

Faster quantum linear system solver beyond the condition number

Alexander M. Dalzell¹, Jianqiang Li², and Yuan Su¹

¹AWS Center for Quantum Computing, Pasadena, CA 91106, USA

²Department of Computer Science, Rice University, Houston, TX 77005, USA

Abstract

The spectral condition number is a widely adopted measure of worst-case cost for quantum linear system solvers. Yet it can significantly overestimate the actual runtime for a typical problem instance.

We present two quantum algorithms that produce the normalized solution of linear system $Ax = |b\rangle$ to accuracy ϵ with complexity independent of the condition number $\kappa = \|A\| \|A^{-1}\|$. We focus on the standard input model where the coefficient matrix A is accessed through a block encoding oracle and the initial vector $|b\rangle$ is prepared by a unitary. But we also introduce an affine dilation model that encodes A and $|b\rangle$ jointly, allowing further refinements of the query complexity.

Our truncation-based solver makes an optimal number of queries to the preparation of $|b\rangle$ and $\mathcal{O}(\kappa_{\text{eff}} \text{polylog}(\frac{\kappa_{\text{eff}}}{\epsilon}))$ queries to the block encoding of A . We prove a family of upper bounds on the effective condition number κ_{eff} , including $\kappa_{\text{eff}} \leq \frac{\|(A^\dagger A)^{-t/2} x\|^{1/t}}{\|x\|^{1/t} \epsilon^{1/t}}$ for positive even integer t and $\kappa_{\text{eff}} \leq \frac{\|A^{-1\dagger} (A^\dagger A)^{-(t-1)/2} x\|^{1/t}}{\|x\|^{1/t} \epsilon^{1/t}}$ for positive odd t , going beyond the naive estimate in terms of κ .

Our filtering-based solver is extremely simple with a favorable runtime prefactor. In particular, the solver has query complexity $6 \frac{\|A^{-1\dagger} x\|}{\|x\| \epsilon} \ln(\frac{1}{\epsilon})$ to leading order when the solution norm $\|x\|$ is known. We then present a similarly simple solution norm estimator with the same asymptotic cost up to logarithmic factors.

Our quantum linear system solvers thus substantially improve a recent algorithm of Li, which has runtime $\mathcal{O}\left(\frac{\|A^{-1\dagger} x\|}{\|x\| \epsilon^2}\right)$ if $\|x\|$ is known and $\mathcal{O}\left(\frac{\|A^{-1\dagger} x\| \|x\|^2}{\epsilon^2}\right)$ otherwise. This enables faster quantum linear system solving beyond the condition number, bringing prohibitively ill-conditioned systems closer to feasibility.

Contents

1	Introduction	3
1.1	Quantum linear system solvers	3
1.2	Main result	4
1.3	Block encoding and affine dilation	6
1.4	Solver based on effective truncation of linear system	7
1.5	Solver based on filtering with effective gap	9
2	Quantum linear system problem and input models	11
2.1	Standard input model	11
2.2	Affine dilation model	12
2.3	Quantum linear system problem	13
3	Beyond-κ solver based on effective truncation of linear system	15
3.1	Effective condition number	15
3.2	Bounds on the effective condition number	18
3.3	Analysis of beyond- κ solver based on effective matrix truncation	21
3.4	Oracle for effective state truncation	23
3.5	Analysis of beyond- κ solver based on effective state truncation	25
4	Beyond-κ solver based on filtering with effective gap	27
4.1	Block encoding and constrained orthogonal decomposition	28
4.2	Effective gap lemma	30
4.3	Eigenstate filtering over unit circle	31
4.4	Constant-prefactor analysis of beyond- κ solver based on filtering	34
4.5	Beyond- κ solution norm estimation	37
5	Discussion	39
A	Analysis of Li's beyond-κ solver in the standard input model	42
B	Analysis of beyond-κ solvers in the affine dilation model	43
B.1	Analysis of truncation-based solvers in the affine dilation model	43
B.2	Analysis of filtering-based solver in the affine dilation model	44
B.3	Complexity comparison	45
C	Generalized truncation property	46
D	Constant-prefactor optimization for filtering with effective gap	47

1 Introduction

1.1 Quantum linear system solvers

Solving systems of linear equations constitutes a foundational task across scientific and engineering disciplines. Classical solvers quickly become computationally prohibitive as the number of constraints and unknown variables increases. This intractability led Harrow, Hassidim, and Lloyd to develop the first quantum linear system algorithm in the seminal work [18], where the BQP-completeness of matrix inversion was established. Remarkably, the query complexity of quantum linear system solvers does not depend explicitly on the problem dimensionality, enabling a potential exponential speedup over classical solvers. It remains challenging to identify end-to-end applications that fully realize this exponential quantum speedup; however, recent studies point to physics-inspired problems such as estimating many-body Green’s functions [40] as promising avenues.

Beyond direct applications, quantum linear system solvers also serve as a key primitive for constructing more advanced quantum algorithms. This was demonstrated in previous work on solving differential equations [5, 6, 20], where a system of linear equations was introduced whose inverse encodes the matrix exponential via a truncated Taylor series. More generally, quantum linear system solvers have found applications in transforming eigenvalues of non-normal operators, achieved through rational generating functions for Faber polynomials [28], contour integration of resolvents [1, 38, 39], and rational approximation [42].

A critical parameter determining the runtime of quantum linear system solvers is the *spectral condition number*. Specifically, consider the linear system $Ax = |b\rangle$, where A is the *coefficient matrix* with spectral norm $\|A\| \leq 1$ accessed through a *block encoding* oracle O_A , and $|b\rangle$ is the *initial vector/state* with unit Euclidean norm $\| |b\rangle \| = 1$ prepared by a unitary O_b from a standard reference state. If A is invertible, then the solution to the linear system is given by

$$x = A^{-1}|b\rangle, \tag{1}$$

and we let the condition number be the largest norm of the solution vector maximized over all possible initial states

$$\kappa = \max_{\| |b\rangle \| = 1} \| A^{-1}|b\rangle \| = \| A^{-1} \|. \tag{2}$$

In practice, the condition number is rarely known exactly; we instead work with an upper bound $\kappa \geq \| A^{-1} \|$. To produce the normalized state $|x\rangle = \frac{x}{\|x\|}$ with accuracy ϵ , the Harrow–Hassidim–Lloyd solver has query complexity $\mathbf{O}(\text{poly}(\kappa, \frac{1}{\epsilon}))$. This was improved through a series of works, ultimately yielding solvers with query complexity $\mathbf{O}(\kappa \log(\frac{1}{\epsilon}))$ [13–15, 29]. This tradeoff is optimal, which is established by an unpublished lower bound of Harrow and Kothari and confirmed recently in [12, Appendix A] and [32].

All known solvers attaining the optimal scaling treat the coefficient matrix and the initial vector on equal footing, and thus make the same asymptotic number of queries to O_A and O_b :

$$\mathbf{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right) \mathbf{Cost}(O_A) + \kappa \log\left(\frac{1}{\epsilon}\right) \mathbf{Cost}(O_b)\right). \tag{3}$$

Alternatively, one may solve the quantum linear system problem using *Variable Time Amplitude Amplification* (VTAA) [2] incurring a cost of [29]

$$\mathbf{O}\left(\kappa \log\left(\frac{\kappa}{\|x\|}\right) \left(\log \log\left(\frac{\kappa}{\|x\|}\right) + \log\left(\frac{1}{\epsilon}\right)\right) \mathbf{Cost}(O_A) + \frac{\kappa}{\|x\|} \mathbf{Cost}(O_b)\right). \tag{4}$$

Here, the solution norm satisfies $1 \leq \|x\| \leq \kappa$, so the scaling for the state preparation cost $\frac{\kappa}{\|x\|}$ can be much better than the worst-case scaling $\kappa \log\left(\frac{1}{\epsilon}\right)$ in terms of the condition number. Moreover, this cost of O_b can be shown to be optimal through a reduction from the quantum search problem. That said, the number of queries to O_A is still constrained by the condition number κ —the block encoding cost is actually logarithmic factors away from the optimal tradeoff and, even with the deterministic amplification schedule [29], VTAA-based solvers are still considerably more complex to implement than those from [13–15, 29].

It is therefore natural to ask whether the condition number barrier can be lifted, especially with respect to the block encoding cost. This question was raised by Harrow, Hassidim, and Lloyd in [18, Section IV] and answered affirmatively in the recent work of Li [24]. Li’s original solver handles only sparse linear systems, though it is fairly straightforward to generalize to the block encoding model—we sketch this generalization in [Appendix A](#). The generalized solver then has the query cost

$$\mathbf{O}\left(\frac{\|A^{-1\dagger}x\|}{\|x\|\epsilon^2}\mathbf{Cost}(O_A) + \frac{\|A^{-1\dagger}x\|}{\|x\|\epsilon^2}\mathbf{Cost}(O_b)\right). \quad (5)$$

Achieving the above scaling requires the nontrivial assumption that the solution norm $\|x\|$ is known up to a constant multiplicative factor. If $\|x\|$ is unknown, then the complexity worsens to

$$\mathbf{O}\left(\frac{\|A^{-1\dagger}x\|\|x\|^2}{\epsilon^2}\mathbf{Cost}(O_A) + \frac{\|A^{-1\dagger}x\|\|x\|^2}{\epsilon^2}\mathbf{Cost}(O_b)\right). \quad (6)$$

Regardless, the cost now depends on vector norms as opposed to the condition number κ , satisfying the inequality $\|x\|^2 \leq \|A^{-1\dagger}x\| \leq \kappa\|x\|$. We shall henceforth call these the *beyond- κ solvers*.

Li demonstrated that for applications such as solving multivariate polynomial systems, these vector norm inequalities are far from saturated, and beyond- κ solvers remain efficient even in the presence of exponentially large condition numbers. This opens the door to solving ill-conditioned linear systems on quantum computers, a task out of reach for all prior approaches. However, at the opposite extreme where $\|A^{-1\dagger}x\| \approx \kappa\|x\| \approx \kappa^2$, the query complexity of Li’s solver is $\mathbf{O}\left(\frac{\kappa}{\epsilon^2}\right)$ with a known $\|x\|$ and degrades to $\mathbf{O}\left(\frac{\kappa^4}{\epsilon^2}\right)$ otherwise, falling far short of optimality. Even in the intermediate regime, these unfavorable exponents could offset any practical benefit of the beyond- κ solvers, undermining their intended advantage over conventional methods.

1.2 Main result

We present two quantum linear system solvers whose complexity goes beyond the condition number barrier: one based on an effective truncation of the linear system, and one based on eigenstate filtering over the unit circle with an effective gap. We summarize our result and compare it against previous results in [Table 1](#).

Our truncation-based solver has query complexity

$$\mathbf{O}\left(\kappa_{\text{eff}} \log\left(\frac{\kappa_{\text{eff}}}{\|x\|}\right)\left(\log\log\left(\frac{\kappa_{\text{eff}}}{\|x\|}\right) + \log\left(\frac{1}{\epsilon}\right)\right)\mathbf{Cost}(O_A) + \frac{\kappa_{\text{eff}}}{\|x\|}\mathbf{Cost}(O_b)\right). \quad (7)$$

Here, the *effective condition number* $1 \leq \kappa_{\text{eff}} \leq \kappa$ is chosen such that the truncated solvers still produce the normalized solution state with the target accuracy ϵ . We show that the following family of bounds hold for the effective condition number

$$\kappa_{\text{eff}} \leq \left(\frac{\|(A)^{-t}_{\text{sv}}x\|}{\|x\|\epsilon}\right)^{\frac{1}{t}}, \quad 0 < t < \infty, \quad (8)$$

Year	Algorithm	Primary innovation	Queries to O_A	Queries to O_b
2008	[18]	Phase estimation + Hamiltonian simulation	$\mathcal{O}\left(\frac{\kappa^2}{\ x\ \epsilon^2}\right)$	$\mathcal{O}\left(\frac{\kappa}{\ x\ }\right)$
2012	[2]	VTAA	$\mathcal{O}\left(\frac{\kappa}{\epsilon^3} \log^3\left(\frac{\kappa}{\epsilon}\right) \log^2\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$
2017	[11]	Linear combination of quantum walks	$\mathcal{O}\left(\frac{\kappa^2}{\ x\ } \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\kappa}{\ x\ } \log\left(\frac{\kappa}{\epsilon}\right)\right)$
2017	[11]	Gapped phase estimation + Quantum walk + VTAA	$\mathcal{O}\left(\kappa \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$
2018	[37]	Randomized adiabatic evolution	$\mathcal{O}\left(\frac{\kappa}{\epsilon} \log(\kappa)\right)$	$\mathcal{O}\left(\frac{\kappa}{\epsilon} \log(\kappa)\right)$
2018	[8]	Detailed analysis of VTAA	$\mathcal{O}\left(\kappa \log(\kappa) \log^2\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\kappa}{\ x\ } \log(\kappa)\right)$
2019	[4]	Continuous time adiabatic evolution	$\mathcal{O}\left(\kappa \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \text{polylog}\left(\frac{\kappa}{\epsilon}\right)\right)$
2019	[25]	Eigenstate filtering	$\mathcal{O}\left(\kappa \log\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \log\left(\frac{\kappa}{\epsilon}\right)\right)$
2022	[13]	Discrete time adiabatic evolution	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$
2023	[9]	QSVT-based GPE + VTAA	$\mathcal{O}\left(\kappa \log(\kappa) \log\left(\frac{\kappa}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\kappa}{\ x\ } \log(\kappa)\right)$
2024	[15]	Kernel reflection	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$
2024	[14]	Poisson-distributed phase randomization	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$
2024	[29]	Tunable VTAA + Discretized inverse state	$\mathcal{O}\left(\kappa \log\left(\frac{\kappa}{\ x\ }\right) \log\left(\frac{\log(\kappa/\ x\)}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\kappa}{\ x\ }\right)$
2024	[29]	Block preconditioning	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$
2025	[24]	Filtering with effective gap	$\mathcal{O}\left(\frac{\ A^{-1\dagger}x\ }{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{\ A^{-1\dagger}x\ }{\epsilon^2}\right)$
2026	This work	Effective truncation of linear system	$\mathcal{O}\left(\kappa_{\text{eff}} \log\left(\frac{\kappa_{\text{eff}}}{\ x\ }\right) \log\left(\frac{\log(\kappa_{\text{eff}}/\ x\)}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\kappa_{\text{eff}}}{\ x\ }\right)$
2026	This work	Improved filtering with effective gap	$\mathcal{O}\left(\frac{\ A^{-1\dagger}x\ }{\ x\ \epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$	$\mathcal{O}\left(\frac{\ A^{-1\dagger}x\ }{\ x\ \epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$

Table 1: Complexity comparison of the new and previous solvers for the quantum linear system problem. Our effective condition number κ_{eff} admits a family of upper bounds indexed by a positive integer t , including $\kappa_{\text{eff}} \leq \frac{\|(A^\dagger A)^{-t/2}x\|^{1/t}}{\|x\|^{1/t}\epsilon^{1/t}}$ for even t and $\kappa_{\text{eff}} \leq \frac{\|A^{-1\dagger}(A^\dagger A)^{-(t-1)/2}x\|^{1/t}}{\|x\|^{1/t}\epsilon^{1/t}}$ for odd t . Our filtering-based solver achieves a constant prefactor of 6 in its leading-order query complexity. The comparison assumes $\|x\|$ is known. Without this assumption, the result of [24] degrades to $\mathcal{O}\left(\frac{\|A^{-1\dagger}x\|\|x\|^2}{\epsilon^2}\right)$, whereas ours remains the same up to logarithmic factors.

where $(A)_{\text{sv}}^{-t}$ denotes the singular value transformation [16] with scalar function $\sigma \mapsto \sigma^{-t}$. For a positive integer t , the singular value transformation is given by the standard matrix inversion as

$$\kappa_{\text{eff}} \leq \begin{cases} \left(\frac{\|A^{-1\dagger}(A^\dagger A)^{-(t-1)/2}x\|}{\|x\|\epsilon}\right)^{1/t}, & t \text{ is odd,} \\ \left(\frac{\|(A^\dagger A)^{-t/2}x\|}{\|x\|\epsilon}\right)^{1/t}, & t \text{ is even.} \end{cases} \quad (9)$$

This reduces further to the succinct bound $\kappa_{\text{eff}} \leq \frac{\|A^{-1\dagger}x\|}{\|x\|\epsilon}$ when $t = 1$, but it also gives alternative bounds with logarithmic scaling in the inverse accuracy when $t = \Theta\left(\frac{\log(\frac{1}{\epsilon})}{\log \log(\frac{1}{\epsilon})}\right)$. As t increases, the vector norm $(\|(A)_{\text{sv}}^{-t}x\|/\|x\|)^{1/t} \nearrow \kappa$ increases monotonically to the condition number while the inverse error $(1/\epsilon)^{1/t} \searrow 1$ decreases to unity, recovering the scaling of conventional VTAA.

Our filtering-based solver is remarkably simple and has a favorable runtime prefactor. In particular, we show that the solver has query complexity

$$6 \frac{\|A^{-1\dagger}x\|}{\|x\|\epsilon} \ln\left(\frac{1}{\epsilon}\right) \mathbf{Cost}(O_A) + 6 \frac{\|A^{-1\dagger}x\|}{\|x\|\epsilon} \ln\left(\frac{1}{\epsilon}\right) \mathbf{Cost}(O_b) \quad (10)$$

to leading order. This is achieved assuming that relative approximations of the solution norm $\|x\|$ are available with an approximation ratio approaching 1. If $\|x\|$ is unknown, then we present a comparably simple solution norm estimator achieving a constant relative precision with query cost

$$\mathbf{O} \left(\frac{\|A^{-1\dagger}x\|}{\|x\|} \log^2 \left(\frac{\|A^{-1\dagger}x\|}{\|x\|} \right) \mathbf{Cost}(O_A) + \frac{\|A^{-1\dagger}x\|}{\|x\|} \log^2 \left(\frac{\|A^{-1\dagger}x\|}{\|x\|} \right) \mathbf{Cost}(O_b) \right). \quad (11)$$

The approximation ratio can then be refined using standard methods such as amplitude estimation.

Our beyond- κ solvers thus substantially improve upon the prior art, while our vector norm scalings degrade to almost linear in κ in the worst case, reproducing the performance of conventional solvers. This constitutes a concrete step toward efficiently solving ill-conditioned linear systems on a quantum computer.

1.3 Block encoding and affine dilation

Quantum linear system solvers can be formulated in an abstract setting where the coefficient matrix A has norm $\|A\| \leq 1$ and is accessed through a block encoding as

$$A = G_1^\dagger O_A G_0, \quad (12)$$

with unitary O_A and isometries G_0, G_1 , while the initial state $|b\rangle$ is prepared by a unitary O_b as

$$|b\rangle = O_b|0\rangle. \quad (13)$$

Assuming the implementation cost of G_0 and G_1 is negligible without loss of generality, the complexity of these solvers is then measured in terms of the number of queries to O_A and O_b respectively. This applies to a range of concrete settings, including sparse linear systems, and extends to cases where A and $|b\rangle$ are constructed indirectly via common linear algebraic operations. We term this the *standard input model* for the remainder of the paper.

The beyond- κ solver of [24] operates on an augmented matrix of the form $\begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}$. Using oracles O_A and O_b , one can construct a block encoding of this augmented matrix with normalization factor $\sqrt{2}$ and thereby reformulate the solver in the standard input model; see [Appendix A](#) for details. However, as shown in [24], the query complexity can be further refined through a direct block encoding of $\begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}$, bypassing the standard input model.

We introduce the *affine dilation model* to facilitate comparison of Li's solver [24] with other quantum linear system solvers while also exploiting this refined block encoding. Specifically, we consider $\begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}$, where the augmented matrix is recovered by setting the scalar $c = 0$. However, if $c \neq 0$, then the affine dilated matrix is invertible and

$$\begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} & \frac{1}{c}A^{-1}|b\rangle \\ 0 & \frac{1}{c} \end{bmatrix}. \quad (14)$$

Consequently, the solution $x = A^{-1}|b\rangle$ can be extracted from the top-right entry of the inverse matrix, while A and $|b\rangle$ are handled jointly in a single block encoding analogous to [24]. We formally introduce the quantum linear system problem and its input models in [Section 2](#). While we primarily focus on the standard input model for simplicity, we also analyze the query complexity of beyond- κ solvers in the affine dilation model in [Appendix B](#).




Figure 1: Illustration of the definition of effective condition number. The step function represents the cumulative distribution of the squared inverse singular values of A weighted by the initial state $|b\rangle$. The effective condition number κ_{eff} is defined so that the total weight below $1/\kappa_{\text{eff}}$ (red region) is at most ϵ^2 .

1.4 Solver based on effective truncation of linear system

We now present our first beyond- κ solver based on an effective truncation of the linear system. To describe this truncation, we define the projections $\Pi_{\text{left},\mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |u_j\rangle\langle u_j|$ and $\Pi_{\text{right},\mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |v_j\rangle\langle v_j|$, for $A = \sum_j \sigma_j |u_j\rangle\langle v_j|$ the singular value decomposition with singular values $\kappa^{-1} \leq \sigma_j \leq 1$ and $\mathcal{S} \subseteq \mathbb{R}$ a subset of real numbers. Then the goal of our truncation-based solver is to produce a quantum state proportional to the truncated solution $\Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]} A^{-1}|b\rangle$. Here, we aim to choose an effective condition number $1 \leq \kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon) \leq \kappa$ as the smallest value satisfying

$$\left\| \frac{\Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]} A^{-1}|b\rangle}{\|\Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]} A^{-1}|b\rangle\|} - \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|} \right\| = \mathbf{O}(\epsilon), \quad (15)$$

so that the solver still outputs the desired solution state up to error $\mathbf{O}(\epsilon)$.

We show how to fulfill this truncation requirement by characterizing κ_{eff} through the generalized inverse (quantile function) of the weighted singular value distribution, where the weights are determined by the squared overlaps of $A^{-1}|b\rangle$ with the singular vector subspaces:

$$\frac{\left\| \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]} A^{-1}|b\rangle \right\|^2}{\|A^{-1}|b\rangle\|^2} \leq \epsilon^2, \quad \frac{\left\| \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]} A^{-1}|b\rangle \right\|^2}{\|A^{-1}|b\rangle\|^2} > \epsilon^2. \quad (16)$$

This definition is illustrated in [Figure 1](#). In practice, determining the exact value of the effective condition number κ_{eff} can be computationally prohibitive. We address this by developing a family of upper bounds on κ_{eff} such as Eq. (9), which can be analyzed using standard linear algebraic techniques. See [Theorem 1](#) for details.

Note that although the singular vector projection $\Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]}$ may be realized to precision ϵ using Quantum Singular Value Transformation (QSVT) [16] with cost $\mathbf{O}(\kappa_{\text{eff}} \log(\frac{1}{\epsilon}))$, we do *not* pursue this approach, as this would incur query overhead and significantly complicate the algorithm. Instead, we invoke a quantum linear system solver with the condition number set to κ_{eff} , without modifying the input oracles O_A and O_b . Since our κ_{eff} is no longer an upper bound on the condition number, this constitutes a somewhat “abnormal” use of the quantum linear system solver, and the resulting behavior is generally unpredictable. To address this, we demand the solver to effectively truncate the coefficient matrix in the following sense.

Definition (Strong truncation property). *Given a linear system $Ax = |b\rangle$ with $\|A\| \leq 1$ and $\| |b\rangle \| = 1$, let O_A and O_b be the oracles for block encoding and initial state preparation respectively.*

A quantum linear system solver is said to satisfy the strong truncation property if, when configured with precision ϵ and condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon) \geq 1$, it produces the truncated solution state $\Pi_{\text{right}, [\kappa_{\text{eff}}^{-1}, 1]} A^{-1} |b\rangle / \left\| \Pi_{\text{right}, [\kappa_{\text{eff}}^{-1}, 1]} A^{-1} |b\rangle \right\|$ to accuracy $\mathbf{O}(\epsilon)$.

Under this definition, any quantum linear system solver with the strong truncation property can produce a truncated solution state, with a query complexity determined by the effective condition number κ_{eff} . Recall from Eq. (8) and Eq. (9) that κ_{eff} can be bounded by vector norms that are independent of the condition number; consequently, we have obtained a beyond- κ solver. We formally state this result in [Theorem 2](#).

It remains to identify solvers that fulfill the strong truncation requirement. One simple solver can be constructed by block encoding A^{-1} using QSVT with condition number set to κ_{eff} , followed by a (fixed-point) amplitude amplification. This QSVT-based solver makes $\mathbf{O}\left(\frac{\kappa_{\text{eff}}^2}{\|x\|} \log\left(\frac{\kappa_{\text{eff}}}{\|x\|\epsilon}\right)\right)$ queries to O_A and $\mathbf{O}\left(\frac{\kappa_{\text{eff}}}{\|x\|}\right)$ queries to O_b . Running VTAA with κ_{eff} yields an improved query complexity, as shown in Eq. (7). However, rigorously establishing the strong truncation property for VTAA is highly nontrivial. The difficulty arises because VTAA recursively amplifies all components of the state that potentially lead to success. When configuring with κ_{eff} , we intend to zero out components corresponding to singular values $\sigma_j < \kappa_{\text{eff}}^{-1}$. However, VTAA still treats these components as potentially successful and amplifies them, requiring an additional postprocessing step. In effect, VTAA amplifies first and then truncates, rather than truncating first and then amplifying. This nuance leads to a nonstandard amplification schedule that is difficult to analyze directly. To circumvent this technical difficulty, we introduce the weak truncation property. Combined with an effective truncation of the initial state, this admits a significantly cleaner proof.

Specifically, with the same effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ as above, we consider the truncated initial state $|b_{\text{eff}}\rangle$ proportional to $\Pi_{\text{left}, [\kappa_{\text{eff}}^{-1}, 1]} |b\rangle$. We then construct an oracle $O_{b_{\text{eff}}}$ that prepares $|b_{\text{eff}}\rangle$ with one query to O_b and approximates O_b , i.e.,

$$O_{b_{\text{eff}}}|0\rangle = \frac{\Pi_{\text{left}, [\kappa_{\text{eff}}^{-1}, 1]} |b\rangle}{\left\| \Pi_{\text{left}, [\kappa_{\text{eff}}^{-1}, 1]} |b\rangle \right\|} = |b_{\text{eff}}\rangle, \quad \|O_{b_{\text{eff}}} - O_b\| = \mathbf{O}\left(\frac{\|x\|\epsilon}{\kappa_{\text{eff}}}\right). \quad (17)$$

Once again, the singular vector projection is *not* performed in the actual implementation of the algorithm. Instead, we run a quantum linear system solver configured with condition number κ_{eff} , while retaining the original input oracles O_A and O_b . Suppose that the solver we use has an *optimal query complexity of initial state preparation*, i.e., it makes $\mathbf{O}\left(\frac{\kappa_{\text{eff}}}{\|x\|}\right)$ queries to O_b . Then replacing every occurrence of O_b by $O_{b_{\text{eff}}}$ introduces an error at most $\mathbf{O}(\epsilon)$. Up to this error tolerance, our solver is conceptually an algorithm with condition number κ_{eff} and oracles O_A and $O_{b_{\text{eff}}}$; in this sense the initial state is effectively truncated.

As before, this is a non-standard use of the quantum linear system solver, and its correct execution requires the following additional assumption on the underlying solver.

Definition (Weak truncation property). *Given a linear system $Ax = |b\rangle$ with $\|A\| \leq 1$ and $\|b\| = 1$, let O_A and O_b be the oracles for block encoding and initial state preparation respectively. A quantum linear system solver is said to satisfy the weak truncation property if, when configured with precision ϵ and condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon) \geq 1$, it produces the truncated solution state $A^{-1}|b\rangle / \|A^{-1}|b\rangle\|$ to accuracy $\mathbf{O}(\epsilon)$, as long as the initial vector is supported in $|b\rangle \in \mathbf{Im}\left(\Pi_{\text{left}, [\kappa_{\text{eff}}^{-1}, 1]}\right)$.*

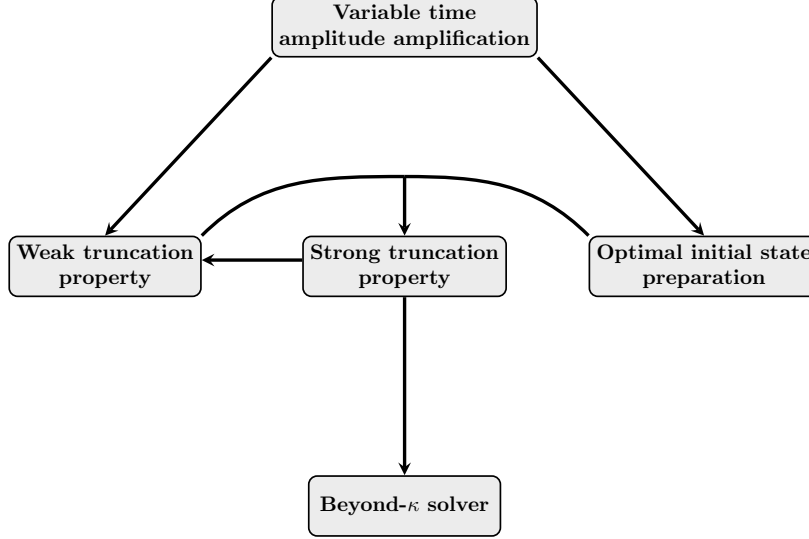


Figure 2: Illustration of the algorithmic template for constructing beyond- κ solvers based on effective truncation of the linear system.

Note that the strong truncation property naturally implies the weak truncation property. This follows from the fact that $A^{-1}|b\rangle = A^{-1}\Pi_{\text{left},[\kappa_{\text{eff}}^{-1},1]}|b\rangle = \Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]}A^{-1}|b\rangle$ whenever $|b\rangle \in \mathbf{Im}\left(\Pi_{\text{left},[\kappa_{\text{eff}}^{-1},1]}\right)$ holds. Conversely, we show in [Theorem 3](#) that a solver satisfying the weak truncation property with optimal query complexity of initial state preparation necessarily fulfills the strong truncation requirement. Given the additional promise that $|b\rangle \in \mathbf{Im}\left(\Pi_{\text{left},[\kappa_{\text{eff}}^{-1},1]}\right)$, it is routine to check that VTAA satisfies the weak truncation property. As VTAA also has an optimal query complexity of initial state preparation, it can be used to construct beyond- κ solvers based on effective state truncation ([Corollary 7](#)). We illustrate this algorithmic template in [Figure 2](#) and present a detailed analysis of the truncation-based solver in [Section 3](#).

To obtain the beyond- κ solvers, we require that VTAA produces the truncated solution state when configured with the effective condition number κ_{eff} . Interestingly, VTAA can be modified to produce the truncated solution state for *any* α supplied as the condition number parameter. This is a more general truncation property of VTAA and is discussed further in [Appendix C](#).

1.5 Solver based on filtering with effective gap

Our second beyond- κ solver is based on an improved eigenstate filtering on the unit circle with an effective gap. Specifically, suppose that the augmented matrix $[\beta A \quad -|b\rangle]$ is block encoded by the overlap of isometries $G_1^\dagger G_0$ with some scalar $\beta > 0$. We then start with the initial state $G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ partitioned conformally to the augmented matrix, which decomposes as

$$\begin{aligned}
 G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix} \\
 &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} G_0 G_0^\dagger G_1 A^{-1\dagger} x.
 \end{aligned} \tag{18}$$

Our goal is to retain the first term while discarding the second, from which a state proportional to the solution $x = A^{-1}|b\rangle$ can be extracted.

A direct calculation shows that the first term belongs to $\mathbf{Ker}(\Pi_1\Pi_0)$ while the second term is supported in $\mathbf{Im}(\Pi_0\Pi_1)$, where $\Pi_0 = G_0G_0^\dagger$ and $\Pi_1 = G_1G_1^\dagger$. As the smallest nonzero singular value of $\Pi_1\Pi_0$ is at least $\Omega(\kappa^{-1})$ [24, Appendix C], one can perform a right singular vector filtering targeting at the singular value 0 with $\Pi_1\Pi_0$ as the underlying operator. This preserves the first term while arbitrarily suppressing the second to magnitude below $\mathbf{O}(\delta)$, although the query cost would be $\mathbf{O}(\kappa \log(\frac{1}{\delta}))$ scaling linearly with the condition number.

In contrast, the beyond- κ solver uses the operator $W = (2\Pi_0 - I)(I - 2\Pi_1)$. A key observation is that the first term from the decomposition is an eigenvector of W with eigenvalue $e^{i0} = +1$, while the second term corresponds to eigenvalues $e^{i\theta}$ with nonzero phase angles $\theta \neq 0$. The effective gap lemma [22, Lemma 4.2] then allows us to reduce the magnitude of the second term below $\mathbf{O}(\delta)$ with a query cost $\mathbf{O}(\frac{1}{\delta})$ independent of the condition number. This suggests performing an eigenstate filtering over the unit circle that preserves the $\theta = 0$ component while suppressing the remaining ones with the help of effective gap lemma.

This eigenstate filtering can be realized in a variety of ways. For instance, one can use generalized quantum signal processing [33, 43] with the target polynomial chosen as the Dolph-Chebyshev polynomial over the unit circle [23]. Prior work [23, Page 30] also proposed a filtering method based on the conventional quantum signal processing [27, 30] and QSVT. However, their method relies on $\sin(\theta)$ and appears incapable of distinguishing between eigenphases $\theta = 0$ and $\theta = \pm\pi$, a critical issue left unresolved. We describe a corrected QSVT-based filtering algorithm with a matching asymptotic complexity and a reduced constant prefactor (Appendix D). Regardless, this improves upon the prior result of [24] based on quantum phase estimation and repeated sampling.

If $\|x\|$ is known a priori, we can run the beyond- κ solver with $\beta = \Theta(\|x\|)$, successfully producing the solution state with a constant probability. This filter-then-measure strategy is exceptionally simple. In fact, we show that its query complexity has a constant prefactor 6 to leading order; see Eq. (10). If $\|x\|$ is unknown, then we develop a similarly simple solution norm estimation algorithm by applying an exponential search starting from $\beta = 1$ and terminating at $\beta = \Theta(\|x\|)$ with a high probability. This yields a constant multiplicative approximation of the solution norm, and the approximation ratio can be further refined using amplitude estimation as in [15]. We present the filtering-based solver in detail in Section 4 with the result formally stated in Theorem 4.

We remark that although both solvers approximate the solution state $|x\rangle$ to an accuracy of $\mathbf{O}(\epsilon)$, their error profiles are structurally different. The truncation-based solver produces a state approximately supported on the right singular vectors of A with singular values of at least κ_{eff}^{-1} , discarding all components corresponding to smaller singular values. The filtering-based solver preserves the full solution but incurs errors from incomplete suppression of an auxiliary component in a constrained orthogonal decomposition induced by the operator W . Despite this structural difference, the two solvers provide equivalent approximation guarantees: both produce a quantum state within Euclidean distance $\mathbf{O}(\epsilon)$ of the normalized $|x\rangle$. Which solver to use then depends on the regime of interest: the truncation-based solver (via VTAA) achieves nearly linear scaling in the effective condition number with an optimal number of queries to the initial state preparation, while the filtering-based solver (with the effective spectral gap) provides a much simpler approach with a more favorable runtime prefactor.

We conclude in Section 5 with a summary of our main contributions and a number of open questions for future investigation.

2 Quantum linear system problem and input models

In this section, we formally introduce the quantum linear system problem and discuss its input models. We first define the standard input model in [Section 2.1](#) based on block encoding, and then introduce the affine dilation model in [Section 2.2](#) which allows further fine-tuning of the query complexity. Finally, we formulate the quantum linear system problem in [Section 2.3](#) and discuss its solvers.

2.1 Standard input model

Block encoding provides a natural input model for quantum linear algebra algorithms [[8](#), [26](#)]. It can be defined in full generality using unitaries and isometries. We say an operator $G : \mathcal{G} \rightarrow \mathcal{H}$ is an *isometry* between spaces \mathcal{G} and \mathcal{H} if $G^\dagger G = I$. By definition, G is injective with image $\mathbf{Im}(G) = \mathbf{Im}(GG^\dagger)$, whereas G^\dagger is surjective with kernel $\mathbf{Ker}(G^\dagger) = \mathbf{Ker}(GG^\dagger)$. We thus obtain the isometric embedding $\mathcal{G} \xrightleftharpoons[G^\dagger]{G} \mathbf{Im}(GG^\dagger) \subseteq \mathcal{H}$. Fixing any bases with respect to this embedding, the operator G can be represented in matrix form as

$$G = \begin{bmatrix} I \\ 0 \end{bmatrix}. \quad (19)$$

Examples of isometries include: (i) quantum state $|\psi\rangle$, where $\|\psi\| = 1$; (ii) unitary operator U , where $UU^\dagger = U^\dagger U = I$; (iii) tensor product $G_0 \otimes G_1$, where G_0 and G_1 are both isometries; (iv) operator composition $G_1 G_0$, where G_0 and G_1 are both isometries and the composition is well defined; and (v) Stinespring dilation $\sum_j |j\rangle \otimes E_j$, where $\sum_j E_j^\dagger E_j = I$ and $\{|j\rangle\}$ are orthonormal.

We say that A is block encoded by isometries G_0 , G_1 and unitary O_A , if

$$A = G_1^\dagger O_A G_0. \quad (20)$$

In bases compatible with the isometric embeddings, we can write O_A in the matrix form

$$O_A = \begin{bmatrix} A & \cdot \\ \cdot & \cdot \end{bmatrix}, \quad (21)$$

where A appears as the top-left block; hence the name *block encoding*. Note that such a block encoding is feasible only when $\|A\| \leq \left\| G_1^\dagger \right\| \|O_A\| \|G_0\| = 1$. Conversely, for any $\|A\| \leq 1$, one can mathematically define a dilation $O_A = \begin{bmatrix} A & -\sqrt{I-AA^\dagger} \\ \sqrt{I-A^\dagger A} & A^\dagger \end{bmatrix}$ as in [[19](#), 2.7.P2] and [[17](#)], although A may need to be further rescaled by some normalization factor for O_A to be efficiently realizable as a quantum circuit.

For a given linear system $Ax = |b\rangle$, we assume without loss of generality that the initial vector $|b\rangle$ is a normalized quantum state, prepared by oracle O_b from a standard reference state as

$$|b\rangle = O_b |0\rangle. \quad (22)$$

We further assume that the coefficient matrix A is accessed through a block encoding, with any normalization factor absorbed into the matrix so that $\|A\| \leq 1$. Depending on the quantum linear system solver under consideration, we adopt one of two simplified forms of block encoding. In the first form, we neglect the implementation cost of the isometries G_0 , G_1 and simply say that A is block encoded by the unitary O_A . This can happen for a block encoding like $A = (\langle 0| \otimes I) O_A (|0\rangle \otimes I)$ where isometries are given by trivial state preparations. In the second form, we absorb O_A into the definition of either G_0 or G_1 and say that A is block encoded by the overlap of isometries $G_1^\dagger G_0$. The query complexity is then measured by the number of calls to either G_0 or G_1 .

2.2 Affine dilation model

In the standard input model, we assumed that the coefficient matrix A has spectral norm $\|A\| \leq 1$ and is accessed through a block encoding $A = G_1^\dagger O_A G_0$. However, this often cannot be efficiently realized by a quantum circuit unless A is rescaled by an additional factor $\alpha > 1$, giving $\frac{A}{\alpha} = G_1^\dagger O_A G_0$. The normalization factor $\alpha > 1$ pushes the singular values of A toward 0, making the linear system harder to solve quantumly; minimizing α is thus critical for reducing the overall query complexity.

To this end, we introduce the affine dilation model, where the input is given by the block encoding of

$$\tilde{A} = \frac{1}{\tilde{\alpha}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix} \quad (23)$$

for some scalar c and normalization factor $\tilde{\alpha} \geq 1$. If $c = 0$, then the affine dilation reduces to the augmented matrix $\frac{1}{\tilde{\alpha}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}$ used as the input model by the beyond- κ solver of [24]. However, we can also choose $c > 0$. Then the affine dilated matrix becomes invertible and the solution $x = A^{-1}|b\rangle$ can be extracted from the top-right entry of the inverse matrix

$$\tilde{A}^{-1} = \left(\frac{1}{\tilde{\alpha}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix} \right)^{-1} = \tilde{\alpha} \begin{bmatrix} A^{-1} & \frac{1}{c} A^{-1}|b\rangle \\ 0 & \frac{1}{c} \end{bmatrix}. \quad (24)$$

This means the affine dilation model can be paired with any quantum linear system solver, allowing the query complexity to be optimized over different choices of A , $|b\rangle$, c , and $\tilde{\alpha}$.

Let us show that, even with the worst-case estimate, it never incurs significant overhead converting from the standard input model to the affine dilation model. To this end, we set $c = 1$ and decompose the affine dilation into

$$\begin{bmatrix} A & -|b\rangle \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & 1 \end{bmatrix} + \frac{1}{2} \begin{bmatrix} 0 & -|b\rangle \\ -\langle b| & 0 \end{bmatrix} + \frac{1}{2} \begin{bmatrix} 0 & -|b\rangle \\ \langle b| & 0 \end{bmatrix}, \quad (25)$$

where each summand can be block encoded respectively as

$$\begin{aligned} \begin{bmatrix} A & 0 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} G_1^\dagger & 0 \\ 0 & \langle 0| \end{bmatrix} \begin{bmatrix} O_A & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} G_0 & 0 \\ 0 & |0\rangle \end{bmatrix}, \\ \begin{bmatrix} 0 & -|b\rangle \\ -\langle b| & 0 \end{bmatrix} &= \begin{bmatrix} I & 0 \\ 0 & \langle 0| \end{bmatrix} \begin{bmatrix} 0 & -O_b \\ -O_b^\dagger & 0 \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & |0\rangle \end{bmatrix}, \\ \begin{bmatrix} 0 & -|b\rangle \\ \langle b| & 0 \end{bmatrix} &= \begin{bmatrix} I & 0 \\ 0 & \langle 0| \end{bmatrix} \begin{bmatrix} 0 & -O_b \\ O_b^\dagger & 0 \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & |0\rangle \end{bmatrix}. \end{aligned} \quad (26)$$

Taking their linear combination [16, Lemma 52] then yields a block encoding of the affine dilation matrix $\tilde{A} = \frac{1}{2} \begin{bmatrix} A & -|b\rangle \\ 0 & 1 \end{bmatrix}$ with normalization factor $\tilde{\alpha} = 2$. This implies that the condition number increases to at most

$$\|\tilde{A}^{-1}\| = \left\| 2 \begin{bmatrix} A^{-1} & A^{-1}|b\rangle \\ 0 & 1 \end{bmatrix} \right\| \leq \left\| 2 \begin{bmatrix} A^{-1} & 0 \\ 0 & 1 \end{bmatrix} \right\| + \left\| 2 \begin{bmatrix} 0 & A^{-1}|b\rangle \\ 0 & 0 \end{bmatrix} \right\| \leq 4\kappa, \quad (27)$$

while the norm of the solution encoded into the upper-right block rescales to $\|\tilde{\alpha} A^{-1}|b\rangle\| = 2\|x\|$. Since the cost of conventional quantum linear system solvers is governed by the condition number and the solution norm, switching from the standard input model to the affine dilation model thus increases the runtime by at most a constant factor. An exception to this conclusion may occur if

$\mathbf{Cost}(O_b)$ is asymptotically larger than $\mathbf{Cost}(O_A)$. In this case, quantum linear system solvers that make asymptotically fewer calls to the initial vector than to the coefficient matrix may have asymptotically better cost in the standard input model than in the affine dilation model, since block-encoding the dilated coefficient matrix A via linear combination requires querying not only O_A (cheap) but also O_b (expensive). As shown in [Table 1](#), this category of quantum linear system solver includes the VTAA method, but excludes all methods with strictly optimal query complexity to O_A .

Li [24] showed that for sparse linear systems, one can adjust the row norms in the affine dilation model, substantially reducing the query complexity. In general, searching over A , $|b\rangle$, c , and $\tilde{\alpha}$ to minimize the query complexity is a nontrivial task. Accordingly, we primarily state our results in the standard input model, though we also analyze the runtime in the affine dilation model in [Appendix B](#). A complete exploration of the affine dilation model—and the potential improvements it may yield—is an interesting direction for future work.

2.3 Quantum linear system problem

We now introduce notation and terminology for the quantum linear system problem. Specifically, consider the linear system $Ax = |b\rangle$ which has the solution

$$x = A^{-1}|b\rangle \tag{28}$$

for an invertible square matrix A . Then the goal of the quantum linear system problem is to produce the normalized solution vector

$$|x\rangle = \frac{x}{\|x\|} = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|} \tag{29}$$

as a quantum state. Since both A^{-1} and $|b\rangle$ appear linearly in the numerator and denominator of the solution state, we assume without loss of generality that the coefficient matrix satisfies $\|A\| \leq 1$ and the initial vector has unit Euclidean norm $\| |b\rangle \| = 1$. We shall use the Dirac notation only for unit vectors.

We consider the abstract setting where a quantum linear system solver is described by a unitary U , acting jointly on the ancilla and system register, with the behavior

$$\| \langle \langle 0| \otimes I \rangle U |00\rangle \rangle \|^2 > \frac{1}{2}, \quad \left\| \frac{\langle \langle 0| \otimes I \rangle U |00\rangle \rangle}{\| \langle \langle 0| \otimes I \rangle U |00\rangle \rangle \|} - \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|} \right\| = \mathbf{O}(\epsilon). \tag{30}$$

Here, the first condition guarantees that the solver succeeds with probability strictly greater than $\frac{1}{2}$, while the second ensures that the normalized output state approximates the desired solution state to within a fixed constant factor of the target accuracy $\epsilon > 0$. We focus primarily on the standard input model where A is block encoded by an oracle O_A and $|b\rangle$ is prepared by a unitary O_b . The complexity of the solver is then measured in terms of the number of queries to O_A and O_b made by U . When making non-asymptotic comparisons between two solvers that both produce solutions to $\mathbf{O}(\epsilon)$ error as in [Eq. \(30\)](#)—but with different prefactors in front of the ϵ —care must be taken to rescale ϵ accordingly.

As discussed in [Section 1.1](#), the spectral condition number is a key parameter determining the runtime of quantum linear system solvers, which is mathematically defined by $\|A\| \|A^{-1}\|$. However, the precise value of the condition number is often inaccessible in practice; instead, we work with a known upper bound

$$\kappa \geq \|A^{-1}\| = \max_{\| |b\rangle \| = 1} \|A^{-1}|b\rangle\|. \tag{31}$$

As this involves a maximization over all initial states $|b\rangle$, the condition number κ serves as a worst-case cost measure for quantum linear system solvers. The main goal of this paper is to develop beyond- κ solvers whose query complexity does not scale with the condition number.

Another parameter of importance is the solution norm $\|x\| = \|A^{-1}|b\rangle\|$. Since $1 \leq \|x\| \leq \|A^{-1}\|$, this quantity reduces to the condition number in the worst case, but typically provides a more refined cost measure for quantum linear system solvers. Like the condition number, the exact solution norm is often unavailable in practice. Instead, we assume knowledge of a value α_x that approximates $\|x\|$ to within a multiplicative factor $\mu \geq 1$, i.e.,

$$\frac{1}{\mu} \leq \frac{\alpha_x}{\|x\|} \leq \mu. \quad (32)$$

We will often assume that the solution norm is known to constant multiplicative accuracy, i.e., $\mu = \mathbf{O}(1)$. When such an estimate is not available a priori, we provide solution norm estimation algorithms to obtain one—this estimation only needs to be performed once as a preprocessing step. Furthermore, we may assume without loss of generality that $\mu \rightarrow 1$, as this refinement can be achieved using amplitude estimation incurring multiplicative complexity overhead $1/(\mu - 1)$ [15].

To summarize, a quantum linear system solver is described by a unitary $U(O_A, O_b, \kappa, \alpha_x, \epsilon)$ depending on the input oracles O_A and O_b , an upper bound on the condition number κ , a constant multiplicative estimate of the solution norm α_x , and the target accuracy ϵ . With this notation, the query complexity of the solver takes the form

$$\mathbf{Cost}(U(O_A, O_b, \kappa, \alpha_x, \epsilon)) = h_A(\kappa, \alpha_x, \epsilon) \mathbf{Cost}(O_A) + h_b(\kappa, \alpha_x, \epsilon) \mathbf{Cost}(O_b), \quad (33)$$

where h_A and h_b are positive-integer-valued functions denoting the number of calls to the respective oracles. To represent them in closed form, we introduce $\mathbf{Floor}(\cdot)$ to denote the largest integer not exceeding a real number and $\mathbf{Ceil}(\cdot)$ for the smallest integer no less than a number.

The setting above is somewhat restricted for clarity of exposition, but the framework readily accommodates several generalizations. First, the solver need not be a unitary—more generally, it can be described by a quantum channel, with the error measured in trace distance rather than Euclidean distance. This generalization is especially relevant when the solver uses classical randomness, in which case it naturally implements a mixed unitary channel. Second, we have specified the solver’s performance through its success probability and output accuracy. It is also possible to characterize the solver instead by its expected behavior and expected query complexity, assuming that the solver is repeated until a successful run is obtained. Third, we have focused on the standard input model, where A and $|b\rangle$ are accessed through separate oracles. One can also consider the affine dilation model, where A and $|b\rangle$ are block encoded jointly by a single oracle, allowing further refinements of the query complexity. Fourth, the solver’s performance may depend on additional parameters beyond κ and $\|x\|$. In particular, we will introduce a family of parameters $\|(A)_{\text{sv}}^{-t}|x\rangle\|$ with $0 < t < \infty$ based on the singular value transformation of A , which upper bounds the query complexity of beyond- κ solvers. Fifth, we have focused on the case where A is an invertible square matrix, but the framework generalizes to non-square and non-invertible matrices. In particular, when the linear system has no solution—i.e., when $|b\rangle \notin \mathbf{Im}(A)$ —a natural output is the least-squares solution, obtained by applying the Moore–Penrose pseudoinverse of A . Our truncation-based solvers can be adapted to achieve this. Finally, we have focused primarily on query complexity. Other resource measures—such as space complexity, time complexity, and query depth—are also of interest and merit further investigation.

3 Beyond- κ solver based on effective truncation of linear system

In this section, we present an algorithmic template to construct beyond- κ solvers based on effective truncation of the linear system. We begin with a formal definition of the effective condition number κ_{eff} in [Section 3.1](#) and present a family of upper bounds on κ_{eff} in [Section 3.2](#), proving [Theorem 1](#). We then describe an algorithm based on effective truncation of the coefficient matrix and analyze its runtime in [Theorem 2](#) in [Section 3.3](#). Next, we construct an oracle in [Section 3.4](#) to realize effective truncation of the initial state. Finally, we propose and analyze a solver based on effective state truncation in [Theorem 3](#) in [Section 3.5](#).

3.1 Effective condition number

Let $A = \sum_j \sigma_j |u_j\rangle\langle v_j|$ be the singular value decomposition of the coefficient matrix with singular values $\kappa^{-1} \leq \sigma_j \leq 1$. Define the singular vector projections $\Pi_{\text{left},\mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |u_j\rangle\langle u_j|$ and $\Pi_{\text{right},\mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |v_j\rangle\langle v_j|$ for $\mathcal{S} \subseteq \mathbb{R}$ a subset of real numbers. Our aim is to choose an effective condition number κ_{eff} such that

$$\frac{\left\| \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]} A^{-1} |b\rangle \right\|}{\|A^{-1}|b\rangle\|} \leq \epsilon, \quad \frac{\left\| \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]} A^{-1} |b\rangle \right\|}{\|A^{-1}|b\rangle\|} > \epsilon. \quad (34)$$

Intuitively, we use κ_{eff}^{-1} to truncate an ϵ -fraction of weight from the solution vector $x = A^{-1}|b\rangle$, starting from the smallest singular values. In cases where multiple κ_{eff} qualify, we shall select the smallest to minimize the cost of our beyond- κ solvers (corresponding to the largest possible κ_{eff}^{-1}). This defines the effective condition number as a function $\kappa_{\text{eff}}(A, |b\rangle, \epsilon)$; we suppress the arguments when clear from context.

Although the effective condition number can be defined and analyzed directly, we find it more natural to work with cumulative distribution functions and their generalized inverses. Specifically, for any non-decreasing function $f : \mathbb{R} \rightarrow \mathbb{R}$, we denote its one-sided limits by

$$f(x-) = \lim_{\epsilon \searrow 0} f(x - \epsilon), \quad f(x+) = \lim_{\epsilon \searrow 0} f(x + \epsilon). \quad (35)$$

Since f is non-decreasing, both limits exist and satisfy $f(x-) \leq f(x) \leq f(x+)$ for all x . We then say that x is a y -quantile of f if $f(x-) \leq y \leq f(x+)$. In general, the y -quantiles are not uniquely determined when f is non-invertible. We will use the largest possible value to define the generalized inverse f^+ .


Definition 1 (Generalized inverse). *For any non-decreasing function $f : \mathbb{R} \rightarrow \mathbb{R}$, its generalized inverse is defined by*

$$f^+(y) = \sup\{x \in \mathbb{R} : f(x-) \leq y \leq f(x+)\},$$


with the convention $\sup \emptyset = -\infty$.

Similarly, one can choose the smallest y -quantile and write $f^-(y) = \inf\{x : f(x-) \leq y \leq f(x+)\}$, which is more commonly used in probability. These two definitions are conceptually related though not exactly identical. For a detailed treatment of this topic, we refer the reader to the excellent note of [\[41\]](#); see also [Figure 3](#). For our analysis, we need the following properties of the generalized inverse.


Proposition 1 (Equivalent definitions of f^+). *Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be any function satisfying $f(x-) \leq g(x) \leq f(x+)$ for all x . Then the following four quantities are all equal to $f^+(y)$:*



(a) A non-decreasing function f with three behaviors: strictly increasing, plateau, and jump (right-continuous at x_0).



(b) The right-continuous generalized inverse f^+ .



(c) The left-continuous generalized inverse f^- .

Figure 3: A non-decreasing function and its generalized inverses. Plateaus of f become jumps, and jumps of f become plateaus. The only difference between f^+ and f^- is the one-sided continuity at discontinuities.

(a) $\sup\{x \in \mathbb{R} : f(x) \leq y\},$

(b) $\sup\{x \in \mathbb{R} : g(x) \leq y\},$

(c) $\inf\{x \in \mathbb{R} : f(x) > y\},$

(d) $\inf\{x \in \mathbb{R} : g(x) > y\}.$

In particular, whenever $f^+(y)$ is finite, the supremum in [Definition 1](#) is attained, so $f^+(y)$ is itself the largest y -quantile:

$$f(f^+(y)-) \leq y \leq f(f^+(y)). \quad (36)$$

Proof. We will proceed in steps.

Step 1: (a) = (c) and (b) = (d). For any function h , the sets $\{x : h(x) \leq y\}$ and $\{x : h(x) > y\}$ partition \mathbb{R} . When h is non-decreasing (or sandwiched between non-decreasing limits), the first set is a left ray and the second is a right ray. Therefore $\sup\{x : h(x) \leq y\} = \inf\{x : h(x) > y\}$. This gives (a) = (c) (taking $h = f$) and (b) = (d) (taking $h = g$). The equivalence (a) = (c) is also discussed in [\[41, Remark 1\]](#).

Step 2: (a) = (b). Since f itself satisfies $f(x-) \leq f(x) \leq f(x+)$, it is a valid choice of g . Now for any g where $f(x-) \leq g(x) \leq f(x+)$, we have the set inclusions $\{x : f(x+) \leq y\} \subseteq \{x : g(x) \leq$

$y\} \subseteq \{x : f(x-) \leq y\}$, and therefore $\sup\{x : f(x+) \leq y\} \leq \sup\{x : g(x) \leq y\} \leq \sup\{x : f(x-) \leq y\}$. It suffices to show that the two extremes are equal. Let $x \in \{x : f(x-) \leq y\}$. For every $\varepsilon > 0$, we have $f((x - \varepsilon)+) \leq f(x-) \leq y$ (since f is non-decreasing), so $x - \varepsilon \in \{x : f(x+) \leq y\}$. Hence $\sup\{x : f(x+) \leq y\} \geq x - \varepsilon$ for all $\varepsilon > 0$, giving $\sup\{x : f(x+) \leq y\} \geq x$. Since this holds for all x with $f(x-) \leq y$, we conclude $\sup\{x : f(x+) \leq y\} \geq \sup\{x : f(x-) \leq y\}$. Combined with the reverse inequality from the set inclusion, the two extremes are equal, and hence (a) = (b). This also recovers [41, Lemma 1(n)] as a special case.

Let q_0 denote the common value of (a)–(d).

Step 3: $f^+(y) \leq q_0$. The set in Definition 1 satisfies $\{x : f(x-) \leq y \leq f(x+)\} \subseteq \{x : f(x-) \leq y\}$, since the former imposes two conditions while the latter imposes only one. Taking $g(x) = f(x-)$ in (b), we have $q_0 = \sup\{x : f(x-) \leq y\}$. By the set inclusion, $f^+(y) \leq q_0$.

Step 4: $q_0 \leq f^+(y)$. We show that $f(q_0-) \leq y \leq f(q_0+)$, i.e., q_0 is a y -quantile, which gives $q_0 \leq \sup\{q : f(q-) \leq y \leq f(q+)\} = f^+(y)$. *Left inequality:* $f(q_0-) \leq y$ is [41, Lemma 1(g).(8)]. *Right inequality:* Suppose for contradiction that $f(q_0+) < y$. Then there exists $\delta > 0$ such that $f(q_0 + \delta) < y$ (by monotonicity and the definition of the right limit). In particular $f(q_0 + \delta) \leq y$, so by [41, Lemma 1(g).(5)], we obtain $q_0 \geq q_0 + \delta$, a contradiction. Hence $y \leq f(q_0+)$.

In conclusion, Steps 3 and 4 give $f^+(y) = q_0$. When $f^+(y)$ is finite, Step 4 shows q_0 is a y -quantile that dominates all others (by Step 3), so the supremum in Definition 1 is attained. \square

The quantile relation $f(f^+(y)-) \leq y \leq f(f^+(y))$ cannot be refined in general as both inequalities are saturated when f is an invertible continuous function. However, it can be further strengthened when f is a step function. Specifically, given n discrete points $x_1 < x_2 < \dots < x_n$ with associated weights $w_k > 0$, define the step function $f(x) = \sum_{x_k \leq x} w_k$. Then it holds $f(f^+(y)-) \leq y < f(f^+(y))$.

Corollary 2. *Given n real numbers $x_1 < x_2 < \dots < x_n$ with associated weights $w_k > 0$, define the step function $f(x) = \sum_{x_k \leq x} w_k$. Then,*

$$f(x-) = \sum_{x_k < x} w_k, \quad f(x+) = f(x) = \sum_{x_k \leq x} w_k. \quad (37)$$

Moreover, for all $0 < y < \sum_{k=1}^n w_k$, $f^+(y)$ is a y -quantile with the right inequality always being strict:

$$f(f^+(y)-) \leq y < f(f^+(y)). \quad (38)$$

Proof. Note that when the sum is finite, the limit commutes with the sum. Specifically, $f(x+) = \lim_{\varepsilon \searrow 0} \sum_{x_k \leq x+\varepsilon} w_k = \sum_{x_k \leq x} w_k = f(x)$, $f(x-) = \lim_{\varepsilon \searrow 0} \sum_{x_k \leq x-\varepsilon} w_k = \sum_{x_k < x} w_k$, where both equalities hold because $x_k \leq x \pm \varepsilon$ for all sufficiently small $\varepsilon > 0$ if and only if $x_k \leq x$ (resp. $x_k < x$), since the points $\{x_k\}$ are finitely many and isolated.

Since f is right-continuous, Proposition 1 gives $f(f^+(y)-) \leq y \leq f(f^+(y))$. It remains to show that the right inequality is strict.

Let $q = f^+(y) = \sup\{x : f(x) \leq y\}$. Since f is a step function with jumps only at $\{x_k\}$, the point $q = x_j$ must be one of the jump locations. Indeed, if q were not a jump location, then f would be constant in a neighborhood of q , and the supremum would extend further to the right, contradicting the definition of q .

Now suppose for contradiction that $f(x_j) \leq y$. Since f is a step function, it is constant on $[x_j, x_{j+1})$ for the next jump point $x_{j+1} > x_j$. Then $f(x) \leq y$ for all $x \in [x_j, x_{j+1})$, which gives $\sup\{x : f(x) \leq y\} \geq x_{j+1} > x_j$, contradicting $f^+(y) = x_j$. Therefore $f(f^+(y)) = f(x_j) > y$. \square

Suppose that the initial vector expands under the left singular vector basis as $|b\rangle = \sum_j w_j |u_j\rangle$. We now consider the function

$$f(\sigma) = \frac{\|\Pi_{\text{right},[0,\sigma]} A^{-1}|b\rangle\|^2}{\|A^{-1}|b\rangle\|^2} = \frac{\sum_{\sigma_j \leq \sigma} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2}. \quad (39)$$

This function can be seen as the cumulative distribution function of the discrete weights $\frac{\sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2}$. Applying the above analysis, we obtain

$$\frac{\sum_{\sigma_j < f^+(\epsilon^2)} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} = f(f^+(\epsilon^2)-) \leq \epsilon^2, \quad \frac{\sum_{\sigma_j \leq f^+(\epsilon^2)} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} = f(f^+(\epsilon^2)) > \epsilon^2. \quad (40)$$

Finally, let us define

$$\kappa_{\text{eff}} = \frac{1}{f^+(\epsilon^2)} = \sup \left\{ \sigma^{-1} \mid \frac{\sum_{\sigma_j \leq \sigma} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} > \epsilon^2 \right\} = \inf \left\{ \sigma^{-1} \mid \frac{\sum_{\sigma_j \leq \sigma} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} \leq \epsilon^2 \right\}. \quad (41)$$

This then gives the desired inequalities

$$\frac{\sum_{\sigma_j < \kappa_{\text{eff}}^{-1}} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} \leq \epsilon^2, \quad \frac{\sum_{\sigma_j \leq \kappa_{\text{eff}}^{-1}} \sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2} > \epsilon^2. \quad (42)$$

3.2 Bounds on the effective condition number

The effective condition number κ_{eff} from the previous subsection serves as a threshold for truncating the linear system, enabling quantum linear system solving beyond the condition number. However, evaluating κ_{eff} directly can be computationally difficult. We now present a family of upper bounds on κ_{eff} that can be used to bound the runtime of truncation-based beyond- κ solvers.

Specifically, for any tunable parameter $t > 0$, we have

$$\sum_j \sigma_j^{-2-2t} |w_j|^2 \geq \sum_{\sigma_j \leq \kappa_{\text{eff}}^{-1}} \sigma_j^{-2-2t} |w_j|^2 \geq \kappa_{\text{eff}}^{2t} \sum_{\sigma_j \leq \kappa_{\text{eff}}^{-1}} \sigma_j^{-2} |w_j|^2 \geq \kappa_{\text{eff}}^{2t} \epsilon^2 \sum_k \sigma_k^{-2} |w_k|^2, \quad (43)$$

which gives a family of upper bounds indexed by t :

$$\kappa_{\text{eff}} \leq \left(\frac{\sum_j \sigma_j^{-2-2t} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2 \epsilon^2} \right)^{\frac{1}{2t}}. \quad (44)$$

The hierarchy of these bounds can be understood through the following weighted power mean inequality.

Lemma 3 (Weighted power mean inequality [36, Problem 8.3] [34]). *Let $a_j > 0$ and $p_j > 0$ be positive real numbers such that $\sum_j p_j = 1$ is ℓ_1 -normalized. For any $0 < s < t < \infty$, the weighted power mean*

$$\left(\sum_j p_j a_j^s \right)^{\frac{1}{s}} \leq \left(\sum_j p_j a_j^t \right)^{\frac{1}{t}} \quad (45)$$

increases monotonically with the exponent. Moreover, we have the limiting behavior

$$\lim_{t \rightarrow \infty} \left(\sum_j p_j a_j^t \right)^{\frac{1}{t}} = \max_j a_j. \quad (46)$$

To apply this inequality, we rewrite our bounds on κ_{eff} as

$$\kappa_{\text{eff}} \leq \frac{\left(\sum_j p_j a_j^{2t}\right)^{\frac{1}{2t}}}{\epsilon^{\frac{1}{t}}}, \quad a_j = \sigma_j^{-1}, \quad p_j = \frac{\sigma_j^{-2} |w_j|^2}{\sum_k \sigma_k^{-2} |w_k|^2}. \quad (47)$$

Here, we have $a_j > 0$ and we can assume without loss of generality that $p_j > 0$ for all j as well, for otherwise we simply drop the zero terms from the summand. By the power mean inequality, we see that the numerator monotonically increases to

$$\left(\sum_j p_j a_j^{2t}\right)^{\frac{1}{2t}} \nearrow \max_j \sigma_j^{-1}. \quad (48)$$

This is at most the condition number κ , with equality when the initial state $|b\rangle$ has nonzero support on the left singular vector of A with the smallest singular value. On the other hand, the denominator monotonically increases to

$$\epsilon^{\frac{1}{t}} \nearrow 1. \quad (49)$$

Recall that $A = \sum_j \sigma_j |u_j\rangle\langle v_j|$, $|b\rangle = \sum_j w_j |u_j\rangle$, and $x = A^{-1}|b\rangle = \sum_j \sigma_j^{-1} w_j |v_j\rangle$. Using the singular value transformation with an odd function extension [16],

$$\left(\sum_j \sigma_j |u_j\rangle\langle v_j|\right)_{\text{sv}}^{-t} = \sum_j \sigma_j^{-t} |u_j\rangle\langle v_j|, \quad (50)$$

we can then represent the bounds compactly as

$$\kappa_{\text{eff}} \leq \left(\frac{\|(A)_{\text{sv}}^{-t} x\|}{\|x\| \epsilon}\right)^{\frac{1}{t}}. \quad (51)$$

For positive integers t , the singular value transformation reduces to the standard matrix inversion as:

$$\kappa_{\text{eff}} \leq \begin{cases} \left(\frac{\|A^{-1\dagger}(A^\dagger A)^{-(t-1)/2} x\|}{\|x\| \epsilon}\right)^{1/t}, & t \text{ is odd,} \\ \left(\frac{\|(A^\dagger A)^{-t/2} x\|}{\|x\| \epsilon}\right)^{1/t}, & t \text{ is even.} \end{cases} \quad (52)$$

These matrix inversions may then be analyzed using standard linear algebraic techniques.

In particular, when $t = 1$, we can bound the effective condition number by

$$\kappa_{\text{eff}} \leq \frac{\|A^{-1\dagger} x\|}{\|x\| \epsilon}. \quad (53)$$

This yields beyond- κ solvers whose runtime is comparable to that of the solver from [Section 4](#) based on filtering with effective gap. However, other choices of t may yield solvers with even lower query complexity. For instance, we can choose

$$t = \frac{\ln\left(\frac{1}{\epsilon}\right)}{\ln \ln\left(\frac{1}{\epsilon}\right)} \quad (54)$$

such that

$$\left(\frac{1}{\epsilon}\right)^{\frac{1}{t}} = \ln\left(\frac{1}{\epsilon}\right). \quad (55)$$

The solver thus achieves a logarithmic dependence on the inverse accuracy without increasing the vector norm factor beyond the condition number. We collect our findings on the effective condition number in the following theorem.

Theorem 1 (Effective condition number). *Let $A = \sum_j \sigma_j |u_j\rangle\langle v_j|$ be the singular value decomposition of matrix A with singular values $0 < \kappa^{-1} \leq \sigma_j \leq 1$. Define the singular vector projections $\Pi_{left, \mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |u_j\rangle\langle u_j|$ and $\Pi_{right, \mathcal{S}} = \sum_{\sigma_j \in \mathcal{S}} |v_j\rangle\langle v_j|$ for $\mathcal{S} \subseteq \mathbb{R}$ a subset of real numbers. Suppose that $|b\rangle$ is a normalized state expanded under the left singular vector basis as $|b\rangle = \sum_j w_j |u_j\rangle$, $x = A^{-1}|b\rangle = \sum_j \sigma_j^{-1} w_j |v_j\rangle$ and $|x\rangle = \frac{x}{\|x\|}$.*

For any $\epsilon > 0$, the effective condition number

$$\kappa_{eff} = \sup \{ \sigma^{-1} \mid \|\Pi_{right, [0, \sigma]} |x\rangle\| > \epsilon \} = \inf \{ \sigma^{-1} \mid \|\Pi_{right, [0, \sigma]} |x\rangle\| \leq \epsilon \} \quad (56)$$

satisfies

$$\left\| \Pi_{right, [0, \kappa_{eff}^{-1}]} |x\rangle \right\| \leq \epsilon, \quad \left\| \Pi_{right, [0, \kappa_{eff}^{-1}]} |x\rangle \right\| > \epsilon. \quad (57)$$

Moreover, for any $t > 0$,

$$\kappa_{eff} \leq \left(\frac{\|(A)_{sv}^{-t} |x\rangle\|}{\epsilon} \right)^{\frac{1}{t}}, \quad (58)$$

where $\left(\sum_j \sigma_j |u_j\rangle\langle v_j|\right)_{sv}^{-t} = \sum_j \sigma_j^{-t} |u_j\rangle\langle v_j|$ is the singular value transformation, which reduces to

$$\kappa_{eff} \leq \begin{cases} \left(\frac{\|A^{-1\dagger} (A^\dagger A)^{-(t-1)/2} |x\rangle\|}{\epsilon} \right)^{1/t}, & t \text{ is odd,} \\ \left(\frac{\|(A^\dagger A)^{-t/2} |x\rangle\|}{\epsilon} \right)^{1/t}, & t \text{ is even,} \end{cases} \quad (59)$$

when t is a positive integer. For $t = 1$, this gives

$$\kappa_{eff} \leq \frac{\|A^{-1\dagger} |x\rangle\|}{\epsilon}, \quad (60)$$

while $t = \Theta\left(\frac{\log(\frac{1}{\epsilon})}{\log \log(\frac{1}{\epsilon})}\right)$ yields bounds with logarithmic dependence on the inverse accuracy. As t increases, the vector norm $\|(A)_{sv}^{-t} |x\rangle\|^{1/t}$ increases monotonically below κ while the inverse error $(1/\epsilon)^{1/t}$ decreases to unity.

Proof. The claims follow from the analysis preceding the theorem statement applied to the normalized state $|x\rangle = \frac{x}{\|x\|}$, together with the observations that

$$\begin{aligned} \|x\|^2 &= \|A^{-1}|b\rangle\|^2 = \sum_j \sigma_j^{-2} |w_j|^2, \\ \|(A)_{sv}^{-t} x\|^2 &= \left\| \left(A^\dagger \right)_{sv}^{-t-1} |b\rangle \right\|^2 = \sum_j \sigma_j^{-2-2t} |w_j|^2, \\ \|\Pi_{right, [0, \sigma]} x\|^2 &= \|\Pi_{right, [0, \sigma]} A^{-1} |b\rangle\|^2 = \sum_{\sigma_j \leq \sigma} \sigma_j^{-2} |w_j|^2, \\ \|\Pi_{right, [0, \sigma)} x\|^2 &= \|\Pi_{right, [0, \sigma)} A^{-1} |b\rangle\|^2 = \sum_{\sigma_j < \sigma} \sigma_j^{-2} |w_j|^2. \end{aligned} \quad (61)$$

□

3.3 Analysis of beyond- κ solver based on effective matrix truncation

We are now ready to present and analyze our beyond- κ quantum linear system solver based on an effective matrix truncation. The output of the solver will approximate a quantum state proportional to the truncated solution $\Pi_{\text{right},[\kappa_{\text{eff}}^{-1},1]}x$. The following corollary confirms that this truncated solution, after normalization, is close to the true solution state.

Corollary 4. *With the same notation and assumptions as [Theorem 1](#),*

$$\left\| \frac{\Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x}{\|\Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x\|} - |x\rangle \right\| \leq 2\epsilon \quad (62)$$

for any upper bound on the effective condition number $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$.

Proof. Applying the estimate

$$\left\| \frac{y}{\|y\|} - \frac{x}{\|x\|} \right\| \leq \left\| \frac{y}{\|y\|} - \frac{y}{\|x\|} \right\| + \left\| \frac{y}{\|x\|} - \frac{x}{\|x\|} \right\| \leq \frac{2\|y-x\|}{\|x\|} \quad (63)$$

with $y = \Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x$, we obtain

$$\left\| \frac{\Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x}{\|\Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x\|} - \frac{x}{\|x\|} \right\| \leq \frac{2\|\Pi_{\text{right},[\alpha_{\kappa_{\text{eff}}^{-1},1]}x - x\|}{\|x\|} = \frac{2\|\Pi_{\text{right},[0,\alpha_{\kappa_{\text{eff}}^{-1}}]x\|}{\|x\|}. \quad (64)$$

Since $\alpha_{\kappa_{\text{eff}}}^{-1} \leq \kappa_{\text{eff}}^{-1}$, the associated projections are partially ordered as $\Pi_{\text{right},[0,\alpha_{\kappa_{\text{eff}}^{-1}}]} \leq \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]}$, which implies

$$\left\| \Pi_{\text{right},[0,\alpha_{\kappa_{\text{eff}}^{-1}}]}x \right\| \leq \left\| \Pi_{\text{right},[0,\kappa_{\text{eff}}^{-1}]}x \right\|. \quad (65)$$

The claimed inequality then follows from [Theorem 1](#) with the normalization $|x\rangle = \frac{x}{\|x\|}$. □

As noted in [Section 1.4](#), we do *not* perform the singular vector projection, but instead invoke a quantum linear system solver with effective condition number κ_{eff} . Specifically, we denote the solver by the unitary $U(O_A, O_b, \kappa, \epsilon)$, where κ is an upper bound on the condition number and ϵ is the target accuracy. The solver then acts on the ancilla and system register jointly and has the behavior

$$\begin{aligned} & \|(\langle 0| \otimes I) U(O_A, O_b, \kappa, \epsilon) |00\rangle\|^2 > \frac{1}{2}, \\ & \left\| \frac{(\langle 0| \otimes I) U(O_A, O_b, \kappa, \epsilon) |00\rangle}{\|(\langle 0| \otimes I) U(O_A, O_b, \kappa, \epsilon) |00\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \end{aligned} \quad (66)$$

guaranteed to succeed with a constant probability and produce a quantum state close to the normalized solution state. Since κ_{eff} is not an upper bound on the true condition number, the solver's output is no longer well defined in the usual sense. Instead, we require it to correctly output the projected solution state supported on terms whose singular values exceed κ_{eff}^{-1} .

Definition 2 (Strong truncation property). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$. Suppose that $U(O_A, O_b, \kappa, \epsilon)$ is a quantum linear system solver such that with success probability $> \frac{1}{2}$,*

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00\rangle}{\|(\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00)\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (67)$$

provided that $\kappa \geq \|A^{-1}\|$ and $\epsilon > 0$.

We say that the solver satisfies the strong truncation property if with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00\rangle}{\|(\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00)\rangle\|} - \frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x}{\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x\|} \right\| = \mathbf{O}(\epsilon), \quad (68)$$

provided that $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$ with the effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ defined by [Theorem 1](#), where $\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}$ is the right singular vector projection of A associated with singular values from $[\alpha_{\kappa_{\text{eff}}}^{-1}, 1]$.

We can now construct a beyond- κ solver by invoking a quantum linear system algorithm with the strong truncation property, using a suitable effective condition number.

Theorem 2 (Beyond- κ solver based on effective matrix truncation). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$.*

Let U be a quantum linear system solver satisfying the strong truncation property from [Definition 2](#). For any $\epsilon > 0$, define the effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ according to [Theorem 1](#). Then, with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00\rangle}{\|(\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00)\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (69)$$

provided that $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$.

Proof. By the strong truncation property, the output of the solver satisfies

$$\begin{aligned} \|\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00\rangle\|^2 &> \frac{1}{2}, \\ \left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00\rangle}{\|(\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \epsilon) |00)\rangle\|} - \frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x}{\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x\|} \right\| &= \mathbf{O}(\epsilon), \end{aligned} \quad (70)$$

Moreover, the choice of effective condition number guarantees that

$$\left\| \frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x}{\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} x\|} - |x\rangle \right\| \leq 2\epsilon. \quad (71)$$

The claimed error bound then follows from the triangle inequality. \square

Theorem 2 applies broadly, converting any conventional quantum linear system solver with the strong truncation property into a beyond- κ solver via truncation. In particular, this applies to VTAA-based solvers, although a direct proof of this claim is somewhat challenging. The difficulty is that, when configured with the effective condition number κ_{eff} and the objective of preparing a state proportional to $\Pi_{\text{right}[\kappa_{\text{eff}}^{-1}, 1]}|x\rangle$, the solver aims to apply matrix inversion on components of the initial vector corresponding to singular values above κ_{eff}^{-1} while suppressing those below κ_{eff}^{-1} . However, both components are amplified in VTAA, transforming $|b\rangle = \sum_j w_j |u_j\rangle$ into a state proportional to

$$\frac{1}{\Theta(\alpha_{\kappa_{\text{eff}}})} |0\rangle \sum_{\sigma_j \geq \alpha_{\kappa_{\text{eff}}^{-1}}} \sigma_j^{-1} w_j |v_j\rangle + |1\rangle \sum_{\sigma_j < \alpha_{\kappa_{\text{eff}}^{-1}}} w_j |v_j\rangle, \quad (72)$$

where we have omitted all irrelevant ancilla registers used by VTAA, and used singular vectors as opposed to eigenvectors of the walk operator [29, Section 4.1] for presentational purpose. It may be possible to show directly that a sufficiently good approximation of $|x\rangle$ can still be extracted from the output of truncated VTAA; however, this nuance may lead to an unconventional amplification schedule that is cumbersome to analyze. To circumvent this, we take a different route: we construct a beyond- κ solver based on an effective truncation of the initial state, which admits a cleaner, more general analysis that applies to VTAA and implies the above matrix truncation result as a corollary.

3.4 Oracle for effective state truncation

We now develop a beyond- κ solver based on an effective truncation of the initial state. Effectively, the solver replaces every instance of oracle O_b with $O_{b_{\text{eff}}}$ that prepares a truncated initial state. The construction of this truncated state preparation oracle relies on the following rotation within the 2D subspace $\mathbf{Span}(|b\rangle, |b_{\text{eff}}\rangle)$ spanned by $|b\rangle$ and $|b_{\text{eff}}\rangle$.

Lemma 5 (Palais matrix [19, 2.1.P16] [21]). *Let $|b_0\rangle$ and $|b_1\rangle$ be linearly independent unit vectors with $\langle b_0 | b_1 \rangle \in \mathbb{R}$ and define $|w\rangle = \frac{|b_0\rangle + |b_1\rangle}{\| |b_0\rangle + |b_1\rangle \|}$. Then, the Palais matrix*

$$V = I - 2|w\rangle\langle w| + 2|b_1\rangle\langle b_0| = (I - 2|w\rangle\langle w|)(I - 2|b_0\rangle\langle b_0|) \quad (73)$$

satisfies:

1. $V|b_0\rangle = |b_1\rangle$, $V|_{\mathbf{Span}^\perp(|b_0\rangle, |b_1\rangle)} = I$; and
2. $VV^\dagger = V^\dagger V = I$, $\mathbf{Det}(V) = +1$.

Conversely, the above constraints uniquely characterize the matrix V .

Proof. We begin with the uniqueness part. Suppose matrices V and W both satisfy the above constraints. Then, $\mathbf{Span}(|b_0\rangle, |b_1\rangle)$ and $\mathbf{Span}^\perp(|b_0\rangle, |b_1\rangle)$ are invariant under both W and V , whereas $W^\dagger V$ satisfies:

1. $(W^\dagger V)|b_0\rangle = |b_0\rangle$, $(W^\dagger V)|_{\mathbf{Span}^\perp(|b_0\rangle, |b_1\rangle)} = I$; and
2. $(W^\dagger V)(W^\dagger V)^\dagger = (W^\dagger V)^\dagger(W^\dagger V) = I$, $\mathbf{Det}(W^\dagger V) = +1$.

This means that $W^\dagger V$ preserves $\mathbf{Span}(|b_0\rangle, |b_1\rangle)$, restricted to which it is unitary, has determinant +1, and one eigenvalue +1. Its other eigenvalue must therefore also be +1. We conclude that $W^\dagger V = I$ is the identity operator, which yields $W = V$.

Since $|b_0\rangle$ and $|b_1\rangle$ are linearly independent, $|b_0\rangle + |b_1\rangle \neq 0$, so $|w\rangle$ is well defined. Then the claimed decomposition of the Palais matrix follows from a direct calculation

$$\begin{aligned} (I - 2|w\rangle\langle w|)(I - 2|b_0\rangle\langle b_0|) &= I - 2|w\rangle\langle w| - 2|b_0\rangle\langle b_0| + 4\langle w|b_0\rangle|w\rangle\langle b_0| \\ &= I - 2|w\rangle\langle w| - 2|b_0\rangle\langle b_0| + 4\frac{1 + \langle b_1|b_0\rangle}{2 + \langle b_0|b_1\rangle + \langle b_1|b_0\rangle}(|b_0\rangle + |b_1\rangle)\langle b_0| \quad (74) \\ &= I - 2|w\rangle\langle w| + 2|b_1\rangle\langle b_0|. \end{aligned}$$

This implies

$$V|b_0\rangle = -(I - 2|w\rangle\langle w|)|b_0\rangle = -|b_0\rangle + 2\frac{1 + \langle b_0|b_1\rangle}{2 + \langle b_0|b_1\rangle + \langle b_1|b_0\rangle}(|b_0\rangle + |b_1\rangle) = |b_1\rangle, \quad (75)$$

while the remaining claims are straightforward to verify. \square

When $|b_0\rangle \approx |b_1\rangle$, we expect that the Palais matrix rotating $|b_0\rangle$ to $|b_1\rangle$ is close to identity. Explicitly,

$$\begin{aligned} \|V - I\| &= \|-2|w\rangle\langle w| + 2|b_1\rangle\langle b_0|\| = \left\| 2 \begin{bmatrix} -|w\rangle & |b_1\rangle \end{bmatrix} \begin{bmatrix} \langle w| \\ \langle b_0| \end{bmatrix} \right\| \\ &= \left\| 2 \begin{bmatrix} \langle w| \\ \langle b_0| \end{bmatrix} \begin{bmatrix} -|w\rangle & |b_1\rangle \end{bmatrix} \right\| = \left\| 2 \begin{bmatrix} -1 & \langle w|b_1\rangle \\ -\langle b_0|w\rangle & \langle b_0|b_1\rangle \end{bmatrix} \right\| \quad (76) \\ &= \left| \langle b_0|b_1\rangle - 1 \pm i\sqrt{1 - \langle b_0|b_1\rangle^2} \right| = \sqrt{2 - 2\langle b_0|b_1\rangle}, \end{aligned}$$

where the third equality uses the fact that $V - I$ is unitarily diagonalizable with rank 2 [19, Example 1.3.23], and the fifth equality holds as $2 \begin{bmatrix} -1 & \langle w|b_1\rangle \\ -\langle b_0|w\rangle & \langle b_0|b_1\rangle \end{bmatrix}$ has eigenvalues $\langle b_0|b_1\rangle - 1 \pm i\sqrt{1 - \langle b_0|b_1\rangle^2}$. When applied to the oracles for state preparation, this gives the following distance bound.

Corollary 6. *With the same notation and assumptions as Theorem 1, define*

$$|b_{\text{eff}}\rangle = \frac{\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}|b\rangle}{\left\| \Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}|b\rangle \right\|} \quad (77)$$

and let $O_{b_{\text{eff}}} = VO_b$, where O_b is the unitary preparing $|b\rangle$ and V is the unique unitary of determinant +1 that rotates $|b\rangle$ to $|b_{\text{eff}}\rangle$ and acts trivially on $\mathbf{Span}^\perp(|b\rangle, |b_{\text{eff}}\rangle)$. Then,

$$O_{b_{\text{eff}}}|0\rangle = |b_{\text{eff}}\rangle, \quad \|O_{b_{\text{eff}}} - O_b\| \leq \sqrt{2 - 2\sqrt{1 - \frac{\epsilon^2 \|x\|^2}{\alpha_{\kappa_{\text{eff}}}^2}}} = \mathbf{O}\left(\frac{\epsilon \|x\|}{\alpha_{\kappa_{\text{eff}}}}\right) \quad (78)$$

for any $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$.

Proof. When $|b_{\text{eff}}\rangle = |b\rangle$, we adopt the convention that $O_{b_{\text{eff}}} = O_b$ so the claim is trivial. Otherwise,

we use [Theorem 1](#) to bound the distance as

$$\begin{aligned}
\|O_{b_{\text{eff}}} - O_b\| &= \sqrt{2 - 2 \frac{\langle b | \Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} | b \rangle}{\|\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} | b \rangle\|}} = \sqrt{2 - 2 \|\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]} | b \rangle\|} \\
&= \sqrt{2 - 2 \sqrt{1 - \sum_{\sigma_j < \alpha_{\kappa_{\text{eff}}}^{-1}} |w_j|^2}} \leq \sqrt{2 - 2 \sqrt{1 - \frac{1}{\alpha_{\kappa_{\text{eff}}}^2} \sum_{\sigma_j < \alpha_{\kappa_{\text{eff}}}^{-1}} \sigma_j^{-2} |w_j|^2}} \\
&\leq \sqrt{2 - 2 \sqrt{1 - \frac{1}{\alpha_{\kappa_{\text{eff}}}^2} \sum_{\sigma_j < \kappa_{\text{eff}}^{-1}} \sigma_j^{-2} |w_j|^2}} = \sqrt{2 - 2 \sqrt{1 - \frac{\|x\|^2}{\alpha_{\kappa_{\text{eff}}}^2} \|\Pi_{\text{right}, [0, \kappa_{\text{eff}}^{-1}]} | x \rangle\|^2}} \\
&\leq \sqrt{2 - 2 \sqrt{1 - \frac{\epsilon^2 \|x\|^2}{\alpha_{\kappa_{\text{eff}}}^2}}} = \mathbf{O}\left(\frac{\epsilon \|x\|}{\alpha_{\kappa_{\text{eff}}}}\right).
\end{aligned} \tag{79}$$

where [Theorem 1](#) was used for the inequality in the final line. \square

3.5 Analysis of beyond- κ solver based on effective state truncation

We are now ready to present and analyze our beyond- κ quantum linear system solver based on an effective state truncation. As mentioned in [Section 1.4](#), we do *not* perform the state truncation in the actual implementation of the algorithm. Instead, we run a quantum linear system solver configured with condition number κ_{eff} , while retaining the original input oracles O_A and O_b as well as the constant multiplicative approximation of the solution norm $\|x\|$. Suppose that the solver has an optimal query complexity of initial state preparation, i.e., it makes $\mathbf{O}\left(\frac{\kappa_{\text{eff}}}{\|x\|}\right)$ queries to O_b . Then replacing every occurrence of O_b by $O_{b_{\text{eff}}}$ introduces an error at most $\mathbf{O}(\epsilon)$. Up to this error, our solver is thus effectively configured with condition number κ_{eff} and oracles O_A and $O_{b_{\text{eff}}}$.

As κ_{eff} is not an upper bound on the true condition number, this is an unconventional application of quantum linear system solvers. We require that the solver outputs the correct solution state, given the promise that the initial vector has been truncated a priori.

Definition 3 (Weak truncation property with optimal state preparation). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$. Suppose that $U(O_A, O_b, \kappa, \alpha_x, \epsilon)$ is a quantum linear system solver such that with success probability $> \frac{1}{2}$,*

$$\left\| \frac{\langle (0| \otimes I) U(O_A, O_b, \kappa, \alpha_x, \epsilon) | 00 \rangle}{\| \langle (0| \otimes I) U(O_A, O_b, \kappa, \alpha_x, \epsilon) | 00 \rangle \|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \tag{80}$$

provided that $\kappa \geq \|A^{-1}\|$, α_x is a constant multiplicative approximation of $\|x\|$ and $\epsilon > 0$.

We say that the solver satisfies the weak truncation property if with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle (0| \otimes I) U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) | 00 \rangle}{\| \langle (0| \otimes I) U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) | 00 \rangle \|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \tag{81}$$

provided that $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$ with the effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ defined by [Theorem 1](#), and $|b\rangle \in \mathbf{Im}\left(\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}\right)$ where $\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}$ is the left singular vector projection of A associated with singular values from $[\alpha_{\kappa_{\text{eff}}}^{-1}, 1]$.

Additionally, we say that the solver has an optimal query complexity of initial state preparation if $U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)$ makes $\mathbf{O}\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\|x\|}\right)$ queries to O_b .

It follows directly from the definition that the strong truncation property implies the weak truncation property. This is because, when given the promise $|b\rangle \in \mathbf{Im}\left(\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}\right)$, we can say that

$$\frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}x}{\left\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}x\right\|} = \frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}A^{-1}|b\rangle}{\left\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}A^{-1}|b\rangle\right\|} = \frac{A^{-1}\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}|b\rangle}{\left\|A^{-1}\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}|b\rangle\right\|} = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|} = \frac{x}{\|x\|} = |x\rangle. \quad (82)$$

Conversely, if the solver has optimal initial state preparation complexity, then the weak truncation property implies the strong truncation property. This gives rise to the following algorithmic template for constructing beyond- κ solvers.

Theorem 3 (Beyond- κ solver based on effective state truncation). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$.*

Let U be a quantum linear system solver satisfying the weak truncation property from Definition 3 with an optimal query complexity of initial state preparation. For any $\epsilon > 0$, define the effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ according to Theorem 1. Then, with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle\langle 0| \otimes I\rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)|00\rangle}{\|\langle\langle 0| \otimes I\rangle U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)|00\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (83)$$

provided that $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$, α_x is a constant multiplicative approximation of $\|x\|$, and $\epsilon > 0$ is sufficiently small.

Proof. Due to the optimality of initial state preparation, the solver $U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)$ makes $\mathbf{O}\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\|x\|}\right)$ queries to O_b . Our analysis from the previous subsection then implies that the error incurred by replacing O_b with $O_{b_{\text{eff}}}$ is at most $\mathbf{O}(\epsilon)$, that is

$$\|U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) - U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)\| = \mathbf{O}\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\|x\|}\right) \mathbf{O}\left(\frac{\epsilon\|x\|}{\alpha_{\kappa_{\text{eff}}}}\right) = \mathbf{O}(\epsilon). \quad (84)$$

Note that the effective initial state is supported in $O_{b_{\text{eff}}}|0\rangle = |b_{\text{eff}}\rangle \in \mathbf{Im}\left(\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}\right)$, and the corresponding solution norm satisfies $\sqrt{1 - \epsilon^2}\|x\| \leq \|A^{-1}|b_{\text{eff}}\rangle\| = \left\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}x\right\| \leq \|x\|$. Hence, α_x is still a constant multiplicative approximation of $\|A^{-1}|b_{\text{eff}}\rangle\|$ and by the weak truncation property,

$$\begin{aligned} \|\langle\langle 0| \otimes I\rangle U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)|00\rangle\|^2 &> \frac{1}{2}, \\ \left\| \frac{\langle\langle 0| \otimes I\rangle U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)|00\rangle}{\|\langle\langle 0| \otimes I\rangle U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon)|00\rangle\|} - \frac{A^{-1}|b_{\text{eff}}\rangle}{\|A^{-1}|b_{\text{eff}}\rangle\|} \right\| &= \mathbf{O}(\epsilon). \end{aligned} \quad (85)$$

Moreover, the choice of effective condition number guarantees that

$$\left\| \frac{A^{-1}|b_{\text{eff}}\rangle}{\|A^{-1}|b_{\text{eff}}\rangle\|} - |x\rangle \right\| = \left\| \frac{\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}x}{\left\|\Pi_{\text{right}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]}x\right\|} - \frac{x}{\|x\|} \right\| \leq 2\epsilon. \quad (86)$$

Let us choose $\epsilon > 0$ sufficiently small so that

$$\|(\langle 0| \otimes I) U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) |00\rangle\|^2 > \frac{1}{2}, \quad (87)$$

fulfilling the requirement on the success probability. Meanwhile,

$$\left\| \frac{\langle \langle 0| \otimes I) U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) |00\rangle}{\|(\langle 0| \otimes I) U(O_A, O_{b_{\text{eff}}}, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) |00\rangle\|} - \frac{\langle \langle 0| \otimes I) U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) |00\rangle}{\|(\langle 0| \otimes I) U(O_A, O_b, \alpha_{\kappa_{\text{eff}}}, \alpha_x, \epsilon) |00\rangle\|} \right\| = \mathbf{O}(\epsilon). \quad (88)$$

The claimed error bound now follows from the triangle inequality. \square

Like [Theorem 2](#), [Theorem 3](#) is broadly applicable: it converts any solver with the weak truncation property and optimal query cost for state preparation into a beyond- κ solver. However, the weak truncation property has an additional promise on the support of the initial state. Thanks to this promise, verifying the weak truncation property for VTAA is straightforward: the proof from [\[29\]](#) carries over line by line within $\mathbf{Im}(\Pi_{\text{left}, [\alpha_{\kappa_{\text{eff}}}^{-1}, 1]})$. We now obtain:

Corollary 7 (Beyond- κ solver based on variable time amplitude amplification). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$.*

Then, the normalized solution state $|x\rangle = \frac{x}{\|x\|}$ can be prepared to accuracy $\epsilon > 0$ and success probability $> \frac{1}{2}$ with query complexity

$$\mathbf{O}\left(\alpha_{\kappa_{\text{eff}}} \log\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x}\right) \left(\log\log\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x}\right) + \log\left(\frac{1}{\epsilon}\right)\right) \mathbf{Cost}(O_A) + \frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x} \mathbf{Cost}(O_b)\right), \quad (89)$$

where $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}$ is an upper bound on the effective condition number $\kappa_{\text{eff}} = \kappa_{\text{eff}}(\epsilon)$ defined by [Theorem 1](#), and α_x is a constant multiplicative approximation of the solution norm, i.e., $\frac{1}{\mu} \leq \frac{\alpha_x}{\|x\|} \leq \mu$ for some constant $\mu \geq 1$.

Otherwise, the solution norm $\|x\|$ can be estimated to a constant multiplicative accuracy with query complexity

$$\mathbf{O}\left(\alpha_{\kappa_{\text{eff}}} \log^2\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x}\right) \log\log^2\left(\frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x}\right) \mathbf{Cost}(O_A) + \frac{\alpha_{\kappa_{\text{eff}}}}{\alpha_x} \mathbf{Cost}(O_b)\right), \quad (90)$$

where $\alpha_x \leq \|x\|$ is a lower bound on the solution norm, and $\alpha_{\kappa_{\text{eff}}} \geq \kappa_{\text{eff}}(\Theta(1))$ is an upper bound on the effective condition number.

We re-emphasize that the complexity of preparing the solution state $|x\rangle$ has an implicit ϵ dependence through κ_{eff} . On the other hand, for estimating the solution norm $\|x\|$ to constant multiplicative error, one may take $\epsilon = \Theta(1)$ when computing κ_{eff} , and there is no implicit ϵ dependence. Additionally, our estimator has a cost that scales with any known lower bound α_x on the solution norm. This may be improved to the exact value of $\|x\|$ using techniques from [\[7, Theorem 3\]](#) and [\[3, Section 3.3\]](#). However, we leave the detailed analysis for future work.

4 Beyond- κ solver based on filtering with effective gap

In this section, we present a beyond- κ solver based on eigenstate filtering with effective gap. We begin by discussing the input model in [Section 4.1](#) and a constrained orthogonal decomposition induced by the given linear system. We then introduce the effective gap lemma in [Section 4.2](#).

Next, we present a QSVT-based algorithm in [Section 4.3](#) for eigenstate filtering over the unit circle. In [Section 4.4](#), we propose the filtering-based beyond- κ solver and analyze the constant prefactor of its query complexity. Finally, we develop a solution norm estimation algorithm in [Section 4.5](#) with a query complexity independent of the condition number, establishing [Theorem 4](#).

4.1 Block encoding and constrained orthogonal decomposition

Let $Ax = |b\rangle$ be a system of linear equations. Recall that the filtering-based solver operates on an augmented matrix of the form $\begin{bmatrix} A & -|b\rangle \end{bmatrix}$.

Starting from the standard input model—where O_A block encodes A with $\|A\| \leq 1$ and O_b prepares $|b\rangle$ —one can set $c = 0$ in the analysis of [Section 2.2](#) to obtain a block encoding of $\frac{1}{2} \begin{bmatrix} A & -|b\rangle \end{bmatrix}$ with normalization factor 2. However, the normalization factor can be improved to $\sqrt{2}$ through the following construction

$$\frac{1}{\sqrt{2}} \begin{bmatrix} A & -|b\rangle \end{bmatrix} = \frac{1}{\sqrt{2}} [\langle 0| \otimes I \quad -\langle 0| \otimes I] \begin{bmatrix} O_A & 0 \\ 0 & I \otimes O_b \end{bmatrix} \begin{bmatrix} |0\rangle \otimes I & 0 \\ 0 & |0\rangle \otimes |0\rangle \end{bmatrix}. \quad (91)$$

We then denote this block encoding in terms of an overlap of isometries as

$$\frac{1}{\sqrt{2}} \begin{bmatrix} A & -|b\rangle \end{bmatrix} = G_1^\dagger G_0, \quad (92)$$

where

$$G_1^\dagger = \frac{1}{\sqrt{2}} [(\langle 0| \otimes I) O_A \quad -\langle 0| \otimes O_b], \quad G_0 = \begin{bmatrix} |0\rangle \otimes I & 0 \\ 0 & |0\rangle \otimes |0\rangle \end{bmatrix} \quad (93)$$

are isometries satisfying $G_0^\dagger G_0 = \begin{bmatrix} I & 0 \\ 0 & 1 \end{bmatrix}$ and $G_1^\dagger G_1 = I$. Within this block encoding of the augmented matrix, the isometry G_1 makes 1 controlled query to O_A and O_b in the standard input model, while G_0 requires no query to implement.

There are two projections $\Pi_0 = G_0 G_0^\dagger$ and $\Pi_1 = G_1 G_1^\dagger$ associated naturally with the block encoding. As Π_0 and Π_1 do not commute in general, these operators are not simultaneously diagonalizable and the orthogonal decomposition $\mathbf{Im}(\Pi_0) = (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus (\mathbf{Im}(\Pi_0) \cap \mathbf{Im}(\Pi_1))$ need not hold. Nevertheless, we can utilize the following pair of decompositions.

Lemma 8 (Constrained orthogonal decomposition). *Any orthogonal projections Π_0 and Π_1 on the same space \mathcal{H} induce the following pair of orthogonal decompositions*

$$\begin{aligned} \mathcal{H} &= \mathbf{Ker}(\Pi_1 \Pi_0) \oplus \mathbf{Im}(\Pi_0 \Pi_1), \\ \mathbf{Im}(\Pi_0) &= (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1))_{\mathbf{Im}(\Pi_0)}^\perp, \end{aligned} \quad (94)$$

where \oplus denotes the orthogonal direct sum and $(\cdot)_{\mathbf{Im}(\Pi_0)}^\perp = \mathbf{Im}(\Pi_0) \cap (\cdot)^\perp$ is the orthogonal complement within $\mathbf{Im}(\Pi_0)$. Moreover, we have the space containment

$$\begin{aligned} \mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1) &\subseteq \mathbf{Ker}(\Pi_1 \Pi_0), \\ \mathbf{Im}(\Pi_0 \Pi_1) &\subseteq (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1))_{\mathbf{Im}(\Pi_0)}^\perp. \end{aligned} \quad (95)$$

Proof. The second decomposition is a standard orthogonal decomposition restricted to the subspace $\mathbf{Im}(\Pi_0)$, whereas the first decomposition follows from

$$\mathcal{H} = \mathbf{Ker}(\Pi_1 \Pi_0) \oplus (\mathbf{Ker}(\Pi_1 \Pi_0))^\perp = \mathbf{Ker}(\Pi_1 \Pi_0) \oplus \mathbf{Im}(\Pi_1 \Pi_0)^\dagger = \mathbf{Ker}(\Pi_1 \Pi_0) \oplus \mathbf{Im}(\Pi_0 \Pi_1). \quad (96)$$

The first containment can be verified directly as

$$|\psi\rangle \in \mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1) \Rightarrow \Pi_1 \Pi_0 |\psi\rangle = \Pi_1 |\psi\rangle = 0 \Rightarrow |\psi\rangle \in \mathbf{Ker}(\Pi_1 \Pi_0), \quad (97)$$

whereas the second containment follows from

$$\mathbf{Im}(\Pi_0 \Pi_1) = (\mathbf{Ker}(\Pi_1 \Pi_0))^\perp \subseteq (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1))^\perp \quad (98)$$

together with the trivial fact that $\mathbf{Im}(\Pi_0 \Pi_1) \subseteq \mathbf{Im}(\Pi_0)$. \square

Taking the smaller subspace from each containment then yields the following constrained orthogonal decomposition

$$(\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus \mathbf{Im}(\Pi_0 \Pi_1). \quad (99)$$

Here, a generic vector from $\mathbf{Im}(\Pi_0)$ has the form $G_0 \begin{bmatrix} y \\ a \end{bmatrix}$ for some y and a , partitioned conformally to the augmented matrix. For it to also lie in $\mathbf{Ker}(\Pi_1)$, we require that $G_1^\dagger G_0 \begin{bmatrix} y \\ a \end{bmatrix} = \frac{1}{\sqrt{2}} [A \quad -|b\rangle] \begin{bmatrix} y \\ a \end{bmatrix} = 0$, which is satisfied by $G_0 \begin{bmatrix} x \\ 1 \end{bmatrix}$.

On the other hand, a generic vector from $\mathbf{Im}(\Pi_0 \Pi_1)$ has the form $G_0 G_0^\dagger G_1 z = G_0 \frac{1}{\sqrt{2}} \begin{bmatrix} A^\dagger z \\ -\langle b|z \rangle \end{bmatrix}$ for some vector z . Suppose that we initialize the algorithm with $G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix}$. To find its constrained orthogonal decomposition, we then choose $A^\dagger z = x$ to match the first component which yields $z = A^{-1\dagger} x$. This analysis can be slightly generalized to give:

Corollary 9. *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$.*

For any $\beta > 0$, the augmented matrix $\frac{1}{\sqrt{\beta^2+1}} [\beta A \quad -|b\rangle] = G_1^\dagger G_0$ can be block encoded as an overlap of isometries with

$$G_1^\dagger = \frac{1}{\sqrt{2}} [(\langle 0| \otimes I) O_A \quad -\langle 0| \otimes O_b], \quad G_0 = \begin{bmatrix} |0\rangle \otimes I & 0 \\ 0 & |0\rangle \otimes |0\rangle \end{bmatrix} \quad (100)$$

using 1 controlled query to O_A and O_b . Moreover,

$$\begin{aligned} G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\sqrt{\|x\|^2 + \beta^2}} \frac{1}{\sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \frac{\beta}{\|x\| \sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix} \\ &= \frac{\beta}{\sqrt{\|x\|^2 + \beta^2}} \frac{1}{\sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \Pi_0 \Pi_1 \frac{\sqrt{\beta^2 + 1}}{\|x\| \sqrt{\|x\|^2 + \beta^2}} G_1 A^{-1\dagger} x \\ &\in (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus \mathbf{Im}(\Pi_0 \Pi_1), \end{aligned} \quad (101)$$

where $\Pi_0 = G_0 G_0^\dagger$ and $\Pi_1 = G_1 G_1^\dagger$, and \oplus denotes the orthogonal direct sum.

4.2 Effective gap lemma

In the above decomposition, the first term encodes the solution $x = A^{-1}|b\rangle$ and is contained in $\mathbf{Ker}\left(G_1 G_1^\dagger G_0 G_0^\dagger\right)$, whereas the second term is enclosed by $\mathbf{Im}\left(G_0 G_0^\dagger G_1 G_1^\dagger\right)$. Using the fact that the smallest nonzero singular value of $G_1^\dagger G_0$ is gapped from zero by $\Omega\left(\frac{1}{\kappa}\right)$, one can apply eigenstate filtering with the underlying operator $G_1^\dagger G_0$ and extract the first term, although doing so incurs a query cost scaling linearly in the condition number κ [24, Appendix C].

In contrast, we now present the following adaption of the effective gap lemma, which can be applied with the constrained orthogonal decomposition to give a filtering with query cost independent of the condition number.

Lemma 10 (Effective gap lemma [22, Lemma 4.2]). *Let Π_0 and Π_1 be orthogonal projections on the same space, and define the unitary $W = (2\Pi_0 - I)(I - 2\Pi_1)$. For any phase function $f : \{e^{i\theta} \mid \theta \in [-\pi, \pi]\} \rightarrow [-1, 1]$, we have:*

1. for any y , $\|f(W)y\| \leq \max_{\theta \in [-\pi, \pi]} |f(e^{i\theta})| \|y\|$;
2. for any $y \in \mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)$, $f(W)y = f(1)y$; and
3. for any $y = \Pi_0 \Pi_1 z \in \mathbf{Im}(\Pi_0 \Pi_1)$, $\|f(W)y\| = \|f(W)\Pi_0 \Pi_1 z\| \leq \max_{\theta \in [-\pi, \pi]} |f(e^{i\theta})| \frac{\theta}{2} \|\Pi_1 z\|$.

Proof. The first two claims follow directly from the definition. We present a proof of the third claim using Hermitian qubitization. Specifically, we define isometries G_0 and G_1 such that $G_0 G_0^\dagger = \Pi_0$ and $G_1 G_1^\dagger = \Pi_1$. Then,

1. $G_0^\dagger (I - 2G_1 G_1^\dagger) G_0$ is Hermitian; and
2. $G_0^\dagger (I - 2G_1 G_1^\dagger)^2 G_0 = I$.

This satisfies all the assumptions of Hermitian qubitization [29, Appendix C].

Suppose that

$$G_0^\dagger (I - 2G_1 G_1^\dagger) G_0 = \bigoplus_{\lambda_u = \pm 1} [\lambda_u]_{|\phi_u\rangle} \bigoplus_{-1 < \lambda_u < 1} [\lambda_u]_{|\phi_u\rangle}. \quad (102)$$

Here, $|\phi_u\rangle$ are eigenvectors of $G_0^\dagger (I - 2G_1 G_1^\dagger) G_0$ which form an orthonormal basis of the underlying space, and $[\cdot]_{|\phi_u\rangle}$ denotes the matrix representation under the basis vector $|\phi_u\rangle$. Then, we have the matrix representations:

$$\begin{aligned} I - 2G_1 G_1^\dagger &= \bigoplus_{\lambda_u = \pm 1} [\lambda_u]_{|\phi_{u,0}\rangle} \bigoplus_{-1 < \lambda_u < 1} \begin{bmatrix} \lambda_u & \sqrt{1 - \lambda_u^2} \\ \sqrt{1 - \lambda_u^2} & -\lambda_u \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle} \bigoplus_{\xi_v = \pm 1} [\xi_v]_{|\phi_{v,0}\rangle}, \\ G_0 G_0^\dagger &= \bigoplus_{\lambda_u = \pm 1} [1]_{|\phi_{u,0}\rangle} \bigoplus_{-1 < \lambda_u < 1} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle} \bigoplus_{\xi_v = \pm 1} [0]_{|\phi_{v,0}\rangle}, \\ W &= \bigoplus_{\lambda_u = \pm 1} [\lambda_u]_{|\phi_{u,0}\rangle} \bigoplus_{-1 < \lambda_u < 1} \begin{bmatrix} \lambda_u & \sqrt{1 - \lambda_u^2} \\ -\sqrt{1 - \lambda_u^2} & \lambda_u \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle} \bigoplus_{\xi_v = \pm 1} [-\xi_v]_{|\phi_{v,0}\rangle}. \end{aligned} \quad (103)$$

Here, $|\phi_{u,0}\rangle = G_0 |\phi_u\rangle$ and $|\phi_{u,1}\rangle = \frac{(I - 2G_1 G_1^\dagger) G_0 |\phi_u\rangle - \lambda_u G_0 |\phi_u\rangle}{\sqrt{1 - \lambda_u^2}}$ provide an orthonormal basis for the 1D/2D subspace, and $|\phi_{v,0}\rangle$ is an orthonormal basis for the (uninteresting) orthogonal complement.

Within a 2D subspace, $I - 2G_1G_1^\dagger$ has eigenpairs

$$+1 : |\phi_{u,ZX+}\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} \sqrt{1+\lambda_u} \\ \sqrt{1-\lambda_u} \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle}, \quad -1 : |\phi_{u,ZX-}\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} \sqrt{1-\lambda_u} \\ -\sqrt{1+\lambda_u} \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle}, \quad (104)$$

whereas $W = (2G_0G_0^\dagger - I)(I - 2G_1G_1^\dagger)$ has eigenpairs

$$e^{i \arccos(\lambda_u)} : |\phi_{u,Y+}\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle}, \quad e^{-i \arccos(\lambda_u)} : |\phi_{u,Y-}\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}_{|\phi_{u,0}\rangle, |\phi_{u,1}\rangle}. \quad (105)$$

Now suppose we start with an arbitrary $+1$ eigenvector of $G_1G_1^\dagger$ (or equivalently, -1 eigenvector of $I - 2G_1G_1^\dagger$):

$$w = \sum_{\lambda_u=-1} \alpha_u |\phi_{u,0}\rangle + \sum_{-1 < \lambda_u < 1} \alpha_u \frac{\sqrt{1-\lambda_u} |\phi_{u,0}\rangle - \sqrt{1+\lambda_u} |\phi_{u,1}\rangle}{\sqrt{2}} + \sum_{\xi_v=-1} \alpha_v |\phi_{v,0}\rangle. \quad (106)$$

Then, applying $G_0G_0^\dagger$ mathematically transforms the above vector into

$$G_0G_0^\dagger w = \sum_{\lambda_u=-1} \alpha_u |\phi_{u,0}\rangle + \sum_{-1 < \lambda_u < 1} \alpha_u \frac{\sqrt{1-\lambda_u} |\phi_{u,0}\rangle}{\sqrt{2}}. \quad (107)$$

Finally, we apply $f(W)$ onto each eigensubspace with eigenvalue $e^{\pm i \arccos(\lambda_u)}$, producing

$$f(W)G_0G_0^\dagger w = \sum_{-1 \leq \lambda_u < 1} \alpha_u \frac{\sqrt{1-\lambda_u}}{\sqrt{2}} \left(\frac{f(e^{i \arccos(\lambda_u)}) |\phi_{u,Y+}\rangle + f(e^{-i \arccos(\lambda_u)}) |\phi_{u,Y-}\rangle}{\sqrt{2}} \right). \quad (108)$$

It is clear that the resulting (unnormalized) vector has length

$$\begin{aligned} \|f(W)G_0G_0^\dagger w\| &= \left\| \sum_{-1 \leq \lambda_u < 1} \alpha_u \frac{\sqrt{1-\lambda_u}}{\sqrt{2}} \left(\frac{f(e^{i \arccos(\lambda_u)}) |\phi_{u,Y+}\rangle + f(e^{-i \arccos(\lambda_u)}) |\phi_{u,Y-}\rangle}{\sqrt{2}} \right) \right\| \\ &= \sqrt{\sum_{-1 \leq \lambda_u < 1} |\alpha_u|^2 \frac{1-\lambda_u}{2} \left(\frac{|f(e^{i \arccos(\lambda_u)})|^2}{2} + \frac{|f(e^{-i \arccos(\lambda_u)})|^2}{2} \right)} \\ &\leq \sqrt{\sum_{-1 \leq \lambda_u < 1} |\alpha_u|^2} \sqrt{\max_{\theta \in [-\pi, \pi]} \frac{1 - \cos(\theta)}{2} |f(e^{i\theta})|^2} \leq \|w\| \max_{\theta \in [-\pi, \pi]} \left| \frac{\theta}{2} f(e^{i\theta}) \right|. \end{aligned} \quad (109)$$

The proof is now complete by setting $w = \Pi_1 z$. \square

4.3 Eigenstate filtering over unit circle

In light of the constrained orthogonal decomposition, our goal is to choose a filter function to preserve the eigenvalue $+1$ of the underlying unitary operator W and suppress eigenvalues with gap δ on the unit circle. That is, we want to perform $\mathbf{F}(W)$ such that

$$\mathbf{F}(1) = 1, \quad \max_{\delta \leq |\theta| \leq \pi} |\mathbf{F}(e^{i\theta})| \leq \xi, \quad \max_{\theta \in [-\pi, \pi]} |\mathbf{F}(e^{i\theta})| = 1, \quad (110)$$




Figure 4: Illustration of QSVT-based eigenstate filtering over the unit circle by shifting.

with circular gap $\delta > 0$ and precision $\xi > 0$. All the remaining eigenvalues with $|\theta| < \delta$ are then suppressed by the effective gap lemma.

In this subsection, we present a QSVT-based algorithm to realize this eigenstate filtering over the unit circle, correcting the phase ambiguity in the method of [23, Page 30]. The core idea is to shift the input unitary operator, so that the eigenvalue of interest is mapped to the origin. See also [31, Section V]. Specifically, starting from the unitary W , we construct the block encoding

$$\frac{W - I}{2} = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes I \right) (|0\rangle\langle 0| \otimes W - |1\rangle\langle 1| \otimes I) \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes I \right). \quad (111)$$

Their eigenvalues correspond as follows.

1. Eigenvalues of W are on the unit circle. The eigenvalue of interest is $+1$, while eigenvalues to be suppressed are δ -gapped away in circular distance.
2. Eigenvalues of $\frac{W-I}{2}$ are on the circle of radius $\frac{1}{2}$ centered at $-\frac{1}{2}$. The eigenvalue of interest is 0, while eigenvalues to be suppressed are $\sin\left(\frac{\delta}{2}\right)$ -gapped away in Euclidean distance.

See Figure 4 for an illustration of this shifting.

After the shifting, $\frac{W-I}{2}$ is still unitarily diagonalizable, and its eigenvalues coincide with singular values in magnitude. The singular value of interest is 0, while the singular values to be suppressed are $\sin\left(\frac{\delta}{2}\right)$ -gapped away. We achieve this filtering using QSVT to perform the Dolph-Chebyshev polynomials, whose relevant properties are summarized in the following lemma.

Lemma 11 (Dolph-Chebyshev polynomials [15, Appendix B.2]). *Given $0 < \eta < 1$, define the degree- $2l$ even polynomials*

$$\mathbf{F}_{\eta,2l}(x) = \frac{\mathbf{T}_l\left(\frac{1+\eta^2-2x^2}{1-\eta^2}\right)}{\mathbf{T}_l\left(\frac{1+\eta^2}{1-\eta^2}\right)}, \quad (112)$$

where \mathbf{T}_l is the degree- l Chebyshev polynomial of the first kind. The polynomial $\mathbf{F}_{\eta,2l}(x)$ has the following properties.

1. $\mathbf{F}_{\eta,2l}(0) = 1$;

$$2. \max_{\eta \leq |x| \leq 1} |\mathbf{F}_{\eta, 2l}(x)| \leq \frac{1}{\mathbf{T}_l\left(\frac{1+\eta^2}{1-\eta^2}\right)}; \text{ and}$$

$$3. \max_{x \in [-1, 1]} |\mathbf{F}_{\eta, 2l}(x)| = 1.$$

Moreover, to achieve $\frac{1}{\mathbf{T}_l\left(\frac{1+\eta^2}{1-\eta^2}\right)} \leq \xi$, it suffices to choose

$$l = \mathbf{Ceil} \left(\frac{\operatorname{arccosh}\left(\frac{1}{\xi}\right)}{\operatorname{arccosh}\left(\frac{1+\eta^2}{1-\eta^2}\right)} \right) \leq \mathbf{Ceil} \left(\frac{1}{2\eta} \ln \left(\frac{2}{\xi} \right) \right). \quad (113)$$

Corollary 12. Let W be a unitary with spectral decomposition $W = \sum_u e^{i\theta_u} |\phi_u\rangle\langle\phi_u|$. Define the Dolph-Chebyshev polynomials as in Lemma 11. For any circular gap $0 < \delta < \pi$ and precision $0 < \xi < 1$, the operator

$$\left(\mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \right)_{\mathbf{sv}} \left(\frac{W - I}{2} \right) = \sum_u \mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \left(\sin \left(\frac{\theta_u}{2} \right) \right) |\phi_u\rangle\langle\phi_u| \quad (114)$$

can be block encoded using

$$2l = 2 \mathbf{Ceil} \left(\frac{\operatorname{arccosh}\left(\frac{1}{\xi}\right)}{2 \operatorname{arccosh}\left(\frac{1}{\cos(\frac{\delta}{2})}\right)} \right) \leq 2 \mathbf{Ceil} \left(\frac{1}{\delta} \ln \left(\frac{2}{\xi} \right) \right). \quad (115)$$

queries to controlled W and W^\dagger , where the circular function

$$\mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \left(\sin \left(\frac{\theta}{2} \right) \right) = \frac{\mathbf{T}_{2l} \left(\frac{\cos(\frac{\theta}{2})}{\cos(\frac{\delta}{2})} \right)}{\mathbf{T}_{2l} \left(\frac{1}{\cos(\frac{\delta}{2})} \right)} \quad (116)$$

satisfies

$$1. \mathbf{F}_{\sin(\frac{\delta}{2}), 2l}(0) = 1;$$

$$2. \max_{\delta \leq |\theta| \leq \pi} \left| \mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \left(\sin \left(\frac{\theta}{2} \right) \right) \right| \leq \xi; \text{ and}$$

$$3. \max_{\theta \in [-\pi, \pi]} \left| \mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \left(\sin \left(\frac{\theta}{2} \right) \right) \right| = 1.$$

Proof. We start with a block encoding of $\frac{W-I}{2}$ as constructed above using one controlled application of W . Given the circular gap δ , we perform QSVT on $\frac{W-I}{2}$ with $\mathbf{F}_{\sin(\frac{\delta}{2}), 2l}$ as the target polynomial. The resulting algorithm makes $2l$ queries to the block encoding oracle for $\frac{W-I}{2}$ and its inverse, which translates to $2l$ controlled W and W^\dagger as claimed.

The output of QSVT is a block encoding of $\left(\mathbf{F}_{\sin(\frac{\delta}{2}), 2l} \right)_{\mathbf{sv}} \left(\frac{W-I}{2} \right)$. We now examine this singular value transformation more closely. First, recall that $\frac{W-I}{2}$ has eigenvalues $\frac{e^{i\theta_u} - 1}{2} = |\sin(\frac{\theta_u}{2})| e^{ig(\theta_u)}$ for some function g . Here, $|\sin(\frac{\theta_u}{2})|$ is the singular value, whereas $e^{ig(\theta_u)}$ is absorbed into the right singular vector. After QSVT, we have $\mathbf{F}_{\sin(\frac{\delta}{2}), 2l}(|\sin(\frac{\theta_u}{2})|)$ in the singular vector basis. Note that $\mathbf{F}_{\sin(\frac{\delta}{2}), 2l}$ is an even function. So the original phase $e^{ig(\theta_u)}$ will be canceled out by QSVT, and

we have $\mathbf{F}_{\sin(\frac{\delta}{2}),2l}(|\sin(\frac{\theta_u}{2})|)$ in the eigenbasis as well. Using again the fact that $\mathbf{F}_{\sin(\frac{\delta}{2}),2l}$ is an even function, we finally obtain $\mathbf{F}_{\sin(\frac{\delta}{2}),2l}(|\sin(\frac{\theta_u}{2})|) = \mathbf{F}_{\sin(\frac{\delta}{2}),2l}(\sin(\frac{\theta_u}{2}))$ in the eigenbasis. This establishes the claimed eigendecomposition.

The remaining claims follow from the representation of Dolph-Chebyshev polynomials

$$\frac{\mathbf{T}_{2l}\left(\frac{\cos(\frac{\theta}{2})}{\cos(\frac{\delta}{2})}\right)}{\mathbf{T}_{2l}\left(\frac{1}{\cos(\frac{\delta}{2})}\right)} = \frac{\mathbf{T}_{2l}\left(\sqrt{\frac{1-x^2}{1-\eta^2}}\right)}{\mathbf{T}_{2l}\left(\sqrt{\frac{1}{1-\eta^2}}\right)} = \frac{\mathbf{T}_l\left(\mathbf{T}_2\left(\sqrt{\frac{1-x^2}{1-\eta^2}}\right)\right)}{\mathbf{T}_l\left(\mathbf{T}_2\left(\sqrt{\frac{1}{1-\eta^2}}\right)\right)} = \frac{\mathbf{T}_l\left(2\frac{1-x^2}{1-\eta^2} - 1\right)}{\mathbf{T}_l\left(2\frac{1}{1-\eta^2} - 1\right)} = \frac{\mathbf{T}_l\left(\frac{1+\eta^2-2x^2}{1-\eta^2}\right)}{\mathbf{T}_l\left(\frac{1+\eta^2}{1-\eta^2}\right)} \quad (117)$$

under the change of variables $x = \sin(\frac{\theta}{2})$ and $\eta = \sin(\frac{\delta}{2})$. In particular, the claimed query complexity follows from [13, Eq. (113)]. \square

4.4 Constant-prefactor analysis of beyond- κ solver based on filtering

We now present our beyond- κ solver based on eigenstate filtering over the unit circle with effective gap. We also estimate the constant prefactor of its query complexity to leading order.

Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$. Applying Corollary 9, we obtain a block encoding of the augmented matrix as an overlap of isometries

$$\frac{1}{\sqrt{\beta^2 + 1}} [\beta A \quad -|b\rangle] = G_1^\dagger G_0 \quad (118)$$

for any scalar $\beta > 0$, using 1 query to O_A and O_b . Moreover, this induces the constrained orthogonal decomposition

$$\begin{aligned} G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\sqrt{\|x\|^2 + \beta^2}} \frac{1}{\sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \frac{\beta}{\|x\| \sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix} \\ &= \frac{\beta}{\sqrt{\|x\|^2 + \beta^2}} \frac{1}{\sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \Pi_0 \Pi_1 \frac{\sqrt{\beta^2 + 1}}{\|x\| \sqrt{\|x\|^2 + \beta^2}} G_1 A^{-1\dagger} x, \end{aligned} \quad (119)$$

where $\Pi_0 = G_0 G_0^\dagger$ and $\Pi_1 = G_1 G_1^\dagger$.

Define the unitary $W = (2\Pi_0 - I)(I - 2\Pi_1) = \sum_u e^{i\theta_u} |\phi_u\rangle\langle\phi_u|$ and perform the filtering $\mathbf{F}_{\text{sv}}\left(\frac{W-I}{2}\right)$ with circular gap $0 < \delta < \pi$ and precision $0 < \xi < 1$, where we have omitted the subscripts for notational convenience. By the effective gap lemma (Lemma 10),

$$\begin{aligned} \mathbf{F}_{\text{sv}}\left(\frac{W-I}{2}\right) G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\beta}{\|x\|^2 + \beta^2} \mathbf{F}_{\text{sv}}\left(\frac{W-I}{2}\right) G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix} \\ &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \mathbf{F}_{\text{sv}}\left(\frac{W-I}{2}\right) \Pi_0 \Pi_1 G_1 A^{-1\dagger} x, \end{aligned} \quad (120)$$

where Corollary 12 shows

$$\mathbf{F}_{\text{sv}}\left(\frac{W-I}{2}\right) = \sum_u \mathbf{F}\left(\sin\left(\frac{\theta_u}{2}\right)\right) |\phi_u\rangle\langle\phi_u|. \quad (121)$$

We now split the second term further into two sub terms. To this end, we introduce the circular indicator function

$$\mathbf{Ind}_{[0, \sin(\frac{\delta}{2})]} \left(\sin \left(\frac{\theta}{2} \right) \right) = \begin{cases} 1, & |\theta| < \delta, \\ 0, & \delta \leq |\theta| \leq \pi. \end{cases} \quad (122)$$

This yields the final decomposition

$$\begin{aligned} \mathbf{F}_{\text{sv}} \left(\frac{W - I}{2} \right) G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} \\ &\quad - \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \sum_u (\mathbf{F} \mathbf{Ind}) \left(\sin \left(\frac{\theta_u}{2} \right) \right) |\phi_u\rangle\langle\phi_u| \Pi_0 \Pi_1 G_1 A^{-1\dagger} x \\ &\quad - \frac{\beta}{\|x\|^2 + \beta^2} \sum_u (\mathbf{F}(1 - \mathbf{Ind})) \left(\sin \left(\frac{\theta_u}{2} \right) \right) |\phi_u\rangle\langle\phi_u| G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix}. \end{aligned} \quad (123)$$

In this new decomposition, the third term can be bounded directly using [Corollary 12](#) as

$$\begin{aligned} &\left\| \frac{\beta}{\|x\|^2 + \beta^2} \sum_u (\mathbf{F}(1 - \mathbf{Ind})) \left(\sin \left(\frac{\theta_u}{2} \right) \right) |\phi_u\rangle\langle\phi_u| G_0 \begin{bmatrix} x \\ -\frac{\|x\|^2}{\beta} \end{bmatrix} \right\| \\ &\leq \frac{\beta}{\|x\|^2 + \beta^2} \sqrt{\|x\|^2 + \frac{\|x\|^4}{\beta^2}} \max_{\theta \in [-\pi, \pi]} \left| \mathbf{F} \left(\sin \left(\frac{\theta}{2} \right) \right) (1 - \mathbf{Ind}) \left(\sin \left(\frac{\theta}{2} \right) \right) \right| \\ &= \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \max_{\delta \leq |\theta| \leq \pi} \left| \mathbf{F} \left(\sin \left(\frac{\theta}{2} \right) \right) \right| \leq \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi, \end{aligned} \quad (124)$$

whereas for the second term we apply the effective gap lemma ([Lemma 10](#))

$$\begin{aligned} &\left\| \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \sum_u (\mathbf{F} \mathbf{Ind}) \left(\sin \left(\frac{\theta_u}{2} \right) \right) |\phi_u\rangle\langle\phi_u| \Pi_0 \Pi_1 G_1 A^{-1\dagger} x \right\| \\ &\leq \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \max_{\theta \in [-\pi, \pi]} \left| \mathbf{F} \left(\sin \left(\frac{\theta}{2} \right) \right) \mathbf{Ind} \left(\sin \left(\frac{\theta}{2} \right) \right) \frac{\theta}{2} \right| \|G_1 A^{-1\dagger} x\| \\ &\leq \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \max_{|\theta| < \delta} \left| \mathbf{F} \left(\sin \left(\frac{\theta}{2} \right) \right) \frac{\theta}{2} \right| \|A^{-1\dagger} x\| \leq \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger} x\| \frac{\delta}{2}. \end{aligned} \quad (125)$$

We thus obtain the following error bound for the eigenstate filtering

$$\left\| \mathbf{F}_{\text{sv}} \left(\frac{W - I}{2} \right) G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} - \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} \right\| \leq \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger} x\| \frac{\delta}{2} + \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi. \quad (126)$$

This filtering uses $2 \mathbf{Ceil} \left(\frac{1}{\delta} \ln \left(\frac{2}{\xi} \right) \right)$ queries to controlled W and W^\dagger . With $W = (2G_0G_0^\dagger - I)(I - 2G_1G_1^\dagger)$, each W or W^\dagger can be realized using 2 (uncontrolled) queries to the block encoding isometries G_0 and G_1 . Each G_1 then requires 1 controlled application of O_A and O_b whereas G_0 can be realized with zero query (Eq. (100)), giving the cost

$$4 \mathbf{Ceil} \left(\frac{1}{\delta} \ln \left(\frac{2}{\xi} \right) \right) \quad (127)$$

in the standard input model.

Let us choose the effective gap δ and filtering accuracy ξ so that the solution vector x has error at most $\|x\| \epsilon$. Equivalently, we require that

$$\frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}x\| \frac{\delta}{2} + \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi \leq \frac{\beta \|x\|}{\|x\|^2 + \beta^2} \epsilon. \quad (128)$$

A standard strategy is to divide the total error budget into two parts of comparable size as in [23, Section 4.1]. Here, we present an alternative approach that matches their asymptotic scaling but improves the constant prefactor. Specifically, we let

$$\frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}x\| \frac{\delta}{2} = \gamma \frac{\beta \|x\|}{\|x\|^2 + \beta^2} \epsilon, \quad \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi = (1 - \gamma) \frac{\beta \|x\|}{\|x\|^2 + \beta^2} \epsilon, \quad (129)$$

for some tunable parameter $0 < \gamma < 1$ to be chosen momentarily. Then the query cost becomes

$$\begin{aligned} \text{Ceil} \left(\frac{4}{\delta} \ln \left(\frac{2}{\xi} \right) \right) &= \text{Ceil} \left(\frac{2\sqrt{\beta^2 + 1} \|A^{-1\dagger}x\|}{\beta \|x\|} \frac{1}{\gamma \epsilon} \ln \left(\frac{2\sqrt{\|x\|^2 + \beta^2}}{\beta} \frac{1}{(1 - \gamma)\epsilon} \right) \right) \\ &= \text{Ceil} \left(\frac{\sqrt{\beta^2 + 1} \|A^{-1\dagger}x\|}{\|x\| \sqrt{\|x\|^2 + \beta^2}} \frac{1}{\gamma \tilde{\epsilon}} \ln \left(\frac{1}{(1 - \gamma)\tilde{\epsilon}} \right) \right), \end{aligned} \quad (130)$$

where

$$\tilde{\epsilon} = \frac{\beta \epsilon}{2\sqrt{\|x\|^2 + \beta^2}}. \quad (131)$$

In [Appendix D](#), we prove that the function

$$g(\gamma) = \frac{1}{\gamma \tilde{\epsilon}} \ln \left(\frac{1}{(1 - \gamma)\tilde{\epsilon}} \right), \quad 0 < \gamma < 1. \quad (132)$$

is minimized at $\gamma_0 = 1 + \frac{1}{\mathbf{W}_{-1}(-\frac{\tilde{\epsilon}}{e})}$, where $\mathbf{W}_{-1}(\cdot)$ is the (-1) -branch of the Lambert-W function. Correspondingly,

$$\min g(\gamma) = g(\gamma_0) = \frac{\mathbf{W}_{-1}(-\frac{\tilde{\epsilon}}{e})}{1 + \mathbf{W}_{-1}(-\frac{\tilde{\epsilon}}{e})} \frac{1}{\tilde{\epsilon}} \ln \left(\frac{-\mathbf{W}_{-1}(-\frac{\tilde{\epsilon}}{e})}{\tilde{\epsilon}} \right) \quad (133)$$

which reduces to

$$\frac{1}{\tilde{\epsilon}} \ln \left(\frac{1}{\tilde{\epsilon}} \right) \quad (134)$$

when higher order terms are dropped. However, the exact values of $\|x\|$ and $\|A^{-1\dagger}x\|$ are typically unavailable in practice. Instead, we assume that $\alpha_{A^{-1\dagger}|x} \geq \|A^{-1\dagger}|x\rangle\|$ is a norm upper bound and $\frac{1}{\mu} \leq \frac{\alpha_x}{\|x\|} \leq \mu$ is a μ -multiplicative approximation of the solution norm $\|x\|$ for some $\mu \geq 1$. Setting

$\beta = \alpha_x$, we can say that the query complexity is at most

$$\begin{aligned} & \text{Ceil} \left(\frac{\sqrt{\alpha_x^2 + 1} \alpha_{A^{-1}\dagger|x}}{\sqrt{\|x\|^2 + \alpha_x^2}} \frac{\mathbf{W}_{-1} \left(-\frac{\alpha_x \epsilon}{2e\sqrt{\|x\|^2 + \alpha_x^2}} \right)}{1 + \mathbf{W}_{-1} \left(-\frac{\alpha_x \epsilon}{2e\sqrt{\|x\|^2 + \alpha_x^2}} \right)} \frac{2\sqrt{\|x\|^2 + \alpha_x^2}}{\alpha_x \epsilon} \ln \left(-\mathbf{W}_{-1} \left(-\frac{\alpha_x \epsilon}{2e\sqrt{\|x\|^2 + \alpha_x^2}} \right) \frac{2\sqrt{\|x\|^2 + \alpha_x^2}}{\alpha_x \epsilon} \right) \right) \\ & \leq \text{Ceil} \left(\frac{\mu\sqrt{\alpha_x^2 + 1} \alpha_{A^{-1}\dagger|x}}{\sqrt{\mu^2 + 1} \alpha_x} \frac{\mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right)}{1 + \mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right)} \frac{2\sqrt{\mu^2 + 1}}{\epsilon} \ln \left(-\mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right) \frac{2\sqrt{\mu^2 + 1}}{\epsilon} \right) \right), \end{aligned} \quad (135)$$

since shrinking the error budget from $\tilde{\epsilon} = \frac{\alpha_x \epsilon}{2\sqrt{\|x\|^2 + \alpha_x^2}}$ to $\tilde{\epsilon} = \frac{\epsilon}{2\sqrt{\mu^2 + 1}}$ never decreases the cost. To leading order, this simplifies to

$$\frac{\mu\sqrt{\alpha_x^2 + 1} \alpha_{A^{-1}\dagger|x}}{\sqrt{\mu^2 + 1} \alpha_x} \frac{2\sqrt{\mu^2 + 1}}{\epsilon} \ln \left(\frac{2\sqrt{\mu^2 + 1}}{\epsilon} \right) \simeq \frac{2\alpha_{A^{-1}\dagger|x}}{\epsilon} \ln \left(\frac{1}{\epsilon} \right). \quad (136)$$

With the above choice of effective gap δ , filtering accuracy ξ and scalar β , the success amplitude is lower bounded by

$$\frac{\beta \|x\| (1 - \epsilon)}{\|x\|^2 + \beta^2} = \frac{\alpha_x \|x\| (1 - \epsilon)}{\|x\|^2 + \alpha_x^2} = \frac{(1 - \epsilon)}{\frac{\|x\|}{\alpha_x} + \frac{\alpha_x}{\|x\|}} \geq \frac{1 - \epsilon}{\mu + \frac{1}{\mu}}. \quad (137)$$

Applying the fixed-point amplitude amplification [15, Page 29], we obtain the solution state $|x\rangle = \frac{x}{\|x\|}$ with accuracy ϵ and failure probability p_{fail} using

$$2 \mathbf{Floor} \left(\frac{1}{2} \text{Ceil} \left(\frac{\mu + \frac{1}{\mu}}{1 - \epsilon} \ln \left(\frac{2}{\sqrt{p_{\text{fail}}}} \right) \right) \right) + 1 \quad (138)$$

repetitions of filtering. Here, we have used $2 \mathbf{Floor} \left(\frac{1}{2} \text{Ceil}(\cdot) \right) + 1$ to denote the smallest odd integer greater than or equal to its argument.

Altogether, to succeed with a probability at least $\frac{1}{2}$, the solver makes

$$\begin{aligned} & \left(2 \mathbf{Floor} \left(\frac{1}{2} \text{Ceil} \left(\frac{\mu + \frac{1}{\mu}}{1 - \epsilon} \ln (2\sqrt{2}) \right) \right) \right) + 1 \\ & \cdot \text{Ceil} \left(\frac{\mu\sqrt{\alpha_x^2 + 1} \alpha_{A^{-1}\dagger|x}}{\sqrt{\mu^2 + 1} \alpha_x} \frac{\mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right)}{1 + \mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right)} \frac{2\sqrt{\mu^2 + 1}}{\epsilon} \ln \left(-\mathbf{W}_{-1} \left(-\frac{\epsilon}{2e\sqrt{\mu^2 + 1}} \right) \frac{2\sqrt{\mu^2 + 1}}{\epsilon} \right) \right) \end{aligned} \quad (139)$$

queries to the oracles O_A and O_b (recall from Eq. (100) that undoing the isometry G_0^\dagger has zero query cost). In the limit $\mu \rightarrow 1$, $\epsilon \rightarrow 0$, $\alpha_x \rightarrow \infty$, $\alpha_{A^{-1}\dagger|x} \rightarrow \infty$, the leading-order query complexity is

$$\left(2 \mathbf{Floor} \left(\frac{1}{2} \text{Ceil} (3 \ln (2)) \right) \right) + 1 \frac{2\alpha_{A^{-1}\dagger|x}}{\epsilon} \ln \left(\frac{1}{\epsilon} \right) = \frac{6\alpha_{A^{-1}\dagger|x}}{\epsilon} \ln \left(\frac{1}{\epsilon} \right). \quad (140)$$

4.5 Beyond- κ solution norm estimation

The simple beyond- κ solver from the previous subsection relies on the assumption that the solution norm $\|x\|$ is known to a constant relative accuracy. If we do not have this prior knowledge, then

we can turn the solver into a solution norm estimation algorithm with essentially the same query complexity up to logarithmic factors.

To start, we recall that the eigenstate filtering introduces an error at most

$$\left\| \mathbf{F}_{\text{sv}} \left(\frac{W - I}{2} \right) G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} - \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} \right\| \leq \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}x\| \frac{\delta}{2} + \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi. \quad (141)$$

Unlike before, we now aim to choose δ and ξ , so that when projected onto the first component, the error in the success amplitude is sufficiently small. Additionally, we perform an exponential search over the rescaling factor $\beta = 1, 2, 4, \dots$, so $\beta \geq 1$ is guaranteed. This means we can choose

$$\delta = \Theta \left(\frac{1}{\alpha_{A^{-1\dagger}|x}} \right), \quad \xi = \Theta(1), \quad (142)$$

where $\alpha_{A^{-1\dagger}|x} \geq \|A^{-1\dagger}|x\rangle\|$ is a known upper bound, so that

$$\begin{aligned} \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}|x\rangle\| \frac{\delta}{2} &\leq \frac{\sqrt{2}\beta \|x\|}{\|x\|^2 + \beta^2} \|A^{-1\dagger}|x\rangle\| \frac{\delta}{2} \leq \frac{1}{8} \frac{\beta \|x\|}{\|x\|^2 + \beta^2}, \\ \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi &\leq \xi \leq \frac{1}{16}. \end{aligned} \quad (143)$$

Then, the filtering has asymptotic query complexity

$$\mathbf{O} \left(\alpha_{A^{-1\dagger}|x} \right). \quad (144)$$

Note that the success amplitude is bounded within the interval

$$\left[\frac{7}{8} \frac{\beta \|x\|}{\|x\|^2 + \beta^2} - \frac{1}{16}, \frac{9}{8} \frac{\beta \|x\|}{\|x\|^2 + \beta^2} + \frac{1}{16} \right] \subseteq \left[\frac{7}{8} \frac{\beta \|x\|}{\|x\|^2 + \beta^2} - \frac{1}{16}, \frac{9}{8} \frac{\beta}{\|x\|} + \frac{1}{16} \right]. \quad (145)$$

Let us choose the rescaling factor $\beta_j = 2^j$ for $j = 0, 1, \dots, \mathbf{Ceil} \left(\log_2 \left(\alpha_{A^{-1\dagger}|x} \right) \right)$ and consider two cases. If $2^j \leq \frac{\|x\|}{16}$, then the success amplitude is at most

$$\frac{9}{8} \frac{2^j}{\|x\|} + \frac{1}{16} \leq \frac{9}{8} \frac{\|x\|}{16} + \frac{1}{16} = \frac{17}{128}. \quad (146)$$

On the other hand, if $\frac{\|x\|}{2} \leq 2^j \leq \|x\|$, then the success amplitude is at least

$$\frac{7}{8} \frac{2^j \|x\|}{\|x\|^2 + 4^j} - \frac{1}{16} \geq \frac{7}{8} \frac{\frac{\|x\|^2}{2}}{\|x\|^2 + \|x\|^2} - \frac{1}{16} = \frac{5}{32} = \frac{20}{128}. \quad (147)$$

To distinguish these two cases, we loop through $j = 0, 1, \dots, \mathbf{Ceil} \left(\log_2 \left(\alpha_{A^{-1\dagger}|x} \right) \right)$, and perform amplitude estimation with a sufficiently small constant precision and failure probability $p_j = \frac{1}{6 \cdot 2^j}$. Then with probability at least $\frac{2}{3}$, the algorithm produces estimates of amplitudes at most $\frac{18}{128}$ for all j such that $2^j \leq \frac{\|x\|}{16}$. On the other hand, if $\frac{\|x\|}{2} \leq 2^j \leq \|x\|$, the algorithm

outputs an estimate at least $\frac{19}{128}$ with probability at least $1 - \frac{1}{6 \cdot 2^j}$. Suppose the first time we find the amplitude above $\frac{19}{128}$ is at j_0 . This means

$$2^{j_0-1} < \frac{\|x\|}{2}, \quad 2^{j_0} > \frac{\|x\|}{16}, \quad (148)$$

which gives a constant multiplicative approximation of solution norm:

$$\frac{1}{4} < \frac{\|x\|}{2^{j_0} \cdot 4} < 4. \quad (149)$$

The total query complexity is then at most

$$\mathbf{O} \left(\sum_{j=0}^{\text{Ceil}(\log_2(\alpha_{A^{-1\dagger}|x\rangle}))} \alpha_{A^{-1\dagger}|x\rangle} \log(6 \cdot 2^j) \right) = \mathbf{O} \left(\alpha_{A^{-1\dagger}|x\rangle} \log^2(\alpha_{A^{-1\dagger}|x\rangle}) \right). \quad (150)$$

The approximation ratio can be refined to μ arbitrarily close to 1 using amplitude estimation incurring multiplicative complexity overhead $1/(\mu - 1)$ [15]. We now summarize the filtering-based beyond- κ solver and the corresponding solution norm estimator in the following theorem.

Theorem 4 (Beyond- κ solver based on filtering with effective gap). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$.*

Then, the normalized solution state $|x\rangle = \frac{x}{\|x\|}$ can be prepared to accuracy $\epsilon > 0$ and success probability $> \frac{1}{2}$ with query complexity

$$\mathbf{O} \left(\frac{\alpha_{A^{-1\dagger}|x\rangle}}{\epsilon} \log \left(\frac{1}{\epsilon} \right) \mathbf{Cost}(O_A) + \frac{\alpha_{A^{-1\dagger}|x\rangle}}{\epsilon} \log \left(\frac{1}{\epsilon} \right) \mathbf{Cost}(O_b) \right), \quad (151)$$

assuming a norm upper bound $\alpha_{A^{-1\dagger}|x\rangle} \geq \|A^{-1\dagger}|x\rangle\|$ and a constant multiplicative approximation α_x of the solution norm, i.e., $\frac{1}{\mu} \leq \frac{\alpha_x}{\|x\|} \leq \mu$ for some constant $\mu \geq 1$.

In the limit where $\mu \rightarrow 1$, $\epsilon \rightarrow 0$, $\alpha_x \rightarrow \infty$, $\alpha_{A^{-1\dagger}|x\rangle} \rightarrow \infty$, the solver has leading-order query complexity

$$\frac{6\alpha_{A^{-1\dagger}|x\rangle}}{\epsilon} \ln \left(\frac{1}{\epsilon} \right). \quad (152)$$

Otherwise, the solution norm $\|x\|$ can be estimated to a constant multiplicative accuracy with query complexity

$$\mathbf{O} \left(\alpha_{A^{-1\dagger}|x\rangle} \log^2(\alpha_{A^{-1\dagger}|x\rangle}) \mathbf{Cost}(O_A) + \alpha_{A^{-1\dagger}|x\rangle} \log^2(\alpha_{A^{-1\dagger}|x\rangle}) \mathbf{Cost}(O_b) \right). \quad (153)$$

5 Discussion

In this work, we have developed algorithms that solve the quantum linear system problem with query complexity independent of the spectral condition number. We have proposed an algorithmic template based on effective truncation of the linear system, which converts conventional quantum linear system solvers into beyond- κ solvers. Our template applies to any solver satisfying the strong truncation property, or the weak truncation property with optimal query complexity for initial state preparation. Additionally, we have presented a beyond- κ solver using eigenstate filtering with

effective gap. We have shown that the filtering-based solver is particularly simple and achieves a favorable runtime prefactor when the solution norm is known. When it is not, we have designed a similarly straightforward solution norm estimator—making our result fully self-contained.

We have introduced the effective condition number κ_{eff} as a measure of complexity for the truncation-based solvers, and established a family of upper bounds $\kappa_{\text{eff}} \leq \left(\frac{\|(A)_{\text{sv}}^{-t}|x\rangle\|}{\epsilon} \right)^{\frac{1}{t}}$ in terms of the singular transformation of the coefficient matrix A , normalized solution vector $|x\rangle$ and adjustable parameter $0 < t < \infty$. The case where t is a positive integer is especially straightforward, as the transformation reduces to $(A)_{\text{sv}}^{-t} = (A^\dagger A)^{-t/2}$ for even t and $(A)_{\text{sv}}^{-t} = A^{-1\dagger}(A^\dagger A)^{-(t-1)/2}$ for odd t and can be analyzed via standard linear algebraic techniques. In particular, choosing $t = 1$ gives $(A)_{\text{sv}}^{-t} = A^{-1\dagger}$. This choice introduces a polynomial dependence on $1/\epsilon$, which can be improved to a polylogarithmic scaling by setting $t = \Theta\left(\frac{\log(\frac{1}{\epsilon})}{\log\log(\frac{1}{\epsilon})}\right)$. Finding efficiently computable bounds on such transformations or alternative tight estimates of κ_{eff} remains an interesting open problem. Our effective condition number κ_{eff} captures both the strength and the limitation of truncation-based solvers. Its reciprocal is the smallest cutoff for which the truncated linear system can still be solved to accuracy ϵ . However, it remains plausible that solvers based on different techniques could achieve scaling below κ_{eff} . An interesting direction is to either construct such solvers or prove a lower bound ruling out their existence. Along similar lines, the filtering-based solver is much simpler but only achieves a scaling with $\frac{\|A^{-1\dagger}|x\rangle\|}{\epsilon} \geq \kappa_{\text{eff}}$. Whether this gap can be closed by a simple algorithm is an interesting question for future work.

In analyzing our solvers, we have focused primarily on the standard input model, where the coefficient matrix A and initial state $|b\rangle$ are accessed through separate oracles. However, we have also defined the affine dilation model, which block encodes A and $|b\rangle$ jointly, allowing further improvements in query complexity. To demonstrate that this improvement can be substantial, we now give a toy example where the separation between the two models diverges as $\kappa \rightarrow \infty$. Specifically, consider $A = \begin{bmatrix} \frac{1}{\kappa} & 0 \\ 0 & 1 \end{bmatrix}$, $|b\rangle = \begin{bmatrix} \frac{1}{\sqrt{1-\frac{1}{\kappa^2}}} \\ 1 \end{bmatrix}$, $x = A^{-1}|b\rangle = \begin{bmatrix} \frac{1}{\sqrt{1-\frac{1}{\kappa^2}}} \\ 1 \end{bmatrix}$, $A^{-1\dagger}x = \begin{bmatrix} \frac{\kappa}{\sqrt{1-\frac{1}{\kappa^2}}} \\ 1 \end{bmatrix}$. In the standard model, we can block encode A with normalization 1 and prepare the normalized state $|b\rangle$. Then the quantity $\|A^{-1\dagger}|x\rangle\|$ which determines the runtime of beyond- κ solvers scales like

$$\|A^{-1\dagger}|x\rangle\| = \Theta\left(\frac{\sqrt{\kappa^2 + 1 - \frac{1}{\kappa^2}}}{\sqrt{1 + 1 - \frac{1}{\kappa^2}}}\right) = \Theta(\kappa). \quad (154)$$

In contrast, we can define $\tilde{A} = \frac{1}{\tilde{\alpha}} \begin{bmatrix} MA & -M|b\rangle \\ 0 & 1 \end{bmatrix}$, $|\tilde{b}\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, $\tilde{x} = \tilde{A}^{-1}|\tilde{b}\rangle = \tilde{\alpha} \begin{bmatrix} x \\ 1 \end{bmatrix}$, and $\tilde{A}^{-1\dagger}\tilde{x} = \tilde{\alpha}^2 \begin{bmatrix} M^{-1\dagger}A^{-1\dagger}x \\ \|x\|^2 + 1 \end{bmatrix}$ in the affine dilation model, where $M = \begin{bmatrix} \kappa & 0 \\ 0 & 1 \end{bmatrix}$ is a diagonal matrix and $\tilde{\alpha} = \Theta(1)$ is the normalization factor for block encoding, giving

$$\|\tilde{A}^{-1\dagger}|\tilde{x}\rangle\| = \tilde{\alpha} \frac{\sqrt{\|M^{-1\dagger}A^{-1\dagger}x\|^2 + (\|x\|^2 + 1)^2}}{\sqrt{\|x\|^2 + 1}} = \Theta(1). \quad (155)$$

This shows that the affine dilation model can offer an arbitrarily large advantage as the condition number κ grows.

The preceding example is in fact a special case of a broader strategy known as preconditioning, which solves the modified linear system $MAx = M|b\rangle$ within the affine dilation model. More

generally, our framework opens a new regime for designing preconditioners for quantum linear system solvers: rather than targeting the condition number $\kappa(A)$ of matrix A , the purpose of this new preconditioning strategy is to redistribute how $|b\rangle$ aligns with the left singular vectors of A , thereby reducing the effective condition number. This can yield substantial complexity improvements even when $\kappa(MA)$ itself remains large, since the gains are captured entirely by the effective condition number.

For sparse linear systems, [24] constructs a preconditioner using diagonal matrices to balance the row norms. Interestingly, the instance-aware parameter $\mathbf{ET} = \sum_i p_i^2 d_i$ introduced in [24] can be recovered within our framework through diagonal preconditioning. Here $p_i = |\langle i|(AA^\dagger)^{-1}|b\rangle|$ is the i -th component of vector $(AA^\dagger)^{-1}|b\rangle$ in absolute value and $d_i = \|a_i\|^2 + |b_i|^2$ is the squared norm of the i -th row of the augmented matrix $[A \quad -|b\rangle]$. The sparse access model normalizes each row by its norm, implicitly block encoding the reweighted system $[MA \quad -M|b\rangle]$ with diagonal matrix $M = \mathbf{Diag}([\dots, 1/\sqrt{d_i}, \dots])$. A direct calculation then gives $(MA)^{-1}x = M^{-1}(AA^\dagger)^{-1}|b\rangle$, so that $\|(MA)^{-1}x\|^2 = \sum_i p_i^2 d_i = \mathbf{ET}$ recovers the result of [24]. Moreover, by designing new input models, one can unlock preconditioning techniques beyond those available in the existing framework. To illustrate, consider a setting where $\| |b\rangle\langle b|A \|$ and $\| (I - |b\rangle\langle b|)A \|$ are substantially different. In this case, one can introduce an input model that block encodes $|b\rangle\langle b|A$ and $(I - |b\rangle\langle b|)A$ separately, each with its own normalization factor. This effectively preconditions the linear system with respect to the orthogonal decomposition $\mathbf{Im}(|b\rangle\langle b|) \oplus \mathbf{Ker}(|b\rangle\langle b|)$. We expect that a systematic exploration of this preconditioning regime, identifying which classical preconditioners can be lifted into suitable input models and what complexity gains they unlock, is a fruitful direction for future work. As another example, [35] showed that for positive-definite matrices, the worst-case runtime lower bound of a quantum linear system solver remains linear in κ , matching the indefinite case. Nevertheless, they identified broad classes of positive-definite systems where this lower bound can be circumvented, achieving a quadratic improvement to $\sqrt{\kappa}$ via more efficient block-encodings or Cholesky-type preconditioners. It would be interesting to investigate whether our beyond- κ techniques can be combined with such structural assumptions to yield further speedups.

Setting aside these technical refinements, we believe the most impactful direction for future work is to identify practical applications where the condition number is prohibitively large for conventional solvers, yet beyond- κ solvers remain efficient. Our filtering-based solver is particularly appealing in this regard—its simplicity makes it a prime candidate for compilation into explicit quantum circuits. A concrete resource estimate for running this algorithm on a fault-tolerant quantum computer, even for a small-scale instance, would be an important step toward demonstrating practical quantum advantage for solving linear systems.

Acknowledgements

A.M.D. and Y.S. thank Fernando Brandão, Oskar Painter, James Hamilton, Nafea Bshara, Peter DeSantis, and Andy Jassy for their involvement and support of the research activities at the AWS Center for Quantum Computing. J.L. was supported by NSF Award FET-2243659 and NSF Career Award FET-2339116, Welch Foundation Grant no. A22-0307, a Microsoft Research Award, an Amazon Research Award, and a Ken Kennedy Research Cluster award (Rice K2I), in part through seed funding from the Ken Kennedy Institute at Rice University, the Rice CS Department, and the George R. Brown School of Engineering and Computing, by Rice University and the Department of Computer Science at Rice University, by the Ken Kennedy Institute and Rice Quantum Initiative, which is part of the Smalley-Curl Institute.

A Analysis of Li's beyond- κ solver in the standard input model

The beyond- κ solver of Li [24] was originally proposed for solving sparse linear systems. In this appendix, we show that it can be extended to work in the standard input model, where the coefficient matrix is accessed through a block encoding oracle and the initial vector is prepared by a unitary. To simplify the discussion, we use the precise values of vector norms such as $\|x\|$ and $\|A^{-1\dagger}x\|$, rather than their known upper and lower bounds.

Specifically, suppose that A is a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and define the solution vector by $x = A^{-1}|b\rangle$. Then we know from Corollary 9 that, for any scalar $\beta > 0$, the augmented matrix

$$\frac{1}{\sqrt{\beta^2 + 1}} [\beta A \quad -|b\rangle] = G_1^\dagger G_0 \quad (156)$$

can be block encoded in terms of an overlap of isometries using 1 query to O_A and O_b . Moreover, this induces a constrained orthogonal decomposition

$$\begin{aligned} G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{\beta}{\|x\|^2 + \beta^2} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix} - \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \Pi_0 \Pi_1 G_1 A^{-1\dagger} x \\ &\in (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus \mathbf{Im}(\Pi_0 \Pi_1), \end{aligned} \quad (157)$$

where $\Pi_0 = G_0 G_0^\dagger$ and $\Pi_1 = G_1 G_1^\dagger$.

Here, the goal is to preserve the first term that encodes the solution vector x while suppressing the second term using the effective gap lemma. In the main text, we show that this can be achieved with QSVT-based eigenstate filtering over the unit circle. In [24], this was instead realized using quantum phase estimation. More concretely, [24, Lemma 3.4] shows that applying quantum phase estimation produces the normalized state $\frac{1}{\sqrt{\|x\|^2 + \beta^2}} G_0 \begin{bmatrix} x \\ \beta \end{bmatrix}$ with accuracy ξ in trace distance, using

$$\mathbf{O} \left(\frac{\|x\|^2 + \beta^2}{\beta^2 \xi^2} \frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}x\| \right) = \mathbf{O} \left(\frac{1}{\beta \xi^2} \|A^{-1\dagger}x\| \right) \quad (158)$$

queries to the oracles O_A and O_b and succeeding with probability

$$\Theta \left(\frac{\beta^2}{\|x\|^2 + \beta^2} \right), \quad (159)$$

assuming $\beta = \Omega(1)$.

To ensure that the normalized solution $|x\rangle = \frac{x}{\|x\|}$ is produced with accuracy ϵ , Li chose $\xi = \Theta \left(\frac{\|x\|^2 \epsilon}{\|x\|^2 + \beta^2} \right)$. Then, he boosted the success probability close to unity by repeating the quantum phase estimation

$$\Theta \left(\frac{\|x\|^2 + \beta^2}{\beta^2} \frac{\|x\|^2 + \beta^2}{\|x\|^2} \right) = \Theta \left(\frac{\|x\|^2}{\beta^2} + \frac{\beta^2}{\|x\|^2} \right) \quad (160)$$

times. This leads to a total query complexity of

$$\mathbf{O} \left(\left(\frac{\|x\|^2}{\beta^2} + \frac{\beta^2}{\|x\|^2} \right) \frac{(\|x\|^2 + \beta^2)^2}{\|x\|^4 \epsilon^2} \frac{\|A^{-1\dagger}x\|}{\beta} \right) = \mathbf{O} \left(\left(\frac{\|x\|^2}{\beta^2} + \frac{\beta^6}{\|x\|^6} \right) \frac{\|A^{-1\dagger}x\|}{\beta \epsilon^2} \right). \quad (161)$$

When the solution norm $\|x\|$ is known, Li set $\beta = \Theta(\|x\|)$, simplifying the query cost to

$$\mathbf{O}\left(\frac{\|A^{-1\dagger}x\|}{\|x\| \epsilon^2}\right). \quad (162)$$

Otherwise, he chose $\beta = 1$. Under the convention that $\|x\| \geq 1$, the query complexity then becomes

$$\mathbf{O}\left(\frac{\|A^{-1\dagger}x\| \|x\|^2}{\epsilon^2}\right). \quad (163)$$

B Analysis of beyond- κ solvers in the affine dilation model

To facilitate a direct comparison with prior quantum linear system solvers, we have focused on the standard input model when presenting our beyond- κ solvers in the main text. In this appendix, we analyze beyond- κ solvers in the affine dilation model, which allows further refinement of the query complexity.

B.1 Analysis of truncation-based solvers in the affine dilation model

In the affine dilation model, the input to a quantum linear system is accessed through the block encoding

$$\tilde{A} = \frac{1}{\tilde{\alpha}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix}. \quad (164)$$

whereas the dilated initial vector is $|\tilde{b}\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$. Here, we assume that $|b\rangle$ is normalized after a potential rescaling of the block encoding. For truncation-based solvers, we take $c > 0$ so that the dilated matrix is invertible. The normalization factor $\tilde{\alpha}$ then satisfies $\tilde{\alpha} \geq \max\{\|A\|, c, 1\}$.

After matrix inversion, we have

$$\tilde{A}^{-1} = \tilde{\alpha} \begin{bmatrix} A^{-1} & \frac{1}{c}x \\ 0 & \frac{1}{c} \end{bmatrix}, \quad (165)$$

which encodes the dilated solution vector

$$\tilde{x} = \tilde{A}^{-1}|\tilde{b}\rangle = \frac{\tilde{\alpha}}{c} \begin{bmatrix} x \\ 1 \end{bmatrix}. \quad (166)$$

The performance of the truncation-based solvers is still determined by the effective condition number from [Theorem 1](#). The only difference is that κ_{eff} is now defined in terms of the dilated \tilde{A} and $|\tilde{b}\rangle$, and we denote it $\tilde{\kappa}_{\text{eff}}$.

It is generally challenging to give a closed-form formula for $\tilde{\kappa}_{\text{eff}}$ analogous to the one in [Theorem 1](#), since the singular value decomposition of the dilated matrix \tilde{A} does not relate simply to that of A . However, the effective condition number admits the succinct upper bound $\tilde{\kappa}_{\text{eff}} \leq \frac{\|\tilde{A}^{-1\dagger}\tilde{x}\|}{\|\tilde{x}\| \epsilon}$ same as in the standard input model, where $\|\tilde{A}^{-1\dagger}\tilde{x}\|$ can be evaluated as

$$\|\tilde{A}^{-1\dagger}\tilde{x}\| = \tilde{\alpha}^2 \left\| \begin{bmatrix} \frac{A^{-1\dagger}x}{c} \\ \frac{\|x\|^2+1}{c^2} \end{bmatrix} \right\| = \tilde{\alpha}^2 \sqrt{\frac{\|A^{-1\dagger}x\|^2}{c^2} + \frac{(\|x\|^2+1)^2}{c^4}}. \quad (167)$$

This implies that the effective condition number for the dilated system is at most

$$\tilde{\kappa}_{\text{eff}} \leq \frac{\tilde{\alpha}}{\epsilon} \frac{1}{\sqrt{\|x\|^2 + 1}} \sqrt{\|A^{-1\dagger}x\|^2 + \frac{(\|x\|^2 + 1)^2}{c^2}}. \quad (168)$$

B.2 Analysis of filtering-based solver in the affine dilation model

For the filtering-based solver, we instead consider the augmented matrix

$$\tilde{A} = \frac{1}{\alpha} [A \quad -|b\rangle]. \quad (169)$$

This corresponds to setting $c = 0$ in the generic affine dilation model and removing the redundant zero blocks at the bottom.

Suppose that this is block encoded by an overlap of isometries as $\tilde{A} = \frac{1}{\alpha} [A \quad -|b\rangle] = G_1^\dagger G_0$. Then by an analysis similar to that in the main text, we have the constrained orthogonal decomposition

$$\begin{aligned} G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} &= \frac{1}{\|x\|^2 + 1} G_0 \begin{bmatrix} x \\ 1 \end{bmatrix} - \frac{\alpha}{\|x\|^2 + 1} \Pi_0 \Pi_1 G_1 A^{-1\dagger} x \\ &\in (\mathbf{Im}(\Pi_0) \cap \mathbf{Ker}(\Pi_1)) \oplus \mathbf{Im}(\Pi_0 \Pi_1), \end{aligned} \quad (170)$$

where $\Pi_0 = G_0 G_0^\dagger$ and $\Pi_1 = G_1 G_1^\dagger$. Performing filtering $\mathbf{F}_{\text{sv}} \left(\frac{W-I}{2} \right)$ with $W = (2\Pi_0 - I)(I - 2\Pi_1)$ then gives

$$\mathbf{F}_{\text{sv}} \left(\frac{W-I}{2} \right) G_0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \frac{1}{\|x\|^2 + 1} G_0 \begin{bmatrix} x \\ 1 \end{bmatrix} - \frac{\alpha}{\|x\|^2 + 1} \mathbf{F}_{\text{sv}} \left(\frac{W-I}{2} \right) \Pi_0 \Pi_1 G_1 A^{-1\dagger} x, \quad (171)$$

where

$$\left\| \mathbf{F}_{\text{sv}} \left(\frac{W-I}{2} \right) \Pi_0 \Pi_1 G_1 A^{-1\dagger} x \right\| \leq \frac{\delta}{2} \|A^{-1\dagger}x\|. \quad (172)$$

Assuming that the solution norm $\|x\| = \|A^{-1}|b\rangle\|$ is estimated to a constant multiplicative accuracy, we set

$$\delta = \Theta \left(\frac{\|x\| \epsilon}{\alpha \|A^{-1\dagger}x\|} \right). \quad (173)$$

Note that $\delta = \mathbf{O}(1)$, since $\epsilon = \mathbf{O}(1)$ and

$$\frac{\alpha \|A^{-1\dagger}x\|}{\|x\|} = \alpha \|A^{-1\dagger}|x\rangle\| \geq \|A\| \|A^{-1\dagger}|x\rangle\| = \|A^\dagger\| \|A^{-1\dagger}|x\rangle\| \geq \|A^\dagger A^{-1\dagger}|x\rangle\| = 1. \quad (174)$$

Finally, we perform $\mathbf{O} \left(\frac{\|x\|^2 + 1}{\|x\|} \right)$ steps of amplitude amplification, so the total query complexity is determined by

$$\frac{\alpha \|x\|^2 + 1}{\epsilon \|x\|^2} \|A^{-1\dagger}x\| \quad (175)$$

up to logarithmic factors.

B.3 Complexity comparison

We now compare the runtimes of the truncation-based and filtering-based solvers in the affine dilation model. In particular, we show that the query complexity of the former is no worse than that of the latter up to logarithmic factors, as in the standard input model.

To this end, suppose we start with a block encoding of the augmented matrix $\frac{1}{\alpha} [A \ -|b\rangle] = G_1^\dagger G_0$. Then, we can construct a block encoding of the affine dilation

$$\tilde{A} = \frac{1}{\sqrt{\alpha^2 + c^2}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix} \quad (176)$$

with normalization factor $\sqrt{\alpha^2 + c^2}$. Indeed, this can be achieved through the following factorization of the dilated matrix as an overlap of isometries

$$\frac{1}{\sqrt{\alpha^2 + c^2}} \begin{bmatrix} A & -|b\rangle \\ 0 & c \end{bmatrix} = \begin{bmatrix} G_1^\dagger & 0 \\ 0 & [1 \ 0 \ 0] \end{bmatrix} \frac{1}{\sqrt{\alpha^2 + c^2}} \begin{bmatrix} \alpha G_0 \\ c \begin{bmatrix} 0 & 1 \\ I & 0 \\ 0 & 0 \end{bmatrix} \end{bmatrix}. \quad (177)$$

Then the effective condition number is bounded by

$$\tilde{\kappa}_{\text{eff}} \leq \frac{\sqrt{\alpha^2 + c^2}}{\epsilon \sqrt{\|x\|^2 + 1}} \sqrt{\|A^{-1\dagger}x\|^2 + \frac{(\|x\|^2 + 1)^2}{c^2}} \leq \frac{\sqrt{\alpha^2 + c^2}}{\epsilon} \frac{\|A^{-1\dagger}x\|}{\sqrt{\|x\|^2 + 1}} + \frac{\sqrt{\alpha^2 + c^2}}{\epsilon} \frac{\sqrt{\|x\|^2 + 1}}{c}. \quad (178)$$

Let us choose

$$c = \alpha \frac{\|x\|^2 + 1}{\|x\|} \geq 2\alpha, \quad (179)$$

so that

$$\sqrt{\alpha^2 + c^2} = \mathbf{O} \left(\alpha \frac{\|x\|^2 + 1}{\|x\|} \right). \quad (180)$$

Meanwhile, by our block encoding convention,

$$1 = \||b\rangle\| = \|Ax\| \leq \alpha \|x\| \leq \alpha^2 \|A^{-1\dagger}x\|, \quad (181)$$

which gives

$$c = \frac{\|x\|^2 + 1}{\frac{\|x\|}{\alpha}} \geq \frac{\|x\|^2 + 1}{\|A^{-1\dagger}x\|}. \quad (182)$$

Altogether,

$$\tilde{\kappa}_{\text{eff}} = \mathbf{O} \left(\frac{\alpha \|x\|^2 + 1}{\epsilon \|x\|} \frac{\|A^{-1\dagger}x\|}{\sqrt{\|x\|^2 + 1}} \right) = \mathbf{O} \left(\frac{\alpha \|x\|^2 + 1}{\epsilon \|x\|^2} \|A^{-1\dagger}x\| \right). \quad (183)$$

Hence, with this specific choice of c in the affine dilation model, the runtime of the truncation-based algorithm matches that of the filtering-based method up to logarithmic factors.

C Generalized truncation property

In the main text, we have defined the strong and weak truncation property (Definition 2 and Definition 3) for a quantum linear system solver. Informally, these capture the requirement that if one supplies α as the condition number parameter, the solver correctly inverts singular value components above $1/\alpha$ and discards those below. To obtain a beyond- κ solver, we have required that the truncation properties hold when α is at least the effective condition number from Theorem 1. More generally, we can define the strong truncation property for an arbitrary α as follows.

Definition 4 (Generalized strong truncation property). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$. Suppose that $U(O_A, O_b, \kappa, \epsilon)$ is a quantum linear system solver such that with success probability $> \frac{1}{2}$,*

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00\rangle}{\|\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (184)$$

provided that $\kappa \geq \|A^{-1}\|$ and $\epsilon > 0$.

We say that the solver satisfies the generalized strong truncation property if with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha, \epsilon) |00\rangle}{\|\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha, \epsilon) |00\rangle\|} - \frac{\Pi_{right, [\alpha^{-1}, 1]} x}{\|\Pi_{right, [\alpha^{-1}, 1]} x\|} \right\| = \mathbf{O}(\epsilon), \quad (185)$$

provided that $1 \leq \alpha \leq \kappa$, where $\Pi_{right, [\alpha^{-1}, 1]}$ is the right singular vector projection of A associated with singular values from $[\alpha^{-1}, 1]$.

When configured with α , VTAA correctly inverts all singular values above $1/\alpha$. Combining this with a standard amplitude amplification yields the truncated solution state supported on singular values above $1/\alpha$. We thus conclude that VTAA satisfies the generalized strong truncation property with a query overhead from amplification. Similarly, one can show that VTAA satisfies the following generalized version of weak truncation property.

Definition 5 (Generalized weak truncation property). *Let A be a matrix with $\|A\| \leq 1$ block encoded by O_A , and $|b\rangle$ be a normalized quantum state prepared by O_b . Assume that A is invertible and denote $x = A^{-1}|b\rangle$ and $|x\rangle = \frac{x}{\|x\|}$. Suppose that $U(O_A, O_b, \kappa, \epsilon)$ is a quantum linear system solver such that with success probability $> \frac{1}{2}$,*

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00\rangle}{\|\langle\langle 0| \otimes I \rangle U(O_A, O_b, \kappa, \epsilon) |00\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (186)$$

provided that $\kappa \geq \|A^{-1}\|$ and $\epsilon > 0$.

We say that the solver satisfies the generalized weak truncation property if with success probability $> \frac{1}{2}$,

$$\left\| \frac{\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha, \epsilon) |00\rangle}{\|\langle\langle 0| \otimes I \rangle U(O_A, O_b, \alpha, \epsilon) |00\rangle\|} - |x\rangle \right\| = \mathbf{O}(\epsilon), \quad (187)$$

provided that $1 \leq \alpha \leq \kappa$, and $|b\rangle \in \mathbf{Im}(\Pi_{left, [\alpha^{-1}, 1]})$ where $\Pi_{left, [\alpha^{-1}, 1]}$ is the left singular vector projection of A associated with singular values from $[\alpha^{-1}, 1]$.




Figure 5: Qualitative behavior of $g(\gamma)$ for constant-prefactor optimization for the filtering-based solver.

D Constant-prefactor optimization for filtering with effective gap

In this appendix, we optimize the constant prefactor of the query complexity for eigenstate filtering with effective gap.

Recall from the main text (Section 4.4) that, in order to prepare the normalized solution state with accuracy ϵ , we choose the effective gap δ and filtering accuracy ξ so that

$$\frac{\sqrt{\beta^2 + 1}}{\|x\|^2 + \beta^2} \|A^{-1\dagger}x\| \frac{\delta}{2} + \frac{\|x\|}{\sqrt{\|x\|^2 + \beta^2}} \xi \leq \frac{\beta \|x\|}{\|x\|^2 + \beta^2} \epsilon. \quad (188)$$

Correspondingly, the query complexity of filtering is

$$\mathbf{Ceil} \left(\frac{4}{\delta} \ln \left(\frac{2}{\xi} \right) \right). \quad (189)$$

A naive strategy is to divide the error budget into equal halves. Here, we propose an alternative approach whose leading-order query complexity has a smaller constant prefactor.

To this end, we consider the function

$$g(\gamma) = \frac{1}{\gamma\epsilon} \ln \left(\frac{1}{(1-\gamma)\epsilon} \right), \quad (190)$$

where ϵ is a fixed positive number and $0 < \gamma < 1$ is a tunable parameter. See Figure 5 for an illustration of its qualitative behavior. This function corresponds to the query cost when we allocate a γ -fraction of the error budget to the effective gap and the remaining fraction to the filtering accuracy parameter. Intuitively, since $\ln(\cdot)$ grows much slower than linearly, we expect the minimum of $g(\gamma)$ to occur at some $\gamma = \gamma_0$ close to 1.

To rigorize this intuition, we differentiate $g(\gamma)$ to get

$$\frac{d}{d\gamma} g(\gamma) = -\frac{(\gamma - 1) \ln \left(\frac{1}{\epsilon(1-\gamma)} \right) + \gamma}{\epsilon(\gamma - 1)\gamma^2}. \quad (191)$$

Thus $g(\gamma)$ is minimized at $\gamma = \gamma_0$ where $\frac{d}{d\gamma}g(\gamma_0) = 0 \Leftrightarrow (\gamma_0 - 1) \ln \left(\frac{1}{\epsilon(1-\gamma_0)} \right) + \gamma_0 = 0$. Defining $v_0 = \gamma_0 - 1$ with $-1 < v_0 < 0$, we have $v_0 \ln \left(\frac{1}{-\epsilon v_0} \right) + v_0 + 1 = 0 \Leftrightarrow \frac{1}{v_0} e^{\frac{1}{v_0}} = -\frac{\epsilon}{e}$. Finally, letting $\zeta_0 = \frac{1}{v_0}$ with $-\infty < \zeta_0 < -1$, we get $\zeta_0 e^{\zeta_0} = -\frac{\epsilon}{e} \Leftrightarrow \zeta_0 = \mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)$, for $\mathbf{W}_{-1}(\cdot)$ the (-1) -branch of the Lambert-W function. This implies

$$\gamma_0 = 1 + \frac{1}{\mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)}. \quad (192)$$

Now let us bound $g(\gamma)$ at $\gamma = \gamma_0$. To this end, we rewrite $\mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right) = \mathbf{W}_{-1} \left(-e^{\ln \left(\frac{\epsilon}{e} \right)} \right) = \mathbf{W}_{-1} \left(-e^{-\ln \left(\frac{1}{\epsilon} \right) - 1} \right)$. Applying bounds on the (-1) -branch of the Lambert-W function [10], we have

$$-1 - \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) - \ln \left(\frac{1}{\epsilon} \right)} < \mathbf{W}_{-1} \left(-e^{-\ln \left(\frac{1}{\epsilon} \right) - 1} \right) < -1 - \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) - \frac{2}{3} \ln \left(\frac{1}{\epsilon} \right)}. \quad (193)$$

This implies

$$g(\gamma_0) < \left(\frac{1 + \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \frac{2}{3} \ln \left(\frac{1}{\epsilon} \right)}}{\sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \frac{2}{3} \ln \left(\frac{1}{\epsilon} \right)}} \right) \frac{1}{\epsilon} \ln \left(\frac{1 + \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \ln \left(\frac{1}{\epsilon} \right)}}{\epsilon} \right), \quad (194)$$

which reduces to

$$\frac{1}{\epsilon} \ln \left(\frac{1}{\epsilon} \right) \quad (195)$$

when higher order terms are dropped.

Proposition 13. *For any $\epsilon > 0$, the function*

$$g(\gamma) = \frac{1}{\gamma \epsilon} \ln \left(\frac{1}{(1-\gamma)\epsilon} \right), \quad 0 < \gamma < 1 \quad (196)$$

is minimized at

$$\operatorname{argmin} g(\gamma) = \gamma_0 = 1 + \frac{1}{\mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)}, \quad (197)$$

where $\mathbf{W}_{-1}(\cdot)$ is the (-1) -branch of the Lambert-W function, with minimum

$$\begin{aligned} \min g(\gamma) = g(\gamma_0) &= \frac{\mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)}{1 + \mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)} \frac{1}{\epsilon} \ln \left(\frac{-\mathbf{W}_{-1} \left(-\frac{\epsilon}{e} \right)}{\epsilon} \right) \\ &< \left(\frac{1 + \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \frac{2}{3} \ln \left(\frac{1}{\epsilon} \right)}}{\sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \frac{2}{3} \ln \left(\frac{1}{\epsilon} \right)}} \right) \frac{1}{\epsilon} \ln \left(\frac{1 + \sqrt{2 \ln \left(\frac{1}{\epsilon} \right) + \ln \left(\frac{1}{\epsilon} \right)}}{\epsilon} \right). \end{aligned} \quad (198)$$

Moreover, the minimum value reduces to

$$\frac{1}{\epsilon} \ln \left(\frac{1}{\epsilon} \right) \quad (199)$$

when higher order terms are omitted.

References

- [1] Abhijeet Alase and Salini Karuvade, *Resolvent-based quantum phase estimation: Towards estimation of parametrized eigenvalues*, 2024, [arXiv:2410.04837](#), [DOI:10.48550/arXiv.2410.04837](#). (Cited on page 3).
- [2] Andris Ambainis, *Variable time amplitude amplification and quantum algorithms for linear algebra problems*, 29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012), vol. 14, pp. 636–647, 2012, [arXiv:1010.4458](#), [DOI:10.4230/LIPIcs.STACS.2012.636](#). (Cited on pages 3 and 5).
- [3] Andris Ambainis, Martins Kokainis, and Jevgēnijs Vihrovs, *Improved algorithm and lower bound for variable time quantum search*, 18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023), vol. 266, pp. 7:1–7:18, 2023, [arXiv:2302.06749](#), [DOI:10.4230/LIPIcs.TQC.2023.7](#). (Cited on page 27).
- [4] 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** (2022), 1–28, [arXiv:1909.05500](#), [DOI:10.1145/3498331](#). (Cited on page 5).
- [5] Dominic W. Berry and Pedro C. S. Costa, *Quantum algorithm for time-dependent differential equations using Dyson series*, Quantum **8** (2024), 1369, [arXiv:2212.03544](#), [DOI:10.22331/q-2024-06-13-1369](#). (Cited on page 3).
- [6] Dominic W. Berry, Andrew M. Childs, Aaron Ostrander, and Guoming Wang, *Quantum algorithm for linear differential equations with exponentially improved dependence on precision*, Communications in Mathematical Physics **356** (2017), 1057–1081, [arXiv:1701.03684](#), [DOI:10.1007/s00220-017-3002-y](#). (Cited on page 3).
- [7] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp, *Quantum amplitude amplification and estimation*, Contemporary Mathematics **305** (2002), 53–74, [arXiv:quant-ph/0005055](#), [DOI:10.1090/conm/305/05215](#). (Cited on page 27).
- [8] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery, *The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation*, 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019), vol. 132, pp. 33:1–33:14, 2019, [arXiv:1804.01973](#), [DOI:10.4230/LIPIcs.ICALP.2019.33](#). (Cited on pages 5 and 11).
- [9] Shantanav Chakraborty, Aditya Morolia, and Anurudh Peduri, *Quantum regularized least squares*, Quantum **7** (2023), 988, [arXiv:2206.13143](#), [DOI:10.22331/q-2023-04-27-988](#). (Cited on page 5).
- [10] Ioannis Chatzigeorgiou, *Bounds on the Lambert function and their application to the outage analysis of user cooperation*, IEEE Communications Letters **17** (2013), no. 8, 1505–1508, [arXiv:1601.04895](#), [DOI:10.1109/LCOMM.2013.070113.130972](#). (Cited on page 48).
- [11] 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** (2017), 1920–1950, [arXiv:1511.02306](#), [DOI:10.1137/16M1087072](#). (Cited on page 5).

- [12] Pedro C.S. Costa, Dong An, Ryan Babbush, and Dominic Berry, *The discrete adiabatic quantum linear system solver has lower constant factors than the randomized adiabatic solver*, Quantum **9** (2025), 1887, [arXiv:2312.07690](https://arxiv.org/abs/2312.07690), DOI:10.22331/q-2025-10-20-1887. (Cited on page 3).
- [13] 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** (2022), 040303, [arXiv:2111.08152](https://arxiv.org/abs/2111.08152), DOI:10.1103/PRXQuantum.3.040303. (Cited on pages 3, 4, 5, and 34).
- [14] Joseph Cunningham and Jérémie Roland, *Eigenpath traversal by Poisson-distributed phase randomisation*, 19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024), vol. 310, pp. 7:1–7:20, 2024, [arXiv:2406.03972](https://arxiv.org/abs/2406.03972), DOI:10.4230/LIPIcs.TQC.2024.7. (Cited on pages 3, 4, and 5).
- [15] Alexander M. Dalzell, *A shortcut to an optimal quantum linear system solver*, 2024, [arXiv:2406.12086](https://arxiv.org/abs/2406.12086), DOI:10.48550/arXiv.2406.12086. (Cited on pages 3, 4, 5, 10, 14, 32, 37, and 39).
- [16] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe, *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, pp. 193–204, 2019, [arXiv:1806.01838](https://arxiv.org/abs/1806.01838), DOI:10.1145/3313276.3316366. (Cited on pages 5, 7, 12, and 19).
- [17] Paul R Halmos, *Normal dilations and extensions of operators*, Summa Brasiliensis Math **2** (1950), 125–134. (Cited on page 11).
- [18] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd, *Quantum algorithm for linear systems of equations*, Physical Review Letters **103** (2009), 150502, [arXiv:0811.3171](https://arxiv.org/abs/0811.3171), DOI:10.1103/PhysRevLett.103.150502. (Cited on pages 3, 4, and 5).
- [19] Roger A. Horn and Charles R. Johnson, *Matrix analysis*, Cambridge University Press, 2012, DOI:10.1017/CBO9781139020411. (Cited on pages 11, 23, and 24).
- [20] Hari Krovi, *Improved quantum algorithms for linear and nonlinear differential equations*, Quantum **7** (2023), 913, [arXiv:2202.01054](https://arxiv.org/abs/2202.01054), DOI:10.22331/q-2023-02-02-913. (Cited on page 3).
- [21] Shuhei Kudo, *Complex Palais matrix and a new unitary transform with bounded component norms*, Special Matrices **12** (2024), no. 1, 20240026, DOI:doi:10.1515/spma-2024-0026. (Cited on page 23).
- [22] Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Špalek, and Mario Szegedy, *Quantum query complexity of state conversion*, 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, pp. 344–353, 2011, [arXiv:1011.3020](https://arxiv.org/abs/1011.3020), DOI:10.1109/FOCS.2011.75. (Cited on pages 10 and 30).
- [23] Guanzhong Li, Lvzhou Li, and Jingquan Luo, *Quantum phase discrimination with applications to quantum search on graphs*, 2025, [arXiv:2504.15194](https://arxiv.org/abs/2504.15194), DOI:10.48550/arXiv.2504.15194. (Cited on pages 10, 32, and 36).
- [24] Jianqiang Li, *A new quantum linear system algorithm beyond the condition number and its application to solving multivariate polynomial systems*, 2025, [arXiv:2510.05588](https://arxiv.org/abs/2510.05588), DOI:10.48550/arXiv.2510.05588. (Cited on pages 4, 5, 6, 10, 12, 13, 30, 41, and 42).

- [25] Lin Lin and Yu Tong, *Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems*, Quantum **4** (2020), 361, [arXiv:1910.14596](https://arxiv.org/abs/1910.14596), [DOI:10.22331/q-2020-11-11-361](https://doi.org/10.22331/q-2020-11-11-361). (Cited on page 5).
- [26] Guang Hao Low and Isaac L. Chuang, *Hamiltonian simulation by uniform spectral amplification*, 2017, [arXiv:1707.05391](https://arxiv.org/abs/1707.05391), [DOI:10.48550/arXiv.1707.05391](https://doi.org/10.48550/arXiv.1707.05391). (Cited on page 11).
- [27] Guang Hao Low and Isaac L. Chuang, *Optimal Hamiltonian simulation by quantum signal processing*, Physical Review Letters **118** (2017), 010501, [arXiv:1606.02685](https://arxiv.org/abs/1606.02685), [DOI:10.1103/PhysRevLett.118.010501](https://doi.org/10.1103/PhysRevLett.118.010501). (Cited on page 10).
- [28] Guang Hao Low and Yuan Su, *Quantum eigenvalue processing*, SIAM Journal on Computing **55** (2026), no. 1, 135–215, [arXiv:2401.06240](https://arxiv.org/abs/2401.06240), [DOI:10.1137/24M1689363](https://doi.org/10.1137/24M1689363). (Cited on page 3).
- [29] Guang Hao Low and Yuan Su, *Quantum linear system algorithm with optimal queries to initial state preparation*, Quantum **10** (2026), 2041, [arXiv:2410.18178](https://arxiv.org/abs/2410.18178), [DOI:10.22331/q-2026-03-23-2041](https://doi.org/10.22331/q-2026-03-23-2041). (Cited on pages 3, 4, 5, 23, 27, and 30).
- [30] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang, *Methodology of resonant equian-gular composite quantum gates*, Physical Review X **6** (2016), 041067, [arXiv:1603.03996](https://arxiv.org/abs/1603.03996), [DOI:10.1103/PhysRevX.6.041067](https://doi.org/10.1103/PhysRevX.6.041067). (Cited on page 10).
- [31] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang, *Grand uni-fication of quantum algorithms*, PRX Quantum **2** (2021), 040203, [arXiv:2105.02859](https://arxiv.org/abs/2105.02859), [DOI:10.1103/PRXQuantum.2.040203](https://doi.org/10.1103/PRXQuantum.2.040203). (Cited on page 32).
- [32] Hitomi Mori, Yuta Kikuchi, Marcello Benedetti, and Matthias Rosenkranz, *Sparsity-dependent complexity lower bound of quantum linear system solvers*, 2026, [arXiv:2601.16697](https://arxiv.org/abs/2601.16697), [DOI:10.48550/arXiv.2601.16697](https://doi.org/10.48550/arXiv.2601.16697). (Cited on page 3).
- [33] Danial Motlagh and Nathan Wiebe, *Generalized quantum signal processing*, PRX Quantum **5** (2024), 020368, [arXiv:2308.01501](https://arxiv.org/abs/2308.01501), [DOI:10.1103/PRXQuantum.5.020368](https://doi.org/10.1103/PRXQuantum.5.020368). (Cited on page 10).
- [34] Art of Problem Solving, *Power mean inequality*, 2024, https://artofproblemsolving.com/wiki/index.php/Power_Mean_Inequality, Last accessed on 2026-05-08. (Cited on page 18).
- [35] Davide Orsucci and Vedran Dunjko, *On solving classes of positive-definite quantum linear systems with quadratically improved runtime in the condition number*, Quantum **5** (2021), 573, [arXiv:2101.11868](https://arxiv.org/abs/2101.11868), [DOI:10.22331/q-2021-11-08-573](https://doi.org/10.22331/q-2021-11-08-573). (Cited on page 41).
- [36] J Michael Steele, *The Cauchy-Schwarz master class: an introduction to the art of mathematical inequalities*, Cambridge University Press, 2004, [DOI:10.1017/CBO9780511817106](https://doi.org/10.1017/CBO9780511817106). (Cited on page 18).
- [37] Yiğit Subaşı, Rolando D. Somma, and Davide Orsucci, *Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing*, Physical Review Letters **122** (2019), 060504, [arXiv:1805.10549](https://arxiv.org/abs/1805.10549), [DOI:10.1103/PhysRevLett.122.060504](https://doi.org/10.1103/PhysRevLett.122.060504). (Cited on page 5).
- [38] Souichi Takahira, Asuka Ohashi, Tomohiro Sogabe, and Tsuyoshi S. Usuda, *Quantum algo-rithm for matrix functions by Cauchy’s integral formula*, Quantum Information and Compu-tation **20** (2020), 14–36, [arXiv:2106.08075](https://arxiv.org/abs/2106.08075), [DOI:10.26421/QIC20.1-2-2](https://doi.org/10.26421/QIC20.1-2-2). (Cited on page 3).

- [39] Souichi Takahira, Asuka Ohashi, Tomohiro Sogabe, and Tsuyoshi Sasaki Usuda, *Quantum algorithms based on the block-encoding framework for matrix functions by contour integrals*, 2021, [arXiv:2106.08076](https://arxiv.org/abs/2106.08076), [DOI:10.48550/arXiv.2106.08076](https://doi.org/10.48550/arXiv.2106.08076). (Cited on page 3).
- [40] Yu Tong, Dong An, Nathan Wiebe, and Lin Lin, *Fast inversion, preconditioned quantum linear system solvers, fast Green's-function computation, and fast evaluation of matrix functions*, *Physical Review A* **104** (2021), 032422, [arXiv:2008.13295](https://arxiv.org/abs/2008.13295), [DOI:10.1103/PhysRevA.104.032422](https://doi.org/10.1103/PhysRevA.104.032422). (Cited on page 3).
- [41] Philipp Wacker, *Please, not another note about generalized inverses*, 2023, [arXiv:2306.06989](https://arxiv.org/abs/2306.06989), [DOI:10.48550/arXiv.2306.06989](https://doi.org/10.48550/arXiv.2306.06989). (Cited on pages 15, 16, and 17).
- [42] Yanqiao Wang and Jin-Peng Liu, *Sign embedding quantum algorithms for matrix equations and matrix functions*, 2026, [arXiv:2604.25333](https://arxiv.org/abs/2604.25333), [DOI:10.48550/arXiv.2604.25333](https://doi.org/10.48550/arXiv.2604.25333). (Cited on page 3).
- [43] Youle Wang, Lei Zhang, Zhan Yu, and Xin Wang, *Quantum phase processing and its applications in estimating phase and entropies*, *Physical Review A* **108** (2023), 062413, [arXiv:2209.14278](https://arxiv.org/abs/2209.14278), [DOI:10.1103/PhysRevA.108.062413](https://doi.org/10.1103/PhysRevA.108.062413). (Cited on page 10).