

A Quantum Computer Amenable Sparse Matrix Equation Solver

by

Christopher D. Phillips

A Thesis submitted to the Faculty of Graduate Studies of
The University of Manitoba
in partial fulfilment of the requirements of the degree of

MASTER OF SCIENCE

Department of Electrical and Computer Engineering
University of Manitoba
Winnipeg

Copyright ©2022 by Christopher D. Phillips

Abstract

Quantum computation offers a promising alternative to classical computing methods in many areas of numerical science, with algorithms that make use of the unique way in which quantum computers store and manipulate data often achieving dramatic improvements in performance over their classical counterparts. The potential efficiency of quantum computers is particularly important for numerical simulations, where the capabilities of classical computing systems are often insufficient for the analysis of real-world problems. In this work, we study problems involving the solution of matrix equations, for which there currently exists no efficient, general quantum procedure. We develop a generalization of the Harrow/Hassidim/Lloyd algorithm by providing an alternative unitary for eigenphase estimation. This unitary, which we have adopted from research in the area of quantum walks, has the advantage of being well defined for any arbitrary matrix equation, thereby allowing the solution procedure to be directly implemented on quantum hardware for any well-conditioned system. The procedure is most useful for sparse matrix equations, as it allows for the inverse of a matrix to be applied with $\mathcal{O}(N_{nz} \log(N))$ complexity, where N is the number of unknowns, and N_{nz} is the total number of nonzero elements in the system matrix. This efficiency is independent of the matrix structure, and hence the quantum procedure can outperform classical methods for many common system types. We show this using the example of sparse approximate inverse (SPAI) preconditioning, which involves the application of matrix inverses for matrices with $N_{nz} = \mathcal{O}(N)$. While these matrices are indeed sparse, it is often found that their inverses are quite dense, and classical methods can require as much as $\mathcal{O}(N^3)$ time to apply an inverse preconditioner.

Acknowledgements

I am grateful to my advisor, Vladimir Okhmatovski, for his patience and invaluable feedback throughout the development of this work. Additionally, I would like to thank my examination committee for taking the time to provide their knowledge and expertise.

I acknowledge the use of IBM Quantum services for this work. The views expressed are those of the author, and do not reflect the official policy or position of IBM or the IBM Quantum team.

Contributions of Authors

The paper which this thesis comprises was a joint effort between myself and my advisor, Dr. Vladimir Okhmatovski. The majority of the work involved was contributed by myself, including nearly all theoretical development of the proposed method, writing of the program used to evaluate the procedure, and the production of all data presented. Dr. Okhmatovski's role was chiefly advisory, with some additional theoretical contributions in the form of discussions at key points in the development.

Contents

Introduction to Thesis 8

References 10

Paper 1: A Quantum Computer Amenable Sparse Matrix Equation Solver 12

 Introduction 12

 The HHL Algorithm 14

 Initialization 15

 Quantum Phase Estimation 15

 Ancilla Rotation and Uncomputation 19

 The Walk Operator 20

 Considerations for Practical Implementation 26

 Treatment of Arbitrary Matrix Equations 30

 Summary 32

 Initial Application of T_0 32

 Quantum Phase Estimation 34

 HHL Ancilla Rotation 36

 Uncomputation 37

 Results 38

 A Fully Quantum Problem 39

 Simulated Solution 42

 Preconditioner Application 42

 Conclusion 45

References 48

Conclusion to Thesis 51

List of Tables

1 Comparison of solution vectors 42

List of Figures

1	Diagram of quantum phase estimation	17
2	Illustration of eigenphase estimation	18
3	Control procedure summary	29
4	Structural outline of full procedure	33
5	Initial application of T_0	34
6	State preparation operator B_j	35
7	Quantum phase estimation using the walk operator	36
8	The walk operator	37
9	HHL rotation	39
10	Gate counts and errors for quantum procedure	41
11	Solved charge distribution	43
12	Classically solved charge distribution	44
13	Nonzero patterns in preconditioner	45
14	Classical preconditioner application time	46
15	Quantum preconditioner application complexity	47
16	Projected results	47
A1	Two-dimensional reflection about the y -axis.	52

Introduction to Thesis

This work is presented in the grouped manuscript style, and comprises a single paper bearing the same title as this thesis. The constituent paper has been submitted to the *IEEE Transactions on Quantum Engineering*, and is currently under review. A preprint version has also been submitted to arXiv (Phillips & Okhmatovski, 2021).

Numerical simulations of physical systems are important tools in research and technological development, as they provide a robust, theoretically transparent framework for the analysis of novel theories and processes. However, the numerical scientist frequently encounters problems which require far greater computational power than can be supplied by available systems. While the development of classical computing systems is an active and highly successful field, the capabilities of these systems continue to struggle in meeting even the basic needs of the computational community. This is one of the chief problems which the field of quantum computing hopes to address, and is the motivation for our current work. With the early successes of quantum algorithms (Grover, 1996; Gerjuoy, 2005; Childs, 2010), it was proven that, at least for certain limited problems, quantum computers could offer dramatic improvements in computational power over classical systems. The question of whether this efficiency can be translated to the analysis of more general problems is then an interesting point of study for computational research, as such reductions in resource requirements would allow for much easier analysis of more practical, complex systems.

Here we consider the application of quantum techniques to the analysis of matrix equations, which are highly relevant in numerical analysis (Chandrupatla & Belegundu, 2012; Gibson, 2008; Özişik, Orlande, Colaço, & Cotta, 2017). This foundation for this work is the Harrow/Hassidim/Lloyd (HHL) algorithm (Harrow, Hassidim, & Lloyd, 2009), which is a well known quantum procedure for the solution of matrix equations. The basic HHL procedure claims a complexity of $\mathcal{O}(\kappa^2 \log(N)/\epsilon)$ to produce a state approximating the solution vector of a given system, where N is the system size, κ is the condition number of the system matrix, and ϵ is the distance between the computed solution vector and the exact solution. This efficiency can be compared with some of the most efficient known classical methods (Chew, Jin, Michielssen, & Song, 2001; Chen, Zheng, & Okhmatovski, 2019; Omar & Jiao, 2015), which offer at best $\mathcal{O}(N)$ or $\mathcal{O}(N \text{polylog}(N))$ scaling, depending on the structure of the problem under analysis. If κ and $1/\epsilon$ scale as $\text{polylog}(N)$, the HHL procedure can then offer an exponential speedup over classical procedures. This is a dramatic result, but unfortunately, it is as yet unrealizable in real world applications. This is due to the statement of the canonical algorithm, which depends upon the efficient implementation of a quantum procedure for Hamiltonian simulation. The original HHL algorithm recommends the procedure of (Childs, 2010; Berry & Childs, 2012), which upon further study is found to depend on a black box model to maintain its complexity. While this does

allow for the query complexity of the HHL procedure to be maintained, practical implementation of the oracle operators inherent to the black box model is nontrivial, and in fact has a significant impact on the overall complexity scaling of the procedure.

This work presents an efficient prescription for implementing the method of (Childs, 2010; Berry & Childs, 2012). This method uses an operator adopted from quantum walk research (Ambainis, 2004; Szegedy, 2004) as a unitary operator for the purposes of quantum phase estimation. The method was originally presented as a technique for Hamiltonian simulation, but has been adapted here for direct application to the HHL procedure. With sufficient considerations for sparse matrices, the complexity of the procedure, in both gate complexity and circuit generation time, is $\mathcal{O}(N_{nz} \log(N))$, where N_{nz} is the total number of nonzeros in the system matrix. The procedure has no dependence on matrix structure, other than the condition that it must be nonsingular, and hence it is capable of outperforming classical solvers when applied to general matrices. To demonstrate this, the example of sparse approximate inverse (SPAI) preconditioning is used (Pan & Sheng, 2014; Takahashi, Coulier, & Darve, 2017; Casacuberta & Kyng, 2021). This method involves the application of the inverse of a matrix with roughly $\mathcal{O}(N)$ nonzero elements to an operand vector. These preconditioning matrices exhibit little reliable structure, and classical methods such as GMRES can require as much as $\mathcal{O}(N^3)$ time to apply the inverse preconditioner. This quantum procedure, however, can apply such matrices in $\mathcal{O}(N \log(N))$ time.

Several improvements over, and alternatives to, the basic HHL algorithm have been proposed (Childs, Kothari, & Somma, 2017; An & Lin, 2022; Lin & Tong, 2020; Costa et al., 2021). These works provide methods for improving the procedure's scaling with κ and ϵ . However, all of these methods are similar to the original HHL procedure in that they require oracles that are impossible to implement in practice. While oracle-based methods are important for the purposes of query-complexity analysis, it is the goal of this work to study an explicit, practically implementable procedure for the purposes of assessing the real-world viability of quantum matrix equation analysis. Note that a constant precision is assumed, as well as a well-conditioned system matrix, such that both ϵ and κ can be omitted from our complexity analysis. Adaptation of one of the above-mentioned works to produce better scaling in ϵ and κ is an interesting route for further study, but as this is a first work along this research direction, this topic has been omitted in favour of a more tractable base procedure.

References

- Ambainis, A. (2004, October). Quantum walk algorithm for element distinctness. In *Proc. 45th annu. IEEE symp. on found. of comput. sci.* (pp. 22–31). doi: 10.1109/FOCS.2004.54
- An, D., & Lin, L. (2022, mar). Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm. *ACM Transactions on Quantum Computing*, 3(2). doi: 10.1145/3498331
- Berry, D. W., & Childs, A. M. (2012, January). Black-box hamiltonian simulation and unitary implementation. *Quantum Info. Comput.*, 12, 29–62.
- Casacuberta, S., & Kyng, R. (2021). *Faster sparse matrix inversion and rank computation in finite fields.* (arXiv:2106.09830)
- Chandrupatla, T. R., & Belegundu, A. D. (2012). *Introduction to finite elements in engineering* (4th ed.). Harlow, England: Pearson.
- Chen, Z., Zheng, S., & Okhmatovski, V. I. (2019, June). Tensor train accelerated solution of volume integral equation for 2-d scattering problems and magneto-quasi-static characterization of multiconductor transmission lines. *IEEE Trans. Microw. Theory Tech.*, 67, 2181–2196. doi: 10.1109/TMTT.2019.2908368
- Chew, W. C., Jin, J. M., Michielssen, E., & Song, J. (2001). *Fast and efficient algorithms in computational electromagnetics* (1st ed.). Norwood, MA, USA: Artech House.
- Childs, A. M. (2010, March). On the relationship between continuous- and discrete-time quantum walk. *Commun. Math. Phys.*, 294, 581–603. doi: 10.1007/s00220-009-0930-1
- Childs, A. M., Kothari, R., & Somma, R. D. (2017). Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6), 1920–1950. doi: 10.1137/16M1087072
- Costa, P. C. S., An, D., Sanders, Y. R., Su, Y., Babbush, R., & Berry, D. W. (2021). *Optimal scaling quantum linear systems solver via discrete adiabatic theorem.* (arXiv:2111.08152)
- Gerjuoy, E. (2005, May). Shor’s factoring algorithm and modern cryptography. an illustration of the capabilities inherent in quantum computers. *Amer. J. Phys.*, 73, 521–540. doi: 10.1119/1.1891170
- Gibson, W. C. (2008). *The method of moments in electromagnetics* (1st ed.). Boca Raton, FL, UAS: Chapman Hall/CRC.
- Grover, L. (1996, July). A fast quantum mechanical algorithm for database search. In *Proc. 28th annu. acm symp. theory comput.* (pp. 212–219). doi: 10.1109/TMTT.2019.2908368
- Harrow, A. W., Hassidim, A., & Lloyd, S. (2009, October). Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103, 150502. doi: 10.1103/PhysRevLett.103.150502
- Lin, L., & Tong, Y. (2020, November). Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems. *Quantum*, 4(1). doi:

10.22331/q-2020-11-11-361

Omar, S., & Jiao, D. (2015, March). A linear complexity direct volume integral equation solver for full-wave 3-d circuit extraction in inhomogeneous materials. *IEEE Trans. Microw. Theory Techn.*, *63*, 897–912.

Özişik, M. N., Orlande, H. R. B., Colaço, M. J., & Cotta, R. M. (2017). *Finite difference methods in heat transfer* (2nd ed.). Boca Raton, FL, USA: CRC Press.

Pan, X., & Sheng, X. (2014, November). Sparse approximate inverse preconditioner for multiscale dynamic electromagnetic problems. *Radio Science*, *49*, 1041–1051. doi: 10.1002/2014RS005387

Phillips, C. D., & Okhmatovski, V. I. (2021). *A quantum computer amenable sparse matrix equation solver*. (arXiv:2112.02600)

Szegedy, M. (2004, October). Quantum speed-up of markov chain based algorithms. In *Proc. 45th annu. IEEE symp. on found. of comput. sci.* (pp. 32–41). doi: 10.1109/FOCS.2004.53

Takahashi, T., Coulier, P., & Darve, E. (2017, July). Application of the inverse fast multipole method as a preconditioner in a 3d helmholtz boundary element method. *J. Comp. Phys.*, *341*, 406–428.

Paper 1: A Quantum Computer Amenable Sparse Matrix Equation Solver

Introduction

Matrix equation formulations are powerful tools in the numerical analysis of physical systems, and their usage is widespread across many areas of engineering (Chandrupatla & Belegundu, 2012; Gibson, 2008; Özişik et al., 2017). Unfortunately, the solution of practical matrix equations often demands an exorbitant supply of computational resources, with many present-day simulation problems already stressing the limits of available computing systems. For a matrix equation involving N unknowns, the most sophisticated algorithms of the current day provide $\mathcal{O}(N)$ or $\mathcal{O}(N \text{polylog}N)$ runtime and memory complexity (Chew et al., 2001; Chen et al., 2019; Omar & Jiao, 2015), and are heavily reliant on the structure of very specific problem classes. Thus, as simulations naturally grow more complex with time, the requisite computational power must grow at least proportionately.

For such computationally intensive problems, quantum computers are growing in popularity as a potential remedy, as they are capable of leveraging procedures of dramatically reduced complexity relative to their classical counterparts. Some of these procedures, such as Shor’s algorithm (Gerjuoy, 2005), promise an impressive exponential speedup over classical procedures, while others, such as Grover’s algorithm (Grover, 1996), are capable of producing a less dramatic, but still useful polynomial speedup. This is chiefly due to the manner in which quantum computers store and operate on data. Values of interest are stored as the amplitudes of possible states for a register of qubits, meaning that a 1-qubit system (using a 2-state qubit) can store two values using the amplitudes of the $|0\rangle$ and $|1\rangle$ states; a 2-qubit system can store four values using the $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$ states; and so on. In general, the storage capacity of a quantum computer scales exponentially with the size of the system. Quantum operators then act directly on the system’s qubits, allowing the effects of operators to also scale exponentially with system size. On the topic of matrix equation solvers, the Harrow/Hassidim/Lloyd (HHL) algorithm (Harrow et al., 2009) has become well established as the leading quantum algorithm. For a system of condition number κ , this algorithm is capable of providing computational complexity $\mathcal{O}(\kappa^2 \log(N)/\epsilon)$ in the calculation of a solution of precision ϵ . Thus, when κ and $1/\epsilon$ scale logarithmically with N , the procedure offers an exponential speedup over classical algorithms. A detailed description of the HHL algorithm is given in section “The HHL Algorithm”, but we provide a brief overview of the procedure here for the sake of motivation. In what follows, note that the notation $|v\rangle$ indicates a vector \bar{v} with its components encoded onto the basis states of a quantum system as $|v\rangle = \sum_{j=0}^{N-1} v_j |j\rangle$, where \bar{v} has N components. We typically suppress normalization factors where their meaning is irrelevant. For simplicity, $|v\rangle$ is understood to also refer to the vector \bar{v} itself. We begin with a matrix equation $A|x\rangle = |b\rangle$, where A is an $N \times N$

matrix, and $|x\rangle$ and $|b\rangle$ are N -dimensional vectors. If a quantum system is initialized to the state $|b\rangle$, then its state will be a superposition of the eigenvectors of A :

$$|b\rangle = \sum_{j=0}^{N-1} \beta_j |u_j\rangle, \quad (1)$$

where $|u_j\rangle$ is an eigenvector of A , and β_j is the component of $|b\rangle$ along $|u_j\rangle$. Note that this statement assumes that A is diagonalizable. Several such assumptions are made by the HHL algorithm, but all necessary properties can be satisfied for general nonsingular matrices through a straightforward preparation procedure. Such system preparation is discussed in depth in section “Treatment of Arbitrary Matrix Equations”. The solution vector $|x\rangle$ can then be obtained by applying the inverse of each eigenvalue λ_j of A to its corresponding eigenvector in the superposition $|b\rangle$:

$$\sum_{j=0}^{N-1} \frac{\beta_j}{\lambda_j} |u_j\rangle = A^{-1} |b\rangle = |x\rangle. \quad (2)$$

This application of inverse eigenvalues, conditioned on the presence of the corresponding eigenvectors, is the basic function of HHL. By using quantum phase estimation (QPE) (Cleve, Ekert, Macchiavello, & Mosca, 1998), it is possible to extract and apply these eigenvalue factors in an efficient manner. However, this phase estimation requires the implementation of a problem-dependent unitary in terms of basic quantum gates. Generally the unitary e^{iA} is employed for this purpose, but the process of decomposing this unitary into a sequence of basic quantum gates is nontrivial. Therefore the process of implementing this operation on real quantum hardware presents a bottleneck for the procedure. In this work we consider a particular method, inspired by research into quantum walks, which produces a unitary with a known decomposition into basic gates.

Quantum walks are a popular area of study in quantum computing, as they have shown success in aiding the construction of efficient algorithms (Venegas-Andraca, 2012). Reference (Kempe, 2004) gives an excellent introduction to the concept of quantum walks. In essence, they are stochastic processes that, when evolved and measured, give information about the system in which they have evolved. The structure of the evolution and the properties measured can be adjusted to change the information obtained. We are interested in the walk procedure developed in (Ambainis, 2004), which considers the problem of element-distinctness. This problem is quite far removed from our current study, but it is important for the fact that it introduces a particular walk operator based on Grover diffusion operators. Following this research, (Szegedy, 2004) further investigated diffusion-based quantum walks defined by probabilistic maps, and considered the spectra of the walk operators. Finally, (Childs, 2010) expanded on these findings from the

perspective of Hamiltonian simulation, and showed that the walk operators could be defined from a Hamiltonian to produce an operator with a spectrum closely related to that of the Hamiltonian itself. Reference (Berry & Childs, 2012) further explored methods by which these walk operators could be implemented in practice.

The method developed in (Childs, 2010) uses QPE to extract the eigenvalues of the walk operator, which can then be related to the eigenvalues of the Hamiltonian. This is the critical component from which we can develop a general procedure for solving matrix equations. By treating a Hermitian system matrix as the Hamiltonian, we can apply QPE using the walk operator to extract the eigenvalues of the system matrix. From there, the HHL algorithm can proceed as usual. An algorithm presented in (Kerenidis & Prakash, 2020) also makes use of the relationship between quantum walk operators and the eigenvalues of the system matrix (in particular, this work considers singular values) to produce an iterative quantum matrix solver. Our procedure differs in that it provides a direct solution method, is more easily generalizable, and does not require explicit computation of any singular values of the system matrix.

The matrix solution procedure we present here is particularly efficient for cases involving sparse matrices, wherein it can easily be seen to outperform classical solvers. This is due to the $\mathcal{O}(N_{nz} \log(N))$ gate complexity of the associated solution circuit, where N_{nz} is the total number of nonzero elements in the system matrix. When system matrices exhibit little reliable structure, it is often impossible for classical procedures to provide performance better than $\mathcal{O}(N^3)$, regardless of the matrix sparsities. A key example of this is the technique of sparse approximate inverse (SPAI) preconditioning (Pan & Sheng, 2014; Takahashi et al., 2017; Casacuberta & Kyng, 2021). This technique typically generates a preconditioning matrix with $\mathcal{O}(N)$ nonzero elements, with a sparsity pattern that can be quite arbitrary, particularly for systems describing complex geometries. This lack of matrix structure makes the development of efficient, general classical solvers extremely difficult, and often impossible. However, our quantum method will always be capable of providing an inversion circuit having $\mathcal{O}(N \log(N))$ complexity (so long as the matrix is well-conditioned).

Using the Qiskit SDK (ANIS et al., 2021), we have developed a program (Phillips & Okhmatovski, 2021) which closely follows the herein described procedure. In the interest of expositional clarity, as well as to aid those seeking to investigate the program itself, we at times make reference to the specific functions of this program. At these points, the program is referred to simply as “the program”.

The HHL Algorithm

Here we describe the basic HHL algorithm which serves as a foundation for our matrix solution procedure. For brevity, we study only a primitive form of the HHL algorithm. More sophisticated

formulations can be used to improve performance by, for example, ignoring ill-conditioned portions of the system matrix.

Initialization

Given a matrix $A \in \mathbb{C}^{N \times N}$ and a vector $|b\rangle \in \mathbb{C}^N$, the HHL algorithm is a quantum procedure to determine $|x\rangle$ such that

$$A|x\rangle = |b\rangle. \quad (3)$$

Due to the nature of quantum computation, several restrictive properties are required of this system. It is important to note, however, that all of these properties can be satisfied for any input system by means of a simple preparation procedure. Section “Treatment of Arbitrary Matrix Equations” considers such preparation in depth. The restrictions are as follows: The vector $|b\rangle$ is to be encoded onto a quantum system, and as such we require that N be a power of two and that $|b\rangle$ be normalized. Additionally, we require that A is Hermitian. This ensures that its eigenvectors form a basis for \mathbb{C}^N (Horn & Johnson, 2013), and also aids us in the definition of the QPE unitary. The first step of the HHL procedure is the preparation of three registers:

$$|b\rangle |0\rangle^{\otimes n_p} |0\rangle. \quad (4)$$

The first register, which we refer to as the vector register, is initialized to $|b\rangle$, and hence it requires $n = \log_2 N$ qubits. The second register is an n_p -qubit phase register that is used to store the phase estimates produced by QPE. The value of n_p is chosen by the user, and it should be large enough to admit sufficient resolution in the phase estimation. The final register is termed the ancilla register, and it consists of a single qubit. Once the eigenvalues of A have been estimated, the state of this register can be rotated to effect the application of the inverse eigenvalue factors.

The initial system state can be expressed in terms of the eigenvectors of A as

$$|b\rangle |0\rangle^{\otimes n_p} |0\rangle = \sum_{j=0}^{N-1} \beta_j |u_j\rangle |0\rangle^{\otimes n_p} |0\rangle, \quad (5)$$

where $|u_j\rangle$ is the j -th eigenvector of A , and β_j is the complex amplitude of $|b\rangle$ along $|u_j\rangle$. The next step of the procedure is to apply QPE to this superposition of eigenvectors to determine the corresponding eigenphases.

Quantum Phase Estimation

Quantum phase estimation is a procedure for the estimation of the eigenphases of a unitary. It begins with the observation that the eigenvalues of unitary operators lie on the complex unit

circle, and hence they can be expressed in the form

$$\mu_j = e^{2\pi i \phi_j}, \quad (6)$$

where μ_j is an eigenvalue of some unitary U , and $\phi_j \in [0, 1)$ is termed the corresponding eigenphase. The fundamental action of QPE is to use one register, initialized to an eigenvector $|u_j\rangle$ of U , to write a binary representation of the corresponding eigenphase onto an external phase register:

$$|u_j\rangle |0\rangle^{\otimes n_p} \xrightarrow{\text{QPE}} |u_j\rangle |\tilde{\phi}_j\rangle. \quad (7)$$

Here $\tilde{\phi}_j$ represents an n_p -bit approximation of ϕ_j : If ϕ_j can be exactly represented by n_p bits, then the phase register will have the final state $|\phi_j\rangle$ exactly¹. Otherwise, the state of the phase register will have maximal amplitude at the best n_p -bit approximation of ϕ_j , with the amplitudes of its other basis states decaying rapidly in the vicinity of this best approximation. That is,

$$|\tilde{\phi}_j\rangle = \sum_{k=1}^{2^{n_p}-1} \rho_{jk} |k\rangle. \quad (8)$$

Where ρ_{jk} spikes sharply in the vicinity of states with $k/2^{n_p} \approx \phi_j$. We do not explore the specifics of the QPE procedure here, but a summary diagram is given in Fig. 1.

Since QPE is itself a linear operator, its effect when applied to a superposition of eigenvectors of U is as follows:

$$\sum_{j=0}^{N-1} \beta_j |u_j\rangle |0\rangle^{\otimes n_p} \xrightarrow{\text{QPE}} \sum_{j=0}^{N-1} \beta_j |u_j\rangle |\tilde{\phi}_j\rangle. \quad (9)$$

That is, the approximated eigenphase states appear on the phase register entangled with their corresponding eigenvectors. Fig. 2 provides an illustration of this behavior.

In the case of our current system, we require a unitary with eigenvalues and eigenvectors which can be easily related to those of A . For the purposes of the HHL algorithm, the QPE unitary is typically defined as the exponential of the system matrix (Hall, 2015),

$$U = e^{iAt}, \quad (10)$$

where t is some constant. This operator is useful due to the close relationships between its eigenvalues and eigenvectors, and those of A . However, it is interesting to note that this operator also corresponds to the simulation of a physical system subjected to a process described by the

¹ Note that ϕ_j is not, in general, an integer. The notation $|\phi_j\rangle$ is used here simply as a label to indicate a precise eigenphase estimation. That is, $|\phi_j\rangle$ is the state $|k\rangle$ such that $k/2^{n_p} = \phi_j$. This notation follows naturally from (8).

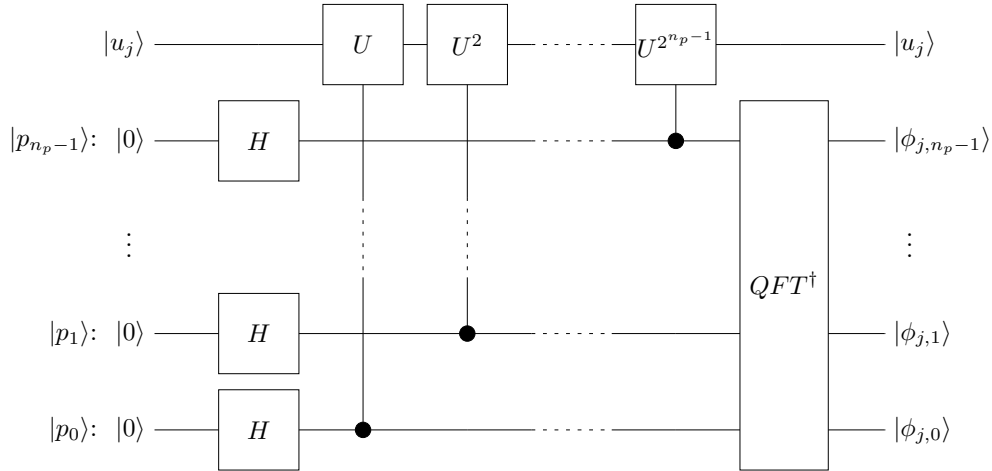


Figure 1

Diagram of quantum phase estimation. To emphasize the action of the eigenphase being read onto the phase register $|p\rangle$, the vector register is shown to have the state of a single eigenvector. In general, it will be prepared in a superposition of all eigenvectors of U . In the state $|\phi_{j,a}\rangle$, the second subscript a indicates the the a -th bit of the eigenphase ϕ_j corresponding to $|u_j\rangle$. This final state assumes ϕ_j can be represented exactly by n_p bits. When ϕ_j cannot be represented thusly, the final state of the phase register will have probabilities which spike sharply in the vicinity of the most accurate approximation to ϕ_j .

Hamiltonian A . The evolution of such a system is given by Schrödinger's equation:

$$i\hbar \frac{\partial |\psi\rangle}{\partial t} = A |\psi\rangle. \quad (11)$$

When evolved from an initial state $|\psi(0)\rangle$ for a time t , the system's final state is given by

$$|\psi(t)\rangle = e^{-iAt/\hbar} |\psi(0)\rangle. \quad (12)$$

Thus, the operator in (10) can be considered to be a time-evolution operator for the Hamiltonian A , up to the factor of $-1/\hbar$.

Different values of t , and even combinations of different values, can be used to improve the performance of the calculation, but from this point on we will use $t = 1$ for notational convenience. Then the unitary $U = e^{iA}$ has eigenvalues $e^{i\lambda_j}$, and hence the eigenvalues of A are related to the eigenphases of U by

$$e^{i\lambda_j} = e^{2\pi i\phi_j} \implies \lambda_j = 2\pi\phi_j + 2\pi l, \quad (13)$$

for some integer l . Recall that ϕ_j is on the range $[0, 1)$. If we assume that A has eigenvalues on the




Figure 2

Illustration of the proportional presence of eigenphase estimations on the phase register after QPE. The unitary used for this calculation was the complex exponential e^{iH} of a randomly generated Hermitian matrix H , and the operand vector was also randomly generated. $|\beta_j|^2$ gives the theoretical probability associated with eigenvector $|u_j\rangle$, and $P(k)$ gives the measured probability of observing the phase register in state $|k\rangle$. The amplitudes of phase states corresponding to accurate eigenphase approximations spike sharply and roughly in proportion to the presence of the eigenvectors, indicating that the state of a given phase estimate is entangled with the corresponding eigenstate. These spikes are not exactly proportional due to the probability being distributed over a nontrivial range of potential approximations.

range $[-\pi, \pi]$ (this can be achieved in general by a simple rescaling of the system matrix) then we see that $\phi_j \in [0, \frac{1}{2}]$ correspond to $\lambda_j \in [0, \pi]$, and $\phi_j \in (\frac{1}{2}, 1)$ correspond to $\lambda_j \in (-\pi, 0)$. Thus, λ_j can be computed as

$$\lambda_j = \begin{cases} 2\pi\phi_j, & \phi_j \in [0, \frac{1}{2}] \\ 2\pi(\phi_j - 1), & \phi_j \in (\frac{1}{2}, 1). \end{cases} \quad (14)$$

Of course, we do not compute the eigenphases exactly. Instead, we transform the phase register to some superposition where the amplitude of a basis state $|k\rangle$ grows large when $k/N_p \approx \phi_j$. Here we have defined $N_p = 2^{n_p}$. Then we in fact need to calculate potential approximations of the eigenvalues as

$$\tilde{\lambda}_k = \begin{cases} 2\pi \frac{k}{N_p}, & k \leq \frac{N_p}{2} \\ 2\pi \left(\frac{k}{N_p} - 1 \right), & k > \frac{N_p}{2}. \end{cases} \quad (15)$$

From the above descriptions, we see that the effect of QPE on the current system state is

$$\sum_{j=0}^{N-1} \beta_j |u_j\rangle |0\rangle^{\otimes n_p} |0\rangle \xrightarrow{\text{QPE}} \sum_{j=0}^{N-1} \sum_{k=0}^{N_p-1} \beta_j \rho_{jk} |u_j\rangle |k\rangle |0\rangle, \quad (16)$$

where ρ_{jk} is the amplitude of the phase state $|k\rangle$ resulting from the action of QPE on eigenvector $|u_j\rangle$. Thus, ρ_{jk} should spike sharply when $|k\rangle$ nears an accurate approximation of the j -th eigenphase of U , and remain near zero otherwise. With the eigenvalue extraction procedure defined, we can apply the third step of the procedure: ancilla rotation.

Ancilla Rotation and Uncomputation

In this step, the ancilla register is used to impose the inverse eigenvalue factors on the system. We use the rotation operator

$$R_y \left(2 \arccos \left(\frac{C}{\tilde{\lambda}_k} \right) \right) |0\rangle = \frac{C}{\tilde{\lambda}_k} |0\rangle + \sqrt{1 - \frac{C^2}{\tilde{\lambda}_k^2}} |1\rangle, \quad (17)$$

where C is a constant chosen to ensure that all arguments of the arccos function obey its domain restrictions. We recommend $C = 2\pi/N_p$, as this corresponds to the minimum possible magnitude of the denominator $\tilde{\lambda}_k$ in (15), excepting $k = 0$. The $k = 0$ case should be explicitly omitted from the calculation, as a nontrivial contribution from this state would indicate an eigenvalue near zero, and therefore an ill-conditioned matrix. By using a controlled version of this rotation operator, with the phase register acting as the control register, it is possible to entangle the $\tilde{\lambda}_k$ rotation with the $|k\rangle$ state of the phase register. Applying this controlled rotation to the ancilla register for all $k \in \{1, \dots, N_p - 1\}$, the system is transformed to the state

$$\sum_{j=0}^{N-1} \sum_{k=1}^{N_p-1} \beta_j \rho_{jk} |u_j\rangle |k\rangle \left(\frac{C}{\tilde{\lambda}_k} |0\rangle + \sqrt{1 - \frac{C^2}{\tilde{\lambda}_k^2}} |1\rangle \right), \quad (18)$$

where we have omitted the $k = 0$ case on the assumption that $\rho_{j0} \approx 0$.

Note that the state (18) is approximately the following ideal state:

$$\sum_{j=0}^{N-1} \sum_{k=0}^{N_p-1} \beta_j \rho_{jk} |u_j\rangle |k\rangle \left(\frac{C}{\lambda_j} |0\rangle + \sqrt{1 - \frac{C^2}{\lambda_j^2}} |1\rangle \right). \quad (19)$$

This is true since ρ_{jk} spikes sharply when $\tilde{\lambda}_k$ approaches λ_j , with the spike narrowing as more accurate approximations become available. Then terms with $\tilde{\lambda}_k$ not approximately equal to λ_j are nullified by their corresponding coefficients ρ_{jk} approaching zero, and we can, up to an

approximation error, replace $\tilde{\lambda}_k$ with λ_j , as we have done in (19). With the knowledge that the QPE error is obscured by this step, we from this point on assume that the system is in state (19) exactly. Acting on this state with an inverse QPE will then “uncompute” the phase register, leaving us in the state

$$\sum_{j=0}^{N-1} \beta_j |u_j\rangle |0\rangle^{\otimes n_p} \left(\frac{C}{\lambda_j} |0\rangle + \sqrt{1 - \frac{C^2}{\lambda_j^2}} |1\rangle \right). \quad (20)$$

This is the final state of the system, from which our solution can be extracted. Selecting those states with $|0\rangle$ in all qubits of the phase register and the ancilla register, we have

$$C \sum_{j=0}^{N-1} \frac{\beta_j}{\lambda_j} |u_j\rangle. \quad (21)$$

This state, when divided by C , is exactly the decomposition of $A^{-1} |b\rangle$ in terms of the eigenvectors of A , and hence (21) gives our final solution $|x\rangle$.

The Walk Operator

A walk operator is a unitary operator that is used during a quantum walk procedure, with each application of the operator representing a “step” made by the system. As mentioned above, we do not here describe a quantum walk procedure per se. Rather, we simply adopt a walk operator from a formerly studied quantum walk procedure (Ambainis, 2004; Szegedy, 2004). The operator chosen is useful because its implementation is well defined for any arbitrary system matrix. The walk operator itself has several constituents which must be defined before the operator can be stated succinctly. To begin, we define the vectors

$$|\phi_j\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \sqrt{\frac{N}{X} A_{jk}^*}, \quad (22)$$

for $j \in \{0, \dots, N-1\}$. Here X is a constant which must satisfy $X \geq N |A_{jk}|_{\max}$, where $|A_{jk}|_{\max} = \max_{j,k} (|A_{jk}|)$. This factor is required for the purposes of ancilla rotation, as discussed in section “Considerations for Practical Implementation”. The $|\phi_j\rangle$ vectors are useful to us as carriers of information about the system matrix, and they form foundational elements of the walk operator. However, they are possessed of some critical shortcomings. First, they are not normalized. While they could be explicitly normalized through appropriate j -dependent factors, subsequent derivations serve to show that this scheme would destroy crucial operator properties. In particular, the result of (34) would no longer hold. The second concern regarding the $|\phi_j\rangle$ vectors is their preparation. Producing unitary operators to generate these states is a nontrivial

problem, and even once computed, these operators would also carry j -dependent normalization factors.

Both of these shortcomings can be remedied with relative ease by introducing an ancillary qubit and studying the set of augmented states

$$|\phi_j^a\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \left[\sqrt{\frac{N}{X}} A_{jk}^* |0\rangle + \sqrt{1 - \frac{N}{X}} |A_{jk}| |1\rangle \right]. \quad (23)$$

These states are normalized, and they can be easily produced through unitary transformations (see section “Considerations for Practical Implementation”). Note that the $|\phi_j\rangle$ states appear when the ancilla is in the $|0\rangle$ state. Hence, the $|1\rangle$ state of the ancilla can be interpreted as a failure state. Storage of each $|\phi_j^a\rangle$ state requires an $(n+1)$ -qubit register. We define two such registers, and refer to them by the names $|r_1\rangle$ and $|r_2\rangle$. The ancilla-augmented basis states for these registers are indicated by a superscript a , e.g. $|j^a\rangle$. Note that there are $2^{n+1} = 2N$ augmented basis states. When necessary, we always assume that the ancillary qubit is appended to the relevant register. In this procedure, there are two basic operators which act on these registers. The first is the swap operator, which swaps the content of $|r_1\rangle$ and $|r_2\rangle$:

$$S = \sum_{j^a=0}^{2N-1} \sum_{k^a=0}^{2N-1} |k^a\rangle |j^a\rangle \langle j^a| \langle k^a|. \quad (24)$$

The second operator is that which produces $|\phi_j^a\rangle$ on $|r_2\rangle$ when $|r_1\rangle$ is in the state $|j\rangle$ with ancilla state $|0\rangle$:

$$T = \sum_{j=0}^{N-1} (|j, 0\rangle |\phi_j^a\rangle \langle j, 0| + |j, 1\rangle |\zeta_j^a\rangle \langle j, 1|). \quad (25)$$

Here the $|\zeta_j^a\rangle$ states correspond to failure states. A forthcoming analysis of the $T^\dagger ST$ operator shows that they must have ancilla state $|1\rangle$, but otherwise their definition is insignificant. These are proper quantum states, and hence they must be normalized.

These prerequisites are sufficient for us to state the walk operator:

$$W = iS(2TT^\dagger - I). \quad (26)$$

The process of implementing W is further explored in section “Considerations for Practical Implementation”. For the remainder of this section, we take it as given that W can be implemented. As stated above, this walk operator is particularly interesting to us because it has eigenvalues and eigenvectors closely related to those of A itself. In deriving these relationships between W and A , we first calculate some useful relationships for the constituent operators T and

S. To begin, consider the product $T^\dagger T$:

$$\begin{aligned} T^\dagger T &= \left[\sum_{j=0}^{N-1} (|j, 0\rangle \langle j, 0| \langle \phi_j^a| + |j, 1\rangle \langle j, 1| \langle \zeta_j^a|) \right] \cdot \left[\sum_{k=0}^{N-1} (|k, 0\rangle \langle \phi_k^a| \langle k, 0| + |k, 1\rangle \langle \zeta_k^a| \langle k, 1|) \right] \\ &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} (|j, 0\rangle \langle \phi_j^a| \langle \phi_k^a| \langle j, 0| \langle k, 0| \langle k, 0| + |j, 1\rangle \langle \zeta_j^a| \langle \zeta_k^a| \langle j, 1| \langle k, 1| \langle k, 1|). \end{aligned} \quad (27)$$

Recalling that the $|\phi_j^a\rangle$ and $|\zeta_j^a\rangle$ states are normalized and that the basis states are orthonormal, we have

$$T^\dagger T = \sum_{j=0}^{N-1} (|j, 0\rangle \langle j, 0| + |j, 1\rangle \langle j, 1|) = I. \quad (28)$$

Note that despite the above, TT^\dagger does not yield the identity operator, and hence T is not a unitary transformation as currently stated. The second relationship of interest is the following:

$$\begin{aligned} T^\dagger ST &= \left[\sum_{j=0}^{N-1} (|j, 0\rangle \langle j, 0| \langle \phi_j^a| + |j, 1\rangle \langle j, 1| \langle \zeta_j^a|) \right] \cdot \left[\sum_{k=0}^{N-1} (|\phi_k^a\rangle \langle k, 0| \langle k, 0| + |\zeta_k^a\rangle \langle k, 1| \langle k, 1|) \right] \\ &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \left(|j, 0\rangle \langle \phi_j^a| \langle k, 0\rangle \langle j, 0| \langle \phi_k^a| \langle k, 0| + |j, 0\rangle \langle \phi_j^a| \langle k, 1\rangle \langle j, 0| \langle \zeta_k^a| \langle k, 1| \right. \\ &\quad \left. + |j, 1\rangle \langle \zeta_j^a| \langle k, 0\rangle \langle j, 1| \langle \phi_k^a| \langle k, 0| + |j, 1\rangle \langle \zeta_j^a| \langle k, 1\rangle \langle j, 1| \langle \zeta_k^a| \langle k, 1| \right). \end{aligned} \quad (29)$$

If this operator acts on an arbitrary input vector $|\psi, 0\rangle$, we have

$$\begin{aligned} T^\dagger ST |\psi, 0\rangle &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \left(|j, 0\rangle \langle \phi_j^a| \langle k, 0\rangle \langle j, 0| \langle \phi_k^a| \langle k, 0| \langle \psi, 0| \right. \\ &\quad \left. + |j, 1\rangle \langle \zeta_j^a| \langle k, 0\rangle \langle j, 1| \langle \phi_k^a| \langle k, 0| \langle \psi, 0| \right). \end{aligned} \quad (30)$$

If each $|\zeta_j^a\rangle$ has ancilla state $|1\rangle$, this expression becomes

$$T^\dagger ST |\psi, 0\rangle = \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} |j, 0\rangle \langle \phi_j^a| \langle k, 0\rangle \langle j, 0| \langle \phi_k^a| \langle k, 0| \langle \psi, 0|. \quad (31)$$

This elimination of the $|\zeta_j^a\rangle$ contributions is a crucial result, and hence we require that these states

have ancilla state $|1\rangle$. The inner products are

$$\begin{aligned}\langle \phi_j^a | k, 0 \rangle &= \left[\frac{1}{\sqrt{N}} \sum_{p=0}^{N-1} \langle p | \left(\sqrt{\frac{N}{X}} \left(\sqrt{A_{jp}^*} \right)^* \langle 0 | + \sqrt{1 - \frac{N}{X}} |A_{jp}| \langle 1 | \right) \right] |k\rangle |0\rangle \\ &= \left(\sqrt{\frac{A_{jk}^*}{X}} \right)^*.\end{aligned}\quad (32)$$

$$\begin{aligned}\langle j, 0 | \phi_k^a \rangle &= \langle 0 | \langle j | \left[\frac{1}{\sqrt{N}} \sum_{p=0}^{N-1} |p\rangle \left(\sqrt{\frac{N}{X}} A_{kp}^* |0\rangle + \sqrt{1 - \frac{N}{X}} |A_{kp}| |1\rangle \right) \right] \\ &= \sqrt{\frac{A_{kj}^*}{X}}.\end{aligned}\quad (33)$$

Then we have

$$\begin{aligned}T^\dagger ST |\psi, 0\rangle &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} |j, 0\rangle \sqrt{\frac{A_{kj}^*}{X}} \left(\sqrt{\frac{A_{jk}^*}{X}} \right)^* \langle k, 0 | \psi, 0 \rangle \\ &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} |j, 0\rangle \frac{A_{jk}}{X} \langle k, 0 | \psi, 0 \rangle \\ &= \left(\frac{1}{X} A |\psi\rangle \right) |0\rangle.\end{aligned}\quad (34)$$

The statement that $\sqrt{A_{kj}^*} \left(\sqrt{A_{jk}^*} \right)^* = A_{jk}$ is true, but it presents an issue for implementation when A_{jk} lies on the negative real axis. This problem is further discussed in section “Considerations for Practical Implementation”. By the above, we see that $T^\dagger ST$, when applied to a vector with ancilla state $|0\rangle$, acts as a rescaled version of A . We introduce the shorthand

$$\hat{A} \equiv \frac{1}{X} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} A_{jk} |j, 0\rangle \langle k, 0|. \quad (35)$$

This transformation is a $2N \times 2N$ matrix, and it can therefore be applied directly to the $2N$ -dimensional augmented states. This simplifies the statement of (34):

$$T^\dagger ST |\psi, 0\rangle = \hat{A} |\psi, 0\rangle. \quad (36)$$

Now we can begin the analysis of the walk operator’s eigenvalues and eigenvectors. We begin with the assertion that the eigenvalues of \hat{A} must lie on the range $[-1, 1]$. By the definition of X , this is in fact always satisfied regardless of the choice of A , as we show in section “Treatment of

Arbitrary Matrix Equations". With this established, consider now the effect of applying W to the vector $T|u_j^a\rangle = T|u_j, 0\rangle$ (recall that $|u_j\rangle$ is an eigenvector of A):

$$WT|u_j^a\rangle = iST|u_j^a\rangle. \quad (37)$$

Note that we have continued to assume that the input vector is provided with ancilla state $|0\rangle$. Meanwhile, if W is applied to $ST|u_j^a\rangle$, we obtain

$$WST|u_j^a\rangle = 2i\hat{\lambda}_jST|u_j^a\rangle - iT|u_j^a\rangle, \quad (38)$$

where $\hat{\lambda}_j = \lambda_j/X$ is the eigenvalue of \hat{A} corresponding to $|u_j^a\rangle$. Thus, if W is applied to any superposition of $T|u_j^a\rangle$ and $ST|u_j^a\rangle$, the result will be a simple superposition of the same two vectors:

$$W(T|u_j^a\rangle + \gamma ST|u_j^a\rangle) = -i\gamma T|u_j^a\rangle + i(1 + 2\hat{\lambda}_j\gamma)ST|u_j^a\rangle, \quad (39)$$

where γ is an arbitrary factor for the relative phase and magnitude of the two contributions. Then $(T|u_j^a\rangle + \gamma ST|u_j^a\rangle)$ is an eigenvector $|v_j\rangle$ of W , with eigenvalue μ_j , if $\gamma = i\mu_j$ and $i\mu_j^2 = i(1 + 2i\hat{\lambda}_j\mu_j)$. Solving this expression for the eigenvalues, we find two solutions:

$$\mu_j^\pm = i\hat{\lambda}_j \pm \sqrt{1 - \hat{\lambda}_j^2}. \quad (40)$$

The eigenvectors $|v_j^\pm\rangle$ corresponding to these eigenvalues can be normalized by computing the inner products of their un-normalized counterparts:

$$\begin{aligned} \langle v_j^\pm | v_j^\pm \rangle &= \left(\langle u_j^a | T^\dagger - i(\mu_j^\pm)^* \langle u_j^a | T^\dagger S^\dagger \right) \left(T|u_j^a\rangle + i\mu_j^\pm ST|u_j^a\rangle \right) \\ &= \langle u_j^a | T^\dagger T |u_j^a\rangle + i\mu_j^\pm \langle u_j^a | T^\dagger ST |u_j^a\rangle - i(\mu_j^\pm)^* \langle u_j^a | T^\dagger S^\dagger T |u_j^a\rangle + |\mu_j^\pm|^2 \langle u_j^a | T^\dagger S^\dagger ST |u_j^a\rangle \\ &= \langle u_j^a | u_j^a \rangle + i\mu_j^\pm \langle u_j^a | \hat{A} |u_j^a\rangle - i(\mu_j^\pm)^* \langle u_j^a | \hat{A}^\dagger |u_j^a\rangle + |\mu_j^\pm|^2 \langle u_j^a | u_j^a \rangle \\ &= 1 + i\mu_j^\pm \hat{\lambda}_j - i(\mu_j^\pm)^* \hat{\lambda}_j + 1 \\ &= 2(1 - \hat{\lambda}_j^2). \end{aligned} \quad (41)$$

Therefore, W has the normalized eigenvectors

$$|v_j^\pm\rangle = \frac{1 + i\mu_j^\pm S}{\sqrt{2(1 - \hat{\lambda}_j^2)}} T|u_j^a\rangle. \quad (42)$$

While the above analysis does not constitute a comprehensive investigation into the eigenvalues and eigenvectors of W , this simplified derivation is, as shown in what follows, sufficient for our purposes. For a detailed analysis, see (Szegedy, 2004). To demonstrate the sufficiency of the above analysis, consider the following superposition of eigenvectors of W :

$$\frac{\left(1 + i\hat{\lambda}_j\mu_j^-\right) |v_j^+\rangle + \left(1 + i\hat{\lambda}_j\mu_j^+\right) |v_j^-\rangle}{\sqrt{2\left(1 - \hat{\lambda}_j^2\right)}}. \quad (43)$$

Expanding the eigenvector expressions and simplifying, we find

$$\begin{aligned} & \frac{1}{2\left(1 - \hat{\lambda}_j^2\right)} \left[\left(1 + i\hat{\lambda}_j\mu_j^-\right) \left(1 + i\mu_j^+ S\right) T |u_j^a\rangle + \left(1 + i\hat{\lambda}_j\mu_j^+\right) \left(1 + i\mu_j^- S\right) T |u_j^a\rangle \right] \\ &= \frac{1}{2\left(1 - \hat{\lambda}_j^2\right)} \left[2 + i\left(\mu_j^+ + \mu_j^-\right) S + i\hat{\lambda}_j\left(\mu_j^+ + \mu_j^-\right) - 2\hat{\lambda}_j\mu_j^+\mu_j^- S \right] T |u_j^a\rangle \\ &= \frac{1}{2\left(1 - \hat{\lambda}_j^2\right)} \left[2 - 2\hat{\lambda}_j S - 2\hat{\lambda}_j^2 + 2\hat{\lambda}_j S \right] T |u_j^a\rangle \\ &= T |u_j^a\rangle. \end{aligned} \quad (44)$$

This result gives us the necessary components to build a matrix equation solver. By applying T to the initial right-hand side vector $|b, 0\rangle$, we can transform the vector into a predictable superposition of the $|v_j^\pm\rangle$'s:

$$\begin{aligned} T |b, 0\rangle &= \sum_{j=0}^{N-1} \beta_j T |u_j^a\rangle \\ &= \sum_{j=0}^{N-1} \frac{\beta_j}{\sqrt{2\left(1 - \hat{\lambda}_j^2\right)}} \left[\left(1 + i\hat{\lambda}_j\mu_j^-\right) |v_j^+\rangle + \left(1 + i\hat{\lambda}_j\mu_j^+\right) |v_j^-\rangle \right]. \end{aligned} \quad (45)$$

Then an application of QPE to this state using W as the requisite unitary yields the phase corresponding to μ_j^+ or μ_j^- when $|u_j^a\rangle$ appears in the initial state. From here it is possible to extract λ_j and apply an HHL-style ancilla rotation. After applying inverse QPE, the final step of the procedure is to uncompute by applying T^\dagger , thereby leaving only the rotated initial state, as in the result of HHL.

Considerations for Practical Implementation

The analysis we have provided so far represents a complete theoretical formulation of the procedure, but practical implementation presents several nontrivial problems. In particular, the implementation of W , the initial application of T , and the eigenvalue extraction require elaboration. Here we present a recommended approach to the resolution of these sticking points. Implementation of W requires, in addition to the straightforward operators i and S , the more involved operator $2TT^\dagger - I$. This operator is a sum of conditional reflectors² about the $|\phi_j^a\rangle$ and $|\zeta_j^a\rangle$ states, as the following calculation shows. We begin by expanding and simplifying the TT^\dagger portion of the operator:

$$\begin{aligned}
TT^\dagger &= \left[\sum_{j=0}^{N-1} (|j,0\rangle |\phi_j^a\rangle \langle j,0| + |j,1\rangle |\zeta_j^a\rangle \langle j,1|) \right] \cdot \left[\sum_{k=0}^{N-1} (|k,0\rangle \langle k,0| \langle \phi_k^a| + |k,1\rangle \langle k,1| \langle \zeta_k^a|) \right] \\
&= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \left[\langle j,0|k,0\rangle (|j,0\rangle \langle k,0| \otimes |\phi_j^a\rangle \langle \phi_k^a|) + \langle j,1|k,1\rangle (|j,1\rangle \langle k,1| \otimes |\zeta_j^a\rangle \langle \zeta_k^a|) \right] \\
&= \sum_{j=0}^{N-1} (|j,0\rangle \langle j,0| \otimes |\phi_j^a\rangle \langle \phi_j^a| + |j,1\rangle \langle j,1| \otimes |\zeta_j^a\rangle \langle \zeta_j^a|). \tag{46}
\end{aligned}$$

Then the full operator $2TT^\dagger - I$ is

$$\begin{aligned}
&2 \sum_{j=0}^{N-1} (|j,0\rangle \langle j,0| \otimes |\phi_j^a\rangle \langle \phi_j^a| + |j,1\rangle \langle j,1| \otimes |\zeta_j^a\rangle \langle \zeta_j^a|) - \sum_{j=0}^{N-1} (|j,0\rangle \langle j,0| \otimes I + |j,1\rangle \langle j,1| \otimes I) \\
&= \sum_{j=0}^{N-1} \left[|j,0\rangle \langle j,0| \otimes (2|\phi_j^a\rangle \langle \phi_j^a| - I) + |j,1\rangle \langle j,1| \otimes (2|\zeta_j^a\rangle \langle \zeta_j^a| - I) \right]. \tag{47}
\end{aligned}$$

That is, the operator reflects $|r_2\rangle$ about $|\phi_j^a\rangle$ when $|r_1\rangle$ is in the state $|j,0\rangle$, or it reflects about $|\zeta_j^a\rangle$ when $|r_1\rangle$ is in the state $|j,1\rangle$. To implement this operator, we first consider a state preparation operator B_j which prepares $|\phi_j^a\rangle$ from the $|0\rangle$ state. By applying, in sequence, the inverse of this state preparation operator, a reflection about the $|0\rangle$ state, and finally the state preparation operator, we find

$$B_j (2|0\rangle \langle 0| - I) B_j^\dagger = 2B_j |0\rangle \langle 0| B_j^\dagger - B_j B_j^\dagger = 2|\phi_j^a\rangle \langle \phi_j^a| - I. \tag{48}$$

Thus, this procedure effectively reflects about the state $|\phi_j^a\rangle$. Likewise, we can reflect about the

² A brief introduction to the concept of reflectors is given in appendix A.

$|\zeta_j^a\rangle$ states using an operator B'_j which prepares $|\zeta_j^a\rangle$ from $|0\rangle$:

$$B'_j (2|0\rangle\langle 0| - I) (B'_j)^\dagger = 2|\zeta_j^a\rangle\langle \zeta_j^a| - I. \quad (49)$$

The full operator can then be implemented as

$$2TT^\dagger - I = \sum_{j=0}^{N-1} \left[|j,0\rangle\langle j,0| \otimes \left(B_j (2|0\rangle\langle 0| - I) B_j^\dagger \right) + |j,1\rangle\langle j,1| \otimes \left(B'_j (2|0\rangle\langle 0| - I) (B'_j)^\dagger \right) \right]. \quad (50)$$

Note that the swap operations involved in the walk operator mean that, at intermediate stages, the system stores essential data on the $|1\rangle$ state of the ancilla, and hence the reflections about the $|\zeta_j\rangle$'s cannot be ignored.

The implementation of the state preparation operators is straightforward. For the B_j operators, we begin by preparing a uniform superposition:

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k,0\rangle. \quad (51)$$

We can then apply a single Pauli X gate to the ancilla bit to obtain

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k,1\rangle. \quad (52)$$

Note that, for $A_{jk} = 0$, the amplitudes of the $|k^a\rangle$ states are now exactly as desired. Then for $|k\rangle$ corresponding to nonzero A_{jk} , we rotate the ancilla to obtain:

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \left[\sqrt{\frac{N}{X}} A_{jk}^* |0\rangle + \sqrt{1 - \frac{N}{X} |A_{jk}|} |1\rangle \right]. \quad (53)$$

It is at this point that the role of X becomes clear. In order to ensure a valid rotation, we select X such that $X/N \geq |A_{jk}|_{\max}$. Note that the calculation of $\sqrt{A_{jk}^*}$ requires some additional consideration. In particular, the result of (34) requires that

$$\sqrt{A_{kj}^*} \left(\sqrt{A_{jk}^*} \right)^* = A_{jk}. \quad (54)$$

In most cases, this can be satisfied by taking

$$\sqrt{A_{jk}^*} = \sqrt{|A_{jk}|} e^{-i \arg(A_{jk})/2}, \quad (55)$$

and using the principal square root for $\sqrt{|A_{jk}|}$. However, if $A_{jk} \in (-\infty, 0)$, then $A_{jk} = A_{kj}$, and we have

$$\sqrt{A_{kj}^*} \left(\sqrt{A_{jk}^*} \right)^* = \left(\sqrt{|A_{jk}|} e^{-i\pi/2} \right) \left(\sqrt{|A_{jk}|} e^{i\pi/2} \right) = |A_{jk}|. \quad (56)$$

For $j \neq k$, we can recover the correct sign by taking the negative square root in (55) whenever $k > j$. That is, if A_{jk} is a negative real number, we use

$$\sqrt{A_{jk}^*} = \text{sign}(j-k) \sqrt{|A_{jk}|} e^{-i \arg(A_{jk})/2}. \quad (57)$$

While this prescription is effective for $j \neq k$, it does not suffice when $j = k$. To address this case, we simply prevent negative elements from appearing on the diagonal of A by adding an appropriate multiple of the identity matrix. Proper treatment of this modification is considered in section ‘‘Treatment of Arbitrary Matrix Equations’’.

The B'_j operators can be implemented by simply switching the ancilla state of the operand register to $|1\rangle$. This gives $|\zeta_j^a\rangle = |0\rangle^{\otimes n} |1\rangle$, which is sufficient for our purposes.

The initial application of T is also of concern. The operation is defined such that, when the $|r_1\rangle$ ancilla qubit is in the $|0\rangle$ state, the $|\phi_j^a\rangle$ states are produced on $|r_2\rangle$ regardless of the initial state of the register. This is problematic to implement in practice, but since the only direct application of T occurs at the beginning of the calculation, we can safely assume that $|r_2\rangle$ begins with all qubits in the $|0\rangle$ state. Then the initial application of T , which we shall distinguish by the term T_0 , can be effectively applied by conditional application of the state preparation operators:

$$T_0 = \sum_{j=0}^{N-1} (|j, 0\rangle \langle j, 0| \otimes B_j + |j, 1\rangle \langle j, 1| \otimes B'_j). \quad (58)$$

Unlike T , this operator is unitary, and hence it is suitable for use in a quantum circuit. Since the initial state must be supplied with ancilla state $|0\rangle$, it is also possible to ignore the applications of B'_j .

While not strictly an issue for a naive implementation, the manner in which controls are applied is important for efficiency. Consider the B_j operator. If row j of A contains N_{nz}^j nonzero elements, then B_j must apply N_{nz}^j ancilla rotations, each with an n -qubit control. The default in-place control scheme provided by, for example, Qiskit requires time exponential in n to apply this control, thereby destroying the efficiency of the procedure. Reference (Barenco et al., 1995) provides a method for implementing such controlled operations with $\mathcal{O}(n^2)$ complexity without including any additional qubits, but this approach requires n recursive square roots of the basic unitary to be computed. Even for unitaries where these square roots can be computed easily, the

complexity of their implementation in terms of basic gates scales poorly as the precision needed to accurately describe the desired operator increases. This could potentially be addressed by providing a constant cutoff precision, but here we instead opt for a control scheme of complexity $\mathcal{O}(n)$ described in (Nielsen & Chuang, 2000). A diagram summarizing the procedure is shown in Fig. 3. This method has the downside that it requires an additional $(n - 1)$ -qubit “work” register to store intermediate calculations related to the control procedure. While this increases the absolute size of system required to apply a solution circuit, it does not affect the overall resource scaling of the procedure, which remains $\mathcal{O}(n + n_p)$. Additionally, it provides an optimal complexity scaling for the control procedure, and applies the desired unitary directly.




Figure 3

Control procedure summary for a 4-qubit control register and arbitrary operator U . This approach provides $\mathcal{O}(n)$ complexity for the implementation of an n -qubit control. This circuit is only applicable for a $|1111\rangle$ control state, but other controls can be easily implemented by applying Pauli X gates to qubits requiring a $|0\rangle$ control. Note that the actual application of U involves only a single control qubit, and hence the complexity of U does not compound with the control complexity.

Using this control scheme, the rotation component of B_j will require $\mathcal{O}\left(N_{nz}^j \log(N)\right)$ basic gates to implement. The initial Hadamard step to produce the uniform superposition can clearly be implemented with $\mathcal{O}(\log(N))$ gates, and the Pauli X gate requires constant complexity.

Therefore, the total gate complexity of each B_j remains $\mathcal{O}\left(N_{nz}^j \log(N)\right)$. Each application of T_0 and W requires N controlled B_j operations, each requiring an $(n + 1)$ -qubit control. However,

note that the actual application of B_j requires only a single control qubit, and hence the complexity of the control procedure is separate from that of B_j . Then a single controlled B_j has $\mathcal{O}(\log(N)) + \mathcal{O}(N_{nz}^j \log(N)) = \mathcal{O}(N_{nz}^j \log(N))$ gate complexity. Note that we have assumed $N_{nz}^j \geq 1$ for all j , which must always be satisfied for a well-conditioned system. For constant precision, the solution circuit will require a constant number of applications of W and T_0 . Then, noting that $\sum_{j=0}^{N-1} N_{nz}^j = N_{nz}$, this gives the final gate complexity of the solution procedure:

$$\mathcal{O}(N_{nz} \log(N)) \quad (59)$$

The last point of concern is the new procedure for eigenvalue extraction. Where the canonical HHL unitary gives the eigenphase relationship $e^{2\pi i \phi_j} = e^{i \lambda_j}$, the walk unitary has $e^{2\pi i \phi_j} = i \hat{\lambda}_j \pm \sqrt{1 - \hat{\lambda}_j^2}$, and hence the extraction procedure must be modified. Recalling that $\hat{\lambda}_j \in [-1, 1]$, this extraction is straightforward if we take the imaginary components:

$$\sin(2\pi \phi_j) = \hat{\lambda}_j \implies \lambda_j = X \sin(2\pi \phi_j). \quad (60)$$

Treatment of Arbitrary Matrix Equations

Recall the following list of restrictions that have been imposed on our system:

1. $|b\rangle$ must be normalized.
2. A must be Hermitian.
3. N must be a power of two.
4. The eigenvalues of A must lie on the interval $[-X, X]$.
5. No negative real numbers may appear on the diagonal of A .

In this section, we develop a method by which an arbitrary matrix equation can be modified to satisfy these constraints. To keep our notation consistent, we begin from the following initial problem: Given $A_0 \in \mathbb{C}^{M \times M}$ and $|b_0\rangle \in \mathbb{C}^M$, find $|x_0\rangle$ such that

$$A_0 |x_0\rangle = |b_0\rangle. \quad (61)$$

This arbitrary input system provides the foundational elements for an appropriately restricted system.

The first restriction is that $|b\rangle$ is normalized. This is easily satisfied by dividing (61) by the magnitude of $|b_0\rangle$:

$$\frac{1}{\|b_0\|_2} A_0 |x_0\rangle = \frac{1}{\|b_0\|_2} |b_0\rangle. \quad (62)$$

Here we assume that $\|b_0\|_2 \neq 0$. Since the zero vector is always a valid solution when $\|b_0\|_2 = 0$, this case is trivial and need not be considered for our purposes.

The second assumption is that the system matrix A must be Hermitian. This is satisfied in general by expanding the problem to the following $2M \times 2M$ system:

$$\frac{1}{\|b_0\|_2} \begin{bmatrix} 0 & A_0 \\ A_0^\dagger & 0 \end{bmatrix} \begin{bmatrix} 0 \\ x_0 \end{bmatrix} = \frac{1}{\|b_0\|_2} \begin{bmatrix} b_0 \\ 0 \end{bmatrix}. \quad (63)$$

Third, we consider that the size of our quantum system must be a power of 2, as the algorithm is constructed for quantum systems using two-state qubits. Therefore, we define $n = \lceil \log(2M) \rceil$ and $N = 2^n$, and once again expand our system, this time to

$$\frac{1}{\|b_0\|_2} \begin{bmatrix} 0 & A_0 & 0 \\ A_0^\dagger & 0 & 0 \\ 0 & 0 & I_{N-2M} \end{bmatrix} \begin{bmatrix} 0 \\ x_0 \\ 0 \end{bmatrix} = \frac{1}{\|b_0\|_2} \begin{bmatrix} b_0 \\ 0 \\ 0 \end{bmatrix}, \quad (64)$$

where I_{N-2M} is the identity matrix of dimension $N - 2M$, and we have appended $N - 2M$ zeros to each vector. Now each vector can be represented by an n -qubit state.

The final problem is the restriction of the eigenvalues of A to the range $[-X, X]$. Let λ_{\max} be the magnitude of the dominant eigenvalue of A . Then we have

$$\lambda_{\max} \leq N |A_{jk}|_{\max} \leq X, \quad (65)$$

by the discussion following (53). Thus $[-\lambda_{\max}, \lambda_{\max}] \subseteq [-X, X]$, and hence the spectrum of A lies entirely on the required range. That is, for this procedure, no rescaling of the system is required to satisfy the eigenvalue restrictions.

Every element on the diagonal of the system matrix is now either 0 or a positive real number, and hence no negative real elements appear on its diagonal. This accounts for all imposed restrictions, leaving us with the following scheme for the preparation of an appropriate system:

$$A = \frac{1}{\|b_0\|_2} \begin{bmatrix} 0 & A_0 & 0 \\ A_0^\dagger & 0 & 0 \\ 0 & 0 & I_{N-2M} \end{bmatrix}, \quad |b\rangle = \frac{1}{\|b_0\|_2} \begin{bmatrix} b_0 \\ 0 \\ 0 \end{bmatrix}, \quad |x\rangle = \begin{bmatrix} 0 \\ x_0 \\ 0 \end{bmatrix}. \quad (66)$$

Of course, it is possible that A_0 is already Hermitian, in which case some efficiency can be gained by not performing the full expansion as stated above. Instead, only the size restriction must be satisfied. To this end, we redefine $n = \lceil \log(M) \rceil$ —keeping the definition $N = 2^n$ —and expand the

system as follows:

$$A = \frac{1}{\|b_0\|_2} \begin{bmatrix} A_0 & 0 \\ 0 & I_{N-M} \end{bmatrix}, \quad |b\rangle = \frac{1}{\|b_0\|_2} \begin{bmatrix} b_0 \\ 0 \end{bmatrix}, \quad |x\rangle = \begin{bmatrix} x_0 \\ 0 \end{bmatrix}. \quad (67)$$

Now it is possible for the system matrix to have negative real values on the diagonal. This can be amended by using a shifted system:

$$(A + dI) |x\rangle = |b\rangle, \quad (68)$$

where d is an upper bound on the magnitude of the offending values on the diagonal of A . For the purposes of the walk operator, $A + dI$ should be used as the system matrix, but note that this has the effect of shifting all eigenvalues of A by d . Hence, in order to directly apply the inverse of A to the initial state, the eigenvalue extraction of (60) must be modified:

$$\lambda_j = X \sin(2\pi\phi_j) - d. \quad (69)$$

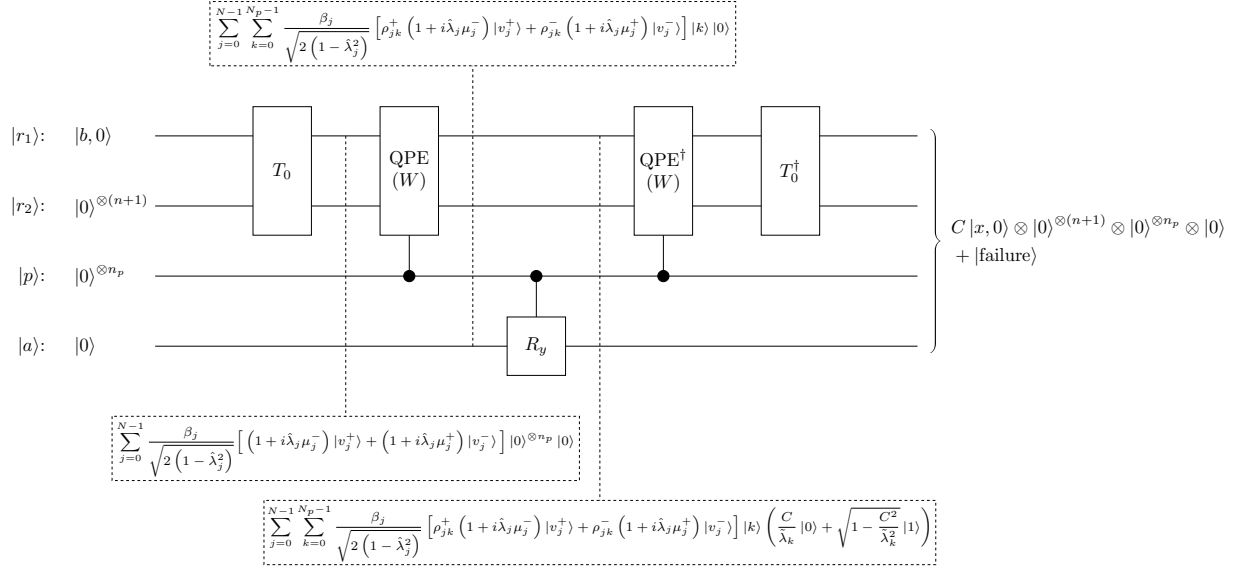
So long as X is calculated from the shifted system matrix, the eigenvalues are still appropriately bounded, and A still requires no additional rescaling to satisfy the eigenvalue restrictions.

Summary

Here we summarize the preceding analysis in the form of a self-contained procedure. We assume that the supplied matrix equation has been properly prepared according to section ‘‘Treatment of Arbitrary Matrix Equations’’. Fig. 4 provides a large-scale pictorial description of the full procedure. Notice that this procedure consists of four main stages: the initial application of T_0 , quantum phase estimation, HHL rotation, and uncomputation.

Initial Application of T_0

The first stage of the procedure is the application of T_0 , which maps $|r_2\rangle$ to a superposition of the $|\phi_j^a\rangle$ and $|\zeta_j^a\rangle$ states defined in section ‘‘The Walk Operator’’. A circuit diagram for our suggested implementation of T_0 is shown in Fig. 5, and a pseudocode description of the operator is given in algorithm 1. In the program, the ‘‘T0’’ function of the ‘‘QWOp’’ module implements this operator. The definition of T_0 depends on the state preparation operators B_j , and as such we have also provided a circuit diagram for an arbitrary B_j in Fig. 6. The corresponding pseudocode description is given in algorithm 2. In the program, B_j is implemented by the ‘‘Bj’’ function of the ‘‘QWOp’’ module. Here we have chosen to forgo any application of the B'_j operator, as we can safely presume that the ancilla state of $|r_1\rangle$ will in practice be initialized to $|0\rangle$.


Figure 4

Structural outline of the full procedure. Dashed boxes indicate the state of the system after every major intermediate stage. The $|r_1\rangle$ and $|r_2\rangle$ registers are the ancilla-augmented vector registers on which the walk operator and its associated constituents act. The phase register $|p\rangle$ provides space for the estimation of the eigenphases of W . The ancilla $|a\rangle$ is used for the application of the HHL rotation described in section “The HHL Algorithm”. We commonly refer to $|a\rangle$ as the “HHL ancilla” to distinguish it from the ancillae of $|r_1\rangle$ and $|r_2\rangle$. The phase amplitudes ρ_{jk}^+ and ρ_{jk}^- are used to distinguish the amplitudes of phase estimates corresponding to each of the two eigenvectors $|v_j^\pm\rangle$. The eigenvalue estimates themselves require no such distinction, as both μ_j^+ and μ_j^- have the same imaginary part.

In this application of T_0 , the system undergoes the transformation

$$\begin{aligned}
 |b, 0\rangle |0 \dots 0, 0\rangle &= \sum_{j=0}^{N-1} \beta_j |u_j, 0\rangle |0 \dots 0, 0\rangle \\
 \xrightarrow{T_0} &\sum_{j=0}^{N-1} \frac{\beta_j}{\sqrt{2(1-\hat{\lambda}_j^2)}} \left[(1+i\hat{\lambda}_j\mu_j^-) |v_j^+\rangle + (1+i\hat{\lambda}_j\mu_j^+) |v_j^-\rangle \right]. \tag{70}
 \end{aligned}$$

Here we have suppressed the phase and HHL ancilla registers, as they are not subject to the effects of T_0 .

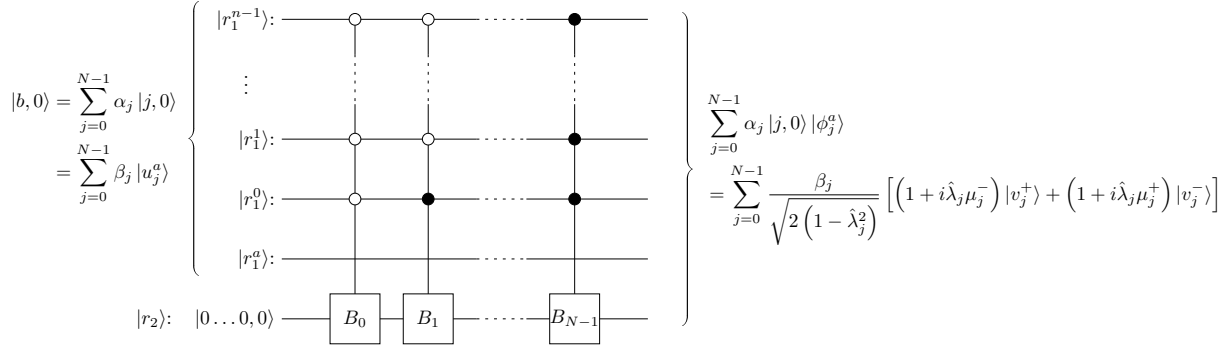

Figure 5

Diagram of the initial application of T_0 . Here state preparation operators are applied to $|r_2\rangle$, conditioned on the state of $|r_1\rangle$. Note that, due to the presumption that $|r_1\rangle$ is initially prepared with an ancilla state of $|0\rangle$, the $|r_1\rangle$ ancilla is ignored in this calculation. Since the resulting system is highly entangled, individual qubit states cannot be separated. As a result, we have only provided a final state for the total system.

Algorithm 1 Implementation of T_0

- 1: **for** $j = 0 \dots N - 1$ **do**
 - 2: Compute the operator B_j
 - 3: Apply B_j to $|r_2\rangle$ with the control condition that $|r_1\rangle$ is in the state $|j, 0\rangle$
 - 4: **end for**
-

Quantum Phase Estimation

With $|r_1\rangle$ and $|r_2\rangle$ prepared in the desired superposition of the walk operator's eigenvectors, an application of QPE using W as the unitary writes the eigenphases of W onto the phase register $|p\rangle$. An expansion of this step, in the typical QPE format, is given in Fig. 7. A circuit diagram for the walk operator itself is shown in Fig. 8, and the corresponding pseudocode is given in algorithm 3. The walk operator is implemented in the program by the “W” function of the “QWOps” module. Note that we have made use of our suggested j -independent form of the $|\zeta_j^a\rangle$ states, wherein $|\zeta_j^a\rangle = |0\rangle^{\otimes n} |1\rangle$. This eliminates the j dependence of the corresponding preparation operator, which we now refer to as simply B' . Since the swap operations can leave important data on the $|1\rangle$ state of $|r_1\rangle$'s ancilla, B' cannot be ignored as in the initial application of T_0 . In the model of our current study, B' can be implemented by simply applying an X gate to the ancilla of $|r_2\rangle$. The program implements B' in the “Bp” function of the “QWOps” module.

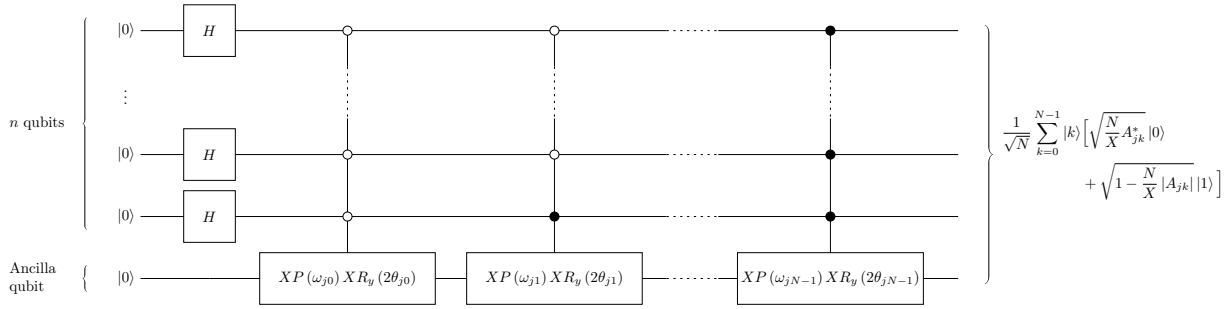
**Figure 6**

Diagram of the state preparation operator B_j applied to a register with all qubits initialized to the $|0\rangle$ state. Here $\theta_{jk} = \arccos\left(\sqrt{|A_{jk}|N/X}\right)$ and $\omega_{jk} = \pm \arg(A_{jk})/2$, where the positive solution for ω_{jk} is chosen iff $\arg(A_{jk}) = \pi$ and $j > k$. Thus, the R_y gates rotate according to the magnitudes of the elements of A , with the factor of XPX subsequently enforcing the correct phase on the $|0\rangle$ state of the ancilla.

Algorithm 2 Implementation of B_j

- 1: **for** $k = 0 \dots n - 1$ **do**
 - 2: Apply H to qubit k
 - 3: **end for**
 - 4: **for** $k = 0 \dots N - 1$ **do**
 - 5: $\theta_{jk} \leftarrow \arccos\left(\sqrt{|A_{jk}|N/X}\right)$
 - 6: $\omega_{jk} \leftarrow -\arg(A_{jk})/2$
 - 7: **if** $\arg(A_{jk}) = \pi$ and $j < k$ **then**
 - 8: $\omega_{jk} \leftarrow -\omega_{jk}$
 - 9: **end if**
 - 10: Apply $XP(\omega_{jk})XR_y(2\theta_{jk})$ to the ancilla qubit with the control condition that the rest of the register is in the state $|k\rangle$
 - 11: **end for**
-

After phase estimation, the system is in the state

$$\sum_{j=0}^{N-1} \sum_{k=0}^{N_p-1} \frac{\beta_j}{\sqrt{2(1-\hat{\lambda}_j^2)}} \left[\rho_{jk}^+ \left(1 + i\hat{\lambda}_j \mu_j^-\right) |v_j^+\rangle + \rho_{jk}^- \left(1 + i\hat{\lambda}_j \mu_j^+\right) |v_j^-\rangle \right] |k\rangle. \quad (71)$$

Here we have reintroduced the phase register, although $|a\rangle$ remains suppressed. As in section “The HHL Algorithm”, ρ_{jk}^\pm extracts accurate eigenvalue approximations by spiking sharply when $\tilde{\lambda}_k$ approaches an actual eigenvalue λ_j .

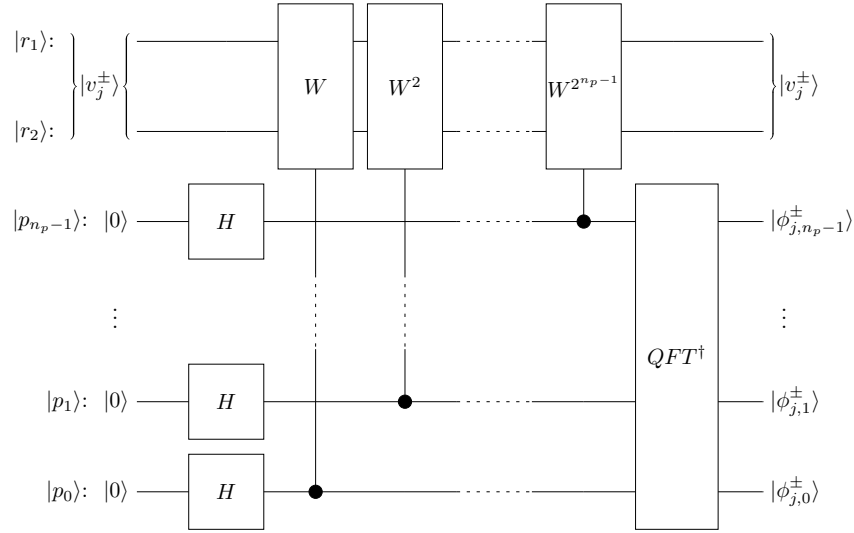


Figure 7

Diagram of phase estimation using the walk operator. Here we use the same simplifications as in Fig. 1.

HHL Ancilla Rotation

With the phase register holding the eigenphases of W , extraction of the corresponding eigenvalues can proceed according to the analysis of sections “Considerations for Practical Implementation” and “Treatment of Arbitrary Matrix Equations”. Specifically, the eigenvalues are calculated according to

$$\lambda_j = X \sin(2\pi\phi_j) - d, \quad (72)$$

where $d = 0$ can be used if the system was expanded to satisfy the Hermitian system matrix property. Of course, the exact values of ϕ_j are not known in general, and in practice approximate eigenvalues must be calculated from the state of the phase register as follows:

$$\tilde{\lambda}_k = X \sin\left(2\pi\frac{k}{N_p}\right) - d. \quad (73)$$

Once these eigenvalue approximations have been calculated, a rotation operator can be used to apply an eigenvalue factor corresponding to the inverse of the system matrix for every possible state of the phase register, as detailed in section “The HHL Algorithm”. A circuit diagram for this procedure is given in Fig. 9, and the relevant pseudocode is provided in algorithm 4. This rotation is implemented in the program by the “HHLRotation” function of the “QAlgs” module.



Figure 8

Diagram of the walk operator W . Here S represents a one-to-one exchange of qubits states between $|r_1\rangle$ and $|r_2\rangle$. The operator iI can be applied to any qubit in $|r_1\rangle$ or $|r_2\rangle$ to produce the desired effect.

Upon completion of this rotation, the system has the state

$$\sum_{j=0}^{N-1} \sum_{k=0}^{N_p-1} \frac{\beta_j}{\sqrt{2(1-\hat{\lambda}_j^2)}} \left[\rho_{jk}^+ (1 + i\hat{\lambda}_j \mu_j^-) |v_j^+\rangle + \rho_{jk}^- (1 + i\hat{\lambda}_j \mu_j^+) |v_j^-\rangle \right] |k\rangle \otimes \left(\frac{C}{\tilde{\lambda}_k} |0\rangle + \sqrt{1 - \frac{C^2}{\tilde{\lambda}_k^2}} |1\rangle \right). \quad (74)$$

Uncomputation

After rotation, $|p\rangle$, $|r_2\rangle$, and the ancilla of $|r_1\rangle$ are uncomputed by applying inverse QPE and T_0 operations. Ideally, these registers are returned to the $|0\rangle$ state, and any other state can be interpreted as a failure of the procedure. This leaves $|r_1\rangle$ entangled with $|a\rangle$. When $|a\rangle$ takes the state $|0\rangle$, the state of $|r_1\rangle$ is proportional (up to the error introduced through QPE) to the solution state $|x\rangle$. Meanwhile, if $|a\rangle$ is in the state $|1\rangle$, the system is in a failure state. That is, the overall final state of the system is

$$C|x, 0\rangle |0\rangle^{\otimes(n+1)} |0\rangle^{\otimes n_p} |0\rangle + |\text{failure}\rangle. \quad (75)$$

This final state is analogous to (20).

From this description, we see that the full procedure is well defined for any input matrix, with all constituent operations having explicit representations. This is in contrast to the canonical HHL algorithm, where the implementation of the QPE unitary presents an ambiguously defined

Algorithm 3 Implementation of W

```

1: for  $j = 0 \dots N - 1$  do
2:   Compute the state preparation operator  $B_j$ 
3:   Apply the reflector  $B_j (2|0\rangle\langle 0| - I) B_j^\dagger$  to  $|r_2\rangle$  with
   the control condition that  $|r_1\rangle$  is in the state  $|j, 0\rangle$ 
4: end for
5: Compute the state preparation operator  $B'$ 
6: Apply the reflector  $B' (2|0\rangle\langle 0| - I) (B')^\dagger$  to  $|r_2\rangle$  with
   the control condition that the  $|r_1\rangle$  ancilla is in the state
    $|1\rangle$ 
7: for  $j = 0 \dots N - 1$  do
8:   Apply a swap operation to the  $j$ -th qubits of  $|r_1\rangle$  and
    $|r_2\rangle$ 
9: end for
10: Apply a swap operation to the ancillae of  $|r_1\rangle$  and  $|r_2\rangle$ 
11: Apply  $iI$  to any qubit of  $|r_1\rangle$  or  $|r_2\rangle$ 

```

Algorithm 4 Implementation of the HHL rotation

```

1: for  $k = 0 \dots N_p - 1$  do
2:    $\tilde{\lambda}_k \leftarrow X \sin(2\pi k/N_p) - d$ 
3:   if  $\tilde{\lambda}_k \neq 0$  then
4:     Apply the rotation  $R_y \left[ 2 \arccos \left( C/\tilde{\lambda}_k \right) \right]$  to  $|a\rangle$ ,
     with the control condition that  $|p\rangle$  is in the state
      $|k\rangle$ 
5:   end if
6: end for

```

bottleneck for the procedure.

Results

In developing suitable tests for the practical performance of this procedure, two important limitations must be considered: First, the quantum systems currently available to the public are quite small, with the largest system we have access to offering only 7 qubits³ (*IBM Quantum*, 2021). Second, errors on these systems can propagate rapidly, with calculations involving more than ~ 800 primitive gates being, in our experience, unlikely to yield the expected states. Therefore, when considering a problem suitable for fully quantum analysis, we are at present limited to very short, small calculations. Section “A Fully Quantum Problem” constructs and analyzes such a problem, but it is possible to construct a more satisfying test using classical

³ This is the *ibmq_jakarta* v1.0.23 system, which is one of the IBM Quantum Falcon Processors.

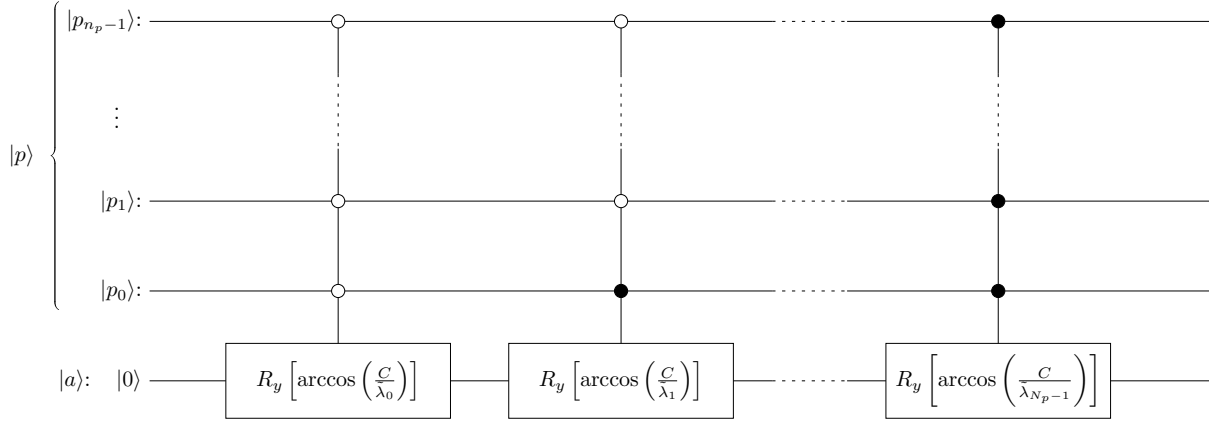


Figure 9

Diagram of the HHL rotation procedure. Here $\tilde{\lambda}_k = X \sin(2\pi k/N_p) - d$, and C is a lower bound on the magnitude of the $\tilde{\lambda}_k$'s that are not zero. Only those rotations for which $\tilde{\lambda}_k \neq 0$ should be applied.

simulation. By using a classical computer to simulate a quantum system, we gain access to more qubits and eliminate gate errors. Such simulation is extremely inefficient, but we have found that as many as 14 qubits can be simulated for this procedure using available classical systems. Section “Simulated Solution” considers a larger, less restricted test problem using classical simulation.

Aside from verifying the accuracy of our method, we must also show that it is efficient compared to classical algorithms. For this task, we use the example of sparse approximate inverse (SPAI) preconditioning. Section “Preconditioner Application” is dedicated analyzing the efficiency of preconditioner application.

A Fully Quantum Problem

Being limited to 7 qubits, we must consider only the smallest of possible problems. Thus we choose a problem of two unknowns, requiring a single qubit for each of the main portions of $|r_1\rangle$ and $|r_2\rangle$. Both of these registers require ancillary qubits, and a third ancilla is also required for the eigenvalue rotation. Then we currently stand at 5 qubits accounted for, with 2 left for the phase register. Since we are most concerned with using a minimal number of qubits, we use the exponential time, in-place control scheme provided by default in Qiskit. Thus, no additional qubits are needed for control considerations. These four possible phase states admit no error in the phase estimation, and hence we must choose a problem such that the eigenphases can be exactly represented by two bits.

By (69), the eigenvalues of the system matrix must satisfy

$$\lambda_j = X \sin(2\pi\phi_j) - d = \begin{cases} -d, & \phi_j = 0, \frac{1}{2} \\ X - d, & \phi_j = \frac{1}{4} \\ -X - d, & \phi_j = \frac{3}{4}. \end{cases} \quad (76)$$

The system

$$A = \begin{bmatrix} -2 & 1 \\ 1 & -2 \end{bmatrix}, \quad d = 3, \quad X = 2, \quad (77)$$

satisfies the constraints of our procedure, and has eigenvalues

$$\lambda_0 = -3 = -d, \quad \lambda_1 = -1 = X - d. \quad (78)$$

Thus, this system provides a suitable candidate for our analysis. Note that the system is Hermitian, and no expansion is required. For a right-hand side vector, we use an equal superposition of both eigenvectors of A :

$$|b\rangle = \frac{1}{2} \left(\begin{bmatrix} -1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right) = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (79)$$

For this problem of known form, some simplification is possible in the various operators of the solution procedure. This allows the size of the final circuit to be significantly reduced. The system matrix defining the walk operator is

$$A + dI = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}. \quad (80)$$

Therefore, all B_j operators apply the same ancilla rotation:

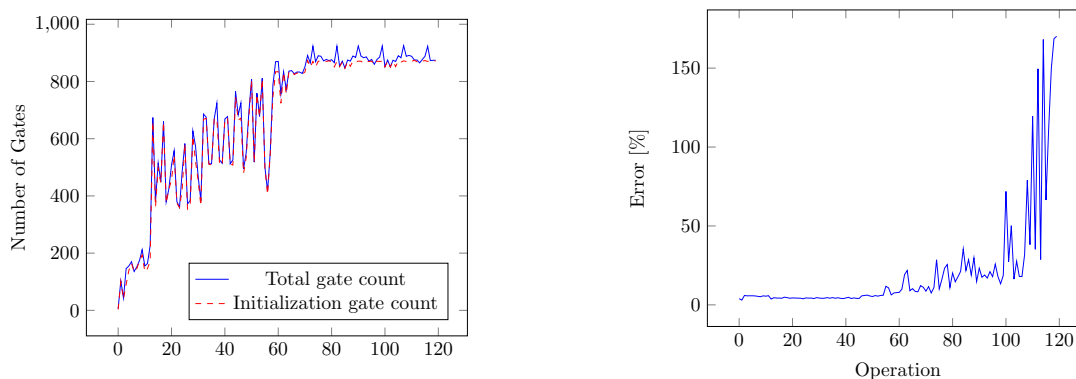
$$R_y \left(2 \arccos \sqrt{1 \cdot \frac{2}{2}} \right) = R_y(0) = I. \quad (81)$$

Since all elements are real, no phase rotation is necessary, and all B_j operators can be implemented by simply applying a Hadamard gate to $|r_2\rangle$. This elimination of the complicated phase and R_y gates, both of which also require additional control considerations, greatly reduces the complexity of each B_j operator.

This simplification of the B_j operators carries through the T_0 and W operators. Without this reduction, the full program produces a circuit consisting of 15,728 basic gates. This figure results

from transpilation using the available CX , I , R_z , \sqrt{X} , and X basis gates. After simplification, the final circuit requires only 2,696 gates. This is a notable reduction, but the circuit is still much too large considering the errors incurred by each gate application. We remedy this by dividing the circuit into many smaller sub-circuits, and running each of these components as its own calculation. By initializing each component based on a classically simulated result from the previous component (rather than continuing from an imprecise result as a direct calculation would), the compounding effect of gate errors can be eliminated. We can then verify the validity of our solution by confirming that the results of each component agree with the simulated results for the same component.

In total, the circuit was divided into 120 components as listed in appendix B. Initializations were performed using Qiskit's built-in initialization function using the state vectors produced by classical simulation of the previous components. The results of the calculation are shown in Fig. 10. We see that good agreement is maintained between the simulated and quantum procedures until the more complex later stages, when gate counts approach and exceed 800. At this point, calculations quickly become unreliable. Note that the sum of the number of gates required for all sub-circuits is significantly larger than the number of gates reported for the direct circuit. This is due to the reinitialization of the system, which adds an unpredictable overhead cost to each operation. The average relative error of the final simulated solution compared to a direct classical solution is 1.12×10^{-10} , which is well within the tolerances prescribed by the program.



(a)

(b)

Figure 10

Gate counts (a) and errors (b) for each circuit component in the solution procedure. Note that, due to the complex intermediate states, most gates are dedicated to component initializations.

Simulated Solution

With space available for a slightly larger system, we consider a practical, though still trivial, problem. We consider a calculation of the charge per unit length on a transmission line consisting of two long, conducting strips of known potentials radiating in free space. This problem can be modeled as a method of moments discretization of a cross-section of the line, yielding a matrix equation for the element charges $|Q\rangle$ as a function of their potentials $|V\rangle$:

$$B|Q\rangle = |V\rangle. \quad (82)$$

Elements of the matrix B are given by (Sadiku, 2001):

$$B_{jk} = \begin{cases} -\frac{l_j}{2\pi\epsilon_0} [\ln(l_j) - 1.5], & j = k \\ -\frac{l_j}{2\pi\epsilon_0} \ln(d_{jk}), & j \neq k. \end{cases} \quad (83)$$

where l_j is the length of the j -th one-dimensional boundary element and d_{jk} is the separation between the centroids of elements j and k . By using elements of uniform size l , we can ensure that B is Hermitian. We give each strip unit width, and orient both strips parallel to each other with unit separation.

The calculated charge density for each element is shown in Fig. 11, and the solution vector and its error are given in table 1. For this calculation, four unknowns and seven phase qubits were used. We see that the quantum results are within $\sim 3\%$ of the correct solution, showing that our procedure produces a reliable approximation. For more accurate solutions, a more complex system involving additional phase qubits would need to be simulated.

Table 1

Comparison of Simulated Quantum and Direct Classical Solution Vectors

Element Index	Quantum [nC]	Classical [nC]	Relative Error
0	0.0383	0.0371	0.0315
1	0.0383	0.0371	0.0315
2	-0.0383	-0.0371	0.0315
3	-0.0383	-0.0371	0.0315

Preconditioner Application

Here we consider the calculation of charge density per unit length on a two-conductor rectangular transmission line using the same method of moments approach as in the previous section. For reference, a classical solution is shown in Fig. 12. This geometry is too complex for our quantum




Figure 11

Solved charge distribution. The left conductor was held at a 1.0V potential, while the right conductor was held at $-1.0V$. Distances are shown in mm.

procedure to solve outright in satisfactory detail, but the larger density of charge elements provides a more enlightening example for preconditioning than the previous two-strip example. To generate the preconditioner, we follow the same procedure as when generating the total system matrix, but only include elements corresponding to positions which lie within a certain distance δ of each other (we use four times the basic element size). Explicitly, for preconditioner P ,

$$P_{jk} = \begin{cases} -\frac{l_j}{2\pi\epsilon_0} [\ln(l_j) - 1.5], & j = k \text{ and } d_{jk} \leq \delta \\ -\frac{l_j}{2\pi\epsilon_0} \ln(d_{jk}), & j \neq k \text{ and } d_{jk} \leq \delta \\ 0, & \text{otherwise.} \end{cases} \quad (84)$$

By applying P^{-1} to the right-hand side vector of the matrix equation prior to performing an iterative solution procedure, the total number of iterations needed to solve the system can be drastically reduced (Pan & Sheng, 2014). However, even though P is generally sparse, P^{-1} can still be quite dense, as illustrated in Fig. 13. This makes efficient classical application of the inverse preconditioner difficult, and often impossible. However, we have already shown that our quantum procedure can apply P in $\mathcal{O}(N_{nz} \log(N))$ time, regardless of the sparsity of P^{-1} .





Figure 12

Classically solved charge distribution. Distances are shown in mm.

Fig. 14 shows the time needed to apply the inverse preconditioner classically using GMRES for various N . For the preconditioning scheme used, matrices contained 5 elements per row, with no significant dependence on N , and hence $N_{nz} = \mathcal{O}(N)$. We see that for large matrices, the time required increases roughly as $\mathcal{O}(N^3)$, showing that no advantage whatsoever is gained from the sparsity of the initial matrix. For the quantum procedure, Fig. 15 shows the size in gates of the final preconditioner inversion circuits, as well as the time taken to compute these circuits. We see that the expected $\mathcal{O}(N \log(N))$ scaling is obtained, and therefore the quantum procedure is, asymptotically, more efficient than the classical procedure.

It is worth noting that the quantum procedure takes a very long time to produce a solution circuit for even modest problem sizes, limiting practical applicability to extremely large problems where the improved scaling can overcome the initial time difference. Fig. 16 gives a first-order extrapolation of our results to provide a rough prediction of the point where the quantum procedure can be expected to outperform a classical system. We predict this to hold true for problems of more than roughly 8×10^5 unknowns, which would certainly indicate a relatively sophisticated problem.

Unfortunately, we are not able to analyze the time complexity of the solution circuit execution due to the limitations of available hardware. The total number of qubits required for a system of

**Figure 13**

Nonzeros in preconditioner P (a) and P^{-1} (b).

size N is


$$2n + 2(n - 1) + 2 + n_p + 1 = 4n + n_p + 1, \quad (85)$$

accounting for the base registers and their ancillae, the control registers, the phase register (which could also require its own work register for controls, but we consider it to have constant size and therefore omit this extra dependence), and the HHL ancilla. Even with a trivial single-qubit phase estimation, a system of 26 qubits would be required to invert our smallest preconditioner ($N = 64$). Since the largest system we can access provides only 7 qubits, this calculation is quite beyond the capabilities of available hardware. Note that, by this calculation, 82 qubits would be required at our estimated crossover point of quantum/classical efficiency (using $n = 20$ as the smallest sufficient power of two).

Conclusion


By combining the known HHL algorithm with a unitary adopted from quantum walk research, we have developed a method for the solution of any well-conditioned matrix equation which is suitable for direct implementation on quantum hardware. The method has demonstrated $\mathcal{O}(N_{nz} \log(N))$ complexity in time and gate count, and for sparse matrices is expected to outperform classical solvers for problems of $N \gtrsim 8 \times 10^5$ unknowns.

While the method itself is suitable for practical implementation, available quantum systems are far too small and unreliable to support the analysis of any nontrivial problem. This situation is expected to improve rapidly, however, with IBM projecting the release of a 1,000 qubit system by the end of 2023 (Cho, 2021). They also anticipate a dramatic decrease in errors (Council, 2021).

**Figure 14**

Time to apply inverse preconditioner classically using GMRES.

Google has also indicated plans to develop a commercial-grade system by 2029 (Castellanos, 2021).



(a)

(b)

Figure 15

Size in gates (a) of quantum circuits, and corresponding time to compute (b). $\mathcal{O}(N \log N)$ line shown for comparison.




Figure 16

Simple extrapolation from computed data sets to estimate quantum/classical efficiency crossover. $\mathcal{O}(N \log N)$ and $\mathcal{O}(N^3)$ trend lines are used for the quantum and classical extrapolation, respectively.

References

- Ambainis, A. (2004, October). Quantum walk algorithm for element distinctness. In *Proc. 45th annu. IEEE symp. on found. of comput. sci.* (pp. 22–31). doi: 10.1109/FOCS.2004.54
- ANIS, M. S., Abby-Mitchell, Abraham, H., AduOffei, Agarwal, R., Agliardi, G., . . . Čepulkovskis, M. (2021). *Qiskit: An open-source framework for quantum computing*. doi: 10.5281/zenodo.2573505
- Barenco, A., Bennett, C. H., Cleve, R., DiVincenzo, D. P., Margolus, N., Shor, P., . . . Weinfurter, H. (1995, Nov). Elementary gates for quantum computation. *Phys. Rev. A*, 52, 3457–3467. doi: 10.1103/PhysRevA.52.3457
- Berry, D. W., & Childs, A. M. (2012, January). Black-box hamiltonian simulation and unitary implementation. *Quantum Info. Comput.*, 12, 29–62.
- Casacuberta, S., & Kyng, R. (2021). *Faster sparse matrix inversion and rank computation in finite fields*. (arXiv:2106.09830)
- Castellanos, S. (2021, May). Google aims for commercial-grade quantum computer by 2029. *WSJ*. (Accessed on: Dec. 2, 2021. [Online]. Available: <https://www.wsj.com/articles/google-aims-for-commercial-grade-quantum-computer-by-2029-11621359156>)
- Chandrupatla, T. R., & Belegundu, A. D. (2012). *Introduction to finite elements in engineering* (4th ed.). Harlow, England: Pearson.
- Chen, Z., Zheng, S., & Okhmatovski, V. I. (2019, June). Tensor train accelerated solution of volume integral equation for 2-d scattering problems and magneto-quasi-static characterization of multiconductor transmission lines. *IEEE Trans. Microw. Theory Tech.*, 67, 2181–2196. doi: 10.1109/TMTT.2019.2908368
- Chew, W. C., Jin, J. M., Michielssen, E., & Song, J. (2001). *Fast and efficient algorithms in computational electromagnetics* (1st ed.). Norwood, MA, USA: Artech House.
- Childs, A. M. (2010, March). On the relationship between continuous- and discrete-time quantum walk. *Commun. Math. Phys.*, 294, 581–603. doi: 10.1007/s00220-009-0930-1
- Cho, A. (2021, September). Ibm promises 1000-qubit quantum computer—a milestone—by 2023. *Science*. (Accessed on: Dec. 2, 2021. [Online]. Available: <https://www.science.org/content/article/ibm-promises-1000-qubit-quantum-computer-milestone-2023>)
- Cleve, R., Ekert, A., Macchiavello, C., & Mosca, M. (1998, January). Quantum algorithms revisited. *Proc. R. Soc. Lond. A.*, 454, 339–354. doi: 10.1098/rspa.1998.0164
- Council, J. (2021, March). The decade of quantum computing is upon us, ibm exec says. *WSJ*. (Accessed on: Dec. 2, 2021. [Online]. Available: <https://www.wsj.com/articles/the-decade-of-quantum-computing-is-upon-us-ibm-exec-says-11614802190>)

- Gerjuoy, E. (2005, May). Shor's factoring algorithm and modern cryptography. an illustration of the capabilities inherent in quantum computers. *Amer. J. Phys.*, *73*, 521–540. doi: 10.1119/1.1891170
- Gibson, W. C. (2008). *The method of moments in electromagnetics* (1st ed.). Boca Raton, FL, USA: Chapman Hall/CRC.
- Grover, L. (1996, July). A fast quantum mechanical algorithm for database search. In *Proc. 28th annu. acm symp. theory comput.* (pp. 212–219). doi: 10.1109/TMTT.2019.2908368
- Hall, B. C. (2015). *Lie groups, lie algebras, and representations: An elementary introduction* (2nd ed.). Cham, Switzerland: Springer.
- Harrow, A. W., Hassidim, A., & Lloyd, S. (2009, October). Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, *103*, 150502. doi: 10.1103/PhysRevLett.103.150502
- Horn, R. A., & Johnson, C. R. (2013). *Matrix analysis* (2nd ed.). New York, NY, USA: CUP.
- IBM Quantum*. (2021). (<https://quantum-computing.ibm.com/>)
- Kempe, J. (2004, October). Quantum random walks: An introductory overview. *Contemporary Physics*, *44*, 307–327. doi: 10.1080/00107151031000110776
- Kerenidis, I., & Prakash, A. (2020, February). Quantum gradient descent for linear systems and least squares. *Phys. Rev. A*, *101*, 022316. doi: 10.1103/PhysRevA.101.022316
- Nielsen, M. A., & Chuang, I. L. (2000). *Quantum computing and quantum information* (1st ed.). Cambridge, UK: CUP.
- Omar, S., & Jiao, D. (2015, March). A linear complexity direct volume integral equation solver for full-wave 3-d circuit extraction in inhomogeneous materials. *IEEE Trans. Microw. Theory Techn.*, *63*, 897–912.
- Özişik, M. N., Orlande, H. R. B., Colaço, M. J., & Cotta, R. M. (2017). *Finite difference methods in heat transfer* (2nd ed.). Boca Raton, FL, USA: CRC Press.
- Pan, X., & Sheng, X. (2014, November). Sparse approximate inverse preconditioner for multiscale dynamic electromagnetic problems. *Radio Science*, *49*, 1041–1051. doi: 10.1002/2014RS005387
- Phillips, C. D., & Okhmatovski, V. I. (2021). *QMES*. (Available: <https://github.com/philli69/QMES>)
- Sadiku, M. N. O. (2001). *Numerical techniques in electromagnetics* (2nd ed.). Boca Raton, FL, USA: CRC Press.
- Szegedy, M. (2004, October). Quantum speed-up of markov chain based algorithms. In *Proc. 45th annu. IEEE symp. on found. of comput. sci.* (pp. 32–41). doi: 10.1109/FOCS.2004.53
- Takahashi, T., Coulier, P., & Darve, E. (2017, July). Application of the inverse fast multipole method as a preconditioner in a 3d helmholtz boundary element method. *J. Comp. Phys.*, *341*, 406–428.

Venegas-Andraca, S. E. (2012, July). Quantum walks: a comprehensive review. *Quantum Inf. Process.*, *11*, 1015–1106. doi: 10.1007/s11128-012-0432-5

Conclusion to Thesis

Here we have presented a closed-form method for the implementation of an HHL-style quantum matrix equation solver. The method does not depend on any oracles or black box operations, and hence it provides an assessment of the practicality of quantum procedures for matrix equation analysis. While this certainly does not prove any lower bound for the complexity of such procedures, the demonstrated efficiency of the method does show that, for arbitrary problems which preclude fast classical algorithms, quantum procedures can provide significantly greater efficiency than their classical counterparts.

As mentioned in the introduction, the future of this research will likely include the analysis of more sophisticated versions of the HHL algorithm. These may be able to provide improved scaling with respect to the condition number of the system matrix as well as the precision of the solution vector. Additionally, for problems having a specific form, it may be possible to develop more sophisticated oracle implementations, thereby providing optimized quantum procedures that are competitive with their specialized classical analogues.

Appendix A

Reflection Operators

Here we briefly describe the nature of reflection operators. In particular, we consider operators of the form $2|v\rangle\langle v| - I$, where $|v\rangle$ is some normalized vector. This operator, when applied to another vector $|w\rangle$ of the same dimension, reflects $|w\rangle$ about $|v\rangle$. As an initial example, consider the two-dimensional case where $|v\rangle$ is the unit vector along the y axis, $|\hat{y}\rangle$, and $|w\rangle = a|\hat{x}\rangle + b|\hat{y}\rangle$ is arbitrary. Then the effect of the reflection operator is as follows:

$$\begin{aligned}
 (2|\hat{y}\rangle\langle\hat{y}| - I)(a|\hat{x}\rangle + b|\hat{y}\rangle) &= (2a|\hat{y}\rangle\langle\hat{y}|\hat{x}\rangle + 2b|\hat{y}\rangle\langle\hat{y}|\hat{y}\rangle) - (a|\hat{x}\rangle + b|\hat{y}\rangle) \\
 &= 2b|\hat{y}\rangle - (a|\hat{x}\rangle + b|\hat{y}\rangle) \\
 &= -a|\hat{x}\rangle + b|\hat{y}\rangle.
 \end{aligned} \tag{86}$$

That is, the x component of $|w\rangle$ is negated, while the y component remains unchanged. Fig. A1 illustrates this behavior, and emphasizes the reflection of the vector.




Figure A1

Two-dimensional reflection about the y -axis.

In general, the operand vector $|w\rangle$ consists of components both perpendicular and parallel to $|v\rangle$. Thus the vector can be written as $|w\rangle = a|w_{\parallel}\rangle + b|w_{\perp}\rangle$, and the action of the reflector is

$$\begin{aligned}
 (2|v\rangle\langle v| - I)(a|w_{\parallel}\rangle + b|w_{\perp}\rangle) &= (2a|v\rangle\langle v|w_{\parallel}\rangle + 2b|v\rangle\langle v|w_{\perp}\rangle) - (a|w_{\parallel}\rangle + b|w_{\perp}\rangle) \\
 &= 2a|w_{\parallel}\rangle - (a|w_{\parallel}\rangle + b|w_{\perp}\rangle) \\
 &= a|w_{\parallel}\rangle - b|w_{\perp}\rangle.
 \end{aligned} \tag{87}$$

Then the arbitrary reflector $2|v\rangle\langle v| - I$ has the effect of negating the perpendicular component of

the operand vector, while leaving the parallel component unchanged.

Appendix B

Quantum Circuit Components

This section lists the individual operations which constitute each component of the divided circuit described in section “A Fully Quantum Problem”. Operations are given in Qiskit-style syntax similar to that used in the program. Operators are listed first, with the operand registers listed on the following line. We assume that C , X , and d have been defined. Note that each of these components also includes the operations necessary to initialize it to the ideal state resulting from the precise application of all previous components. These operations, applied in sequence to an appropriately initialized system, constitute the solution procedure for the relevant system.

```

1  # initial T0
2  HGate()
3      reg_r2
4
5  # QPE
6  HGate()
7      [reg_phase[0]]
8  HGate()
9      [reg_phase[1]]
10
11 HGate().control(2,None,'01')
12     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2[0]]
13 XGate().control(2,None,'11')
14     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2a[0]]
15 ZGate().control(1,None,'1')
16     [reg_phase[0]]+[reg_r2[0]]
17 XGate().control(1,None,'1')
18     [reg_phase[0]]+[reg_r2[0]]
19 ZGate().control(1,None,'1')
20     [reg_phase[0]]+[reg_r2[0]]
21 XGate().control(1,None,'1')
22     [reg_phase[0]]+[reg_r2[0]]
23 XGate().control(2,None,'01')
24     [reg_phase[0]]+reg_r2[:]+reg_r2a[:]
25 ZGate().control(2,None,'01')
26     [reg_phase[0]]+reg_r2[:]+reg_r2a[:]
27 XGate().control(2,None,'01')
28     [reg_phase[0]]+reg_r2[:]+reg_r2a[:]
29 HGate().control(2,None,'01')
30     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2[0]]
31 XGate().control(2,None,'11')
32     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2a[0]]
33 SwapGate().control(1,None,'1')
34     [reg_phase[0]]+reg_r1[:]+reg_r2[:]
35 SwapGate().control(1,None,'1')
36     [reg_phase[0]]+reg_r1a[:]+reg_r2a[:]
37 SGate().control(1,None,'1')
38     [reg_phase[0]]+[reg_r1[0]]
39 XGate().control(1,None,'1')
40     [reg_phase[0]]+[reg_r1[0]]
41 SGate().control(1,None,'1')
42     [reg_phase[0]]+[reg_r1[0]]
43 XGate().control(1,None,'1')
44     [reg_phase[0]]+[reg_r1[0]]
45
46 HGate().control(2,None,'01')
47     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
48 XGate().control(2,None,'11')
49     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
50 ZGate().control(1,None,'1')
51     [reg_phase[1]]+[reg_r2[0]]
52 XGate().control(1,None,'1')
53     [reg_phase[1]]+[reg_r2[0]]
54 ZGate().control(1,None,'1')
55     [reg_phase[1]]+[reg_r2[0]]

```

```

56 XGate().control(1,None,'1')
57     [reg_phase[1]]+[reg_r2[0]]
58 XGate().control(2,None,'01')
59     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
60 ZGate().control(2,None,'01')
61     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
62 XGate().control(2,None,'01')
63     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
64 HGate().control(2,None,'01')
65     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
66 XGate().control(2,None,'11')
67     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
68 SwapGate().control(1,None,'1')
69     [reg_phase[1]]+reg_r1[:]+reg_r2[:]:]
70 SwapGate().control(1,None,'1')
71     [reg_phase[1]]+reg_r1a[:]+reg_r2a[:]:]
72 SGate().control(1,None,'1')
73     [reg_phase[1]]+[reg_r1[0]]
74 XGate().control(1,None,'1')
75     [reg_phase[1]]+[reg_r1[0]]
76 SGate().control(1,None,'1')
77     [reg_phase[1]]+[reg_r1[0]]
78 XGate().control(1,None,'1')
79     [reg_phase[1]]+[reg_r1[0]]
80
81 HGate().control(2,None,'01')
82     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
83 XGate().control(2,None,'11')
84     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
85 ZGate().control(1,None,'1')
86     [reg_phase[1]]+[reg_r2[0]]
87 XGate().control(1,None,'1')
88     [reg_phase[1]]+[reg_r2[0]]
89 ZGate().control(1,None,'1')
90     [reg_phase[1]]+[reg_r2[0]]
91 XGate().control(1,None,'1')
92     [reg_phase[1]]+[reg_r2[0]]
93 XGate().control(2,None,'01')
94     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
95 ZGate().control(2,None,'01')
96     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
97 XGate().control(2,None,'01')
98     [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:]
99 HGate().control(2,None,'01')
100     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
101 XGate().control(2,None,'11')
102     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
103 SwapGate().control(1,None,'1')
104     [reg_phase[1]]+reg_r1[:]+reg_r2[:]:]
105 SwapGate().control(1,None,'1')
106     [reg_phase[1]]+reg_r1a[:]+reg_r2a[:]:]
107 SGate().control(1,None,'1')
108     [reg_phase[1]]+[reg_r1[0]]
109 XGate().control(1,None,'1')
110     [reg_phase[1]]+[reg_r1[0]]
111 SGate().control(1,None,'1')
112     [reg_phase[1]]+[reg_r1[0]]
113 XGate().control(1,None,'1')
114     [reg_phase[1]]+[reg_r1[0]]
115
116 SwapGate()
117     reg_phase
118 HGate()
119     [reg_phase[0]]
120 CZGate().power(0.5).inverse()
121     reg_phase
122 HGate()
123     [reg_phase[1]]
124
125 # HHL ancilla rotation
126 RYGate(2.0*math.acos(C/(-d))).control(nq_phase,None,'00')
127     reg_phase[:]+reg_a_hhl[:]:]
128 RYGate(2.0*math.acos(C/(X-d))).control(nq_phase,None,'01')
129     reg_phase[:]+reg_a_hhl[:]:]

```

```

130 RYGate(2.0*math.acos(C/(-d))).control(nq_phase,None,'10')
131   reg_phase[:]+reg_a_hhl[:]:
132 RYGate(2.0*math.acos(C/(-X-d))).control(nq_phase,None,'11')
133   reg_phase[:]+reg_a_hhl[:]:
134
135 # inverse QPE
136 HGate()
137   [reg_phase[1]]
138 CZGate().power(0.5)
139   reg_phase
140 HGate()
141   [reg_phase[0]]
142 SwapGate()
143   reg_phase
144
145 XGate().control(1,None,'1')
146   [reg_phase[1]]+[reg_r1[0]]
147 SGate().inverse().control(1,None,'1')
148   [reg_phase[1]]+[reg_r1[0]]
149 XGate().control(1,None,'1')
150   [reg_phase[1]]+[reg_r1[0]]
151 SGate().inverse().control(1,None,'1')
152   [reg_phase[1]]+[reg_r1[0]]
153 SwapGate().control(1,None,'1')
154   [reg_phase[1]]+reg_r1a[:]+reg_r2a[:]:
155 SwapGate().control(1,None,'1')
156   [reg_phase[1]]+reg_r1[:]+reg_r2[:]:
157 XGate().control(2,None,'11')
158   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
159 HGate().control(2,None,'01')
160   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
161 XGate().control(2,None,'01')
162   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
163 ZGate().control(2,None,'01')
164   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
165 XGate().control(2,None,'01')
166   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
167 XGate().control(1,None,'1')
168   [reg_phase[1]]+[reg_r2[0]]
169 ZGate().control(1,None,'1')
170   [reg_phase[1]]+[reg_r2[0]]
171 XGate().control(1,None,'1')
172   [reg_phase[1]]+[reg_r2[0]]
173 ZGate().control(1,None,'1')
174   [reg_phase[1]]+[reg_r2[0]]
175 XGate().control(2,None,'11')
176   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
177 HGate().control(2,None,'01')
178   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
179
180 XGate().control(1,None,'1')
181   [reg_phase[1]]+[reg_r1[0]]
182 SGate().inverse().control(1,None,'1')
183   [reg_phase[1]]+[reg_r1[0]]
184 XGate().control(1,None,'1')
185   [reg_phase[1]]+[reg_r1[0]]
186 SGate().inverse().control(1,None,'1')
187   [reg_phase[1]]+[reg_r1[0]]
188 SwapGate().control(1,None,'1')
189   [reg_phase[1]]+reg_r1a[:]+reg_r2a[:]:
190 SwapGate().control(1,None,'1')
191   [reg_phase[1]]+reg_r1[:]+reg_r2[:]:
192 XGate().control(2,None,'11')
193   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
194 HGate().control(2,None,'01')
195   [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
196 XGate().control(2,None,'01')
197   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
198 ZGate().control(2,None,'01')
199   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
200 XGate().control(2,None,'01')
201   [reg_phase[1]]+reg_r2[:]+reg_r2a[:]:
202 XGate().control(1,None,'1')
203   [reg_phase[1]]+[reg_r2[0]]

```



```

204 ZGate().control(1,None,'1')
205     [reg_phase[1]]+[reg_r2[0]]
206 XGate().control(1,None,'1')
207     [reg_phase[1]]+[reg_r2[0]]
208 ZGate().control(1,None,'1')
209     [reg_phase[1]]+[reg_r2[0]]
210 XGate().control(2,None,'11')
211     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2a[0]]
212 HGate().control(2,None,'01')
213     [reg_phase[1]]+[reg_r1a[0]]+[reg_r2[0]]
214
215 XGate().control(1,None,'1')
216     [reg_phase[0]]+[reg_r1[0]]
217 SGate().inverse().control(1,None,'1')
218     [reg_phase[0]]+[reg_r1[0]]
219 XGate().control(1,None,'1')
220     [reg_phase[0]]+[reg_r1[0]]
221 SGate().inverse().control(1,None,'1')
222     [reg_phase[0]]+[reg_r1[0]]
223 SwapGate().control(1,None,'1')
224     [reg_phase[0]]+[reg_r1a[:]]+[reg_r2a[:]]
225 SwapGate().control(1,None,'1')
226     [reg_phase[0]]+[reg_r1[:]]+[reg_r2[:]]
227 XGate().control(2,None,'11')
228     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2a[0]]
229 HGate().control(2,None,'01')
230     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2[0]]
231 XGate().control(2,None,'01')
232     [reg_phase[0]]+[reg_r2[:]]+[reg_r2a[:]]
233 ZGate().control(2,None,'01')
234     [reg_phase[0]]+[reg_r2[:]]+[reg_r2a[:]]
235 XGate().control(2,None,'01')
236     [reg_phase[0]]+[reg_r2[:]]+[reg_r2a[:]]
237 XGate().control(1,None,'1')
238     [reg_phase[0]]+[reg_r2[0]]
239 ZGate().control(1,None,'1')
240     [reg_phase[0]]+[reg_r2[0]]
241 XGate().control(1,None,'1')
242     [reg_phase[0]]+[reg_r2[0]]
243 ZGate().control(1,None,'1')
244     [reg_phase[0]]+[reg_r2[0]]
245 XGate().control(2,None,'11')
246     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2a[0]]
247 HGate().control(2,None,'01')
248     [reg_phase[0]]+[reg_r1a[0]]+[reg_r2[0]]
249
250 HGate()
251     [reg_phase[1]]
252 HGate()
253     [reg_phase[0]]
254
255 # inverse T0
256 HGate()
257     reg_r2

```