (1)$$

where $\mathbf{A_i} \in \mathbb{R}^n$ correspond to the rows of the matrix A and $\mathbf{A}_F \in \mathbb{R}^m$ is a vector whose entries are the ℓ_2 norms of the rows, i.e., $(\mathbf{A}_F)_i = \|A_i\|$. (ii) The time required to store a new entry A_{ij} is $\mathcal{O}(\log^2(mn))$ and data structure size is $O(w\log mn)$, where w is the number of nonzero entries in A.

We provide a brief description of a possible realization of this data structure with a schematic diagram in the Supplemental Material [31]; more details can be found in Ref. [29].

The QSVE algorithm is a quantum walk based algorithm that leverages the connection between the singular values σ_i of the target matrix A and the principal angles θ_i between

certain subspaces associated with A. It uses a factorization $[A/(\|A\|_F)]=\mathcal{M}^\dagger\mathcal{N}$, where $\mathcal{M}\in\mathbb{R}^{mn\times m}$ and $\mathcal{N}\in\mathbb{R}^{mn\times n}$ are isometries with column spaces $\mathcal{C}(\mathcal{M})$ and $\mathcal{C}(\mathcal{N})$ such that the unitary operator $W=(2\mathcal{M}\mathcal{M}^\dagger-I_{mn})\times(2\mathcal{N}\mathcal{N}^\dagger-I_{mn})$, where I_{mn} denotes the identity matrix, can be implemented efficiently using the data structure in Lemma 1.

The operators $2\mathcal{M}\mathcal{M}^\dagger-I_{mn}$ and $2\mathcal{N}\mathcal{N}^\dagger-I_{mn}$ represent reflections on subspaces $\mathcal{C}(\mathcal{M})$ and $\mathcal{C}(\mathcal{N})$, respectively. The application of W thus amounts to two sequential reflections, first on $\mathcal{C}(\mathcal{N})$, and then on $\mathcal{C}(\mathcal{M})$. As such it can be interpreted as taking a step in the bipartite quantum walk [32], where the singular values of the normalized target matrix $A/\|A\|_F$ for QSVE correspond to the principal angles between the subspaces $\mathcal{C}(\mathcal{M})$ and $\mathcal{C}(\mathcal{N})$. The relation between quantum walks and eigenvalues has been well-known in the literature and has been used in several previous results [32–34]. However, the quantum walk defined by the QSVE algorithm is particularly interesting for linear systems, as instead of the sparsity $s(A)$ it depends on the Frobenius norm $\|A\|_F$. Algorithm 1 describes the QSVE algorithm, the analysis is contained in the following lemma:

Algorithm 1. Quantum singular value estimation. [29].

1. Create the arbitrary input state $|\alpha\rangle=\sum_i\alpha_i|\mathbf{v}_i\rangle$
2. Append a register $|0^{\lceil\log m\rceil}\rangle$ and create the state $|\mathcal{N}\alpha\rangle=\sum_i\alpha_i|\mathcal{N}\mathbf{v}_i\rangle=\sum_i\alpha_i(\omega_i^+|\mathbf{w}_i^+\rangle+\omega_i^-|\mathbf{w}_i^-\rangle)$.
3. Perform phase estimation [30] with precision $2\delta>0$ on input $|\mathcal{N}\alpha\rangle$ for $W=(2\mathcal{M}\mathcal{M}^\dagger-I_{mn})(2\mathcal{N}\mathcal{N}^\dagger-I_{mn})$ and obtain $\sum_i\alpha_i(\omega_i^+|\mathbf{w}_i^+,\bar{\theta}_i\rangle+\omega_i^-|\mathbf{w}_i^-,-\bar{\theta}_i\rangle)$, where $\overline{\theta_i}$ is the estimated phase θ_i in binary bit strings.
4. Compute $\bar{\sigma}_i=\cos(\pm\overline{\theta_i}/2)\|A\|_F$.
5. Uncompute the output of the phase estimation and apply the inverse transformation of step (2) to obtain

$$\sum_i\alpha_i|\mathbf{v}_i\rangle|\overline{\sigma_i}\rangle. \tag{2}$$

Lemma 2.—(Preparation of the isometries [29]).

Let $A\in\mathbb{R}^{m\times n}$ be a matrix with singular value decomposition $A=\sum_i\sigma_i\mathbf{u}_i\mathbf{v}_i^\dagger$ stored in the data structure described in Lemma 1. Then there exist matrices $\mathcal{M}\in\mathbb{R}^{mn\times m}$, and $\mathcal{N}\in\mathbb{R}^{mn\times n}$, such that (1) $\mathcal{M}$, $\mathcal{N}$ are isometries, that is $\mathcal{M}^\dagger\mathcal{M}=I_m$ and $\mathcal{N}^\dagger\mathcal{N}=I_n$ such that A can be factorized as $A/\|A\|_F=\mathcal{M}^\dagger\mathcal{N}$. Multiplication by $\mathcal{M}$, $\mathcal{N}$, i.e., the mappings $|\alpha\rangle\to|\mathcal{M}\alpha\rangle$ and $|\beta\rangle\to|\mathcal{N}\beta\rangle$ can be performed in time $\mathcal{O}(\mathrm{polylog}(mn))$. (2) The reflections $2\mathcal{M}\mathcal{M}^\dagger-I_{mn}$, $2\mathcal{N}\mathcal{N}^\dagger-I_{mn}$, hence, the unitary $W=(2\mathcal{M}\mathcal{M}^\dagger-I_{mn})(2\mathcal{N}\mathcal{N}^\dagger-I_{mn})$ can be implemented in time $\mathcal{O}(\mathrm{polylog}(mn))$. (3) The unitary W acts as rotation by θ_i on the two dimensional invariant subspace $\{\mathcal{M}\mathbf{u}_i,\mathcal{N}\mathbf{v}_i\}$ plane, such that $\sigma_i=\cos(\theta_i/2)\|A\|_F$, where σ_i is the ith singular value for A.

We outline the ideas involved in the analysis of the QSVE algorithm and refer to Ref. [29] for further details. The map $\mathcal{M}$ appends to an arbitrary input state vector $|\alpha\rangle$ a register that encodes the row vectors $\mathbf{A_i}$ of A, such that

$$\mathcal{M}\colon\ |\alpha\rangle=\sum_{i=1}^m\alpha_i|i\rangle\to\sum_{i=1}^m\alpha_i|i,\mathbf{A_i}\rangle=|\mathcal{M}\alpha\rangle.$$

The map $\mathcal{N}$ similarly appends to an arbitrary input state vector $|\alpha\rangle$ a register that encodes the vector $\mathbf{A_F}$ whose entries are the ℓ_2 norms $\|\mathbf{A_i}\|$ of the rows of A,

$$\mathcal{N}\colon\ |\alpha\rangle=\sum_{j=1}^n\alpha_j|j\rangle\to\sum_{j=1}^n\alpha_j|\mathbf{A_F},j\rangle=|\mathcal{N}\alpha\rangle.$$

The above maps can be efficiently implemented given the memory structure described by Lemma 1.

The factorization of A follows from the amplitude encoding of $\mathbf{A_i}$ and $\mathbf{A}_F$. We have $|i,\mathbf{A_i}\rangle=(1/\|\mathbf{A_i}\|)\times\sum_{j=1}^nA_{ij}|i,j\rangle$ and $|\mathbf{A}_F,j\rangle=(1/\|A\|_F)\sum_{i=1}^m\|\mathbf{A_i}\||i,j\rangle$, implying that $(\mathcal{M}^\dagger\mathcal{N})_{ij}=\langle i,\mathbf{A_i}|\mathbf{A}_F,j\rangle=[A_{ij}/(\|A\|_F)]$. Similarly, it follows that $\mathcal{M}$ and $\mathcal{N}$ have orthonormal columns and thus $\mathcal{M}^\dagger\mathcal{M}=I_m$ and $\mathcal{N}^\dagger\mathcal{N}=I_n$. The relation between the eigenvalues of W and the singular values of A follows from the observation that W acts on $|\mathcal{N}\mathbf{v}_i\rangle$ as a rotation in the plane of $\{\mathcal{M}\mathbf{u}_i,\mathcal{N}\mathbf{v}_i\}$ by θ_i, such that $\cos(\theta_i/2)=[\sigma_i/(\|A\|_F)]$. The detailed calculation leading to this fact is included in the Supplemental Material [31]. The two dimensional subspace spanned by $\{\mathcal{M}\mathbf{u}_i,\mathcal{N}\mathbf{v}_i\}$ is, therefore, invariant under the action of W which acts on it as a rotation by angle θ_i. The corresponding eigenvectors of W, $|\mathbf{w}_i^\pm\rangle$ hence have eigenvalues $\exp(\pm i\theta_i)$. We have $|\mathcal{N}\mathbf{v}_i\rangle=\omega_i^+|\mathbf{w}_i^+\rangle+\omega_i^-|\mathbf{w}_i^-\rangle$, with $|\omega_i^-|^2+|\omega_i^+|^2=1$, and phase estimation can be performed to estimate $\pm\overline{\theta_i}$. Computing $\overline{\sigma_i}=\cos(\overline{\theta_i}/2)\|A\|_F$ then leads to estimation of the singular values. We have shown the correctness of Algorithm 1, which is summarized in the following theorem.

Theorem 2.—(Quantum singular value estimation [29]) Let $A\in\mathbb{R}^{m\times n}$ be a matrix with singular value decomposition $A=\sum_i\sigma_i\mathbf{u}_i\mathbf{v}_i^\dagger$ stored in the data structure in Lemma 1. Further, let $\delta>0$ be the precision number. There is an algorithm that runs in $\mathcal{O}(\mathrm{polylog}(mn)/\delta)$ and performs the mapping $\sum_i\alpha_i|\mathbf{v}_i\rangle|0\rangle\to\sum_i\alpha_i|\mathbf{v}_i\rangle|\overline{\sigma_i}\rangle$, where $\overline{\sigma_i}\in\sigma_i\pm\delta\|A\|_F$ for all i with probability at least $1-1/\mathrm{poly}(n)$.

The runtime of QSVE is dominated by the phase estimation procedure which returns a δ-close estimate of θ_i, s.t. $|\bar{\theta}_i-\theta_i|\le2\delta$, which translates into the estimated singular value via $\bar{\sigma}_i=\cos(\bar{\theta}_i/2)\|A\|_F$. The error in σ_i can then be bounded from above by $|\bar{\sigma}_i-\sigma_i|\le\delta\|A\|_F$. The unitary W can be implemented in time $\mathcal{O}(\mathrm{polylog}(mn))$ by Lemma 2. Therefore, by Theorem 1 the running time for estimating of the singular values with additive error $\delta\|A\|_F$ in $\mathcal{O}(\mathrm{polylog}(mn)/\delta)$.

Quantum linear system algorithm.—We now give an algorithm for solving linear systems of equations in $\tilde{\mathcal{O}}(\sqrt{n})$ time for arbitrary matrices with bounded spectral norm. A QSVE algorithm immediately yields a linear system solver for positive definite matrices as the estimated singular values and eigenvalues are related via $\overline{\sigma_i} = |\overline{\lambda_i}|$. In order to solve general linear systems we need to recover the sign of each $\overline{\lambda_i}$. We provide a simple algorithm that recovers the signs using the QSVE procedure as a black box incurring only a constant overhead. We assume that A has been rescaled so that its eigenvalues lie within the interval $[-1, -1/\kappa] \cup [1/\kappa, 1]$. The QLSA [3] also makes the same assumption; this is also indicated in the review [35]. The main result of this Letter is the following theorem:

Algorithm 2. Quantum linear system solver.

1. Create the state $|\mathbf{b}\rangle = \sum_i \beta_i |\mathbf{v}_i\rangle$ with $\mathbf{v}_i$ being the singular vectors of A.
2. Perform two QSVEs as in Algorithm 1 for matrices A, $A + \mu I$ with $\delta < 1/2\kappa$ and $\mu = 1/\kappa$ to obtain
$$\sum_i \beta_i |\mathbf{v}_i\rangle_A |\bar{\lambda}_i|\rangle_B |\bar{\lambda}_i + \mu|\rangle_C.$$
3. Add an auxiliary register and set it to 1 if the value in register B is greater than that in register C and apply a conditional phase gate:
$$\sum_i (-1)^{f_i} \beta_i |\mathbf{v}_i\rangle_A |\bar{\lambda}_i|\rangle_B |\bar{\lambda}_i + \mu|\rangle_C |f_i\rangle_D.$$
4. Add an ancilla register and apply a rotation conditioned on register B with $\gamma = O(1/\kappa)$. Then uncompute the registers B, C, D, to obtain
$$\sum_i (-1)^{f_i} \beta_i |\mathbf{v}_i\rangle \left(\frac{\gamma}{|\bar{\lambda}_i|} |0\rangle + \sqrt{1 - \left(\frac{\gamma}{|\bar{\lambda}_i|} \right)^2} |1\rangle \right).$$

Postselect on the ancilla register being in state $|0\rangle$.

Theorem 3.—Let $A \in \mathbb{R}^{n \times n}$ be a Hermitian matrix with spectral decomposition $A = \sum_i \lambda_i \mathbf{u_i} \mathbf{u_i}^\dagger$ stored in the data structure in Lemma 1. Further, let κ be the condition number A, and $\|A\|_F$ the Frobenius norm, and $\epsilon > 0$ be a precision parameter. Then Algorithm 2 has runtime $\mathcal{O}(\kappa^2 \text{polylog}(n) \|A\|_F / \epsilon)$ that outputs state $|\overline{A^{-1}\mathbf{b}}\rangle$ such that $\| |\overline{A^{-1}\mathbf{b}}\rangle - |A^{-1}\mathbf{b}\rangle \| \leq \epsilon$.

Proof.—We first argue that Algorithm 2 correctly recovers the sign of the λ_i. The algorithm compares the estimates obtained by performing QSVE for A and for $A' = A + \mu I_n$, where μ is a positive scalar to be chosen later. The matrix A' has the same eigenvectors as A, but has eigenvalues $\lambda_i + \mu$. Note that for $\lambda_i \geq 0$, we have $|\lambda_i + \mu| = |\lambda_i| + |\mu| \geq |\lambda_i|$. However, if $\lambda_i < -\mu/2$, then $|\lambda_i + \mu| < |\lambda_i|$. Thus, if the estimates were perfect, then choosing $\mu < 2/\kappa$ would recover the sign correctly as the eigenvalues of A lie in the interval $[-1, -1/\kappa] \cup [1/\kappa, 1]$. With the choice $\mu = 1/\kappa$ and $\delta < 1/2\kappa$ we find that the signs are still correct for all λ_i.

The analysis of the error bounds appears in the Supplemental Material [31], where it is shown that the error ϵ for the linear system solver is related to the QSVE precision parameter δ via $\delta = \mathcal{O}[\epsilon/(\kappa \|A\|_F)]$. Considering the success probability of the postselection step, we require on average $\mathcal{O}(\kappa^2)$ repetitions of the coherent computation. This can be reduced to $\mathcal{O}(\kappa)$ using amplitude amplification [36]. Therefore, an upper-bound of the runtime of our algorithm is given by $\mathcal{O}(\kappa^2 \text{polylog}(n) \|A\|_F / \epsilon)$. ∎

The error dependence on the Frobenius norm suggests that our algorithm is most accurate when the $\|A\|_F$ is bounded by some constant, in which case the algorithm returns the output state with a constant ϵ error in polylogarithmic time even if the matrix is nonsparse. More generally, as in the QLSA we can assume that the spectral norm $\|A\|_*$ is bounded by a constant, although the Frobenius norm may scale with the dimensionality of the matrix. In such cases we have $\|A\|_F = \mathcal{O}(\sqrt{n})$. In these scenarios the proposed algorithm runs in $\mathcal{O}(\kappa^2 \sqrt{n} \text{polylog}(n)/\epsilon)$ and returns the output with a constant ϵ error. Note that since $\|A\|_F \leq \sqrt{r} \|A\|_*$, where r denotes the rank of A, we may also write the runtime as $\mathcal{O}(\kappa^2 \sqrt{r} \text{polylog}(n)/\epsilon)$. Hence a particularly advantageous runtime is achievable if the rank of A is polylogarithmic in n.

We note that the memory model of Lemma 1 is stronger than the black-box model; therefore, a direct comparison of our results with the black-box approach of Ref. [23] is not appropriate. The QSVE-based linear system solver achieves a $\tilde{\mathcal{O}}(\sqrt{n})$-scaling in this stronger memory model, and it is an interesting open problem to prove a similar scaling in the black-box model. We also note that for practical implementations, the constant runtime overhead with respect to a given set of elementary fault-tolerant quantum gates is an important consideration. It has been shown by Scherer *et al.* [37] that current approaches to the QLSA potentially suffer from a large constant overhead, hindering prospects of near-term applications. Whether our proposed QSVE-based algorithm exhibits a more advantageous constant overhead due to the absence of Hamiltonian simulation, remains an open question.

The authors thank Simon Benjamin, Joseph Fitzsimons, Patrick Rebentrost, and Nathan Wiebe for helpful comments on the manuscript, and Andrew Childs for his feedback on the earlier version and Danial Dervovic for help with the figure. The authors also acknowledge support from Singapore's Ministry of Education and National Research Foundation. This material is based on research funded in part by the Singapore National Research Foundation under NRF Grants No. NRF-NRFF2013-01 and No. NRF-NRFF2013-13.

L. W., Z. Z., and A. P. contributed equally to this work.

*Corresponding author.
zhikuan_zhao@mymail.sutd.edu.sg

[1] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2004), Vol. 14, ISBN: .
[2] C. M. Bishop, *Pattern Recognition and Machine Learning*, Vol. 4 (Springer, New York, 2006), ISBN: .
[3] A. W. Harrow, A. Hassidim, and S. Lloyd, Phys. Rev. Lett. **103**, 150502 (2009).
[4] P. Rebentrost, M. Mohseni, and S. Lloyd, Phys. Rev. Lett. **113**, 130503 (2014).
[5] M. Schuld, I. Sinayskiy, and F. Petruccione, Phys. Rev. A **94**, 022342 (2016).
[6] N. Wiebe, D. Braun, and S. Lloyd, Phys. Rev. Lett. **109**, 050505 (2012).
[7] N. Wiebe and C. Granade, arXiv:1512.03145.
[8] N. Wiebe, C. Granade, and D. G. Cory, New J. Phys. **17**, 022005 (2015).
[9] N. Wiebe, A. Kapoor, and K. M. Svore, Quantum Inf. Comput. **15**, 318 (2015).
[10] N. Wiebe, A. Kapoor, and K. M. Svore, arXiv:1412.3489.
[11] N. Wiebe, A. Kapoor, and K. M. Svore, arXiv:1602.04799.
[12] W. Zeng and B. Coecke, *Proceedings of the 2016 Workshop on Semantic Spaces at the Intersection of NLP, Physics and Cognitive Science, Glasgow, Scotland, 2016* (2016), Vol. 221.
[13] Z. Zhao, J. K. Fitzsimons, and J. F. Fitzsimons, arXiv:1512.03929.
[14] D. Coppersmith and S. Winograd, J. Symb. Comput. **9**, 251 (1990).
[15] F. Le Gall, in *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation* (ACM, New York, 2014), p. 296.
[16] A. M. Childs, R. Kothari, and R. D. Somma, SIAM J. Comput. **46**, 1920 (2017).
[17] A. Ambainis, in *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, Paris, France* (LIPIcs, 2012).
[18] B. D. Clader, B. C. Jacobs, and C. R. Sprouse, Phys. Rev. Lett. **110**, 250504 (2013).
[19] P. Rebentrost, M. Schuld, F. Petruccione, and S. Lloyd, arXiv:1612.01789.
[20] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, arXiv:1707.08561.
[21] A. M. Childs, Nat. Phys. **5**, 861 (2009).
[22] S. Aaronson, Nat. Phys. **11**, 291 (2015).
[23] D. W. Berry and A. M. Childs, Quantum Inf. Comput. **12**, 29 (2012).
[24] A. G. Wilson, C. Dann, and H. Nickisch, arXiv:1511.01870.
[25] J. Yangqing, Ph. D. thesis, University of California, Berkeley, 2014.
[26] S. Chetlur and C. Woolley, arXiv:1410.0759.
[27] V. Giovannetti, S. Lloyd, and L. Maccone, Phys. Rev. Lett. **100**, 160501 (2008).
[28] V. Giovannetti, S. Lloyd, and L. Maccone, Phys. Rev. A **78**, 052310 (2008).
[29] I. Kerenidis and A. Prakash, *Proceedings of 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017).
[30] A. Y. Kitaev, arXiv: 9511026.
[31] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.120.050502 for detailed calculations and a schematic diagram for the data structure.
[32] M. Szegedy, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (IEEE, Bellingham, WA, 2004), p. 32.
[33] A. M. Childs, Commun. Math. Phys. **294**, 581 (2010).
[34] M. Santha, *Theory Applications Models Computation*, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2008), pp. 31–46.
[35] A. W. Harrow, Quantum Algorithms for Systems of Linear Equations, in *Encyclopedia of Algorithms*, edited by M. Y. Kao (Springer, New York, 2016), p. 1680.
[36] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, in *AMS Contemporary Mathematics* (American Mathematical Society, Providence, Rhode Island, 2002).
[37] A. Scherer, B. Valiron, S.-C. Mau, S. Alexander, E. van den Berg, and T. E. Chapuran, Quantum Inf. Process. **16**, 60 (2017).