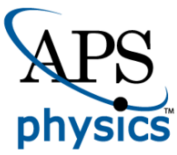

Quantum Algorithms for Systems of Linear Equations Inspired by Adiabatic Quantum Computing

Yiğit Subaşı, Rolando D. Somma, and Davide Orsucci

Phys. Rev. Lett. **122**, 060504 — Published 14 February 2019

DOI: 10.1103/PhysRevLett.122.060504

Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing

Yiğit Subaşı and Rolando D. Somma
Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, US.

Davide Orsucci
Department of Theoretical Physics, University of Innsbruck, Innsbruck, Austria.
(Dated: January 8, 2019)

We present two quantum algorithms based on evolution randomization, a simple variant of adiabatic quantum computing, to prepare a quantum state $|x\rangle$ that is proportional to the solution of the system of linear equations $A\vec{x} = \vec{b}$. The time complexities of our algorithms are $O(\kappa^2 \log(\kappa)/\epsilon)$ and $O(\kappa \log(\kappa)/\epsilon)$, where κ is the condition number of A and ϵ is the precision. Both algorithms are constructed using families of Hamiltonians that are linear combinations of products of A, the projector onto the initial state $|b\rangle$, and single-qubit Pauli operators. The algorithms are conceptually simple and easy to implement. They are not obtained from equivalences between the gate model and adiabatic quantum computing. They do not use phase estimation or variable-time amplitude amplification, and do not require large ancillary systems. We discuss a gate-based implementation via Hamiltonian simulation and prove that our second algorithm is almost optimal in terms of κ. Like previous methods, our techniques yield an exponential quantum speedup under some assumptions. Our results emphasize the role of Hamiltonian-based models of quantum computing for the discovery of important algorithms.

PACS numbers: 03.67.Ac, 03.67.Lx, 03.65.Xp, 89.70.Eg

Introduction. Recently, there has been significant interest in quantum algorithms to solve various linear algebra problems [1–5], as quantum computers can implement certain linear transformations more efficiently than their classical counterparts. Such algorithms may find applications in a wide range of topics, including machine learning [6–8], graph problems [9], solving differential equations [10], and physics problems [11, 12]. A main example is the algorithm of Harrow, Hassidim, and Lloyd (HHL) of Ref. [1] for the so-called quantum linear systems problem (QLSP), where the goal is to prepare a quantum state $|x\rangle$ that is proportional to the solution of a system of linear equations $A\vec{x} = \vec{b}$. If the $N \times N$ matrix A and N-dimensional vector $\vec{b}$ are sparse, and for constant precision, the complexity of the algorithm in Ref. [1] is polynomial in $\log N$ and κ, where κ is the condition number of A. In contrast, classical algorithms to invert matrices are of complexity polynomial in N, suggesting that quantum computers would be able to solve certain problems related to systems of linear equations exponentially faster than classical computers. Improvements of the HHL algorithm can be found in Refs. [3–5].

The referenced algorithms are described in the standard gate-based model of quantum computing, where quantum states are prepared by applying a sequence of elementary (e.g., two-qubit) gates to some initial state. However, Hamiltonian-based alternatives to the gate-based model exist, such as adiabatic quantum computing (AQC) [13]. One advantage of considering these other alternatives is that new and simple quantum algorithms can be found, even if such algorithms will ultimately be implemented using a different but equivalent model.

In AQC, for example, the computation is performed by smoothly changing the parameters of a Hamiltonian that evolves a quantum system. The adiabatic theorem asserts that if the continuously related eigenstates remain non-degenerate and the Hamiltonians change sufficiently slowly, then the evolved state is sufficiently close to the eigenstate of the final Hamiltonian [14]. Such an eigenstate encodes information about the solution to a problem; in our case the final eigenstate would be $|x\rangle$ (or $|\phi\rangle \otimes |x\rangle$ if ancillas are used). A closely related method is the randomization method (RM) described in Ref. [15]. Both, AQC and RM are examples of eigenpath traversal [16]. Nevertheless, an advantage of the RM with respect to AQC is that better convergence guarantees can sometimes be obtained, as shown in Refs. [17, 18].

In this paper, we develop two simple quantum algorithms that solve the QLSP based on the RM. To this end, we construct families of Hamiltonians whose continuously related eigenstates connect $|b\rangle$, the quantum state proportional to $\vec{b}$, with $|x\rangle$. The average evolution times of our algorithms, i.e. the time complexities, are nearly order κ^2 and κ, respectively. Here κ is the condition number of A. Additionally, the time complexities of both algorithms are linear in $1/\epsilon$, where ϵ is a precision parameter. In contrast to previous approaches, our algorithms do not use any form of phase estimation, amplitude amplification, or function approximation, thus reducing the number of ancillary qubits significantly.

Our first quantum algorithm solves the QLSP by preparing the lowest-energy states of a family of Hamil-

tonians, whereas our second algorithm achieves this by preparing energy states that lie exactly at the middle of the spectrum, i.e., excited states. Our second algorithm is noteworthy in that it is almost optimal, having time complexity almost linear in κ.

The Hamiltonians involved in our algorithms are easily described in terms of the inputs of the problem. They may not correspond to a physical model and actual implementations of our algorithms on analog quantum computing devices may be impractical. However, the quantum algorithms could still be efficiently implemented in the gate-based model by using the Hamiltonian simulation results of Refs. [19–21]. This will require oracle access to the matrix A as well as a procedure to prepare the state $|b\rangle$. A resulting gate-model algorithm for the QLSP following this procedure will be nearly optimal according to Ref. [1]. That is, like Refs. [3, 4], the query complexity is almost linear in κ, a quadratic improvement over that of the HHL algorithm. We give more specifics below.

Quantum linear systems problem. The QLSP in Refs. [1, 3, 4] is stated as follows. We are given an $N \times N$ Hermitian matrix A and a vector $\vec{b} = (b_1, \ldots, b_N)^T$, with $N = 2^n$. The goal is to prepare an ϵ-approximation of a quantum state

$$|x\rangle := \frac{\sum_{j=1}^{N} x_j |j\rangle}{\sqrt{\sum_{j=1}^{N} |x_j|^2}} = \frac{(1/A)|b\rangle}{\|(1/A)|b\rangle\|} , \tag{1}$$

where $\vec{x} = (x_1, \ldots, x_N)^T$ is the solution to the linear system $A\vec{x} = \vec{b}$, $|b\rangle \propto \sum_{j=1}^{N} b_j |j\rangle$, and $0 < \epsilon < 1$ is a precision parameter. We assume that A is invertible, having condition number $\kappa < \infty$, and $\|A\| \leq 1$. The approximated state $|\tilde{x}\rangle$ satisfies $\| |\tilde{x}\rangle - |x\rangle \| \leq \epsilon$. Here, we consider a slightly modified version of this problem where the goal is to prepare a mixed state ρ_x such that the trace distance satisfies

$$\frac{1}{2}\mathrm{Tr}\,|\rho_x - |x\rangle\langle x|| \leq \epsilon . \tag{2}$$

Note that this modified version is adequate since the ultimate purpose of the QLSP is for obtaining expectation values of observables. Thus, both $|\tilde{x}\rangle$ and ρ_x will provide same-order approximations for such calculations.

Algorithm evolving on ground states. We first define the family of Hamiltonians

$$H(s) := A(s) P_{\bar{b}}^{\perp} A(s) . \tag{3}$$

Here, $A(s) := (1-s)Z \otimes \mathbb{1} + sX \otimes A$, $|\bar{b}\rangle := |+, b\rangle$, $P_{\bar{b}}^{\perp} := \mathbb{1} - |\bar{b}\rangle\langle\bar{b}|$, and $s \in [0, 1]$ is a parameter. X and Z are single-qubit Pauli operators. These Hamiltonians act on a Hilbert space of dimension $2N$, i.e., the space of A plus one ancilla qubit. The reason for using an ancilla is to guarantee that $A(s)$ is invertible for all s. We introduce the family of states

$$|x(s)\rangle := \frac{1/(A(s))|\bar{b}\rangle}{\|1/(A(s))|\bar{b}\rangle\|} , \tag{4}$$

which satisfy $H(s)|x(s)\rangle = 0$. In Supp. Mat. we show that $|x(s)\rangle$ is the unique ground state of $H(s)$ and the energy gap satisfies $\Delta(s) \geq \Delta^*(s) := (1-s)^2 + (s/\kappa)^2$. As s is increased from 0 to 1, the ground state continuously changes from $|x(0)\rangle = |-, b\rangle$ to $|x(1)\rangle = |+, x\rangle$. Exact preparation of $|x(1)\rangle$ implies exact preparation of the target state $|x\rangle$ by discarding the ancillary qubit.

In the case $A > 0$, we can opt for the simpler choice $A(s) := (1-s)\mathbb{1} + sA$, and still have $A(s)$ non-singular for all s. Then, $|x(s)\rangle \propto A(s)^{-1}|b\rangle$ is the unique ground state of $H(s)$. The following analysis is for general A.

Randomization method. The details of the RM as well as its complexity analysis can be found in Ref. [15]. Here, we mainly study and describe how to use the RM to solve the QLSP. Roughly, the method can be viewed as a version of AQC, where the parameter s is changed discretely rather than continuously, and the Hamiltonian evolution is for a random time. This process effectively simulates an approximate projective measurement of the desired ground state (or any other eigenstate). It then allows to make transformations within the ground states (eigenstates) of the Hamiltonians. The time complexity of the RM in general is $O(L^2/(\epsilon\Delta))$, where L is the so-called path length (which we define later), and Δ is the minimum gap of the Hamiltonians. We observe that the dependence on Δ is optimal [18], while general bounds for AQC provide a worse time complexity of $O(1/\Delta^3)$ [22]. This observation is key to achieve our results. Then, obtaining the actual time complexity for the QLSP requires studying the properties of the Hamiltonians $H(s)$ and eigenstates $|x(s)\rangle$. With this information, we can find discrete values of s as well as values for the evolution times needed to implement the RM.

The full complexity analysis for the QLSP is given in Supp. Mat.. According to Refs. [15, 16, 18], to obtain the discrete values of s, it is convenient to work with a "natural" parametrization $s(v)$. This is defined so that the norm of the rate of change of the eigenstate with respect to v can be bounded by a constant. We find that a natural parametrization for this case is

$$s(v) := \frac{e^{v\frac{\sqrt{1+\kappa^2}}{\sqrt{2}\kappa}} + 2\kappa^2 - \kappa^2 e^{-v\frac{\sqrt{1+\kappa^2}}{\sqrt{2}\kappa}}}{2(1+\kappa^2)} . \tag{5}$$

Here, $v_a \leq v \leq v_b$, with

$$v_a := \frac{\sqrt{2}\kappa}{\sqrt{1+\kappa^2}} \log(\kappa\sqrt{1+\kappa^2} - \kappa^2) , \tag{6}$$

$$v_b := \frac{\sqrt{2}\kappa}{\sqrt{1+\kappa^2}} \log(\sqrt{1+\kappa^2} + 1) . \tag{7}$$

The discrete values $s^j = s(v^j)$ are obtained from discrete values of v, which are evenly distributed in q points as $v_a < v^1 < v^2 < \ldots < v^q = v_b$. Here, $v^j = v_a + j\delta$ $(j = 1, \ldots, q)$ and $s^0 = s(v_a) = 0$, $s^q = s(v_b) = 1$. The number of steps of the RM is $q = \Theta(\log^2(\kappa)/\epsilon)$, and $\delta = (v_b - v_a)/q$. The choice of q implies

$$1 - |\langle x(s^j)|x(s^{j+1})\rangle|^2 = O(\epsilon/q) \ . \tag{8}$$

That is, a sequence of q projective measurements of $|x(s^j)\rangle$, starting from $|x(0)\rangle$, will produce $|x(1)\rangle$ with probability $1 - O(\epsilon)$. These measurements are simulated by evolution randomization.

Our algorithm is as follows. At each step $j = 1, \ldots, q$, we evolve with the Hamiltonian $H(s^j)$ for a random time t^j. The evolution time can be sampled from the uniform distribution $t^j \in [0, 2\pi/\Delta^*(s^j)]$ [15, 18] and satisfies $\langle t^j \rangle = \pi/(\Delta^*(s^j))$. The time complexity of this algorithm is $T := \sum_{j=1}^{q} \langle t^j \rangle$ and in Supp. Mat. we show

$$T = O\left(\kappa^2 \log(\kappa)/\epsilon\right) . \tag{9}$$

Note that, in each run, the overall evolution time is always bounded by $2T$.

Our first algorithm then uses the RM to prepare a mixed state ρ_x that satisfies Eq. (2), after discarding the ancilla. The time complexity is almost quadratic in κ. The pseudocode for the algorithm is shown below.

Algorithm

Given condition number κ and precision ϵ:
– Compute v_a and v_b. Set $q = \Theta(\log^2(\kappa)/\epsilon)$, $\delta = (v_b - v_a)/q$
– For $j = 1, \ldots, q$, let $v^j = v_a + j\delta$, $s^j = s(v^j)$, and t^j be sampled from the uniform distribution $\left[0, 2\pi/\Delta^*(s^j)\right]$
– Apply $e^{-it^q H(s^q)} \ldots e^{-it^1 H(s^1)}$ to $|\bar{b}\rangle$, discard the ancilla

Spectral gap amplification. One way to improve the time complexity of the first algorithm is by considering other families of Hamiltonians where the relevant spectral gap is larger than that of $H(s)$. This idea was considered in Ref. [23] and resulted in various polynomial quantum speedups for quantum state preparation. A quadratic spectral gap amplification is indeed possible when the Hamiltonians satisfy a so-called frustration free property. Very roughly, a possible Hamiltonian with an amplified gap can be interpreted as the square root of the frustration-free Hamiltonian. A zero eigenvalue remains zero and an eigenvalue $\lambda > 0$ is transformed into eigenvalues $\pm\sqrt{\lambda}$. ($\sqrt{\lambda} \gg \lambda$ if $\lambda \ll 1$.) To avoid additional complexity overheads, the Hamiltonians with an amplified gap must satisfy certain constrains related to the difficulty of their simulation. We refer to [23] for details.

Motivated by these results, we now consider another family of Hamiltonians for solving the QLSP using the RM. This family is given by

$$H'(s) := \sigma^+ \otimes A(s) P_{\bar{b}}^{\perp} + \sigma^- \otimes P_{\bar{b}}^{\perp} A(s) \ , \tag{10}$$

where $\sigma^{\pm} = (X \pm iY)/2$ are single-qubit (raising and lowering) operators, and $s \in [0, 1]$. We note that $H'(s)$ acts on a Hilbert space of dimension $4N$. Then

$$(H'(s))^2 = \begin{pmatrix} H(s) & 0 \\ 0 & P_{\bar{b}}^{\perp}(A(s))^2 P_{\bar{b}}^{\perp} \end{pmatrix} \ , \tag{11}$$

where each block of the matrix is of dimension $2N \times 2N$. Using $B(s) := A(s)P_{\bar{b}}^{\perp}$, the blocks on the diagonal of Eq. (11) can be written as $B(s)^\dagger B(s)$ and $B(s)B(s)^\dagger$, and thus have the same spectrum. Consequently, the eigenvalues of $H'(s)$ are $\{0, 0, \pm\sqrt{\gamma_1(s)}, \ldots, \pm\sqrt{\gamma_{2N-1}(s)}\}$, where $\gamma_j(s) > 0$ are the nonzero eigenvalues of $H(s)$. The subspace of $H'(s)$ of eigenvalue zero is spanned by $\{|0\rangle \otimes |x(s)\rangle, |1\rangle \otimes |\bar{b}\rangle\}$.

In contrast to the first algorithm that aimed at preparing the ground state of $H(s)$, we now aim at preparing one of the two eigenstates of zero eigenvalue of $H'(s)$ that lies exactly at the middle of the spectrum. Nevertheless, the RM can be used to prepare any eigenstate as long as it is separated by a nonzero spectral gap from the other eigenstates. One may wonder if the double degeneracy is a problem for this case. The answer is negative as the Hamiltonian does not allow for transitions between the two eigenstates, that is, $\langle 0| \otimes \langle x(s)|H'(s')\,|1\rangle \otimes |\bar{b}\rangle = 0$. If we initialize our quantum computer in $|0\rangle \otimes |x(0)\rangle$, a sequence of perfect projective measurements of the eigenstates of $H'(s)$ at sufficiently close points will allow us to prepare $|0\rangle \otimes |x(1)\rangle$ with sufficiently high probability. The relevant spectral gap is now bounded by $\sqrt{\Delta^*(s)} > 0$.

The eigenstate $|0\rangle \otimes |x(s)\rangle$ has similar properties as $|x(s)\rangle$: the path length and norm of the rate of change are the same. Then, our second algorithm can be constructed by using the same discretization points s^j that were used for the first algorithm. At each step, we now need to evolve with the Hamiltonian $H'(s^j)$ for a random time t^j. This time can be sampled from the uniform distribution $t^j \in [0, 2\pi/\sqrt{\Delta^*(s^j)}]$. The time complexity of this algorithm is $T := \sum_{j=1}^{q} \langle t^j \rangle$ and, in Supp. Mat., we show

$$T = O\left(\kappa \log(\kappa)/\epsilon\right) \ . \tag{12}$$

After discarding the two ancilla qubits, the final state is ρ_x and satisfies Eq. (2). The time complexity of our second algorithm is then almost linear in κ. The pseudocode for this algorithm follows from the previous one by replacing $\Delta^*(s^j)$ with $\sqrt{\Delta^*(s^j)}$, $H(s)$ with $H'(s)$, and $|\bar{b}\rangle$ with $|0\rangle \otimes |\bar{b}\rangle$, in the second and third lines.

Simulation results. We tested the validity of our quantum algorithms by performing numerical simulations. For this purpose, we randomly generated Hermitian matrices A of dimension $N = 16$ that are 4-sparse

and $N = 32$ that are 5-sparse, both satisfying $\|A\| = 1$. The generated matrices result in a range of values for the condition number κ. We post-selected matrices for which $\kappa \approx 10$ and $\kappa \approx 50$ (to within absolute error 10^{-3}), for $N = 16$ and $N = 32$ respectively. Similarly, we randomly generated 4-sparse and 5-sparse vectors for $\vec{b}$. The parameters s^j and t^j were chosen according to the previous discussion and depend on κ and ϵ (or q). In each execution, we prepare a pure quantum state that is not guaranteed to be ϵ-close to the pure eigenstate of the final Hamiltonian. However, the expected error of the prepared pure states from many repeated executions of the algorithms is indeed bounded by ϵ.

We ran simulations for which the number of repetitions of our algorithms were $n_{rep} = 50$ and $n_{rep} = 200$, respectively. For each case, we first construct a finite-sampling density matrix $(1/n_{rep}) \sum_{i=1}^{n_{rep}} |\psi_i\rangle\langle\psi_i|$. Here, $|\psi_i\rangle$ is the pure state output at the i'th repetition. Tracing out the ancilla qubits, we get a density matrix $\tilde{\rho}_x$ that describes the state of the system only. Note that $\tilde{\rho}_x$ is, in general, slightly different from ρ_x of Eq. (2). However, $\tilde{\rho}_x \to \rho_x$ in the limit of $n_{rep} \to \infty$. The error computed in our numerical simulations is then the trace distance between $\tilde{\rho}_x$ and $|x\rangle\langle x|$.

In Fig. 1, we show the dependence of the inverse of the error on the number of steps q. While the results are for two particular matrices A with $\kappa \approx 10$ and $\kappa \approx 50$, other matrices show similar results. We observe that the inverse of the error for the two quantum algorithms, denoted by ϵ_Q and ϵ_L respectively, scales almost linearly with q. The dispersion around the linear fit is smaller for larger n_{rep}. The results are then in accordance with our theoretical analysis.

Gate-based implementations. Our algorithms are based on Hamiltonian evolutions and can be implemented on a gate-based quantum computer using a Hamiltonian simulation method. We focus on the method of Ref. [20], which implements the truncated Taylor series of the evolution operator. It requires the Hamiltonian to be given as a linear combination $\sum_l \alpha_l V_l$, where the V_l are unitaries that are easy to apply and $\alpha_l > 0$. The V_l are applied $\tilde{O}(\tau)$ times, where τ is the product of the evolution time and $\sum_l \alpha_l$. The $\tilde{O}$ notation hides logarithmic factors in τ.

Our second algorithm applies the evolution under $H'(s^j)$ for time t^j. The main challenge is then to find a decomposition of the Hamiltonian in terms of unitaries. For technical reasons, we consider another Hamiltonian $\tilde{H}'(s^j)$, but whose evolution operator mimics that of $H'(s^j)$. This Hamiltonian is discussed in Supp. Mat.. It turns out that $\tilde{H}'(s) = \frac{d+1}{16} \sum_{l=1}^{32} V_l(s)$, where $V_l(s)$ are unitaries. As previous approaches for the QLSP [4], we assume access to a quantum oracle $\mathcal{O}_A$ for the matrix A. This oracle outputs the nonzero matrix elements and their indices, for any row of A. We also assume access to a (controlled) unitary U_b that prepares the state $|b\rangle$ and

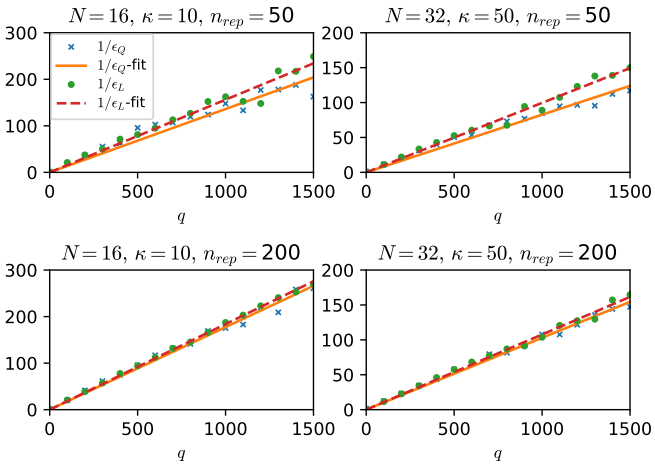

FIG. 1. The inverse of the error for the two quantum algorithms as a function of q, the number of steps in the RM. Subscript Q refers to the quantum algorithm with complexity that is almost quadratic in κ and L to the quantum algorithm with complexity almost linear in κ. n_{rep} is the number of repetitions of the of our algorithm. The results are for two randomly generated matrices A with $N = 16$, $\kappa \approx 10$, and $N = 32$, $\kappa \approx 50$.

the (controlled) $U_b^\dagger$, as in Refs. [1, 3, 4]. Each unitary $V_l(s)$ can be applied using, at most, a constant number of $\mathcal{O}_A$ and (controlled) U_b and $U_b^\dagger$. In addition, it may require $O(n)$ two-qubit gates – see Supp. Mat..

In our construction, we have $\tau = O(t^j d)$ if the evolution time is t^j. Since our algorithm implements evolutions with q Hamiltonians, the total number of uses of $\mathcal{O}_A$ and (controlled) U_b and $U_b^\dagger$, or query complexity, is then $\tilde{O}(Td)$, where T is the total evolution time. The number of additional two-qubit gates is a multiplicative factor of order n away from the query complexity.

Substituting T from Eq. (12) gives the query complexity of our approach as $\tilde{O}(\kappa d/\epsilon)$. In Ref. [1], it was shown that quantum algorithms for the QLSP must have a query complexity that is, at least, linear in κ. Then, the gate-based implementation following Ref. [20] is almost optimal. Note that the query complexity of evolving with $\tilde{H}'(s)$ is of the same order as that of evolving directly with A, which is needed for the HHL algorithm.

Discussion. We presented simple quantum algorithms for solving the QLSP that were motivated by AQC and not the standard gate-based model. A nice feature about AQC and related models, such as the RM or general eigenpath traversal methods [16], is that the time complexity is typically dominated by a single quantity, i.e., the inverse of the minimum spectral gap of the Hamiltonians. Then, the root of the quantum speedup is more clear in this case than in the gate-based model, allowing for algorithmic improvements by considering different Hamiltonians with larger spectral gaps. Another

nice feature is that some problems are naturally reduced to preparing the eigenstate of a Hamiltonian, and eigenpath traversal methods are useful in that context. We showed that this is the case for the QLSP. Our results emphasize the importance of considering models of quantum computing, which go beyond the gate-based model, for discovering novel quantum algorithms – see Ref. [24] for another example.

The further significance of our results is as follows. Previous algorithms for the QLSP [1, 3, 4] use three main subroutines: (i) Hamiltonian simulation, (ii) phase estimation or function approximation, and (iii) some form of amplitude amplification. The method of "variable-time amplitude amplification" is used in Refs. [3, 4] to achieve near-optimal complexity in terms of κ. That method alone requires $\Omega(\log(1/\epsilon)\log(\kappa/\epsilon)/\epsilon^2)$ and $\Omega(\log(\kappa)\log(\kappa/\epsilon))$ ancillary qubits, respectively, which become excessively large for large κ. In contrast, our algorithms use only Hamiltonian simulation (which has the same query complexity as in previous methods) thereby reducing the number of ancillary qubits significantly. Our result additionally implies a significant reduction in the number of conditional operations to solve the QLSP, making our algorithms more attractive for implementations on quantum computers of smaller size. To this point, our algorithm has already been used in Ref. [25] to solve an 8-dimensional linear system on a 4-qubit NMR device, the largest dimension up to date.

The time complexity of our methods is linear in $1/\epsilon$. This complexity can be improved to polylogarithmic in $1/\epsilon$ using the fast methods for eigenpath traversal of Ref. [16]. These methods will provide a different way of obtaining an exponential improvement in terms of precision with respect to the HHL algorithm, as in Ref. [4]. They, however, require repeated uses of phase estimation and thus many additional ancillary qubits.

Last, it would be interesting to study if our results can also impact classical methods for solving systems of linear equations.

Acknowledgements. We thank Anirban Chowdhury for discussions. Part of this material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams program. We also thank the Laboratory Directed Research and Development Program at LANL for support. DO acknowledges the support by the Austrian Science Fund (FWF) through the DK-ALM W1259-N27, the SFB FoQuS F4012, and the Templeton World Charity Foundation grant TWCF0078/AB46.

[1] A. W. Harrow, A. Hassidim, and S. Lloyd, Phys. Rev. Lett. **103**, 150502 (2009).
[2] A. Prakash, *Quantum Algorithms for Linear Algebra and Machine Learning*, Ph.D. thesis, University of California, Berkeley (2008).
[3] A. Ambainis, in *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science* (2012) pp. 636–647.
[4] A. Childs, R. Kothari, and R. D. Somma, SIAM J. Comp. **46**, 1920 (2017).
[5] S. Chakraborty, A. Gilyén, and S. Jeffery, arXiv:1804.01973 (2018).
[6] N. Wiebe, D. Braun, and S. Lloyd, Phys. Rev. Lett. **109**, 050505 (2012).
[7] S. Lloyd, M. Mohseni, and P. Rebentrost, arXiv:1307.0411 (2013).
[8] S. Lloyd, M. Mohseni, and P. Rebentrost, Nature Physics **10(9)**, 631 (2014).
[9] G. Wang, Quantum Information & Computation **17**, 987 (2017).
[10] D. W. Berry, J. Phys. A: Math. Theor. **47**, 105301 (2014).
[11] B. D. Clader, B. C. Jacobs, and C. R. Sprouse, Phys. Rev. Lett. **110**, 250504 (2013).
[12] A. Chowdhury and R. D. Somma, Quant. Inf. Comp. **17**, 0041 (2016).
[13] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, arXiv:quant-ph/0001106 (2000).
[14] A. Messiah, *Quantum Mechanics* (Dover Publications, 1999).
[15] S. Boixo, E. Knill, and R. Somma, Q. Inf. Comp. **9**, 0833 (2009).
[16] S. Boixo, E. Knill, and R. Somma, arXiv:1005.3034 (2010).
[17] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, Phys. Rev. Lett. **101**, 130504 (2008).
[18] H.-T. Chiang, G. Xu, and R. D. Somma, Phys. Rev. A **89**, 012314 (2014).
[19] D. Berry, G. Ahokas, R. Cleve, and B. Sanders, Comm. Math. Phys. **270**, 359 (2007).
[20] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Phys. Rev. Lett. **114**, 090502 (2015).
[21] G. Low and I. Chuang, Phys. Rev. Lett. **118**, 010501 (2017).
[22] S. Jansen, M. Ruskai, and R. Seiler, J. Math. Phys. **48**, 102111 (2007).
[23] R. D. Somma and S. Boixo, SIAM J. Comp **42**, 593 (2013).
[24] R. Somma, D. Nagaj, and M. Kieferova, Phys. Rev. Lett. **109**, 050501 (2012).
[25] J. Wen, X. Kong, S. Wei, T. Xin, B. Wang, K. Li, Y. Zhu, and G. Long, arXiv:1806.03295 (2018).