Research Square

Preprints are preliminary reports that have not undergone peer review. They should not be considered conclusive, used to inform clinical practice, or referenced by the media as validated information.

Quantum oblique eigenprojection

Yuan Su

buptsuyuan@gmail.com

Amazon https://orcid.org/0000-0003-1144-3563

Alexander Dalzell

AWS Center for Quantum Computing https://orcid.org/0000-0002-3756-8500

Article

Keywords:

Posted Date: October 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11354688/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: There is **NO** Competing Interest.

Quantum oblique eigenprojection

Alexander M. Dalzell and Yuan Su*

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

Abstract

Every square matrix decomposes its underlying Hilbert space into generalized, nonorthogonal eigensubspaces. We show that a quantum computer can perform such an oblique eigenprojection Π given block encoding access to the input matrix. Our approach has a query complexity nearly linear in the inverse gap and a normalization factor close to $\|\Pi\|$ under a spectral-set condition on the input. This covers common assumptions on the numerical range or diagonalizability and matches known results for orthogonal eigenprojections. We achieve this with a two-sided block preconditioning that uses a discrete Fourier transform of the matrix resolvent. As the oblique eigenprojection serves as a fundamental primitive in the design of quantum algorithms, our result can be used to extend and improve a variety of prior results. Specifically, we describe applications to: (i) preparing eigenstates of matrices with complex eigenvalues, extending the quantum eigenvalue transformation algorithm of Low and Su beyond real spectra; (ii) solving continuous-time algebraic Riccati equations, cubically speeding up a prior solver of Rodenas-Ruiz, Zhao, and Lee; and (iii) solving ordinary Sylvester equations, quadratically improving a direct augmented method of Wang and Liu. Our result suggests a promising route to applying nonanalytic matrix functions on quantum computers.

Introduction

Quantum computers hold the promise of solving numerous computational problems faster than their classical counterparts [17, 39]. A prominent example is the ground state projection problem, whose solution gives access to the static properties of many-body quantum systems relevant to application areas like chemistry, nuclear physics, and materials science. In this problem, we are given a description of a many-body Hamiltonian A along with an assumption on the gap δ between the ground energy and the first excited energy, and the goal is to efficiently perform the eigenprojection Π_{λ_0} onto the eigenstate associated with the ground energy λ_0. One solution to this problem—known already in the early days of quantum algorithms research—applies the phase estimation algorithm to the time evolution unitary e^{iAt}, effecting the projection Π_{λ_0} [1, 10, 22]. To resolve the ground and excited states this way, the evolution must be performed out to time $t = \mathbf{\Theta}(1/\delta)$—a fundamental lower bound related to the energy–time uncertainty principle [5].

Modern approaches to eigenprojection for Hermitian matrices indeed achieve optimal total query cost $\mathbf{\Theta}(1/\delta)$ [15, 26], but they are better viewed through the lens of linear algebra: one starts with a description of the input matrix A and aims to perform a particular matrix function $f(A)$ on a quantum computer. When A is Hermitian, its singular values agree with eigenvalues in magnitude, whereas its eigenstates form an orthonormal basis. Then the task can be accomplished by the

*Corresponding author: buptsuyuan@gmail.com

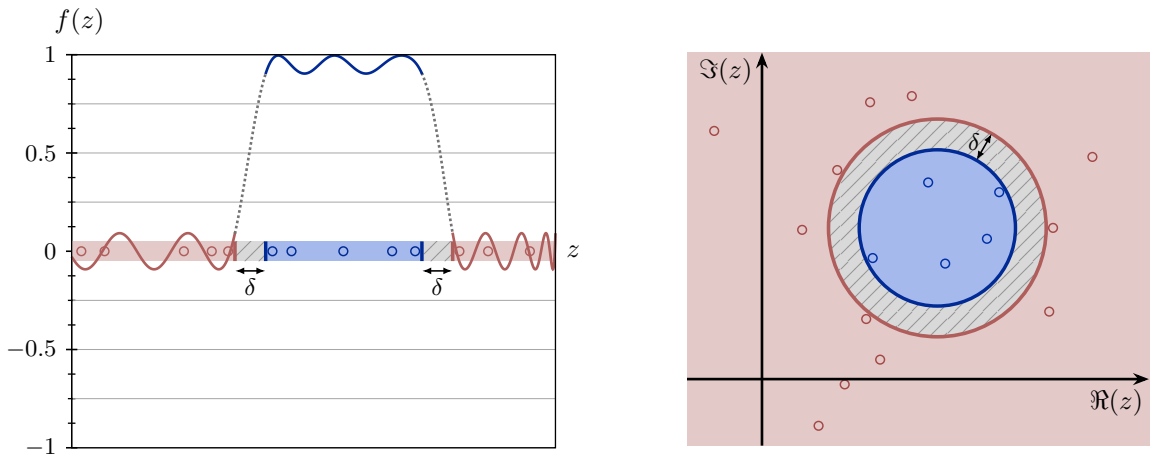

(a) Real spectrum

(b) Complex spectrum

Figure 1: The eigenprojection problem asks to project onto generalized eigensubspaces associated with eigenvalues (open circles) lying within a particular target region (blue region), promised that all other eigenvalues lie at least δ-far (red region) from the target region. We illustrate the target region as a disk, but we also consider more general regions. (a) If all eigenvalues are real, there exist suitable polynomials f with degree linear in $1/\delta$ which approximate the indicator function, that is, a function which is 1 on the blue region, 0 on the red region, and arbitrary in the gray region where there are no eigenvalues. This enables optimal (in $1/\delta$) eigenprojection via QSVT or QEVT. (b) The same does not hold when the spectrum is complex. It is unclear how to construct a polynomial (or even analytic) function which approximates the analogous indicator function on the blue and red regions.

Quantum Singular Value Transformation (QSVT) [16]. Crucially, QSVT can perform polynomial transformations to Hermitian input matrices with a query complexity linear in the degree of the polynomial, and thus can remain efficient even when the underlying Hilbert space has an exponentially large dimension. For example, QSVT optimally solves the orthogonal eigenprojection problem by using a degree-$\mathbf{O}(1/\delta)$ polynomial f that approximates the indicator function on the target part of the spectrum, so that $f(A) \approx \Pi_{\lambda_0}$ [16, Lemma 29]. But QSVT applies far more widely, providing a unified framework for developing many quantum algorithms [31] with nearly optimal query complexity [33], including Hamiltonian simulation [27], Gibbs state preparation, and solving linear systems [16], which correspond to different choices of the function f.

Performing $f(A)$ for a general, non-normal matrix A is considerably harder, as the eigenvalues and singular values of A no longer relate and the associated generalized eigensubspaces are typically oblique. Nonetheless, the ability to implement Quantum EigenValue Transformation (QEVT) for non-normal operators is broadly useful: it arises in simulating non-Hermitian physics [4], in transcorrelated approaches to quantum chemistry [32], in solving differential equations [7, 8, 24], and in studying irreversible Markov chains [6, 9]. When the target function f is analytic, several approaches have been developed to transform non-normal matrices on a quantum computer, which can be broadly grouped into three methodologies: those based on resolvent integration [20, 41–43, 46], those based on solving a system of linear equations built from a matrix generating function [28], and those based on a linear combination of Hamiltonian simulations or a linear combination of Hermitian matrices [3, 34, 47].

An even more challenging task is to perform *nonanalytic* functions of general matrices. A simple yet fundamental example is the eigenprojection Π associated with a target region of the

spectrum in the complex plane, corresponding to the indicator function of that region—a natural generalization of the influential ground state projection problem discussed above. As illustrated in Figure 1, a polynomial can approximate this indicator function when the spectrum is real, but doing so is fundamentally obstructed when the spectrum is complex [23, Corollary B.24][19, Theorem 11.5.5], where it is unclear how to polynomially approximate the indicator function in every direction [11]. Such oblique eigenprojections are nonetheless useful in a variety of settings, including preparing eigenstates of non-normal matrices, solving Sylvester equations, and solving Riccati equations arising in quantum chemistry [35, 48].

Resolvent integrals offer a way around this obstruction of nonanalyticity. The strategy is to select a contour enclosing the target spectrum and integrate the resolvent along it, which recovers the oblique eigenprojection through Cauchy's integral formula for matrices. This resolvent-integration approach was first used to estimate the eigenvalue density in a target region [14], and was subsequently developed within the block encoding framework and analyzed in detail in [35]. In the special case where the target region is a half-plane, one can instead integrate the resolvent along the axis [48]. However, a drawback with this approach is that it introduces a normalization factor $\alpha = \mathbf{\Theta}(\|\Pi\| / \delta)$ with δ the gap between the spectrum of A and boundary of the target region, which is much larger than the desired normalization factor $\|\Pi\|$. This is in stark contrast to the QSVT and QEVT frameworks for analytic functions, where $f(A)$ can be implemented with a normalization factor as small as $\|f(A)\|$. The inflated normalization factor significantly increases the cost of the resulting quantum algorithms when the eigenprojection is used as a subroutine within a larger computation—as in the solution of Riccati equations that we discuss in this paper.

In this work, we show that a quantum computer can perform the oblique eigenprojection Π given block encoding access to the input matrix A. Our approach has a query complexity almost linear in $1/\delta$, and it produces a block encoding with a normalization factor close to $\alpha = \mathbf{\Theta}(\|\Pi\|)$ under a spectral-set condition on A, covering common assumptions on its numerical range or diagonalizability. As $\|\Pi\|$ is the smallest factor achievable by any block encoding of Π, this normalization is necessarily nearly optimal. Compared with the standard resolvent integration approach, our result quadratically improves the cost of oblique eigenprojection as a subroutine in a larger computation, and it matches analogous results previously established only for the orthogonal case [26] and real spectrum [28].

Our main technical contribution is a two-sided block preconditioning that uses a discrete Fourier transform of the matrix resolvent. Whereas the standard resolvent integral extracts only the top-left corner carrying the projector, we characterize every block through the discrete Fourier transform of the resolvent. This exposes a circulant structure of the block encoding, and preconditioning against it suppresses the inverse-gap growth of the normalization factor down to nearly $\|\Pi\|$, while maintaining the same query complexity as the standard approach.

As the oblique eigenprojection serves as a fundamental primitive in the design of quantum algorithms, our result can be used to extend and improve a variety of prior results. We describe three such applications. First, we prepare eigenstates of matrices with complex eigenvalues, extending the QEVT algorithm of Low and Su [28] beyond real spectra. Second, we solve continuous-time algebraic Riccati equations, cubically improving a prior solver of Rodenas-Ruiz, Zhao, and Lee [35]. Third, we solve ordinary Sylvester equations, quadratically improving a direct augmented method of Wang and Liu [48]. These applications suggest a promising route toward applying more general nonanalytic matrix functions on quantum computers.

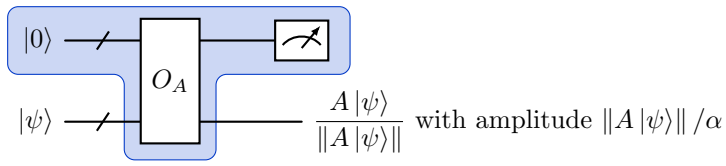

Figure 2: Block encoding as a unitary access model: initialize the ancilla register in $|0\rangle$, apply the joint unitary O_A, and measure the ancilla. Suppose that $A/\alpha = (\langle 0| \otimes I)\, O_A\, (|0\rangle \otimes I)$. Conditioned on the desired outcome $|0\rangle$, the input register $|\psi\rangle$ is transformed by the matrix A with success amplitude $\|A|\psi\rangle\|/\alpha$.

Results

The standard resolvent-integral block encoding

A common framework for developing and analyzing quantum algorithms in linear algebra is the *block encoding* framework [16], which provides a unitary model of access to an input matrix A. To describe it, consider a typical quantum computation consisting of three steps: we first initialize an ancilla register in a fixed state $|0\rangle$, then apply a joint unitary O_A to the ancilla and input register, and finally measure the ancilla with the desired measurement outcome flagged by $|0\rangle$. Conditioned on this outcome, the computation enacts a transformation on the input register. Bundling these steps together, we may abstract away the ancilla and refocus on how the overall procedure transforms the input state, which is precisely the action of the matrix A we wish to access.

Formally, a unitary O_A is said to be a block encoding of a matrix A if

$$A = (\langle 0| \otimes I)\, O_A\, (|0\rangle \otimes I)\,, \tag{1}$$

where $|0\rangle$ denotes the fixed initial state of the ancilla register. Equivalently, using the matrix representation $|0\rangle = \left[\begin{smallmatrix}1\\0\end{smallmatrix}\right]$ and $\langle 0| = \left[\begin{smallmatrix}1 & 0\end{smallmatrix}\right]$, A appears as the top-left block of O_A,

$$O_A = \begin{bmatrix} A & \cdot \\ \cdot & \cdot \end{bmatrix}, \tag{2}$$

with the remaining blocks left unspecified, subject only to the unitarity of O_A. A key feature of this framework is that common matrix arithmetic operations can be performed directly at the level of block encodings: given block encodings of matrices A and B, one can construct block encodings of their linear combination, product AB, Hermitian conjugate $A^\dagger$, and tensor product $A \otimes B$, as well as the inverse A^{-1} [16]. More generally, the quantum singular value transformation allows one to block encode a polynomial transformation of the singular values of A, with query complexity linear in the polynomial degree.

For a general matrix A, the above definition must be relaxed to allow for a normalization factor α, so that the unitary O_A block encodes the rescaled matrix A/α rather than A itself:

$$A/\alpha = (\langle 0| \otimes I)\, O_A\, (|0\rangle \otimes I)\,. \tag{3}$$

Since the left-hand side is a submatrix of a unitary, its spectral norm is at most one, and hence $\alpha \geq \|A\|$ is necessary for such a block encoding to exist. This lower bound is purely a property of A; in practice, the normalization factor achieved by a specific construction depends on the algorithm used and can be larger than $\|A\|$. The normalization factor directly governs the cost of the resulting quantum algorithm. Indeed, when O_A is applied to an input state $|\psi\rangle$, the desired outcome is heralded with success amplitude $\|A|\psi\rangle\|/\alpha$, so a larger α means a smaller success probability

and a correspondingly higher cost to prepare the output near deterministically. Reducing the normalization factor of a block encoding is therefore essential to reducing the overall cost of quantum algorithms. See Figure 2 for a circuit illustration of this point.

We now introduce the central problem of this work. Given an arbitrary square matrix A, the underlying Hilbert space decomposes into the generalized eigensubspaces of A, associated with its complex eigenvalues $\mathbf{Spec}(A)$. Any partition of the spectrum into two disjoint groups $\mathbf{Spec}_{\text{in}}(A)$ and $\mathbf{Spec}_{\text{out}}(A)$ then induces a corresponding decomposition of the space into a pair of complementary invariant subspaces, together with an associated eigenprojection Π onto the first and $I-\Pi$ onto the second. For presentational purpose, we first consider the case in which the eigenvalues are separated by the unit circle, so that the spectrum splits into those lying inside the circle and those lying outside. The problem is then to perform the oblique eigenprojection Π onto the generalized eigensubspaces whose eigenvalues $\mathbf{Spec}_{\text{in}}(A)$ lie inside the unit circle.

The eigenprojection Π admits a classical representation as a contour integral of the resolvent $(\zeta I - A)^{-1}$, known as the Riesz projector [19]. Taking the target region to be the unit disk, the contour is the unit circle, which we parametrize as $\zeta = \mathrm{e}^{\mathrm{i}\theta}$ with $\mathrm{d}\zeta = \mathrm{i}\,\mathrm{e}^{\mathrm{i}\theta}\,\mathrm{d}\theta$. Factoring $\mathrm{e}^{\mathrm{i}\theta}$ out of the resolvent then yields the simplified expression

$$\Pi = \frac{1}{2\pi\mathrm{i}} \int_{|\zeta|=1} (\zeta I - A)^{-1}\,\mathrm{d}\zeta = \frac{1}{2\pi} \int_0^{2\pi} \left(I - \mathrm{e}^{-\mathrm{i}\theta} A\right)^{-1} \mathrm{d}\theta, \tag{4}$$

where the integrand depends on A only through the shifted resolvent $(I - \mathrm{e}^{-\mathrm{i}\theta}A)^{-1}$ along the unit circle.

We now review the standard approach to implementing this oblique eigenprojection on a quantum computer [35]. We first discretize the integral using the n-point trapezoidal rule at the equally spaced nodes $\theta_j = 2\pi j/n$, giving $\Pi \approx \frac{1}{n}\sum_{j=0}^{n-1}\left(I - \omega_n^{-j}A\right)^{-1}$, $\omega_n = \mathrm{e}^{2\pi\mathrm{i}/n}$, which converges rapidly as n increases. To implement the right-hand side, we assemble the block-diagonal matrix $M = \sum_{j=0}^{n-1} |j\rangle\langle j| \otimes \left(I - \omega_n^{-j}A\right)$, each of whose blocks is a shifted copy of A and hence admits a block encoding through linear combination. Applying QSVT to invert M yields a block encoding of

$$M^{-1} = \begin{bmatrix} (I-A)^{-1} & & & \\ & \left(I - \omega_n^{-1}A\right)^{-1} & & \\ & & \ddots & \\ & & & \left(I - \omega_n^{-(n-1)}A\right)^{-1} \end{bmatrix}. \tag{5}$$

Preparing the uniform superposition $\mathbf{Had}\,|0\rangle = \frac{1}{\sqrt{n}}\sum_{j=0}^{n-1}|j\rangle$ by a Hadamard operation on each side and extracting the average of the diagonal blocks, we obtain

$$\left(\langle 0|\,\mathbf{Had} \otimes I\right) M^{-1} \left(\mathbf{Had}\,|0\rangle \otimes I\right) = \frac{1}{n}\sum_{j=0}^{n-1}\left(I - \omega_n^{-j}A\right)^{-1} \approx \Pi, \tag{6}$$

so that the composite operation block encodes the Riesz projector. Equivalently, in the matrix representation, the Hadamard-conjugated inverse is an $n \times n$ array of blocks,

$$\left(\mathbf{Had} \otimes I\right) M^{-1} \left(\mathbf{Had} \otimes I\right) \approx \begin{bmatrix} \Pi & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ \vdots & \vdots & \ddots & \vdots \\ \cdot & \cdot & \cdots & \cdot \end{bmatrix}, \tag{7}$$

with the eigenprojection Π appearing in the top-left block and all remaining blocks left unspecified.

Let us examine the normalization factor incurred by this approach. Using QSVT, each shifted resolvent can be block encoded with normalization factor $\mathbf{\Theta}\left(\left\|(I-\omega_n^{-j}A)^{-1}\right\|\right)$ to accuracy ε, using $\mathbf{\Theta}\left(\left\|(I-\omega_n^{-j}A)^{-1}\right\|\log(1/\varepsilon)\right)$ queries to the input block encoding of A. Taking the linear combination then yields a block encoding of Π with normalization factor $\mathbf{\Theta}(\alpha_{\mathrm{res}})$ at the same query complexity, where

$$\alpha_{\mathrm{res}} = \max_{\theta\in[0,2\pi)} \left\|\left(I - \mathrm{e}^{-\mathrm{i}\theta}A\right)^{-1}\right\|. \tag{8}$$

This can be far larger than the desired value $\|\Pi\|$: already when A is normal, $\alpha_{\mathrm{res}} = 1/\delta$ scales as the inverse distance between the spectrum of A and the unit circle, whereas $\|\Pi\| = 1$. As noted earlier, this inflated normalization factor increases the cost of the resulting quantum algorithm, especially when the eigenprojection is used as a subroutine within a larger computation, such as in solving the Riccati equations discussed in this paper.

A discrete Fourier transform of the matrix resolvent

The inefficiency of the standard approach can be traced to how little it uses of the matrix M^{-1}: the Hadamard conjugation (7) exposes only the top-left block Π, while every remaining block is left undetermined. A valid block encoding must nonetheless accommodate all of these blocks, and normalizing the entire matrix forces the large factor α_{res}, even though only the top-left block is of interest.

Our key insight is that these blocks are highly structured. Rather than conjugating M^{-1} by the Hadamard operation, we conjugate by the quantum Fourier transform. This resolves M^{-1} into a complete array of blocks,

$$(\mathbf{QFT}\otimes I)\, M^{-1}\left(\mathbf{QFT}^{\dagger}\otimes I\right) = \begin{bmatrix} S_0 & S_{n-1} & \cdots & S_1 \\ S_1 & S_0 & \cdots & S_2 \\ \vdots & \vdots & \ddots & \vdots \\ S_{n-1} & S_{n-2} & \cdots & S_0 \end{bmatrix}, \qquad S_m = \frac{1}{n}\sum_{j=0}^{n-1} \omega_n^{jm}\left(I-\omega_n^{-j}A\right)^{-1}, \tag{9}$$

in which every block is determined: the (k,l) block depends only on the difference $k-l$ and is given by the discrete Fourier transform S_m of the shifted resolvent. In particular, the top-left block S_0 approximately recovers the eigenprojection Π, exactly as in the Hadamard construction, but now the surrounding blocks carry structure we can exploit rather than merely accommodate. The modes admit a closed form in terms of the two restrictions of A to the invariant subspaces inside and outside the unit circle.

Theorem 1 (Discrete Fourier transform of the matrix resolvent)**.** *Let A be a square matrix whose spectrum is disjoint from the unit circle, with oblique eigenprojection Π onto the invariant subspace inside the circle. Write $A_{\mathrm{in}} = A\Pi$ and $A_{\mathrm{out}} = A(I-\Pi)$ for the restrictions of A to the inside and outside subspaces, with A_{out}^{-1} the inverse of A_{out} on the outside subspace. Then, for $0 \le m \le n-1$,*

$$S_m = A_{\mathrm{in}}^{m}(I-A_{\mathrm{in}}^{n})^{-1}\,\Pi - A_{\mathrm{out}}^{m-n}(I-A_{\mathrm{out}}^{-n})^{-1}\,(I-\Pi). \tag{10}$$

Both A_{in} (inside the circle) and A_{out}^{-1} (outside) have spectral radius below one, so their powers decay geometrically under a spectral-set condition on the input. We say a compact set $\mathcal{K}$ containing the spectrum of a matrix B is a *spectral set* with constant c if $\|g(B)\| \le c\max_{z\in\mathcal{K}}|g(z)|$ for every g analytic near $\mathcal{K}$. For a normal matrix the spectrum itself is a spectral set with $c = 1$; for

a non-normal matrix one enlarges the set or pays a larger constant. By the Crouzeix–Palencia theorem [12], the numerical range is always a spectral set with $c = 1+\sqrt{2}$, and for a diagonalizable matrix the spectrum is a spectral set with c equal to the eigenbasis condition number. Assuming A_{in} and A_{out}^{-1} admit spectral sets inside $\{|z| \leq 1-\delta\}$ with constant c, taking $g(z) = z^j$ gives $\left\|A_{\text{in}}^j\right\|, \left\|A_{\text{out}}^{-j}\right\| \leq c(1-\delta)^j$. Hence the corner block $S_0 = \Pi + \mathbf{O}(c(1-\delta)^n \|\Pi\|)$ is exponentially close to the eigenprojection, while every other block satisfies $S_m = \mathbf{O}(c\,\|\Pi\|)$.

Two-sided block preconditioning on the unit disk

We reduce the normalization factor of the oblique eigenprojection Π through a two-sided block preconditioning in the Fourier basis. Let $T = t|0\rangle\langle 0| + \sum_{k=1}^{n-1} |k\rangle\langle k|$ be a scaling operator with $0 < t < 1$, which can be block encoded with normalization factor 1 at no additional query cost. Applying QSVT to the preconditioned matrix $(T\otimes I)(\mathbf{QFT}\otimes I)M(\mathbf{QFT}^\dagger\otimes I)(T\otimes I)$ yields a block encoding of its inverse

$$\left(T^{-1}\otimes I\right)\left(\mathbf{QFT}\otimes I\right)M^{-1}\left(\mathbf{QFT}^\dagger\otimes I\right)\left(T^{-1}\otimes I\right) = \begin{bmatrix} \frac{S_0}{t^2} & \frac{S_{n-1}}{t} & \cdots & \frac{S_1}{t} \\ \frac{S_1}{t} & S_0 & \cdots & S_2 \\ \vdots & \vdots & \ddots & \vdots \\ \frac{S_{n-1}}{t} & S_{n-2} & \cdots & S_0 \end{bmatrix}. \tag{11}$$

As t decreases, the first block row and column grow, so the condition number of the inversion increases; but the normalization factor of the extracted top-left block decreases at the quadratically higher rate t^2. Balancing the two effects gives the main result.

Theorem 2 (Quantum oblique eigenprojection on the unit disk)**.** *Let A be a square matrix accessed through a block encoding with normalization factor α_A, whose spectrum is disjoint from the unit circle, with oblique eigenprojection Π onto the invariant subspace inside the circle. Suppose the restrictions $A_{\text{in}} = A\Pi$ and A_{out}^{-1} admit spectral sets inside $\{|z| \leq 1-\delta\}$ with common constant c. Then for any $\varepsilon > 0$, two-sided block preconditioning block encodes Π with normalization factor*

$$\alpha_\Pi = \mathbf{O}(c\,\|\Pi\|), \tag{12}$$

near its optimal value $\|\Pi\|$, and accuracy ε, using

$$\mathbf{O}\left(\frac{c\,\alpha_A\,\|\Pi\|}{\delta}\log\frac{c\,\|\Pi\|}{\varepsilon}\log\frac{1}{\varepsilon}\right) \tag{13}$$

queries to the block encoding of A. In particular, this holds with $c = 1+\sqrt{2}$ under a numerical-range assumption on the two restrictions, and with $c = \kappa_V$ the eigenbasis condition number under a diagonalizability assumption.

The normalization factor $\mathbf{O}(c\,\|\Pi\|)$ is near its optimal value: $\|\Pi\|$ is the smallest factor achievable by any block encoding of Π. Compared with the unpreconditioned algorithm, whose condition number is $\mathbf{O}(\|\Pi\|/\delta)$, the preconditioned inversion has a condition number larger only by a logarithmic factor, so the query complexity remains almost the same as the standard approach while the normalization factor is reduced to nearly optimal. The balanced choice $t = \mathbf{\Theta}(\sqrt{\delta})$ realizes this tradeoff; a different choice $t = 1/\sqrt{n}$ removes the logarithmic factor from the normalization at the same overall query cost.

Quantum oblique eigenprojection for general regions

The same approach adapts to more general regions of the complex plane, whose boundary separates the spectrum of A into an inside part and an outside part. Suppose f is analytic on disjoint open sets $\mathcal{S}_{\text{in}}$ and $\mathcal{S}_{\text{out}}$ covering the inside and outside parts of the spectrum, and separates them in modulus: $|f(w)| < 1$ for $w \in \mathcal{S}_{\text{in}}$ and $|f(w)| > 1$ for $w \in \mathcal{S}_{\text{out}}$. Then the oblique eigenprojection onto the inside part is

$$\Pi = \frac{1}{2\pi \mathrm{i}} \int_{\mathcal{C}} (wI - A)^{-1} \, \mathrm{d}w = \frac{1}{2\pi \mathrm{i}} \int_{|z|=1} (zI - f(A))^{-1} \, \mathrm{d}z, \tag{14}$$

for a curve $\mathcal{C}$ enclosing the inside part, reducing the general-region problem to the unit-disk construction applied to $f(A)$ in place of A. The map f only needs to separate the two parts in modulus. It need not be injective or conformal, so we are free to choose a low-degree rational function.

Theorem 3 (Quantum oblique eigenprojection for general regions)**.** *Let A be a square matrix with oblique eigenprojection Π onto the invariant subspace associated with a target region, and let f be an analytic separating map for that region, with $f(A)$ accessed through a block encoding with normalization factor $\alpha_{f(A)}$. Suppose the restrictions $A_{\text{in}} = A\Pi$ and $A_{\text{out}} = A(I - \Pi)$ admit spectral sets $\mathcal{K}_{\text{in}}$ and $\mathcal{K}_{\text{out}}$ with common constant c, on neighborhoods of which f is analytic, and that f has gap δ_f,*

$$\max_{z \in \mathcal{K}_{\text{in}}} |f(z)| \leq 1 - \delta_f, \qquad \min_{z \in \mathcal{K}_{\text{out}}} |f(z)| \geq 1 + \delta_f. \tag{15}$$

Then for any $\varepsilon > 0$, the two-sided block preconditioning of Theorem 2 applied to $f(A)$ block encodes Π with normalization factor $\alpha_\Pi = \mathbf{O}(c \, \|\Pi\|)$ and accuracy ε, using

$$\mathbf{O}\left(\frac{c \, \alpha_{f(A)} \, \|\Pi\|}{\delta_f} \log \frac{c \, \|\Pi\|}{\varepsilon} \log \frac{1}{\varepsilon} \right) \tag{16}$$

queries to the block encoding of $f(A)$.

For a simply connected region a conformal separating map exists by the Riemann mapping theorem, and a polynomial approximation implemented by QEVT [28] suffices; in practice the regions arising in applications are bounded by lines and circular arcs, for which a separating map is built directly from elementary Möbius factors. In particular, after rescaling A so that $\|A\| \leq \frac{1}{2}$, the half-planes are separated by

$$f(A) = \frac{(A - I)^2}{(A + I)^2}, \tag{17}$$

which sends the right half-plane into the unit disk and the imaginary axis onto the unit circle. Here $A + I$ is well conditioned, so $f(A)$ is block encoded at constant normalization with only a logarithmic query overhead, while the gap is preserved up to a constant factor. This construction underlies our improvements to the Riccati and Sylvester equations below. See Figure 3.

Applications

As the oblique eigenprojection serves as a fundamental primitive in the design of quantum algorithms, our result can be used to extend and improve a range of prior results.

- **Eigenstate preparation for complex spectra.** Preparing an eigenstate of a given operator is among the most fundamental tasks in quantum computation, which is relevant for studying static properties of the underlying physical system. For a Hermitian Hamiltonian the eigenstates

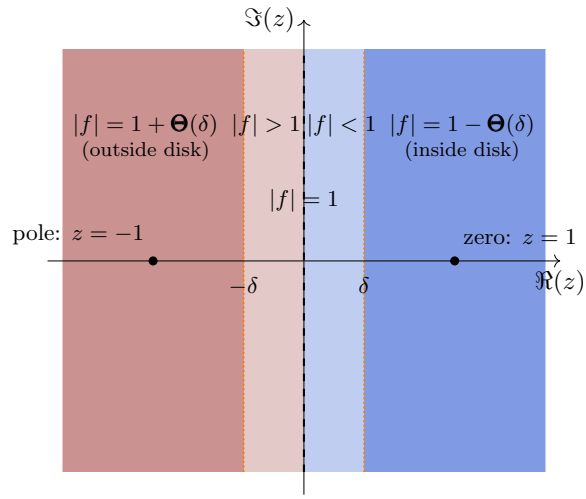

Figure 3: The rational map $f(z) = (z-1)^2/(z+1)^2$ separates the complex plane by the imaginary axis. The right half-plane (blue) is mapped inside the unit circle ($|f| < 1$) and the left half-plane (red) outside ($|f| > 1$), with the imaginary axis ($|f| = 1$) mapped onto the unit circle. The double zero at $z = 1$ and double pole at $z = -1$ lie in the right and left half-planes respectively. Darker shading indicates the gapped regions $|\Re(z)| \geq \delta$, on which $|f| = 1 - \mathbf{\Theta}(\delta)$ (right) and $|f| = 1 + \mathbf{\Theta}(\delta)$ (left), so the gap to the unit circle is preserved, $\delta_f = \mathbf{\Theta}(\delta)$.

are orthogonal, and the task is well understood. Indeed, given an initial state with nonnegligible overlap on the target eigenspace and a spectral gap δ separating it from the rest of the spectrum, one can prepare the target eigenstate efficiently [26]. The same task is natural for an arbitrary, non-normal matrix A, whose eigenvalues may lie anywhere in the complex plane and whose eigenspaces are oblique. Such matrices arise in simulating non-Hermitian physics [4], in transcorrelated approaches to quantum chemistry [32], and in studying irreversible Markov chains [6, 9], among other settings. Given an initial state $|\psi\rangle$ with overlap $\eta = \|\Pi|\psi\rangle\|$, we prepare $\Pi|\psi\rangle/\eta$ for a target disk in the complex plane. This extends the QEVT algorithm of Low and Su [28], whose polynomial filters are tied to real spectra. Amplitude amplification applies the projector $\mathbf{O}(\alpha/\eta)$ times, where α is its normalization factor. For the resolvent integral, $\alpha_{\text{res}} = \mathbf{O}(\|\Pi\|/\delta)$ contributes an extra $1/\delta$, whereas our factor is close to $\alpha_\Pi \approx \|\Pi\|$. As a result, our initial state is prepared about $\mathbf{O}(\|\Pi\|/\eta)$ times independent of the gap, whereas our total number of queries to the block encoding is close to $\mathbf{O}(\|\Pi\|^2/(\eta\delta))$, in contrast to the scaling $\mathbf{O}(\|\Pi\|^2/(\eta\delta^2))$ for the standard resolvent integral approach.

- **Continuous-time algebraic Riccati equations.** Riccati equations are central to optimal control, filtering, and stability theory, and in quantum chemistry they arise directly in ring coupled-cluster doubles and the random-phase approximation [35]. The stabilizing solution X of the continuous-time algebraic Riccati equation $XQX - XP - P^\dagger X - R = 0$ is characterized through an invariant subspace of the associated Hamiltonian matrix whose spectrum is symmetric about the imaginary axis, so solving it reduces to projecting onto an invariant subspace of a non-normal matrix split by the imaginary axis—an oblique eigenprojection of exactly the kind our primitive provides. The solution is encoded in the antistable invariant subspace; writing the projector onto it as $\Pi = [\,\Pi_1\ \Pi_2\,]$, one has $X = -\Pi_2^+\Pi_1$ [35]. Implementing the pseudoinverse from a block encoding of Π with normalization factor α requires $\mathbf{O}(\alpha)$ uses of the projector, each costing about $1/\delta$ queries, and yields X with normalization factor $\mathbf{O}(\alpha)$. The resolvent integral

has $\alpha_{\text{res}} = \mathbf{O}(\|\Pi\| / \delta)$, so the inverse gap $1/\delta$ enters the final query complexity three times. Our normalization factor $\alpha_\Pi \approx \|\Pi\|$ removes two of these, cubically improving the gap dependence of the solver of Rodenas-Ruiz, Zhao, and Lee.

- **Ordinary Sylvester equations.** Sylvester equations arise throughout control theory, model reduction, and the numerical solution of matrix and differential equations. As with the Riccati equation, the solution is encoded in an invariant subspace of an associated block matrix, now with the spectra of A and $-B$ lying in opposite open half-planes separated, after a shift and rotation, by the imaginary axis. The solution X of the ordinary Sylvester equation $AX + XB = C$ appears directly as the off-diagonal block of the projector $\left[\begin{smallmatrix} I & X \\ 0 & 0 \end{smallmatrix}\right]$ onto the right-half-plane invariant subspace of $\left[\begin{smallmatrix} A & C \\ 0 & -B \end{smallmatrix}\right]$. So no pseudoinverse is needed, and the normalization factor of the projector becomes that of X. Reducing it quadratically improves the gap dependence of the direct augmented method of Wang and Liu [48, Remark 6.4].

We give the detailed cost analysis of all three applications in Methods.

Discussion

We have shown that a quantum computer can perform an oblique eigenprojection Π given block encoding access to a non-normal input matrix. Our method has a query complexity almost linear in the inverse gap and a normalization factor close to its optimal value $\|\Pi\|$. This improves the standard resolvent-integral approach and matches known results for the orthogonal case. The improvement carries through to concrete applications: preparing eigenstates of matrices with complex eigenvalues, and solving the continuous-time algebraic Riccati and Sylvester equations, where it sharpens the gap dependence of recent quantum algorithms by a quadratic to cubic factor.

Underlying the result is a complete characterization of the block encoding produced by the standard resolvent integral. Rather than tracking only the top left corner that carries the projector, we express every block through the discrete Fourier transform of the resolvent, and it is this full description of the otherwise uninteresting entries that makes the two-sided block preconditioning possible. A similar attention to the entire structure of a block encoding, rather than a single distinguished entry, appears in several works, for reducing the ancilla overhead of block encodings through uncomputing [45], for maintaining matrix encoding in the off-diagonal blocks of a Hamiltonian to compose arithmetic operations [21], and for constraining the powers of a block encoding to reproduce those of the encoded matrix [18, 30]. We expect that a further understanding of this kind will find use in the design of quantum algorithms.

Our preconditioning rests on the closed-form expression for the discrete Fourier transform of the matrix resolvent, and it does not directly apply to every construction. In particular, the main approach of Wang and Liu [48] represents the matrix sign through a product of two resolvents rather than a single one, for which we do not have an analogous closed form for the discrete Fourier transform, so the two-sided block preconditioning does not carry over. Adapting it to reduce the normalization factor in that setting is an interesting open direction.

It would also be valuable to better understand the case of more general spectral regions. Our construction there first block encodes $f(A)$ for a separating map f and then applies the unit-disk algorithm, so the cost and the spectral-set hypotheses are inherited through f. It would be desirable to avoid the explicit block encoding of $f(A)$, and instead work with the resolvent of A itself, integrated over a general curve enclosing the target spectrum. Such a construction could considerably simplify the algorithm, but would require preconditioning directly against the geometry of the region.

A different route to the eigenprojection is through quantum eigenvalue estimation: one estimates the eigenvalues of the input matrix in an ancillary register and then projects by selecting those in the target region. Beyond the QEVT technique from [28], a number of algorithms estimate eigenvalues of non-normal or non-Hermitian matrices [2, 37, 38, 49]. This approach, however, requires resolving individual eigenvalues to high precision and it typically imposes additional structural assumptions on spectrum of the input matrix and properties of the initial state. Our construction instead projects by location in the complex plane, requiring only that the input admit a spectral set separated from the region boundary, and never estimates an eigenvalue.

A further contrast is worth noting. For a Hermitian operator, one can transform it by QSVT and its cost is solely determined by the function alone: implementing a window function, for instance, costs inversely in its transition width independently of the input matrix. Our oblique eigenprojection algorithm does not admit such a feature. Because the indicator of a spectral region is nonanalytic, it cannot be approximated uniformly and applied blindly; our construction instead controls the resolvent on the boundary of the region, which requires the spectrum of the input to be separated from that boundary by a gap. Our cost is therefore governed by a property of the input matrix—its gap—rather than by a property of the target function. It would be interesting to know whether an eigenprojection can be implemented at a cost depending only on the geometry of the target region, as in the Hermitian case, rather than on the spectral gap of the input.

Conceptually, the oblique eigenprojection is not well served by QSVT and QEVT, because the required function is necessarily nonanalytic, which rules out a direct implementation through polynomial approximation. Prior to our work, it was unclear whether the normalization factor for block encoding could be kept close to optimal for such nonanalytic functions. Our resolvent-based construction answers this question affirmatively for the important case of oblique eigenprojection, achieving complexity analogous to that of QSVT and QEVT through quite different techniques. Whether the same discrete Fourier transform and preconditioning ideas extend to other nonanalytic functions of a non-normal matrix is an intriguing direction for future work.

Methods

We provide the detailed constructions and cost analyses behind the three theorems of the Results. Throughout the Results we wrote Π for the oblique eigenprojection; in the constructions below it is convenient to display the complementary projector explicitly, so we write $\Pi_{\text{in}} = \Pi$ for the eigenprojection onto the inside subspace and $\Pi_{\text{out}} = I - \Pi_{\text{in}}$ for its complement. Full proofs of all statements are given in the Supplementary Information.

Discrete Fourier transform of the resolvent

Conjugating M^{-1} by the quantum Fourier transform resolves it into the block-circulant array (9), whose (k, l) block depends only on $k - l$ and equals the discrete Fourier mode S_m of the shifted resolvent. To evaluate the modes in closed form, we split A across the two complementary invariant subspaces. Writing $\Pi_{\text{out}} = I - \Pi_{\text{in}}$, we define

$$A_{\text{in}} \;=\; A\,\Pi_{\text{in}}, \qquad A_{\text{out}} \;=\; A\,\Pi_{\text{out}}, \tag{18}$$

so that $A = A_{\text{in}} + A_{\text{out}}$, with A_{in} supported on $\mathcal{H}_{\text{in}} = \mathbf{Im}(\Pi_{\text{in}})$ and A_{out} on $\mathcal{H}_{\text{out}} = \mathbf{Im}(\Pi_{\text{out}})$. By construction, A_{in} has all its eigenvalues inside the unit circle, hence spectral radius less than one. The restriction of A to $\mathcal{H}_{\text{out}}$ has all its eigenvalues outside the unit circle and is thus invertible on that subspace; we write A_{out}^{-1} for the inverse of this restriction, so that A_{out}^{-1} has spectral radius less than one on $\mathcal{H}_{\text{out}}$, extended by zero on $\mathcal{H}_{\text{in}}$.

On each subspace, the shifted resolvent expands as a geometric series: in nonnegative powers of A_{in} on $\mathcal{H}_{\text{in}}$, and in negative powers of A_{out} on $\mathcal{H}_{\text{out}}$. Carrying out the discrete Fourier sum yields the closed form (10) of Theorem 1. Since Π_{in} acts as the identity on $\mathcal{H}_{\text{in}}$ and as zero on $\mathcal{H}_{\text{out}}$, the deviation of the zeroth mode from the oblique eigenprojection is

$$S_0 - \Pi_{\text{in}} \;=\; \begin{cases} A_{\text{in}}^{n}\,(I - A_{\text{in}}^{n})^{-1} & \text{on } \mathcal{H}_{\text{in}}, \\ -\,A_{\text{out}}^{-n}\left(I - A_{\text{out}}^{-n}\right)^{-1} & \text{on } \mathcal{H}_{\text{out}}. \end{cases} \tag{19}$$

The correction terms A_{in}^{n} and A_{out}^{-n} are exponentially small in n, since both A_{in} and A_{out}^{-1} have spectral radius below one. Thus $S_0 = \Pi_{\text{in}} + \mathbf{O}(\|A_{\text{in}}^{n}\|\,\|\Pi_{\text{in}}\| + \left\|A_{\text{out}}^{-n}\right\|\,\|\Pi_{\text{out}}\|)$, recovering $S_0 \approx \Pi_{\text{in}}$, while the remaining modes are uniformly bounded as $S_m = \mathbf{O}(\alpha_{\text{in}}\,\|\Pi_{\text{in}}\| + \alpha_{\text{out}}\,\|\Pi_{\text{out}}\|)$, where $\alpha_{\text{in}} = \max_{0\le j\le n-1}\left\|A_{\text{in}}^{j}\right\|$ and $\alpha_{\text{out}} = \max_{1\le k\le n}\left\|A_{\text{out}}^{-k}\right\|$ collect the largest powers of A_{in} and A_{out}^{-1}.

Two-sided block preconditioning

The block preconditioning technique we use was introduced in recent work on quantum linear system solvers, where it was employed to obtain an optimal query complexity of block encoding and improved complexity of initial state preparation [29], and subsequently to develop solvers whose cost is governed by instance-dependent quantities beyond the worst-case condition number [13, 25]. At a high level, the idea is to block encode a *scaling operator* and bundle it together with the input matrix, so that the composite block encoding represents a rescaled version of the original. A judicious choice of scaling operator leaves the condition number of the matrix unchanged, while amplifying selected blocks—equivalently, reducing their normalization factors. When applied to the Fourier modes of the resolvent, it drives the normalization of the Π_{in} block encoding down from α_{res} toward its optimal value $\|\Pi_{\text{in}}\|$.

We use a two-sided variant of the block preconditioning. Let T_1 and T_2 be scaling operators, diagonal in the computational basis, that act nontrivially only on the top-left entry,

$$T_1 = \begin{bmatrix} t_1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{bmatrix}, \qquad T_2 = \begin{bmatrix} t_2 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{bmatrix}, \qquad 0 < t_1, t_2 < 1. \tag{20}$$

Rather than inverting M directly, we assemble a block encoding of the preconditioned matrix $(T_1 \otimes I)\,(\mathbf{QFT} \otimes I)\,M\,\left(\mathbf{QFT}^{\dagger} \otimes I\right)(T_2 \otimes I)$, which is available at no additional query cost, since the quantum Fourier transform $\mathbf{QFT}$ and the diagonal scalings T_1, T_2 can all be block encoded with normalization factor one. Applying QSVT to invert it then yields a block encoding of

$$\left(T_2^{-1} \otimes I\right)(\mathbf{QFT} \otimes I)\,M^{-1}\left(\mathbf{QFT}^{\dagger} \otimes I\right)\left(T_1^{-1} \otimes I\right) \;=\; \begin{bmatrix} \frac{1}{t_1 t_2} S_0 & \frac{1}{t_2} S_{n-1} & \frac{1}{t_2} S_{n-2} & \cdots & \frac{1}{t_2} S_1 \\ \frac{1}{t_1} S_1 & S_0 & S_{n-1} & \cdots & S_2 \\ \frac{1}{t_1} S_2 & S_1 & S_0 & \cdots & S_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{t_1} S_{n-1} & S_{n-2} & S_{n-3} & \cdots & S_0 \end{bmatrix}. \tag{21}$$

This implies

$$\begin{aligned}\alpha_{\text{precond}} &= \left\| \left(T_2^{-1} \otimes I\right) \left(\mathbf{QFT} \otimes I\right) M^{-1} \left(\mathbf{QFT}^\dagger \otimes I\right) \left(T_1^{-1} \otimes I\right) \right\| \\ &\leq \frac{1}{t_1 t_2} \|S_0\| + \frac{\sqrt{n-1}}{t_2} \max_m \|S_m\| + \frac{\sqrt{n-1}}{t_1} \max_m \|S_m\| + \left\|M^{-1}\right\| \\ &= \mathbf{O}\left(\frac{\|\Pi_{\text{in}}\|}{t_1 t_2} + \left(\frac{\sqrt{n}}{t_1} + \frac{\sqrt{n}}{t_2}\right) \left(\alpha_{\text{in}} \|\Pi_{\text{in}}\| + \alpha_{\text{out}} \|\Pi_{\text{out}}\|\right) + \alpha_{\text{res}} \right).\end{aligned} \tag{22}$$

Here, the four terms bound the four structural parts of the inverse matrix: the top-left entry $\frac{1}{t_1 t_2} S_0$, the remainder of the top row (scaled by $\frac{1}{t_2}$), the remainder of the left column (scaled by $\frac{1}{t_1}$), and the bottom-right $(n-1) \times (n-1)$ block of unscaled modes. The bottom-right block has norm at most $\left\|M^{-1}\right\| = \max_{0 \leq j \leq n-1} \left\|(I - \omega_n^{-j} A)^{-1}\right\| \leq \alpha_{\text{res}}$ due to the resolvent bound. The top row and left column are controlled by the individual modes, which satisfy $\max_m \|S_m\| = \mathbf{O}(\alpha_{\text{in}} \|\Pi_{\text{in}}\| + \alpha_{\text{out}} \|\Pi_{\text{out}}\|)$, while the top-left entry is exponentially close to $\|\Pi_{\text{in}}\|$.

Under the spectral-set assumption of the Results, the Crouzeix–Palencia bound [12] applied on each subspace gives $\left\|A_{\text{in}}^j\right\| \leq (1+\sqrt{2})(1-\delta)^j$ and $\left\|A_{\text{out}}^{-k}\right\| \leq (1+\sqrt{2})(1-\delta)^k$, so that $\alpha_{\text{in}}, \alpha_{\text{out}} = \mathbf{O}(1)$. Summing the geometric series bounds the resolvent likewise,

$$\alpha_{\text{res}} = \mathbf{O}\left(\frac{\|\Pi_{\text{in}}\|}{\delta}\right), \tag{23}$$

and the discretization error $\|S_0 - \Pi_{\text{in}}\| = \mathbf{O}((1-\delta)^n \|\Pi_{\text{in}}\|)$ falls below ε as long as $n = \mathbf{O}\left(\frac{1}{\delta} \log \frac{\|\Pi_{\text{in}}\|}{\varepsilon}\right)$. Substituting these bounds with the balanced choice $t_1 = t_2 = \mathbf{\Theta}\left(\sqrt{\delta}\right)$, the preconditioned inverse has norm at most

$$\alpha_{\text{precond}} = \mathbf{O}\left(\frac{\|\Pi_{\text{in}}\|}{\delta} \sqrt{\log \frac{\|\Pi_{\text{in}}\|}{\varepsilon}}\right), \tag{24}$$

which sets the condition number for the QSVT inversion. Extracting the top-left entry and rescaling, the resulting block encoding of Π_{in} has reduced normalization factor

$$\alpha_\Pi = t_1 t_2 \alpha_{\text{precond}} = \mathbf{O}\left(\|\Pi_{\text{in}}\| \sqrt{\log \frac{\|\Pi_{\text{in}}\|}{\varepsilon}}\right). \tag{25}$$

The full treatment, including the general spectral-set formulation and the diagonalizable case, is given in the Supplementary Information.

General regions and the half-plane construction

The reduction behind Theorem 3 relies on the holomorphic functional calculus: for a closed curve $\mathcal{C}$ enclosing $\mathbf{Spec}_{\text{in}}(A)$ but not $\mathbf{Spec}_{\text{out}}(A)$ and disjoint from $\mathbf{Spec}(A)$, the change-of-variables identity (14) holds because A and $f(A)$ share the same invariant subspaces while f carries $\mathbf{Spec}_{\text{in}}(A)$ inside the unit circle and $\mathbf{Spec}_{\text{out}}(A)$ outside, by the spectral mapping theorem [19, Theorem 12.7.1][36, Theorem 8.3]. The general-region problem thus reduces to the unit-disk construction applied to $f(A)$ in place of A.

Because f need not be injective, conformal, or boundary-preserving, we may separate the design of the algorithm from the construction of f. Treating $f(A)$ as an oracle, the construction above applies verbatim with A replaced by $f(A)$, yielding a block encoding of Π_{in} with normalization factor close to $\|\Pi_{\text{in}}\|$ and query complexity governed by the gap of $f(A)$ to the unit circle. The

oracle model requires only a weak analyticity hypothesis: because $f(A)$ is given as input rather than approximated, f need only be analytic on neighborhoods of the two block numerical ranges $\mathcal{W}(A_{\text{in}})$ and $\mathcal{W}(A_{\text{out}})$ *separately*, and not on the full numerical range $\mathcal{W}(A)$. Since $\mathcal{W}(A)$ is convex and generally strictly larger than the two block ranges, a singularity of f may lie inside $\mathcal{W}(A)$ while remaining outside each block range. The oracle model therefore admits any separating map whose singularities fall between the blocks, including maps that an explicit polynomial approximation of $f(A)$ over the whole numerical range could not accommodate.

We now construct a block encoding of $f(A)$ from a block encoding of A. Let $\mathcal{K}$ be a spectral set for A with constant c, and let f be analytic on $\mathcal{K}$ and separate the spectrum,

$$|f(w)| < 1 \quad (w \in \mathbf{Spec}_{\text{in}}(A)), \qquad |f(w)| > 1 \quad (w \in \mathbf{Spec}_{\text{out}}(A)). \tag{26}$$

Approximating f by a polynomial p and applying the spectral-set definition to $g = f - p$ gives $\|f(A) - p(A)\| \leq c \sup_{z\in\mathcal{K}} |f(z) - p(z)|$, so a polynomial that approximates f uniformly on $\mathcal{K}$ yields a block encoding of $f(A)$ to matching accuracy. A common choice is the numerical range $\mathcal{K} = \mathcal{W}(A)$, a spectral set for every A with the dimension-independent constant $c = 1 + \sqrt{2}$ by the Crouzeix–Palencia theorem [12]. The polynomial p is applied by the QEVT algorithm [28, Section 8 and Appendix B] at a cost of $\mathbf{O}(d)$ queries to the block encoding of A, where d is its degree. The required degree follows from the Bernstein–Walsh theorem [40, Section III.3]: if f is analytic on a neighborhood of $\mathcal{W}(A)$, the best degree-d polynomial approximation on $\mathcal{W}(A)$ decays geometrically, so $d = \mathbf{O}(\log(1/\varepsilon))$ suffices for accuracy ε, with the rate set by how far the nearest singularity of f lies beyond $\mathcal{W}(A)$. This generalizes the approximation of a function on $[-1, 1]$ through its analyticity within a Bernstein ellipse [44, Theorem 8.1 and 8.2].

For the half-plane, suppose $\mathbf{Spec}_{\text{in}}(A)$ lies in the open right half-plane and $\mathbf{Spec}_{\text{out}}(A)$ in the open left half-plane, separated by the imaginary axis. The separating map (17), $f(z) = (z-1)^2/(z+1)^2$, has a double zero at $z = 1$ and a double pole at $z = -1$, and sends the right half-plane into the unit disk, the imaginary axis onto the unit circle, and the left half-plane outside. Assume a block encoding of A is given with $\|A\| \leq \frac{1}{2}$. Then $A + I$ is well conditioned: its singular values lie in $[\frac{1}{2}, \frac{3}{2}]$, since $\|A + I\| \leq \frac{3}{2}$ and $\left\|(A + I)^{-1}\right\| \leq (1 - \|A\|)^{-1} \leq 2$, a condition number at most 3. QSVT-based matrix inversion therefore block encodes $(A+I)^{-1}$ to error ε with normalization $\mathbf{O}(1)$ using $\mathbf{O}(\log(1/\varepsilon))$ queries to the block encoding of A. Applying $(A+I)^{-1}$ twice and multiplying by $(A - I)^2$ yields a block encoding of $f(A) = (A - I)^2(A + I)^{-2}$ with normalization $\|f(A)\| = \mathbf{O}(1)$, again at $\mathbf{O}(\log(1/\varepsilon))$ query cost. Thus the map f is applied directly, with only a constant-factor overhead in normalization and a logarithmic query cost.

The map must also preserve the gap. Suppose the spectrum is separated from the imaginary axis by δ, in the sense that $|\Re(\lambda)| \geq \delta$ for every $\lambda \in \mathbf{Spec}(A)$. Under $\|A\| \leq \frac{1}{2}$, the image spectrum $\mathbf{Spec}(f(A))$ is separated from the unit circle by a gap $\delta_f = \mathbf{\Theta}(\delta)$, so the map preserves the spectral gap up to a constant factor. The same holds for the numerical range: if the per-block numerical ranges $\mathcal{W}(A_{\text{in}})$ and $\mathcal{W}(A_{\text{out}})$ are separated from the imaginary axis by δ, then $|f|$ is at most $1 - \mathbf{\Theta}(\delta)$ on $\mathcal{W}(A_{\text{in}})$ and at least $1 + \mathbf{\Theta}(\delta)$ on $\mathcal{W}(A_{\text{out}})$, so applying the Crouzeix–Palencia bound to the composite function over each per-block numerical range yields the desired estimates with a comparable gap. The same argument applies to any spectral sets for A_{in} and A_{out} separated from the imaginary axis by δ, with the Crouzeix–Palencia constant replaced by the spectral-set constant c. Consequently the cost bounds apply to $f(A)$ with δ replaced by $\delta_f = \mathbf{\Theta}(\delta)$. The detailed treatment is given in the Supplementary Information.

Applications

The oblique eigenprojection is a general-purpose primitive, which we now instantiate in three settings. In each, two-sided block preconditioning reduces the normalization factor of the block-encoded projector to near its optimal value, and this reduction carries through to a lower overall cost, sharpening a recent quantum algorithm.

Eigenstate preparation for complex spectra

When the spectrum of A lies on the real line, eigenstate preparation can be accomplished by the QEVT algorithm of Low and Su [28], which applies a polynomial transformation to the eigenvalues of A through the Chebyshev and Faber history states. Choosing the polynomial to be close to one near the target eigenvalue and close to zero elsewhere prepares the corresponding eigenstate at a cost scaling as $1/\delta$ in the spectral gap, recovering the nearly optimal scaling of ground-state preparation for Hermitian Hamiltonians. This filtering is effective because the target eigenvalues lie on a line, where a polynomial can approximate an indicator. For eigenvalues distributed in the complex plane there is an intrinsic obstruction: when the target region is surrounded by the rest of the spectrum, the maximum modulus principle forces any polynomial—indeed any analytic function—to attain its maximum magnitude on the enclosing boundary, so it cannot be close to one inside the region and close to zero just outside it [11]. Preparing an eigenstate of a matrix with complex eigenvalues therefore calls for a different mechanism. One may instead form the Riesz projector through the resolvent integral, as done by Rodenas-Ruiz, Zhao, and Lee [35]; this bypasses the analyticity requirement and applies to complex spectra. Its cost, however, scales as $1/\delta^2$: the block-encoded projector carries a normalization factor $\alpha_{\text{res}} = \mathbf{O}(\|\Pi_{\text{in}}\| / \delta)$ set by the resolvent norm on the contour, and this leads the $1/\delta$ factor to enter the eigenstate-preparation cost twice—once in the $\mathbf{O}(1/\delta)$ queries needed to block encode the projector, and again in the $\mathbf{O}(1/\delta)$ rounds of amplitude amplification needed to project the initial state onto the target subspace.

Our algorithm applies oblique eigenprojection at a cost matching the result of Low and Su while handling complex spectra. Suppose the target eigenvalue or cluster is enclosed by a disk $\{z : |z - z_0| < \rho\}$ that separates it from the rest of the spectrum with a gap δ, where z_0 and ρ are considered constants, independent of δ. The separating map is the linear function $f(z) = (z - z_0)/\rho$, so $f(A) = (A - z_0 I)/\rho$ is obtained from the block encoding of A by a shift and a rescaling, at only a constant-factor overhead in normalization for a center $|z_0| = \mathbf{O}(1)$. Feeding $f(A)$ into the unit-disk construction block encodes the eigenprojection Π_{in} onto the enclosed invariant subspace with a query complexity scaling as $1/\delta$ and, after two-sided block preconditioning, a normalization factor close to its optimal value $\|\Pi_{\text{in}}\|$. Because the normalization no longer grows with the inverse gap, the amplitude amplification that projects the initial state onto the target subspace needs only $\widetilde{\mathbf{O}}(1)$ rounds, and the eigenstate is prepared in $\widetilde{\mathbf{O}}(1/\delta)$ total queries. This quadratically improves the resolvent-integral approach and matches the scaling in $1/\delta$ of Low and Su for nonnormal matrices with real spectra—preparing an eigenstate associated with a complex eigenvalue is no more expensive, in δ, than filtering a real one. The full analysis is given in the Supplementary Information.

Continuous-time algebraic Riccati equations

The stabilizing solution of the Riccati equation is characterized through an invariant subspace of the associated Hamiltonian matrix

$$J = \begin{bmatrix} P & -Q \\ -R & -P^\dagger \end{bmatrix}, \tag{27}$$

whose spectrum is symmetric about the imaginary axis. Splitting the spectrum along the imaginary axis into a stable part ($\Re(\lambda) < 0$) and an antistable part ($\Re(\lambda) > 0$), the stable invariant subspace is spanned by the columns of $\left[\begin{smallmatrix} U \\ V \end{smallmatrix}\right]$, and the stabilizing solution is recovered as $X = VU^{-1}$. Solving the Riccati equation thus reduces to projecting onto an invariant subspace of a non-normal matrix split by the imaginary axis—an oblique eigenprojection of exactly the kind our primitive provides.

Rodenas-Ruiz, Zhao, and Lee [35] solve the Riccati equation on a quantum computer along these lines. The stable invariant subspace spanned by $\left[\begin{smallmatrix} I \\ X \end{smallmatrix}\right]$ is the kernel of the projector Π_{in} onto the antistable subspace, so writing $\Pi_{\text{in}} = \left[\begin{smallmatrix} \Pi_1 & \Pi_2 \end{smallmatrix}\right]$ gives $\Pi_1 + \Pi_2 X = 0$ and hence $X = -\Pi_2^+ \Pi_1$, where $(\cdot)^+$ denotes the pseudoinverse. They therefore block encode the antistable projector Π_{in} through the resolvent integral, and then extract X from its sub-blocks by implementing Π_2^+ through QSVT. Writing $\delta = \min_\lambda |\Re(\lambda)|$ for the gap from the imaginary axis, the block-encoded projector carries a normalization factor $\alpha_{\text{res}} = \mathbf{O}(\|\Pi_{\text{in}}\| / \delta)$ set by the resolvent norm on the contour. Beyond the $\mathbf{O}(1/\delta)$ queries needed to block encode the projector, this normalization enters the extraction twice more—in the condition number of the pseudoinverse that forms $X = -\Pi_2^+ \Pi_1$, and in the normalization of the resulting solution—so that preparing the Riccati solution costs $\widetilde{\mathbf{O}}(1/\delta^3)$ queries in the gap, ignoring polylogarithmic factors.

Applying our oblique eigenprojection removes the two normalization-driven factors. Since the spectrum is split by the imaginary axis, we use the half-plane construction, with a degree-two rational map sending the antistable half-plane into the unit disk and the imaginary axis onto the unit circle, and the per-block numerical-range hypotheses inherited with a transformed gap $\delta_f = \mathbf{\Theta}(\delta)$. Two-sided block preconditioning then block encodes the projector with a normalization factor close to its optimal value $\|\Pi_{\text{in}}\|$, in place of $\alpha_{\text{res}} = \mathbf{O}(\|\Pi_{\text{in}}\| / \delta)$. The pseudoinverse condition number and the solution normalization then no longer grow with the inverse gap, so the two extra factors collapse to $\widetilde{\mathbf{O}}(1)$ and the Riccati solution is prepared in $\widetilde{\mathbf{O}}(1/\delta)$ queries—a cubic improvement over the standard resolvent-integral solver of [35]. The detailed analysis is given in the Supplementary Information.

Ordinary Sylvester equations

As with the Riccati equation, the solution of the Sylvester equation $AX + XB = C$ is encoded in an invariant subspace of an associated block matrix:

$$K = \begin{bmatrix} A & C \\ 0 & -B \end{bmatrix}. \tag{28}$$

Its spectrum is the union of the spectra of A and $-B$, and it has two invariant subspaces: the A branch is spanned by $\left[\begin{smallmatrix} I \\ 0 \end{smallmatrix}\right]$, while the $-B$ branch is spanned by $\left[\begin{smallmatrix} -X \\ I \end{smallmatrix}\right]$, since the invariance condition $K \left[\begin{smallmatrix} -X \\ I \end{smallmatrix}\right] = \left[\begin{smallmatrix} -X \\ I \end{smallmatrix}\right](-B)$ is precisely the Sylvester equation $AX + XB = C$. Stacking these bases into $S = \left[\begin{smallmatrix} I & -X \\ 0 & I \end{smallmatrix}\right]$, the oblique projector onto the A branch along the B branch is $\Pi_{\text{in}} = S \left[\begin{smallmatrix} I & 0 \\ 0 & 0 \end{smallmatrix}\right] S^{-1} = \left[\begin{smallmatrix} I & X \\ 0 & 0 \end{smallmatrix}\right]$, so the solution X appears directly in its off-diagonal block. When the spectra of A and $-B$ lie in opposite open half-planes—separated, after a shift and rotation, by the imaginary axis—this is again an oblique eigenprojection handled by our algorithm, with X read off from a single block of the projector. Unlike the Riccati equation, the extraction requires no further pseudoinverse.

Wang and Liu [48] give a quantum algorithm for the Sylvester equation along these lines. We compare against their direct augmented method [48, Remark 6.4], which block encodes the projector through the resolvent integral applied to the augmented matrix K, so that the block-encoded projector carries a normalization factor $\alpha_{\text{res}} = \mathbf{O}(\|\Pi_{\text{in}}\| / \delta)$ set by the resolvent norm on the contour, with δ the gap from the imaginary axis. Since the solution is the off-diagonal block

of the projector, no pseudoinverse is needed, and this normalization enters the cost only once, for amplifying the block encoding, on top of the $\mathbf{O}(1/\delta)$ queries needed to block encode the projector, so that preparing the solution costs $\widetilde{\mathbf{O}}(1/\delta^2)$ queries in the gap.

Our oblique eigenprojection applies directly and improves on this complexity. Using the half-plane construction to project onto the invariant subspace of the A branch, two-sided block preconditioning replaces the normalization factor $\alpha_{\text{res}} = \mathbf{O}(\|\Pi_{\text{in}}\|/\delta)$ with one close to its optimal value $\|\Pi_{\text{in}}\| = \sqrt{\|X\|^2 + 1}$, and the extra inverse-gap factor collapses to $\widetilde{\mathbf{O}}(1)$. One directly obtains a near-normalized block encoding of X, and the $\mathbf{O}(1/\delta)$ cost of amplifying it is avoided. The total query complexity to obtain the solution is $\widetilde{\mathbf{O}}(1/\delta)$—a quadratic improvement over the direct augmented method of [48, Remark 6.4]. The detailed analysis is given in the Supplementary Information.

References

[1] Daniel S. Abrams and Seth Lloyd, *Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors*, Physical Review Letters **83** (1999), 5162–5165, arXiv:quant-ph/9807070, DOI:10.1103/PhysRevLett.83.5162.

[2] 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.

[3] Dong An, Andrew M. Childs, Lin Lin, and Lexing Ying, *Laplace transform–based quantum eigenvalue transformation via linear combination of Hamiltonian simulation*, SIAM Journal on Computing **55** (2026), no. 2, 376–409, arXiv:2411.04010, DOI:10.1137/24M1720160.

[4] Yuto Ashida, Zongping Gong, and Masahito Ueda, *Non-Hermitian physics*, Advances in Physics **69** (2020), no. 3, 249–435, arXiv:2006.01837, DOI:10.1080/00018732.2021.1876991.

[5] Yosi Atia and Dorit Aharonov, *Fast-forwarding of Hamiltonians and exponentially precise measurements*, Nature Communications **8** (2017), no. 1, 1572, arXiv:1610.09619, DOI:10.1038/s41467-017-01637-7.

[6] Avah Banerjee, Asim Sharma, and Sooraj Soman, *Quantum fast-forwarding beyond reversibility: The α-perturbed n-cycle*, 2026, arXiv:2606.26584, DOI:10.48550/arXiv.2606.26584.

[7] 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.

[8] 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.

[9] Baptiste Claudon, Jean-Philip Piquemal, and Pierre Monmarché, *Quantum speedup for non-reversible Markov chains*, Nature Communications **16** (2025), no. 1, 10732, arXiv:2501.05868, DOI:10.1038/s41467-025-65761-5.

[10] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca, *Quantum algorithms revisited*, Proceedings of the Royal Society A **454** (1998), no. 1969, 339–354, arXiv:quant-ph/9708016, DOI:10.1098/rspa.1998.0164.

[11] John D. Cook, *Not even close*, https://www.johndcook.com/blog/2020/07/24/not-even-close/, 2020, Accessed 2026-09-20.

[12] M. Crouzeix and C. Palencia, *The numerical range is a* $(1+\sqrt{2})$*-spectral set*, SIAM Journal on Matrix Analysis and Applications **38** (2017), no. 2, 649–655, arXiv:1702.00668, DOI:10.1137/17M1116672.

[13] Alexander M. Dalzell, Jianqiang Li, and Yuan Su, *Faster quantum linear system solver beyond the condition number*, 2026, arXiv:2607.07691, DOI:10.48550/arXiv.2607.07691.

[14] Yasunori Futamura, Xiucai Ye, and Tetsuya Sakurai, *Contour integral-based quantum algorithm for estimating matrix eigenvalue density*, 2021, arXiv:2112.05395, DOI:10.48550/arXiv.2112.05395.

[15] Yimin Ge, Jordi Tura, and J. Ignacio Cirac, *Faster ground state preparation and high-precision ground energy estimation with fewer qubits*, Journal of Mathematical Physics **60** (2019), no. 2, 022202, arXiv:1712.03193, DOI:10.1063/1.5027484.

[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, DOI:10.1145/3313276.3316366.

[17] Lov K. Grover, *A fast quantum mechanical algorithm for database search*, Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing (STOC), pp. 212–219, 1996, arXiv:quant-ph/9605043, DOI:10.1145/237814.237866.

[18] Xabier Gutiérrez, Lorenzo Laneve, and Mikel Sanz, *Quantum eigenvalue transformations for arbitrary matrices*, 2026, arXiv:2604.19688, DOI:10.48550/arXiv.2604.19688.

[19] Jeffrey Humpherys, Tyler J. Jarvis, and Emily J. Evans, *Foundations of applied mathematics, volume 1: Mathematical analysis*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017, DOI:10.1137/1.9781611974904.

[20] Shan Jiang and Dong An, *Contour-integral based quantum eigenvalue transformation: analysis and applications*, 2026, arXiv:2601.11959, DOI:10.48550/arXiv.2601.11959.

[21] Christopher Kang and Yuan Su, *Quantum matrix arithmetics with Hamiltonian evolution*, ACM Transactions on Quantum Computing **7** (2026), no. 4, 1–86, arXiv:2510.06316, DOI:10.1145/3816436.

[22] A Yu Kitaev, *Quantum measurements and the abelian stabilizer problem*, 1995, arXiv:quant-ph/9511026, DOI:10.48550/arXiv.quant-ph/9511026.

[23] Anthony W. Knapp, *Basic real analysis*, digital second ed., Birkhäuser, 2016, DOI:10.3792/euclid/9781429799997.

[24] Hari Krovi, *Improved quantum algorithms for linear and nonlinear differential equations*, Quantum **7** (2023), 913, arXiv:2202.01054, DOI:10.22331/q-2023-02-02-913.

[25] Jianqiang Li, *A new quantum linear system algorithm beyond the condition number and its application to solving multivariate polynomial systems*, 2025, arXiv:2510.05588, DOI:10.48550/arXiv.2510.05588.

[26] Lin Lin and Yu Tong, *Near-optimal ground state preparation*, Quantum **4** (2020), 372, arXiv:2002.12508, DOI:10.22331/q-2020-12-14-372.

[27] Guang Hao Low and Isaac L. Chuang, *Optimal Hamiltonian simulation by quantum signal processing*, Physical Review Letters **118** (2017), 010501, arXiv:1606.02685, DOI:10.1103/PhysRevLett.118.010501.

[28] Guang Hao Low and Yuan Su, *Quantum eigenvalue processing*, SIAM Journal on Computing **55** (2026), no. 1, 135–215, arXiv:2401.06240, DOI:10.1137/24M1689363.

[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, DOI:10.22331/q-2026-03-23-2041.

[30] Guang Hao Low and Nathan Wiebe, *Hamiltonian simulation in the interaction picture*, 2018, arXiv:1805.00675, DOI:10.48550/arXiv.1805.00675.

[31] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang, *Grand unification of quantum algorithms*, PRX Quantum **2** (2021), 040203, arXiv:2105.02859, DOI:10.1103/PRXQuantum.2.040203.

[32] Sam McArdle and David P. Tew, *Improving the accuracy of quantum computational chemistry using the transcorrelated method*, 2020, arXiv:2006.11181, DOI:10.48550/arXiv.2006.11181.

[33] Ashley Montanaro and Changpeng Shao, *Quantum and classical query complexities of functions of matrices*, Proceedings of the 56th Annual ACM Symposium on Theory of Computing, pp. 573–584, 2024, arXiv:2311.06999, DOI:10.1145/3618260.3649665.

[34] Hongkang Ni and Lexing Ying, *Quantum eigenvalue transformation via linear combination of Hamiltonian simulation: A Weyl calculus approach*, 2026, arXiv:2606.29848, DOI:10.48550/arXiv.2606.29848.

[35] Pablo Rodenas-Ruiz, Andrew Zhao, and Joonho Lee, *Quantum solvers for nonlinear matrix equations in quantum chemistry*, 2026, arXiv:2605.16189, DOI:10.48550/arXiv.2605.16189.

[36] Steven Roman, *Advanced linear algebra*, 3 ed., Graduate Texts in Mathematics, vol. 135, Springer, New York, 2008, DOI:10.1007/978-0-387-72831-5.

[37] Changpeng Shao, *Computing eigenvalues of diagonalizable matrices on a quantum computer*, ACM Transactions on Quantum Computing **3** (2022), no. 4, 1–20, arXiv:1912.08015, DOI:10.1145/3527845.

[38] Changpeng Shao and Jin-Peng Liu, *Solving generalized eigenvalue problems by ordinary differential equations on a quantum computer*, Proceedings: Mathematical, Physical and Engineering Sciences **478** (2022), no. 2262, 1–25, arXiv:2010.15027, DOI:10.1098/rspa.2021.0797.

[39] Peter W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM Journal on Computing **26** (1997), no. 5, 1484–1509, arXiv:quant-ph/9508027, DOI:10.1137/S0097539795293172.

[40] P. K. Suetin and EV Pankratiev, *Series of Faber polynomials*, vol. 1, CRC Press, 1998.

[41] Souichi Takahira, Asuka Ohashi, Tomohiro Sogabe, and Tsuyoshi S. Usuda, *Quantum algorithm for matrix functions by Cauchy's integral formula*, Quantum Information and Computation **20** (2020), 14–36, arXiv:2106.08075, DOI:10.26421/QIC20.1-2-2.

[42] Souichi Takahira, Asuka Ohashi, Tomohiro Sogabe, and Tsuyoshi S Usuda, *Quantum algorithms based on the block-encoding framework for matrix functions by contour integrals*, Quantum Information and Computation **22** (2022), no. 11&12, 965–979, arXiv:2106.08076, DOI:10.26421/QIC22.11-12-4.

[43] 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, DOI:10.1103/PhysRevA.104.032422.

[44] Lloyd N Trefethen, *Approximation theory and approximation practice, extended edition*, Society for Industrial and Applied Mathematics, 2019, DOI:10.1137/1.9781611975949.

[45] Francisca Vasconcelos and András Gilyén, *Methods for reducing ancilla-overhead in block encodings*, 2025, arXiv:2507.07900, DOI:10.48550/arXiv.2507.07900.

[46] Chao Wang, Xi-Ning Zhuang, Menghan Dou, Zhao-Yun Chen, and Guo-Ping Guo, *A unified Poisson summation framework for generalized quantum matrix transformations*, 2026, arXiv:2604.02874, DOI:10.48550/arXiv.2604.02874.

[47] Yanqiao Wang, Yixuan Liang, Hongjia Chen, and Jin-Peng Liu, *Optimal quantum eigenvalue transformation via linear combinations of Hermitian matrices*, 2026, arXiv:2607.25812, DOI:10.48550/arXiv.2607.25812.

[48] Yanqiao Wang and Jin-Peng Liu, *Sign embedding quantum algorithms for matrix equations and matrix functions*, 2026, arXiv:2604.25333, DOI:10.48550/arXiv.2604.25333.

[49] Xiao-Ming Zhang, Yukun Zhang, Wenhao He, and Xiao Yuan, *Exponential quantum advantages for practical non-Hermitian eigenproblems*, Physical Review Letters **135** (2025), 140601, arXiv:2401.12091, DOI:10.1103/3n8f-k8pl.

Data availability

No datasets were generated or analysed during the current study.

Code availability

This study did not involve custom code or software central to its conclusions.

Use of large language models

The authors used a large language model (Anthropic Claude, accessed through Amazon Quick) as a writing and drafting aid during the preparation of this manuscript: to draft and revise expository text, reorganize sections, and check internal consistency of notation and cross-references. The model was used under the authors' direction and supervision. All mathematical results, proofs,

and claims were conceived, derived, and verified by the authors, who take full responsibility for the content of the work.

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, and Lin Lin for reviewing an earlier draft.

Author contributions

A.M.D. contributed technical ideas to improve the complexity, gave feedback on the draft, and verified the results for correctness. Y.S. conceived the research question, carried out the analysis, and wrote the manuscript.

Competing interests

The authors declare no competing interests.

Funding

This work received no external funding.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- obliqueprojectionsi.pdf