

# The discrete adiabatic quantum linear system solver has lower constant factors than the randomised adiabatic solver

Pedro C. S. Costa<sup>1,2</sup>, Dong An<sup>3</sup>, Ryan Babbush<sup>4</sup>, and Dominic W. Berry<sup>2</sup>

<sup>1</sup>ContinoQuantum, Sydney, NSW 2093, AU

<sup>2</sup>School of Mathematical and Physical Sciences, Macquarie University, Sydney, NSW 2109, AU

<sup>3</sup>Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, USA

<sup>4</sup>Google Quantum AI, Venice, CA 90291, USA

The solution of linear systems of equations is the basis of many other quantum algorithms, and recent results provided an algorithm with optimal scaling in both the condition number  $\kappa$  and the allowable error  $\epsilon$  [PRX Quantum 3, 040303 (2022)]. That work was based on the discrete adiabatic theorem, and worked out an explicit constant factor for an upper bound on the complexity. Here we show via numerical testing on random matrices that the constant factor is in practice about 1,200 times smaller than the upper bound found numerically in the previous results. That means that this approach is far more efficient than might naively be expected from the upper bound. In particular, it is about an order of magnitude more efficient than using a randomised approach from [arXiv:2305.11352] that claimed to be more efficient.

## 1 Introduction

Finding solutions to a system of linear equations is of great importance in many areas of science and technology, and is the foundation for many quantum algorithms. When we consider classical approaches for solving linear system problems, we have in the best cases a total number of steps proportional to the number of equations, which can make the solution prohibitively costly when we have a large linear systems problem. On the other hand, quantum computers allow us to deal with the linear systems problems in these prohibitive regions by providing the output as a quantum state. The first proposal for the quantum linear system problem (QLSP) was given by Harrow, Hassidim, and Lloyd (HHL) [1], which was subsequently refined in later works [2, 3, 4, 5] with the optimal approach provided by Costa et al. in [6].

The goal of the QLSP is to efficiently produce a quantum state  $|x\rangle$  proportional to the solution of  $Ax = b$ , where  $A \in \mathbb{C}^{N \times N}$ . For simplicity we assume  $N = 2^n$  and  $\|A\| = 1$ , where  $\|\cdot\|$  denotes the spectral norm. The ground-breaking algorithm from

Pedro C. S. Costa: [pcs.costa@protonmail.com](mailto:pcs.costa@protonmail.com)

HHL produces the state  $|x\rangle$  with cost  $\mathcal{O}(\text{poly}(n)\kappa^2/\epsilon)$ , where  $\kappa = \|A\|\|A^{-1}\|$  is the condition number of  $A$ , and  $\epsilon$  is the maximum allowable error. Despite the exponential improvement given by HHL over the classical results on the dimension of the problem, e.g.,  $\mathcal{O}(N\sqrt{\kappa}\log(1/\epsilon))$  achieved by the conjugate gradient descent method, there was a plenty of scope to improve on the other parameters in the complexity.

Quantum algorithms aim to provide complexity polynomial in  $n$ , an exponential speedup over the classical methods. It is known [7] that the lower bound for scaling of the complexity of quantum algorithms in  $\kappa$  and  $\epsilon$  is  $\Omega(\kappa \log(1/\epsilon))$  (see Section A for details). This bound is in terms of calls to a block encoding for  $A$ , and the gate complexity will have a factor corresponding to the cost of block encoding  $A$ , usually taken to be  $\text{poly}(n)$  for efficient block encodings. That shows that the optimal scaling is linear in  $\kappa$  when aiming for the exponential quantum speedup in  $n$ , as opposed to the  $\sqrt{\kappa}$  classical complexity. This lower bound on the quantum complexity was achieved in [6], thereby providing the asymptotically optimal quantum algorithm.

The most significant improvements for the QLSP after [1] are based on adiabatic quantum computing (AQC) [2, 8, 9], as opposed to the quantum phase estimation and amplitude amplification approach of HHL. Reference [6] is based upon AQC, but the main new feature is a discrete formulation of the adiabatic theorem, where the time evolution is based on the quantum walk operator  $W_T(s)$ ; that enables the optimal complexity for the QLSP. We refer to that as the quantum walk (QW) method. In addition to the asymptotically scaling results for the QLSP in [6], the authors also provided strict upper bounds, which can be used to estimate the total number of quantum walk steps for a given solution error  $\epsilon$ .

The upper bound derived for the discrete adiabatic theorem (DAT) given in [6] results from a combination of several approximations from different operators. The approximations allow them to show that the total number of walk steps grows linearly in  $\kappa$ , i.e.,  $T = \alpha\kappa/\Delta$  for allowable error  $\Delta$  when the parameters of the QLSP are plugged in the DAT. However, due to the loose upper bound from DAT, the value for  $\alpha$

is very large. The constant factor  $\alpha = 2305$ , given on page 16 of [6], can be used only for a rough estimate of the total number of quantum walk steps needed to yield a desired error.

An earlier proposal for solving the QLSP is from Ref. [9], which uses time evolution for randomised times in order to emulate measurements for the quantum Zeno effect. Here, we refer to that as the Randomized method (RM). The complexity of that approach scales as  $\mathcal{O}(\kappa \log(\kappa/\epsilon))$ , so has an extra factor of  $\log \kappa$  in the complexity over the QW method. A later analysis in Ref. [10] showed a relatively tight upper bound for the RM, and argued that it would be faster in practice than the QW method due to the smaller constant factor.

Here, we test the actual performance of the QW method and RM by performing numerical simulation of the solution with many samples of randomly generated matrices. Our numerical tests show that in practice, the constant factor for the QW method is  $\alpha = 1.84$ , which is about 1,200 times smaller than what is given in [6]. Moreover, our tests demonstrate that the quantum walk method is, on average, approximately 7 times more efficient than the RM. This results in an overall improvement of at least a factor of 3 in total complexity, including the filtering step for both methods, for realistic target errors  $\epsilon$  of the solution even without accounting for the additional overhead associated with simulating time evolution.

The background of the discrete quantum walk and Randomized method are presented in Section 2. Details about the numerical tests for the quantum walk and Randomized method are given in Section 3. In Section 4, we discuss the cost of implementation of the filtering method, which is applied after the walk steps, to estimate the appropriate target precision for the adiabatic evolution. We then conclude in Section 5. In appendices we provide technical details about the lower bound on the complexity, circuit construction of the walk operators, and further numerical tests.

## 2 Background

In this section we summarise the background of how the discrete quantum walk method and Randomized method are performed. Readers who are familiar with the background may wish to skip to the results in Section 3.

### 2.1 Quantum walk method

We start this section by introducing the basic theory of adiabatic quantum computation which is independent of the application. In AQC, we have a time-dependent system constructed from a linear interpolation between two time-independent Hamiltonians as follows

$$H(s) = (1 - f(s))H_0 + f(s)H_1, \quad (1)$$

where the function  $f(s) : [0, 1] \rightarrow [0, 1]$  is called the schedule function. Normally we have an eigenpath between these Hamiltonians where the starting eigenstate of  $H_0$  is easy to prepare, and the final eigenstate given by  $H_1$  is the one that encodes the solution to the problem that we are trying to determine.

For linear systems solvers, the final eigenstate given by  $H(1)$  should encode the normalised solution for a linear system. In other words, for  $A \in \mathbb{C}^{N \times N}$  an invertible matrix with  $\|A\| = 1$ , and a normalised vector  $|b\rangle \in \mathbb{C}^N$  the goal is to prepare a normalised state  $|\tilde{x}\rangle$ , which is an approximation of  $|x\rangle = A^{-1}|b\rangle / \|A^{-1}|b\rangle\|$ . For precision  $\epsilon$  of the approximation, we require  $\| |\tilde{x}\rangle - |x\rangle \| \leq \epsilon$ .

In the following subsections, we show how we should encode both  $H_0$  and  $H_1$  when our goal is to use them for solving QLSP when  $A$  is a Hermitian and positive-definite matrix, and also the case where  $A$  does not have to be a Hermitian matrix. In both cases, we also look at a family of scheduling functions for the problem that slows down the evolution according to the gap from the interpolated Hamiltonian given by Eq. (1).

After we set up the time-dependent Hamiltonian in both scenarios of  $A$ , we apply the method from [6] to transform the Hamiltonian into a discrete quantum walk, which is used as the evolution operator for the discrete adiabatic evolution. This discrete quantum walk is to obtain an approximate solution of the linear system, then a filter is used to ensure that the solution is obtained to high accuracy.

#### 2.1.1 Positive definite and Hermitian matrix $A$

As explained in [6] we have  $|\tilde{x}\rangle$  as the state achieved from the steps of the walk, and  $|x\rangle$  would be obtained from the ideal adiabatic evolution. Starting with the simplest case, where  $A$  is Hermitian and positive definite, one takes the Hamiltonians [2]

$$H_0 := \begin{pmatrix} 0 & Q_b \\ Q_b & 0 \end{pmatrix}, \quad (2)$$

and

$$H_1 := \begin{pmatrix} 0 & AQ_b \\ Q_b A & 0 \end{pmatrix}, \quad (3)$$

where  $Q_b = I_N - |b\rangle\langle b|$ . The state  $|0, b\rangle$  is an eigenstate of  $H_0$  with eigenvalue 0, and one would aim for this to evolve adiabatically to eigenstate  $|0, A^{-1}b\rangle$  of  $H_1$ .

Denoting the condition number of the matrix as  $\kappa$ , a lower bound for the gap of  $H(s)$  is [2]

$$\Delta_0(s) = 1 - f(s) + f(s)/\kappa. \quad (4)$$

As explained in [2, 6] the goal is to provide a schedule function that slows down the evolution as the gap becomes small, according to the condition given in [11]

$$\dot{f}(s) = d_p \Delta_0^p(s), \quad (5)$$

where  $f(0) = 0$ ,  $p > 0$  and  $d_p = \int_0^1 \Delta_0^{-p}(u) du$  is a normalisation constant chosen so that  $f(1) = 1$ . According to [2] the following schedule function

$$f(s) = \frac{\kappa}{\kappa - 1} \left[ 1 - (1 + s(\kappa^{p-1} - 1))^{\frac{1}{1-p}} \right], \quad (6)$$

satisfies Eq. (5).

Distinct from the continuous version of the adiabatic theorem, in the discrete version of the theorem, we have to take into account the smallest gap between the different groups of eigenvalues of  $W(s)$  for  $s, s + 1/T$ , and  $s + 2/T$ , which are given by

$$\Delta_k(s) = 1 - f(s + k/T) + f(s + k/T)/\kappa, \quad k = 0, 1, 2. \quad (7)$$

The methods for block encoding the walk operator are explained in Appendix B.

### 2.1.2 Non-Hermitian matrix $A$

To address the case where  $A$  is not positive definite or Hermitian, we pay an extra space cost by doubling its dimension to transform it into a Hermitian matrix before solving it in the quantum algorithm, i.e.,

$$\mathbf{A} := \begin{pmatrix} 0 & A \\ A^\dagger & 0 \end{pmatrix}. \quad (8)$$

As described in Ref. [2] the Hamiltonian is then constructed as

$$H(s) = \begin{pmatrix} 0 & A(f(s))Q_{\mathbf{b}} \\ Q_{\mathbf{b}}A(f(s)) & 0 \end{pmatrix}, \quad (9)$$

where

$$A(f) := (1-f)\sigma_z \otimes I_N + f\mathbf{A} = \begin{pmatrix} (1-f)I & fA \\ fA^\dagger & -(1-f)I \end{pmatrix}, \quad (10)$$

and  $Q_{\mathbf{b}} = I_{2N} - |0, b\rangle\langle 0, b|$  as the projection for this Hamiltonian formulation. As discussed in [6] it is equivalent to taking  $H(s) = (1 - f(s))H_0 + f(s)H_1$  with

$$\begin{aligned} H_0 &= \sigma_+ \otimes [(\sigma_z \otimes I_N)Q_{\mathbf{b}}] + \sigma_- \otimes [Q_{\mathbf{b}}(\sigma_z \otimes I_N)] \\ H_1 &= \sigma_+ \otimes [\mathbf{A}Q_{\mathbf{b}}] + \sigma_- \otimes [Q_{\mathbf{b}}\mathbf{A}], \end{aligned} \quad (11)$$

where  $\sigma_+ = |1\rangle\langle 0|$ ,  $\sigma_- = |0\rangle\langle 1|$ .

As in the positive-definite case, the schedule function can be determined from gap via the differential equation (5). In this case the gap is lower bounded as  $\sqrt{(1 - f(s))^2 + (f(s)/\kappa)^2}$  so there is not a closed-form solution of Eq. (5).

Instead, one can use the relation that for  $0 \leq f(s) \leq 1$ ,

$$\sqrt{(1 - f(s))^2 + (f(s)/\kappa)^2} \geq (1 - f(s) + f(s)/\kappa)/\sqrt{2}. \quad (12)$$

As a result of using this formula, and the way we use the rotation operator  $R(s)$ , we can use the same schedule function when  $A$  is not positive definite.

## 2.2 Randomized method

A randomised QLSP algorithm inspired by AQC was first proposed in Ref. [9] and further improved in recent work [10]. The method uses the same  $H_0$ ,  $H_1$ , and the same input state as the adiabatic quantum walk method, in order to approximate the desired eigenstate of  $H_1$  with zero eigenvalue encoding the solution.

Let  $H(v) = (1 - f(v))H_0 + f(v)H_1$  be the interpolating time-dependent Hamiltonian. Here we change the notation of the dimensionless time from  $s$  to  $v$  to avoid possible confusion, as in the Randomized method  $v$  does not range from 0 to 1. Specifically, let

$$f(v) = \frac{-\kappa^2 \exp\left(-v \frac{\sqrt{\kappa^2 + 1}}{\sqrt{2\kappa}}\right) + \exp\left(v \frac{\sqrt{\kappa^2 + 1}}{\sqrt{2\kappa}}\right) + 2\kappa^2}{2(\kappa^2 + 1)}, \quad (13)$$

and

$$v_a = \frac{\sqrt{2\kappa}}{\sqrt{\kappa^2 + 1}} \log\left(\kappa\sqrt{1 + \kappa^2} - \kappa^2\right), \quad (14)$$

$$v_b = \frac{\sqrt{2\kappa}}{\sqrt{\kappa^2 + 1}} \log\left(\sqrt{1 + \kappa^2} + 1\right). \quad (15)$$

Such a choice ensures that  $f(v_a) = 0$  and  $f(v_b) = 1$ . The Randomized method then implements the operation

$$e^{-it_q H(v_q)} \dots e^{-it_2 H(v_2)} e^{-it_1 H(v_1)} \quad (16)$$

on the input state. Here  $v_j = v_a + j(v_b - v_a)/q$ , and  $t_j$  are independent random variables.

In the original work [9],  $t_j$  was chosen to be uniformly distributed in  $[0, 2\pi/\Delta(v_j)]$ . Recent work [10] suggests to choose  $t_j$  according to the probability density function

$$p_j(t) \propto \left( \frac{J_p(\Delta(v_j)|t|/2)}{\Delta^{p-1}(v_j)|t|^p} \right)^2, \quad (17)$$

where  $J_p$  is a Bessel function of first kind and  $p = 1.165$ . It has the average time  $\langle |t_j| \rangle = 2.32132/\Delta(v_j)$ . We remark that the way of constructing the optimal probability distribution has been proposed in [12]. It uses

$$p_j(t) = \left( \frac{1}{\sqrt{2\pi}} \sum_{l \geq 0} a_l \int_{-1}^1 \cos(xt\Delta(v_j)/2) (1 - x^2)x^{2l} dx \right)^2, \quad (18)$$

where the coefficients  $a_l$  can be numerically computed by minimising  $\langle |t_j| \rangle$ . This optimal choice gives  $\langle |t_j| \rangle = 2.3160/\Delta(v_j)$ , which is slightly better than Eq. (17). The corresponding probability distributions (Eq. (17) and Eq. (18)) are almost the same (see Fig. 1 for a comparison), so in this work we still use Eq. (17).

Applying Eq. (16) to the input state produces the final state. It is shown in [10] that, in order to achieve  $1 - \delta$  fidelity, it suffices to choose  $q$  such that

$$\left( 1 - \frac{(v_b - v_a)^2}{q^2} \right)^q \geq 1 - \delta. \quad (19)$$




Figure 1: A comparison between two probability distribution functions used in the Randomized method. Here the legend “JLPSS” refers to the one in [10] and “Optimal” to the one in [12]. The gap  $\Delta(v_j)$  is chosen to be 0.5 in this test.

Asymptotically  $q \sim (v_b - v_a)^2/\delta \sim (\log \kappa)^2/\delta$ . In practice,  $\delta$  can be chosen to be  $\mathcal{O}(1)$ , and the state produced by the Randomized method will be passed into the filter to improve its fidelity.

To implement the randomised evolution operator in Eq. (16), Ref. [10] uses quantum signal processing (QSP) together with the Jacobi-Anger expansion. The cost of implementing an exponential  $e^{-itH(v)}$  is given as

$$3(e\alpha_H|t|/2 + \log(2c/\gamma)) \quad (20)$$

queries to the block encoding of  $H(v)$ , with  $c = 1.47762$  and  $\gamma$  the allowed error in each exponential. Therefore the overall number of queries to the block encoding of  $H(v)$  becomes

$$3 \left( \frac{e\alpha_H}{2} \sum_{j=1}^q |t_j| + q \log(2c/\gamma) \right). \quad (21)$$

That complexity analysis gives a non-asymptotic expression suitable for small times. In contrast, Ref. [13] gives an asymptotic analysis of the cost of implementing  $e^{-itH(v)}$  (with the same algorithm) to give the leading-order terms

$$2\alpha_H|t| + 3^{2/3}(\alpha_H|t|)^{1/3}(\log(1/\gamma))^{2/3} \quad (22)$$

for the number of queries to the block encoding of  $H(v)$ . The constant factor on  $\alpha_H|t|$  is only 2 here, which is better than the factor  $3e/2$  in the bound of Eq. (20). Recent work [14] removes a factor of 2 in the cost of quantum signal processing for Hamiltonian simulation. That would remove a factor of 2 from both complexity estimates. The leading-order term for the analysis of Ref. [13] would then be  $\alpha_H|t|$ . Therefore, with  $\alpha_H = 1$  the number of calls to the block encoding of  $H(v)$  should be approximately (and no less than)  $|t|$ . It is therefore reasonable to compare

the total time in the Randomized method with the number of steps in the quantum walk approach.

### 3 Numerical Tests

Our main results are given in Fig. 2, which includes the case where the matrix is  $A$  positive definite (PD) and Hermitian as well as the case where  $A$  is a general non-Hermitian matrix. For the comparison we are quantifying the complexity for the discrete walk approach by the number of calls to the block encoding of  $A$ , whereas for the Randomized method we measure the cost by the total evolution time. As discussed above, the total evolution time for the Randomized method enables us to lower bound the complexity in terms of block encodings, and thereby show that it is more costly than the discrete walk approach.

In Fig. 2 we can see an average complexity about 7 times worse for the randomised method than the QW method. This is comparing the number of calls to a block encoding of  $A$  for the QW to a Hamiltonian evolution time for the RM. In practice the simulation of the Hamiltonian evolution will incur extra overhead, so the advantage of the QW method will be even greater.

Moreover, we consider an improved probability distribution for the RM from Ref. [12], so the performance of the RM as simulated is better than that presented in Ref. [10]. We also stress that the quantum walk performance is not only better on average over the thousands of tested instances of problems, but also for each instance tested. With our results, we can avoid future misunderstanding of the performance of the method of Ref. [6].

In both methods for solving the QLSP, we tested the cases where the matrices  $A \in \mathbb{R}^{16 \times 16}$  are positive definite and non-Hermitian. As shown in Fig. 2, we consider matrices with a range of condition numbers  $\kappa = 10, 20, 30, 40, 50$ , where for each condition number we generate 100 independent samples. We also consider matrices with lower dimensions in Section C.

The plots show the geometric mean of the complexity over the 100 instances for each condition number considered. We also included the error bars, with both the interquartile range and the complete data range. The large size of the error bars results in part from the large value used as the threshold for the target error in the QLSP. In particular, for the quantum walk, the principle is that we start with a small value for the total number of steps. We then check the error between the actual and the ideal state. If it is above the target error, we keep increasing the total number of steps until we reach this error. However, the error between the ideal and the target state can vary rapidly with a small increment in the total number of steps, when we are not in the adiabatic regime. Consequently, we can have large variation in complexity between different instances of the matrices.



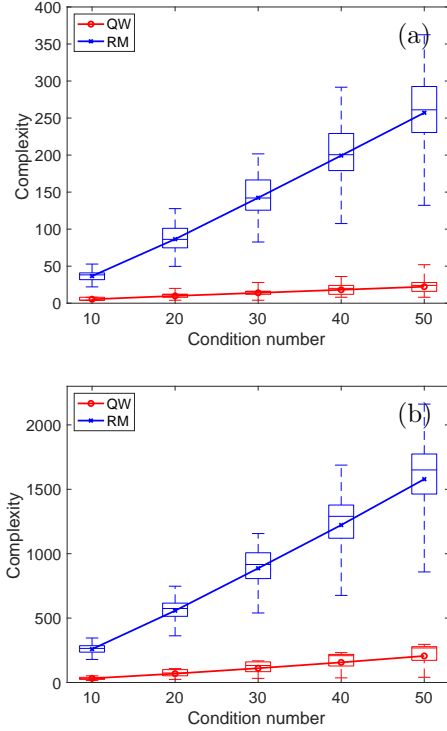


Figure 2: Comparison between the quantum walk method and the Randomized method. The complexity is quantified by the number of calls to the block encoding of the matrix for the quantum walk case and by the total evolution time for the Randomized approach until we get the 2-norm error,  $\Delta = 0.4$  in the QW method and the infidelity error  $\delta = 0.4$  in for the Randomized method. The complexities are computed by the geometric mean of the 100 instances of 16-dimensional matrices (solid line), and the box plots show their distributions. Part (a): positive-definite case. Part (b): general non-Hermitian case.

### 3.1 Quantum Walk method

We start testing the AQC in the discrete framework for solving QLSPs. For the positive-definite case the quantum circuit to construct the qubitised walk is shown in Fig. 7 of Section B. The results for dimension  $16 \times 16$  are given in the first plot of Fig. 2, and more detail and other matrix sizes are given in Section C. The results are given for 100 independent random matrices for each condition number. We computed the norm of the difference between the ideal and the actual state, aiming for  $\|\tilde{x} - |x\rangle\| \leq \Delta$  with  $\Delta = 0.4$  as the target error. The complete QLSP algorithm uses an adiabatic component followed by filtering, both for the QW method and RM, and we denote the error for the adiabatic part of the algorithm as  $\Delta$  to distinguish it from the final error  $\epsilon$  obtained after filtering. We use the schedule function as in Ref. [6] with  $p = 1.4$  (see Section 2.1). In the first plot of Fig. 2 (see also Fig. 9 for more details and other dimensions) we observe a roughly linear growth of walk steps in condition number, as theoretically proven in [6].

In the second plot of Fig. 2, given with further details in Fig. 10, we present the results for non-Hermitian matrices. See Section B for explanation of the block encoding of the qubitised walk and the source code to generate these results is accessible at [15]. Similarly to the PD case, we generate 100 independent random matrices for each condition number. In the non-Hermitian case, we again see the approximately linear growth of the walk steps in condition number.

We also provide results in Table 1 for non-Hermitian matrices, where rather than slowly increasing the total number of steps until we obtain the desired threshold, we use the same number of total steps for all instances. The total number is chosen to obtain root-mean-square (RMS) error (for 100 instances) as close as possible to  $\Delta = 0.4$  but never higher. We also include a table for PD matrices in Section C. From the results in Table 1, we can infer the constant factor  $\alpha$  in  $T = \alpha\kappa/\Delta$  which we will find in practice for the QLSP. In particular, for  $\kappa = 50$ , we find  $\alpha = 1.84$ , which gives a constant factor about 1,250 smaller than what is presented in Ref. [6].

In addition to the randomly tested matrices without a specific application, we included two matrices from the collection presented in [16], which provides a variety of sparse matrices used for benchmarking across different applications. The results for these practical cases are reported in Table 2.

### 3.2 Randomized method

In the Randomized method, we measure the cost by the total time, defined as  $\sum_{j=1}^q |t_j|$ , to achieve desired accuracy. Hamiltonian evolution may be simulated using calls to block encodings of the Hamiltonian via quantum signal processing. With  $\|H\| = 1$ , the simulation via the most efficient known scheme requires at least as many calls to the Hamiltonian as the time [14], though the cost can be larger (as a ratio to time) for shorter simulation times, as analysed in Ref. [10]. It is therefore reasonable to compare the total time here to the number of walk steps; it enables us to lower bound the true cost of the RM, to show that it is more costly than the QW. See Section 2.2 for more details, and the source code to generate these results are accessible at [17].

In our tests, we record the averaged minimal total time to reduce the RMS error below 0.4 over the 100 independent samples for each dimension and condition number, with the same sample matrices as for the QW method. For each sample we perform 200 repetitions with new randomly selected times. Such a choice of RMS error is comparable to the error threshold in the QW tests. We remark that the actual numerical RMS error is slightly smaller than 0.4, since in the Randomized method only the number of exponentials  $q$  is tunable and the resulting errors are

$\kappa$	Quantum Walk		Randomized Method		Ratio
	Number of Steps	$\Delta$	Averaged Time	$\Delta$	
10	36	0.348	281	0.395	7.81
20	76	0.381	604	0.397	7.95
30	120	0.400	963	0.398	8.03
40	176	0.399	1330	0.397	7.56
50	232	0.397	1722	0.399	7.42

Table 1: Comparison between the Quantum Walk and Randomized Methods for a range of condition numbers applied to non-Hermitian matrices of dimension  $16 \times 16$ . For the discrete adiabatic method, the RMS error, measured as the norm of the difference in quantum states, is computed over 100 instances for each condition number. For the randomized method, the average evolution time required to achieve an average error  $\Delta$  of approximately 0.4 is reported. The RMS error is also computed over 100 instances, with 200 repetitions per instance. The Ratio column shows the factor of improvement of the QW over RM.

Matrix	$\kappa$	Quantum Walk		Randomized Method		Ratio
		Number of Steps	$\Delta$	Averaged Time	$\Delta$	
Directed Graph – ID=168	$4.041 \times 10^2$	$2.416 \times 10^3$	0.398	$2.315 \times 10^4$	0.406	9.58
Circuit Simulation Problem – ID=1199	$6.302 \times 10^5$	$3.980 \times 10^4$	0.399	$1.817 \times 10^7$	0.402	457

Table 2: Performance comparison between the Quantum Walk and RM for two non-Hermitian matrices of dimension  $32 \times 32$  from [16]. For each matrix, we report its application, ID, condition number, and performance metrics including number of steps and error for QW, and averaged time and error for RM. The Ratio column shows the factor of improvement of the QW over RM.

changing discontinuously. As an illustration, results for 16-dimensional non-Hermitian matrices are shown in Table 4, and more results in PD cases can be found in Section C.

From Fig. 2, which can be seen with further details in Figs. 11 and 12, we observe roughly linear growth of time in condition number, as theoretically proven. We also observe that the total time for the non-Hermitian case is about 6.4 times larger than that for the PD case. In both PD and non-Hermitian cases, the overall complexity of the RM is significantly higher than that of the QW method. Because we are comparing evolution time for the RM to the number of steps for the QW, the advantage will be even greater when the complexity of implementing the time evolution is taken into account.

Same as the QW method, we also perform a numerical test of RM using the same two practical matrices from [16]. The results are reported in Table 2 and show that RM can be 9.2 times worse than QW in the directed graph example, and 4.0 times worse in the circuit simulation problem.

## 4 Cost with filter

Typically, we would aim to perform the discrete adiabatic algorithm followed by a filtering step to ensure the correct eigenvalue has been obtained; this filtering procedure is a common step in both the discrete adiabatic and randomized methods, and incurs similar cost in each case. First assume that the adiabatic

algorithm gives an infidelity error  $\delta$ , in the sense that the probability of measuring the correct eigenstate is  $1 - \delta$ . The relation between  $\Delta$  (2-norm error) and infidelity error  $\delta$  is  $\delta = \Delta^2 - \Delta^4/4$ . Then if we were to perform filtering as described in [6] (see Eq. (113) of that work), the complexity of the filter would be  $\kappa \ln(2/\epsilon_f)$ , with  $\epsilon_f$  the parameter for the filter. In that work it is shown that the 2-norm error at the end is upper bounded by  $\epsilon_f$  (for  $\delta > 1/2$ ), and  $\epsilon_f^2$  is also the maximum probability for failure (accepting an incorrect state as correct).

If the state fails the filtering step, then the adiabatic algorithm and filter are repeated. It should also be taken into account that the error probability will be increased if there are multiple failures and filters. The set of possible outcomes is as follows.

1. With probability  $1 - \delta$  the correct state is obtained the first time, which then passes the filtering test.
2. With probability  $\delta$  there is an incorrect state. With probability no greater than  $\delta\epsilon_f^2$  this state will be accepted regardless, and otherwise it will be flagged as a failure. The probability of this occurring is at least  $\delta(1 - \epsilon_f^2)$ .
3. To upper bound the probability of error we should take the probability of performing the procedure a second time as  $\delta(1 - \epsilon_f^2)$ . The probability of incorrectly accepting the wrong state is then upper bounded as  $\delta^2\epsilon_f^2(1 - \epsilon_f^2)$ .

In general, the probability of flagging a failure and

repeating the measurement  $n$  times is  $[\delta(1 - \epsilon_f^2)]^n$ . For the average number of repetitions, it will be, in this worst-case scenario for error,

$$\sum_{n=0}^{\infty} (n+1)[\delta(1 - \epsilon_f^2)]^n = \frac{1}{[1 - \delta(1 - \epsilon_f^2)]^2}. \quad (23)$$

For the costing, we need to bear in mind that for the cases where the filter fails the LCU approach will on average flag a failure after half the maximum number of operations. This means that the cost of the filter can be given as

$$\begin{aligned} & \kappa \ln(2/\epsilon_f) + \frac{1}{2} \kappa \ln(2/\epsilon_f) \left( \frac{1}{[1 - \delta(1 - \epsilon_f^2)]^2} - 1 \right) \\ &= \frac{1}{2} \kappa \ln(2/\epsilon_f) \left( \frac{1}{[1 - \delta(1 - \epsilon_f^2)]^2} + 1 \right) \end{aligned} \quad (24)$$

If the number of steps in the adiabatic algorithm to obtain error  $\Delta$  is  $\alpha\kappa/\Delta$ , then the cost of the adiabatic repetitions is

$$\frac{\alpha\kappa}{\Delta[1 - \delta(1 - \epsilon_f^2)]^2}, \quad (25)$$

for a total complexity

$$\frac{1}{2} \kappa \ln(2/\epsilon_f) \left( \frac{1}{[1 - \delta(1 - \epsilon_f^2)]^2} + 1 \right) + \frac{\alpha\kappa}{\Delta[1 - \delta(1 - \epsilon_f^2)]^2}. \quad (26)$$

Here we can use  $\delta = \Delta^2 - \Delta^4/4$  or  $\Delta = \sqrt{2 - 2\sqrt{1 - \delta}}$  to give the expression in terms of a single variable ( $\delta$  or  $\Delta$ ). The value of  $\kappa$  factors out, and given  $\alpha$  one can optimise  $\delta$  or  $\Delta$  for a given  $\epsilon_f$  to minimise the cost. In the case where the cost scales with  $1/\delta$  instead of  $1/\Delta$ , then the expression for the total complexity is changed slightly to

$$\frac{1}{2} \kappa \ln(2/\epsilon_f) \left( \frac{1}{[1 - \delta(1 - \epsilon_f^2)]^2} + 1 \right) + \frac{\beta\kappa}{\delta[1 - \delta(1 - \epsilon_f^2)]^2}, \quad (27)$$

where we are using  $\beta$  for the constant factor, and a similar minimisation is possible. It is important to stress that, although the filtering step is in common between the methods, its cost depends on  $\Delta$ . Because the different methods have different optimal values of  $\Delta$ , the cost of the filtering step is different between the methods.

A further adjustment to the complexity can be made by noting that  $\epsilon_f$  here is just an upper bound on the 2-norm error  $\epsilon$ . As per Eq. (J3) of [6], the fidelity in the filtered state can be lower bounded as

$$f \geq \frac{1 - \delta}{1 - \delta + \delta\epsilon_f^2}. \quad (28)$$

Then the 2-norm error is  $\epsilon = \sqrt{2 - 2\sqrt{f}}$ , so solving for  $\epsilon_f$  gives

$$\epsilon_f = \epsilon \frac{\sqrt{(1 - \delta)(1 - \epsilon^2/4)}}{(1 - \epsilon^2/2)\sqrt{\delta}}. \quad (29)$$

We are using the notation  $\epsilon_f$  for the parameter used in the filter, and  $\epsilon$  for the final allowable error. We can use this expression for  $\epsilon_f$  in Eq. (26) or Eq. (27) to give a slightly tighter bound for the complexity.

The results for numerically solving for the optimal  $\Delta$  are shown in Fig. 3 for a range of values of  $\alpha$ . For typical values of  $\alpha$ , optimal values of  $\Delta$  are around 0.4. When  $\alpha$  is reduced to 0.17 (as for the PD case for the quantum walk), the optimal  $\Delta$  is around 0.15. For the Randomized method with cost scaling as  $\beta/\delta$  instead, the optimal values of  $\delta$  are around 0.5. In the preceding section, we have used the same  $\Delta$  for the different methods despite the different optimal values to give a fair comparison of the approaches.




Figure 3: The values of  $\Delta$  as a function of  $\epsilon$  that minimise the cost. The results are for  $\alpha = 0.17$  (green) and  $\alpha = 2$  (blue) using the cost in Eq. (26). Results using Eq. (27) are shown for  $\beta = 0.8$  (red) and  $\beta = 5.3$  (orange).

In this result it is assumed that the scaling holds exactly, even for larger values of  $\Delta$ . In reality, this scaling only holds in the asymptotic limit of small  $\Delta$ , and determining the scaling constant with larger values of  $\Delta$  can give larger values. In Table 3 we show the results for each case calculated with a range of values of  $\Delta$ . In the case of the QW with general matrices, the performance is notably better for smaller values of  $\Delta$ , and we estimate  $\alpha = 1.84, 1.66$  and  $1.37$ , for  $\Delta \approx 0.4, 0.3$  and  $0.2$ , respectively. To more accurately estimate the optimal values of  $\Delta$ , in each case we have used the results for three values of  $\Delta$  (those shown in Table 3) to interpolate the scaling constants  $\alpha$  and  $\beta$  as functions of  $\Delta$ .

The resulting optimal values of  $\Delta$  are shown in Fig. 4. In the case of the QW with general matrices, the predicted optimal values of  $\Delta$  are reduced by about 0.1 as compared to those in Fig. 3. There are more minor differences in the other cases. The value  $\Delta = 0.4$  used in Section 3 is in the middle of the range for the optimal values, which is the reason it was chosen as a common value for the numerical comparisons above.

The division of the costs between the adiabatic and filtering steps is shown in Fig. 5 for the interpolated values of  $\alpha$  for the PD case using the quantum walk.

Case	Target $\Delta$					
	0.1	0.2	0.3	0.4	0.5	0.6
QW, positive	108 (0.0971)	56 (0.1990)	32 (0.2821)			
QW, general		344 (0.1986)	288 (0.2878)	232 (0.3967)		
RM, positive			421 (0.3012)	271 (0.3916)	195 (0.4959)	
RM, general				1722 (0.3991)	1100 (0.4987)	771 (0.5989)

Table 3: The number of steps or averaged time for a range of values of  $\Delta$  that are used for the interpolation. We show both the QW and RM, both for positive definite and general matrices, with  $16 \times 16$  matrices and  $\kappa = 50$  for all. The values in brackets are the actual values of  $\Delta$  obtained.




Figure 4: The values of  $\Delta$  as a function of  $\epsilon$  that minimise the cost. We include results assuming various values of  $\alpha$ , as well as with interpolation of the value of  $\alpha$  as a function of  $\Delta$ .

The contribution to the cost from the adiabatic part of the algorithm is relatively independent of  $\epsilon$ , whereas the filtering has a clear  $\ln(1/\epsilon)$  dependence as would be expected. For values of  $\epsilon$  above about 0.0004, the majority of the contribution to the cost is from the adiabatic part of the algorithm. For smaller values of  $\epsilon$  the cost of the filtering step is larger, but these are smaller values of  $\epsilon$  than would be expected in applications. Because QLSP outputs the solution encoded in a quantum state, classical results can only be obtained from sampling. It would require millions of samples to provide an accuracy better than 0.0004, and there would not be an advantage to solving the QLSP to significantly higher accuracy than the sampling.

The Randomized method with  $\beta$  interpolated is also shown in Fig. 5. The adiabatic part of the algorithm is significantly more costly than the filtering part for the entire range of  $\epsilon$  shown. It takes unreasonably small values of  $\epsilon$  below about  $10^{-12}$  for the filtering to be more costly. Both the adiabatic and filtering parts of the algorithm are more costly than for the quantum walk, with the adiabatic part being much more costly. This means that the *total* cost (including both the adiabatic component and filtering component) for the randomised walk is substantially larger than for the discrete adiabatic approach, as shown in Fig. 6.

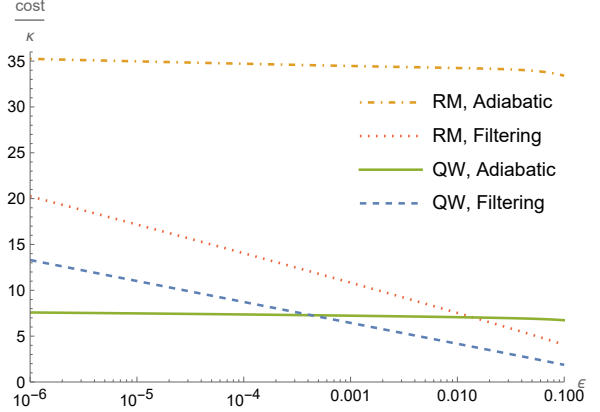


Figure 5: The two contributions to the minimised costs as a function of  $\epsilon$  for the discrete adiabatic approach and for the randomised walk. In both cases the scaling constant is fitted as a function of  $\Delta$ .

## 5 Conclusion

The optimal quantum linear systems algorithm of [6] works by first applying a quantum walk to provide a result with low accuracy, then filtering the resulting state to improve the precision to the desired level. The initial quantum walk has an upper bound on its error with a large constant factor in [6] which would suggest it is a costly algorithm despite having optimal asymptotic complexity.

Here we have numerically tested this initial quantum walk with random matrices and found that the constant factor is about 1,200 times smaller. This means that the algorithm will be highly efficient in practice, with the cost of the quantum walk comparable to the cost of the filtering. The cost of the quantum walk is typically only a little larger than the cost of the filtering, but in special cases, such as positive-definite matrices or  $\epsilon \ll 1$ , its cost is even smaller, as would be expected since the filtering step has the factor of  $\log(1/\epsilon)$  in the complexity.

Reference [10] provided a tighter upper bound on the complexity for an approach with randomised evolution times for initial adiabatic evolution, followed by a filtering step that is similar to [6]. In our numerical testing of that approach, we find that it is about an order of magnitude more costly in practice, despite






Figure 6: The total costs as a function of  $\epsilon$  for the discrete adiabatic approach with  $\alpha$  fitted as a function of  $\Delta$  (blue) and the randomised walk with  $\beta$  fitted as a function of  $\delta$  (orange).

having a tighter analytic bound on the complexity, where we have  $1/\Delta$  in the cost of the adiabatic part vs  $1/\delta$  in the cost of the randomized method. It is somewhat ambiguous to compare the complexity because it is in terms of an evolution time rather than the number of steps of a quantum walk. The evolution time is about 7 times the number of steps of the quantum walk needed for the discrete adiabatic approach, and in practice the cost to simulate that Hamiltonian evolution would be even larger.

### Acknowledgements

DWB worked on this project under a sponsored research agreement with Google Quantum AI. DWB is also supported by Australian Research Council Discovery Projects DP210101367 and DP220101602. DA acknowledges the support by the Department of Defense through the Hartree Postdoctoral Fellowship at QuICS, and the seed grant at the NSF Quantum Leap Challenge Institute for Robust Quantum Simulation (QLCI grant OMA-2120757).

### References

- [1] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. “Quantum algorithm for linear systems of equations”. *Physical Review Letters* **103**, 150502 (2009).
- [2] Dong An and Lin Lin. “Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm”. *ACM Transactions on Quantum Computing* **3**, 5 (2022).
- [3] Andris Ambainis. “Variable time amplitude amplification and a faster quantum algorithm for solving systems of linear equations” (2010). url: [arxiv.org/abs/1010.4458](https://arxiv.org/abs/1010.4458).
- [4] Lin Lin and Yu Tong. “Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems”. *Quantum* **4**, 361 (2020).
- [5] Andrew M. Childs, Robin Kothari, and Rolando D. Somma. “Quantum algorithm for systems of linear equations with exponentially improved dependence on precision”. *SIAM Journal on Computing* **46**, 1920–1950 (2017).
- [6] Pedro C.S. Costa, Dong An, Yuval R. Sanders, Yuan Su, Ryan Babbush, and Dominic W. Berry. “Optimal scaling quantum linear-systems solver via discrete adiabatic theorem”. *PRX Quantum* **3**, 040303 (2022).
- [7] Aram W. Harrow and Robin Kothari. “”. In preparation (2025).
- [8] Lin Lin and Yu Tong. “Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems”. *Quantum* **4**, 361 (2020).
- [9] Yiğit Subaşı, Rolando D Somma, and Davide Orsucci. “Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing”. *Physical Review Letters* **122**, 060504 (2019).
- [10] David Jennings, Matteo Lostaglio, Sam Pallister, Andrew T Sornborger, and Yiğit Subaşı. “Efficient quantum linear solver algorithm with detailed running costs” (2023). url: [arxiv.org/abs/2305.11352](https://arxiv.org/abs/2305.11352).
- [11] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. “Bounds for the adiabatic approximation with applications to quantum computation”. *Journal of Mathematical Physics* **48**, 102111 (2007).
- [12] Yuval R. Sanders, Dominic W. Berry, Pedro C.S. Costa, Louis W. Tessler, Nathan Wiebe, Craig Gidney, Hartmut Neven, and Ryan Babbush. “Compilation of fault-tolerant quantum heuristics for combinatorial optimization”. *PRX Quantum* **1**, 020312 (2020).
- [13] Ryan Babbush, Dominic W. Berry, and Hartmut Neven. “Quantum simulation of the sachdev-ye-kitaev model by asymmetric qubitization”. *Physical Review A* **99**, 040301 (2019).
- [14] Dominic W. Berry, Danial Motlagh, Giacomo Pantaleoni, and Nathan Wiebe. “Doubling the efficiency of hamiltonian simulation via generalized quantum signal processing”. *Phys. Rev. A* **110**, 012612 (2024).
- [15] Pedro C. S. Costa. “Qlsp via discrete adiabatic method - source code”. <https://github.com/PcostaQuantum/QLSP-via-discrete-adiabatic-method/>

[blob/main/Walk\\_error\\_Herm.m](#) (2025). Accessed: 2025-01-24.

- [16] Tim Davis and Yifan Hu. “Suitesparse matrix collection”. <https://sparse.tamu.edu/> (2024). Accessed: 2024-10-23.
- [17] Pedro C. S. Costa. “Qlsp via randomisation method - source code”. <https://github.com/PcostaQuantum/QLSP-via-randomisation-method> (2024). Accessed: 2025-01-24.
- [18] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. “Limit on the speed of quantum computation in determining parity”. *Phys. Rev. Lett.* **81**, 5442–5444 (1998).
- [19] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. “Quantum lower bounds by polynomials”. *J. ACM* **48**, 778–797 (2001).

## A Lower bound

The  $\Omega(\kappa \log(1/\epsilon))$  part of the lower bound from [7] can be obtained via a minor modification of the proof from HHL [1]. Adopting the notation of [1], in Eq. (4) of that work a unitary  $U$  is constructed such that for  $T + 1 \leq t \leq 2T$ , applying  $U^t$  to  $|1\rangle |\psi\rangle$  yields  $|t + 1\rangle \otimes U_T \cdots U_1 |\psi\rangle$  for some sequence of unitary operators  $U_j$ . Now define  $A = I - U\epsilon^{1/T}$  so the condition number  $\kappa = \mathcal{O}(T/\log(1/\epsilon))$  and we can expand

$$A^{-1} = \sum_{k=0}^{\infty} U^k \epsilon^{k/T}. \quad (30)$$

After applying  $A^{-1}$  one would obtain  $t$  in the range  $T + 1$  to  $2T$  with  $\Theta(\epsilon)$  probability. Therefore, solving the linear system with precision  $\epsilon$  provides an unbounded-error algorithm for implementing  $T$  steps  $U_T \cdots U_1$ .

It is straightforward to choose the operations  $U_j$  so that implementing them in succession gives the parity of  $T$  bits, and each corresponds to a single query to these bits. An unbounded-error algorithm for determining the parity of  $T$  bits requires  $\Omega(T)$  queries [18, 19]. That means there must be a lower bound  $\Omega(T)$  for the number of queries to the operators  $U_j$ . Now  $A$  can be block encoded with a single query to  $U_j$ , there is similarly a lower bound  $\Omega(T)$  to the number of calls to the block encoding of  $A$ . Since  $\kappa = \mathcal{O}(T/\log(1/\epsilon))$ , we have  $T = \Omega(\kappa \log(1/\epsilon))$ , so the complexity of the linear systems algorithm is  $\Omega(\kappa \log(1/\epsilon))$ .

## B Block encoding

### B.1 Positive definite $A$

Now we show how to block encoding  $H(s)$  from Eq. (1), when we have  $H_0$  and  $H_1$  given by Eq. (2) and Eq. (3) respectively. The general principle for the block encoding of our Hamiltonian is to block encode both  $H_0$  and  $H_1$ , and apply the following controlled operation

$$\text{SEL} = |0\rangle \langle 0| \otimes U_{H_0} + |1\rangle \langle 1| \otimes U_{H_1}, \quad (31)$$

to select between  $H_0$  and  $H_1$ . To give the linear combination of  $H_0$  and  $H_1$  we apply the one-qubit rotation on the ancilla qubit

$$R(s) = \frac{1}{\sqrt{(1-f(s))^2 + f(s)^2}} \begin{pmatrix} 1-f(s) & f(s) \\ f(s) & -(1-f(s)) \end{pmatrix}. \quad (32)$$

In the following we describe the details of this implementation.

We denote the unitary for the block encoding of  $A$  as  $U_A$ , which acts on an ancilla denoted with subscript  $a$  and the system such that

$${}_a \langle 0| U_A |0\rangle_a = A. \quad (33)$$

We also denote the unitary oracle for preparing  $|b\rangle$  as  $U_b$  such that

$$U_b |0\rangle = |b\rangle. \quad (34)$$

As well as the ancilla system used for the block encoding of  $A$ , we use three ancilla qubits. These ancilla qubits are used as follows.

1. The first selects between the blocks in  $H_0$  and  $H_1$ .
2. The next is used for preparing the combination of  $H_0$  and  $H_1$ .
3. The third is used in implementing  $Q_b$ .

We use the subscripts  $a_h$ ,  $a_1$  and  $a_2$  for these three qubits. Note that although we have called  $a_h$  an ancilla qubit, it part of the system that the Hamiltonian acts upon, not an ancilla used in the block encoding. This means that we do not apply the reflection operator on that ancilla to construct the walk operator.

To make the block encoding  $H(s)$  simpler let us rewrite Eq. (3) as

$$H_1 = \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix} \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix}, \quad (35)$$

and Eq. (2) as

$$H_0 = \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix} \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix}. \quad (36)$$

This new way to express the Hamiltonians will then facilitate the construction of the select operator since the total Hamiltonian can be expressed as

$$H(s) = \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix} \left[ (1-f(s)) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} + f(s) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix} \right] \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix}. \quad (37)$$

For the first block matrix inside the bracket in the equation above, we have

$$\begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} = \sigma_{a_h}^x \otimes I, \quad (38)$$

where  $\sigma_{a_h}^x$  indicates that we have the Pauli  $X$  operator acting on the register  $a_h$ . Now, for the other block matrix which comes from  $H_1$ , we have

$$\begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix} = I_{a_h} \otimes A, \quad (39)$$

where we have included  $I_{a_h}$  to indicate that the operation acts as the identity on the first qubit. Moreover, we have  $U_A$  as the unitary for the block encoding of  $A$ , with ancilla qubit  $a$ . Thus when we combine Eq. (38) and Eq. (39), the two block matrices for  $H_1$ , we have

$$\begin{aligned} & (\sigma_{a_h}^x \otimes I_N \otimes I_a) (I_{a_h} \otimes U_A) \\ &= |0\rangle \langle 1|_{a_h} \otimes U_A + |1\rangle \langle 0|_{a_h} \otimes U_A. \end{aligned} \quad (40)$$

However, the encoding above does not guarantee the hermiticity of  $H_1$ , which can be fixed by the following expression

$$|0\rangle \langle 1|_{a_h} \otimes U_A + |1\rangle \langle 0|_{a_h} \otimes U_A^\dagger. \quad (41)$$

Therefore, we have Eq. (41) as the block encoding of the two middle operators of  $H_1$  and  $\sigma_{a_h}^x \otimes I_N \otimes I_a$  as the block encoding of the two middle operators of  $H_0$ . We can therefore select between  $H_0$  and  $H_1$  using the operator

$$\begin{aligned} U_{A(f)} &= |0\rangle \langle 0|_{a_1} \otimes \sigma_{a_h}^x \otimes I_N \otimes I_a \\ &+ |1\rangle \langle 1|_{a_1} \otimes \left( |0\rangle \langle 1|_{a_h} \otimes U_A + |1\rangle \langle 0|_{a_h} \otimes U_A^\dagger \right), \end{aligned} \quad (42)$$

together with the operators involving  $Q_b$  in Eq. (37) and rotation on the ancilla qubit  $a_1$ . The operator involving  $Q_b$  in Eq. (37) can be rewritten as

$$\begin{aligned} \begin{pmatrix} I & 0 \\ 0 & Q_b \end{pmatrix} &= |0\rangle \langle 0|_{a_h} \otimes I_N + |1\rangle \langle 1|_{a_h} \otimes Q_b \\ &= |0\rangle \langle 0|_{a_h} \otimes I_N + |1\rangle \langle 1|_{a_h} \otimes (I_N - |b\rangle \langle b|) \\ &= |0\rangle \langle 0|_{a_h} \otimes I_N \\ &+ \frac{1}{2} |1\rangle \langle 1|_{a_h} \otimes (I_N + (I_N - 2|b\rangle \langle b|)) \end{aligned}$$

$$\begin{aligned} &= |0\rangle \langle 0|_{a_h} \otimes I_N \\ &+ \frac{1}{2} |1\rangle \langle 1|_{a_h} \otimes \left[ I_N + (I_N - 2U_b |0\rangle \langle 0| U_b^\dagger) \right]. \end{aligned} \quad (43)$$

In the second equality, we have written the projection in terms of a sum of the identity and a projector, and in the last equality, we used the oracle for preparing  $|b\rangle$  as given in Eq. (34). We can block encode this projector using the ancilla  $a_2$ . We simply need to create this ancilla in an equal superposition and use the following unitary operation

$$\begin{aligned} CU_{Q_b} &= \left( I_{a_h} \otimes I_{a_2} \otimes U_b^\dagger \right) \left\{ |0\rangle \langle 0|_{a_h} \otimes I_{a_2} \otimes I_N \right. \\ &+ |0\rangle \langle 0|_{a_h} \otimes [|0\rangle \langle 0|_{a_2} \otimes I_N + |1\rangle \langle 1|_{a_2} \\ &\left. \otimes (I_N - 2|0\rangle \langle 0|)] \right\} (I_{a_h} \otimes I_{a_2} \otimes U_b). \end{aligned} \quad (44)$$

To block encode  $H(s)$  as in Eq. (37), the block encoding above is applied before and after the select operator  $U_{A(f)}$ .

We also need to perform the rotation  $R(s)$  in order to interpolate between the Hamiltonians. The interpolation operator can then act on the ancilla qubit used in the select operator. However, instead of applying the inverse of  $R(s)$ , we simply perform a Hadamard and to produce the right weights for the Hamiltonians we need to apply the rotation operator controlled by ancilla qubit  $a_h$  as

$$CR(s) = |0\rangle \langle 0|_{a_h} \otimes R(s)_{a_1} + |1\rangle \langle 1|_{a_h} \otimes \mathcal{H}_{a_1} \quad (45)$$

where  $\mathcal{H}$  denotes the Hadamard operation. This means that our block encoding gives an extra factor of  $1/\sqrt{2[(1-f(s))^2 + f(s)^2]}$ . This prefactor is between  $1/\sqrt{2}$  and 1, and will reduce the gap.

Then the complete circuit diagram can be given as in Fig. 7. The walk operator  $W_T(s)$  can be completed by a reflection about zero on all ancilla registers except  $a_h$  which is part of the target system for the Hamiltonian.

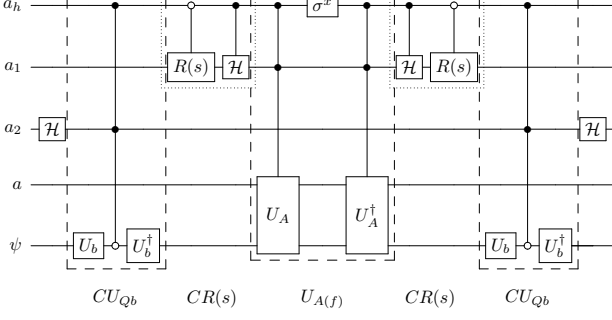


Figure 7: The block encoding of the Hamiltonian  $H(s)$ , where the target system is labeled as  $|\psi\rangle$ , and  $a_h$  as the extra qubit used to construct the Hamiltonian for the QLSP. The ancillas resulting from the block-encoded construction of  $H(s)$  are labeled by  $a_1$ ,  $a_2$ , and  $a$  (where  $a$  is the ancilla needed in the block encoding of  $A$ ). The dashed boxes show  $CU_{Qb}$  and  $U_{A(f)}$ , whereas the dotted boxes show  $CR(s)$ .

## B.2 Non-Hermitian matrix $A$

For the general case, we briefly review what is given in [6] for the ancilla qubits needed, since we enlarge the space to address the non-Hermitian property of  $A$ .

We have the ancilla qubit  $a$  for the block encoding of  $A$ , as given for the PD case, and the remainder of the ancilla qubits are used as follows.

1. The first selects between the blocks in  $A(f)$ .
2. The next is used for preparing the combination of  $\sigma_z \otimes I$  and  $A$ .
3. The third is used in implementing  $Q_b$ .
4. The fourth selects between the blocks in  $H(s)$ .

The first and fourth ancillas above are part of the target system for the Hamiltonian  $H(s)$ , and are denoted  $a_{h1}$  and  $a_{h2}$ , respectively. The second and third qubits in the list above are denoted  $a_1$  and  $a_2$ , respectively. The circuit diagram for the block encoding for the general case is presented in Fig. 8. See Ref. [6] for its complete derivation.

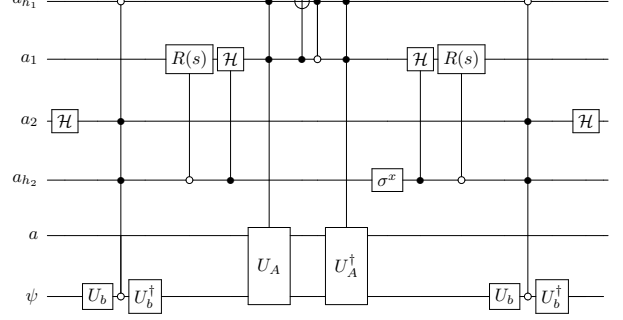


Figure 8: The block encoding of the Hamiltonian  $H(s)$  for the non-Hermitian case, where again the target system is labeled as  $|\psi\rangle$ , and the ancilla are labeled by  $a_{h1}$ ,  $a_1$ ,  $a_2$ ,  $a_{h2}$  and  $a$  (where  $a_{h1}$  and  $a_{h2}$  are part of the system acted upon by the Hamiltonian).

## C Numerical Tests - Supporting Material

### C.1 QW method

In Fig. 9 (positive definite) and Fig. 10 (non-Hermitian) we show the results from Fig. 2 for dimension  $16 \times 16$  matrices, as well as results for  $4 \times 4$  and  $8 \times 8$  matrices. In these figures, we show the results for the three dimensions in subfigures as well and find that the dimension does not significantly alter the complexity of the QLSP, as we expect from theory.

The table for the PD for a fixed total number of steps for all the 100 instances tested for each condition number is included in Table 4. The worst relative performance of the QW method compared to the RM is for matrices with condition number equal to 10, but the QW performance is still 13 times better.

Moreover, we observed a large amount of outlier instances in the quantum walk for the non-Hermitian matrices where the solution is obtained after a small number of steps. This occurs due to the relatively large allowed error  $\Delta = 0.4$  because when taking sufficiently large number of samples, one can be sufficiently close to the solution purely by chance. Such outliers are expected when sampling randomly, especially at low step counts, and do not contradict the broader trend that the discrete adiabatic solver consistently requires fewer resources across all instances.

### C.2 Randomized Method

In Fig. 11 (positive definite) and Fig. 12 (non-Hermitian) we show results from Fig. 2 for dimension  $16 \times 16$ , and also dimension  $4 \times 4$  and  $8 \times 8$  matrices. There is even less dependence of the complexity on the dimension than for the QW method. The dimension dependence in the complexity comes from the matrix encoding, which does not influence the numerical

Condition Number	Number of steps	$\Delta$
10	4	0.378
20	12	0.310
30	16	0.339
40	20	0.353
50	24	0.362

not overlap.

Table 4: Number of steps for average error  $\Delta \sim 0.4$  with a range of condition numbers for PD matrices of dimension  $16 \times 16$ . The RMS of the error (as measured by the norm of the difference in states), is computed over 100 instances for each condition number considered.

analysis.

We also include the table (Table 5) that gives the RMS for the target error for the PD matrices, when for each condition number we fix the same averaged time for all instances of matrices tested.

Condition Number	Averaged time	$\Delta$
10	37.2	0.379
20	88.1	0.395
30	144.5	0.396
40	202.7	0.398
50	270.9	0.392

Table 5: Average evolution time for average error  $\Delta \sim 0.4$  via RM with a range of condition numbers for PD matrices of dimension  $16 \times 16$ . The RMS error is computed over 100 instances for each condition number considered and 200 repetitions per instance.

### C.3 QW vs RM - scatter plots

In this section, we compare the performance between the RM and the QW method of each instance tested for  $A \in \mathbb{R}^{8 \times 8}$  for PD and Hermitian matrices in Fig. 13 and non-Hermitian matrices in Fig. 14. The vertical axis shows the number of walk steps of the instances to obtain the desired solution error  $\Delta = 0.4$ . In the horizontal axis, we have the average minimal time of the Randomized method to reduce the RMS error below  $\Delta = 0.4$ .

In the numerical results, there were no tested cases where the Randomized method surpassed the QW performance. This can be seen from the fact that the complexity for all instances is below the dashed line in Fig. 13 and Fig. 14. We even find that the QW still wins if we compare different problem instances with the same condition number. That is, the maximum cost of the QW is below the minimum cost of the RM (provided we select matrices of the same type, either both PD or both non-Hermitian).

The same can be concluded from the cases of other dimensions by looking at the error bars, since they do



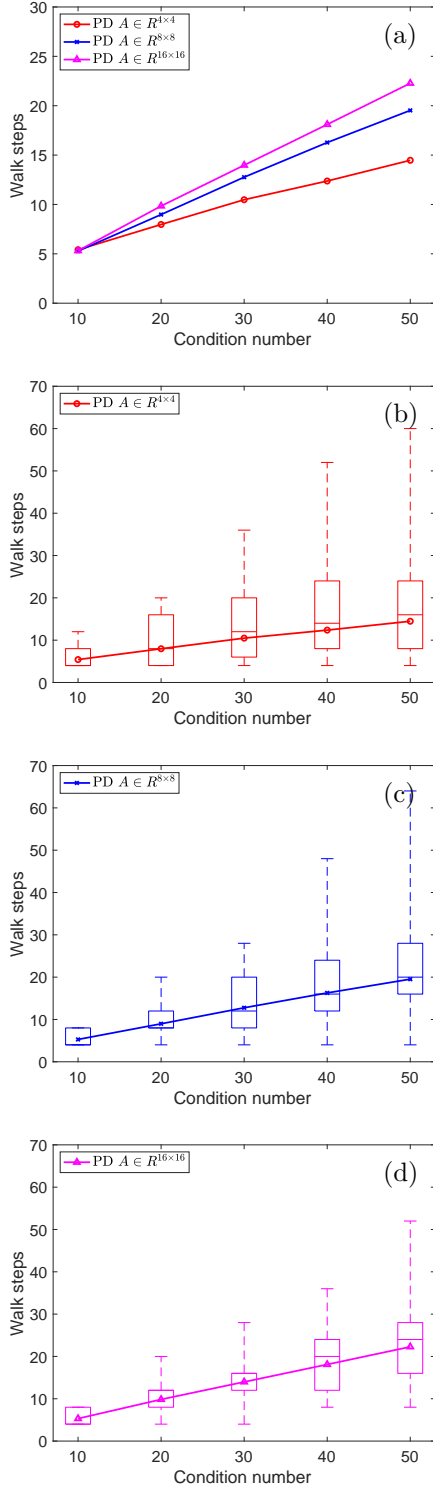


Figure 9: Numerical tests of the QW method with positive definite matrices of dimension  $4 \times 4$ ,  $8 \times 8$ , and  $16 \times 16$ . The vertical axis shows the required total walk steps with  $\Delta = 0.4$  as the threshold error. Part (a) shows the averaged results of 100 randomly generated instances, and parts (b) to (d) show the box plots.

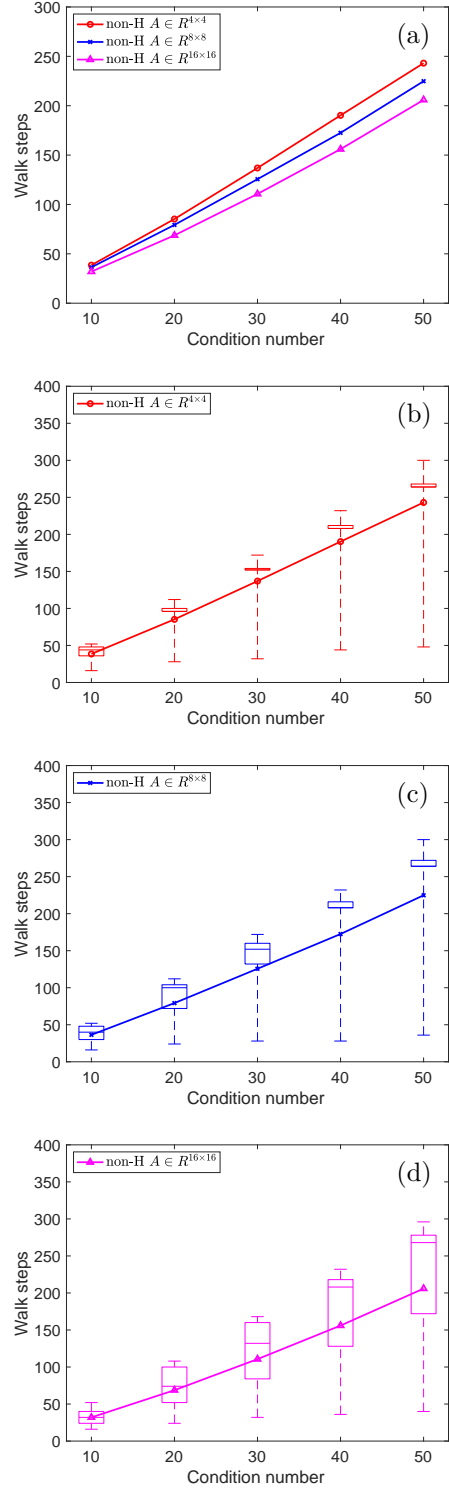


Figure 10: Numerical tests of the QW method with non-Hermitian matrices of dimension  $4 \times 4$ ,  $8 \times 8$ , and  $16 \times 16$ . The vertical axis shows the required total walk steps with  $\Delta = 0.4$  as the threshold error. Part (a) shows the averaged results of 100 randomly generated instances, and parts (b) to (d) show the box plots.

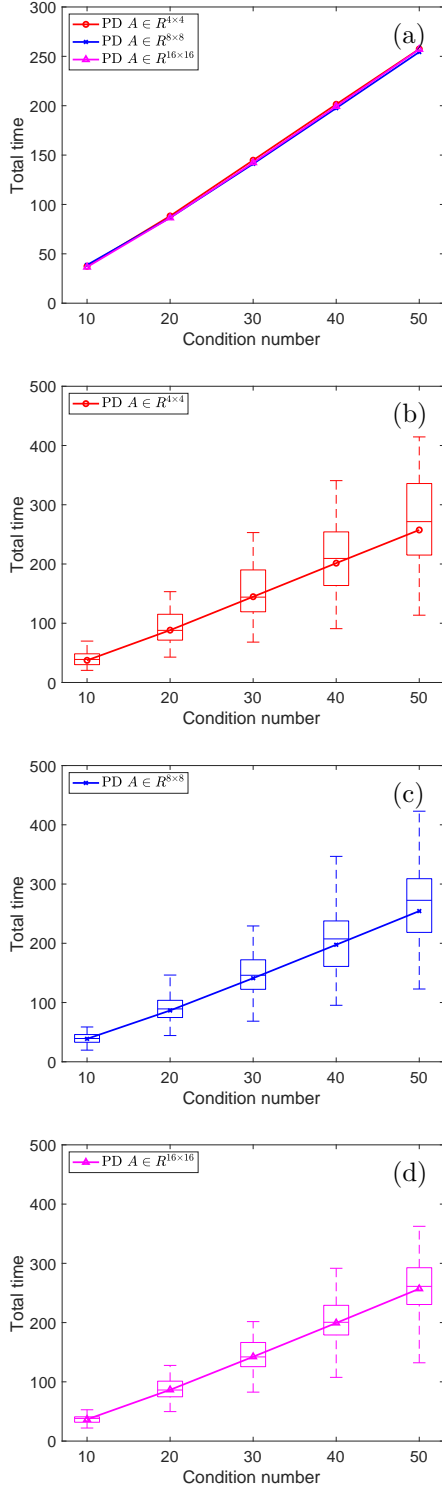


Figure 11: Numerical tests of the Randomized method with positive definite matrices of dimension  $4 \times 4$ ,  $8 \times 8$ , and  $16 \times 16$ . The vertical axis shows the required total time (defined as  $\sum_{j=1}^q |t_j|$ ) to reduce the RMS below 0.4. Part (a) shows the averaged results of 100 randomly generated instances, and parts (b) to (d) show the box plots.

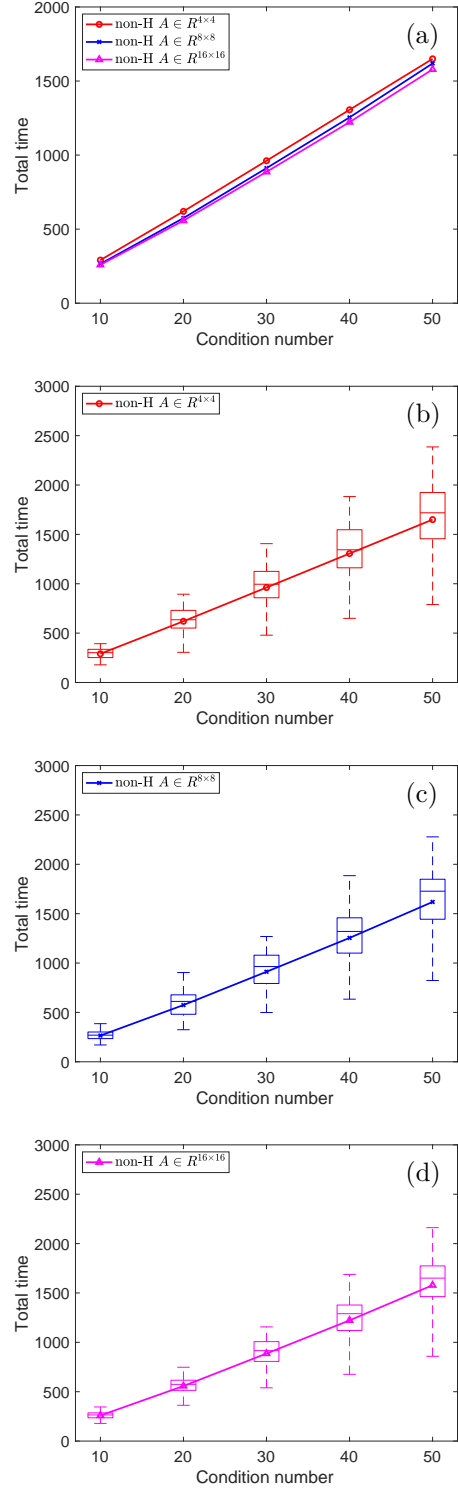


Figure 12: Numerical tests of the Randomized method with non-Hermitian matrices of dimension  $4 \times 4$ ,  $8 \times 8$ , and  $16 \times 16$ . The vertical axis shows the required total time (defined as  $\sum_{j=1}^q |t_j|$ ) to reduce the RMS below 0.4. Part (a) shows the averaged results of 100 randomly generated instances, and parts (b) to (d) show the box plots.




Figure 13: Scatter plots for the positive definite instances used in testing QW and RM. In the vertical axis, we have the complexity (number of walk steps) for the quantum walk method, and in the horizontal axis the complexity (evolution time) for the Randomized method.




Figure 14: Scatter plots for the non-Hermitian instances used in testing QW and RM. The vertical axis is the complexity (number of walk steps) for the quantum walk method, and the horizontal axis is the complexity (evolution time) for the Randomized method.