Letter

Quantum eigenvalue transformations for arbitrary matrices

Xabier Gutiérrez *

Department of Physical Chemistry, University of the Basque Country UPV/EHU, Apartado 644, 48080 Bilbao, Spain and EHU Quantum Center, University of the Basque Country UPV/EHU, Apartado 644, 48080 Bilbao, Spain

Lorenzo Laneve

Faculty of Informatics—Università della Svizzera Italiana, 6900 Lugano, Switzerland

Mikel Sanz

Department of Physical Chemistry, University of the Basque Country UPV/EHU, Apartado 644, 48080 Bilbao, Spain; EHU Quantum Center, University of the Basque Country UPV/EHU, Apartado 644, 48080 Bilbao, Spain; and Basque Center for Applied Mathematics (BCAM), Alameda de Mazarredo 14, 48009 Bilbao, Spain

(Received 24 April 2026; revised 24 July 2026; accepted 2 October 2026; published 16 November 2026)

Quantum signal processing (QSP) and quantum singular value transformation (QSVT) provide an efficient framework for implementing polynomials of block-encoded matrices, and thus offer a systematic approach to quantum algorithm design. However, despite a number of recent advances, important limitations remain. In particular, QSP can only transform unitary matrices, by applying a polynomial to their eigenvalues, while QSVT is a singular value transformation and thus one can only obtain the polynomial of Hermitian matrices. As a consequence, these techniques do not directly apply to an arbitrary non-Hermitian matrix that is not diagonalizable. In this work, we propose a simple yet powerful method to extend these ideas to arbitrary square matrices by acting on their eigenvalues. To this end, we introduce the notion of an n-regular block encoding, namely, a block encoding whose kth power reproduces the kth power of the encoded matrix for every $0 \leqslant k \leqslant n$. We show that applying QSP to any unitary with this property is equivalent to applying a polynomial of degree at most n to the block-encoded matrix, independently of its internal structure. Moreover, we provide a simple construction that transforms any block encoding into an n-regular one using only $\mathcal{O}(\log n)$ ancillary qubits and $\mathcal{O}(n \log(n/\epsilon))$ operations. Finally, we show that this construction induces the desired transformation on the eigenvalues associated with the Jordan normal form of the matrix; for nondiagonalizable matrices, this yields the full matrix function $P(A)$, whose action on each Jordan block involves the derivatives of P, and not merely the map $\lambda \mapsto P(\lambda)$.

DOI: 10.1103/sm6w-5ttl

Introduction. Quantum algorithms have the potential to deliver substantial speedups over their classical counterparts, yet until very recently their design had no truly systematic framework and instead relied largely on case-by-case ingenuity. In this context, the method of quantum signal processing (QSP) [1,2] has emerged as a unifying paradigm, linking algorithm design with techniques from numerical analysis and, in particular, with the implementation of polynomials of matrices encoded through unitary transformations. This perspective enabled a more abstract and systematic approach to algorithm design, and it has already been applied in a variety of settings, including Hamiltonian simulation [1,3], quantum search [2], quantum linear systems [4], and quantum state preparation [5–8].

In QSP, one has access to a *signal* z—typically a complex number of unit modulus or, depending on the convention, a real number in $[-1, 1]$—encoded in a single-qubit unitary known as the *signal* operator. By interleaving applications of this operator with a sequence of operations that is independent of z—the *processing* operators—one can construct essentially any polynomial transformation $P(z) = \sum_{k=0}^{n} a_k z^k$ of the signal, which may, for instance, appear in the upper-left entry of the resulting unitary [9–11]. Moreover, efficient and numerically stable methods are known for constructing a single-qubit protocol realizing a prescribed polynomial P [12–15]. By leveraging the spectral theorem, this construction can be lifted so that the scalar signal z is replaced by a unitary operator U [11]. The resulting circuit then implements the nonunitary transformation $P(U) = \sum_{k=0}^{n} a_k U^k$, thereby applying P directly to the eigenvalues of U in a *black-box* fashion, that is, without the need to explicitly construct or estimate the corresponding eigenstates. Remarkably, this requires only $n + 1$ single-qubit operations and n controlled uses of U.

On the other hand, when a matrix A, not necessarily square, appears in the upper-left block of a unitary U_A that we can implement—a so-called *block encoding*—one can

*Contact author: xabier.gutierrez@ehu.eus

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

similarly apply a QSP polynomial to its singular values by a closely related construction, known as quantum singular value transformation (QSVT) [10,16]. This method has received particular attention because of its practical usefulness. It is important, however, to note that a polynomial of a square matrix does not, in general, coincide with the polynomial transformation of its singular values.

Despite their rapid development, these techniques still suffer from important limitations. In particular, the existing frameworks can perform either an eigenvalue transformation of a unitary matrix or a singular value transformation of a block-encoded matrix A, but they do not directly provide an eigenvalue transformation of a block-encoded matrix, except in the Hermitian case, where singular values and eigenvalues coincide up to sign. Consequently, for a general non-Hermitian matrix, neither method can be used directly to implement its polynomial transformation.

Several approaches beyond the QSP framework have been proposed to partially bridge this gap, especially the construction of Chebyshev and Faber history states [17], and the method of linear combination of Hamiltonian simulations (LCHS) [18], which has been shown to be useful for the simulation of nonunitary dynamics [19,20]. In particular, the latter approach has been used to block encode eigenvalue transformations by expressing them as inverse Laplace transforms [21]. Recently, a Weyl calculus approach to LCHS has been proposed, extending the class of admissible matrices [22], an approach based on Cauchy's integral formula [23], and a linear combination of Hermitian matrices approach [24].

However, these methods rely on the linear combination of unitaries (LCU) technique [25], which predates QSP and exhibits several limitations. First, the complexity of implementing a target transformation $f(z)$ depends strongly on its specific form and typically requires a nontrivial synthesis of the corresponding quantum circuit. Moreover, a linear combination of N terms generally requires an N-dimensional register, so the qubit overhead can grow rapidly, especially in the aforementioned approaches, where LCU is used to approximate complicated integrals to a prescribed precision. Another drawback of LCU-based methods is that the resulting block-encoded matrices can suffer from severe subnormalization—that is, instead of obtaining a block encoding of A, one obtains a block encoding of A/α for some potentially very large α. This leads to small success probabilities and, consequently, to further overhead in the overall complexity.

In this Letter, we propose a method to carry out transformations of eigenvalues, understood as solutions to the characteristic polynomials of a matrix, of an arbitrary square matrix, even nondiagonalizable ones. In order to do this, we introduce the concept of *n-regular* block encoding as a block encoding whose kth power reproduces the kth power of the encoded matrix for every $0 \leqslant k \leqslant n$, and prove that this kind of block encoding is always possible by adding only $\mathcal{O}(\log n)$ ancilla qubits and $\mathcal{O}(n \log(n/\epsilon))$ extra two-qubit gates. We then show that a standard application of QSP yields the eigenvalue transformation when this property holds. This allows us to obtain matrix transformations by remaining in this established framework.

This technique extends the QSP formalism to several applications, such as simulation of general inhomogeneous differential equations and nonunitary dynamics, which arise in a broad range of settings, including dissipative dynamics treated via Carleman linearization [26,27], partial differential equations [28], and open quantum systems [29], among others. In particular, since the cost of our construction for a matrix A is independent of the eigenvector conditioning $\kappa(S)$, where $A = SJS^{-1}$ is the Jordan decomposition of the matrix, it applies without penalty to defective generators (i.e., with high $\kappa(S)$) at exceptional points, a regime of intense current interest in non-Hermitian photonics, *PT*-symmetric and open quantum systems.

Block encodings and (generalized) quantum signal processing. Here, we summarize the main definitions and results that we are going to use in the rest of the work. Let us start by clarifying the notation: We will use $\mathbb{D}$ to denote the complex unit disk $\mathbb{D} = \{z \in \mathbb{C} : |z| < 1\}$, while $\mathbb{T} = \partial\mathbb{D}$ will be the unit circle. Additionally, if not said otherwise, $\|\cdot\|$ refers to the operator norm. Let us now provide the definition of a block encoding:

Definition 1. A unitary U is said to (a, ϵ) *block encode* a matrix A if

$$\|(\langle 0|^{\otimes a} \otimes I)U(|0\rangle^{\otimes a} \otimes I) - A\| \leqslant \epsilon, \tag{1}$$

with a the number of ancillary qubits needed for the block encoding. In other words, the unitary is of the form

$$U = \begin{pmatrix} \tilde{A} & \cdot \\ \cdot & \cdot \end{pmatrix} \tag{2}$$

for some $\tilde{A}$ satisfying $\|\tilde{A} - A\| \leqslant \epsilon$.

Notice that any unitary (0, 0) block encodes itself. Moreover, any block-encoded matrix A must necessarily satisfy $\|A\| \leqslant 1$ in order for the unitarity condition to hold for U.

Within the block-encoding framework, we aim to identify a set of operations that can be used to compose block encodings and thereby obtain more *useful* matrices from those that are easier to construct. A powerful tool for this purpose is generalized quantum signal processing (GQSP). Throughout, $\tilde{w}$ denotes the signal operator, defined by $\tilde{w} = \mathrm{diag}(1, z)$, which encodes a complex number $z \in \mathbb{T}$.

Theorem 1 (GQSP theorem, adapted from Ref. [11]). Let $P(z)$ be a polynomial in z of degree n satisfying $|P(z)| \leqslant 1$ on $\mathbb{T}$. There exists a sequence of single-qubit operations $R_0, R_1, \ldots, R_n \in SU(2)$ such that

$$\langle 0|R_0 \tilde{w} R_1 \tilde{w} \cdots \tilde{w} R_n|0\rangle = P(z). \tag{3}$$

Moreover, it is always possible, without loss of generality, to choose $R_1, \ldots, R_n$ to be a product of only two Pauli rotations (e.g., X and Z), whose angles are now easily computed with numerical algorithms [12–15].

Although previous works have only considered GQSP on the unit circle $\mathbb{T}$, any analytic function $f(z)$ satisfying $|f(z)| \leqslant 1$ on $\mathbb{T}$ automatically satisfies the same bound on $\mathbb{D}$, by virtue of the maximum modulus principle. This shows that, using existing techniques, we are already able to approximate any function $f(z)$ that is analytic and bounded by 1 on the unit *disk*.

A particularly simple construction enabled by GQSP is the transformation of the eigenvalues of a unitary. Specifically, if U is a $d \times d$ unitary, then its controlled version has the block-diagonal form

$$C(U) = \mathrm{diag}(\mathrm{I_d}, \mathrm{U}),$$

which closely resembles $\tilde{w}$. In particular, by applying the R_k to the control qubit and replacing $\tilde{w}$ with $C(U)$, the resulting circuit would simply block encode $P(U)$, that is, the matrix obtained by mapping each eigenvalue of U through the transformation $z \mapsto P(z)$.

Regular block encodings and regularization by "dumping branches". Current techniques based on quantum signal processing allow one to transform the eigenvalues of a unitary (as in the construction above) or the singular values of general block-encoded matrices (using QSVT), but no construction based on GQSP was previously known to successfully carry out *eigenvalue* transformations of generic block-encoded square matrices (except in the case of Hermitian matrices, where eigenvalue and singular value decompositions coincide and QSVT can be applied).

Unfortunately, directly performing an eigenvalue transformation $U \mapsto P(U)$ as we did above does not yield a block encoding of $P(A)$, because, in simple terms, U^k does not generally block encode A^k. Let us formally define this property:

Definition 2. An (a, ϵ) block encoding U of A is called n-regular if U^k $(a, k\epsilon)$ block encodes A^k for every $0 \leqslant k \leqslant n$.

The bound on the error implied by the definition simply comes from the bound $\|A^k - B^k\| \leqslant k\|A - B\|$. These n-regular block encodings show the following crucial property:

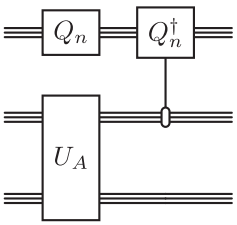
Lemma 1. Let $P(z)$ be any polynomial of degree n satisfying $|P(z)| \leqslant 1$ on $\mathbb{D}$. If U is a n-regular (a, ϵ) block encoding of A, then $P(U)$ is an $(a, \sqrt{n^3/3}\epsilon)$ block encoding of $P(A)$.

The proof of this lemma is given in the Supplemental Material [30], while a worked example on a 2-regular block encoding is found in the Supplemental Material [31]. This regularity notion is not particularly difficult to achieve, in the sense that any block encoding can be easily transformed into an n-regular block encoding via the simple construction shown in the following theorem:

Theorem 2. Let U_A be an (a, ϵ) block encoding of A, and suppose $n = 2^b$. Let Q_n denote the quantum incrementer, defined by

$$Q_n|x\rangle = |(x + 1) \bmod n\rangle. \tag{4}$$

Then, the following circuit yields an n-regular $(b + a, \epsilon)$ block encoding of A:



In other words, the circuit first applies A and then increments the newly introduced b-qubit register by 1 whenever the ancillary qubits are not in the all-zero state. The underlying idea is the following: When we apply a block encoding, we effectively implement the mapping

$$|0\rangle^{\otimes a}|\psi\rangle \mapsto |0\rangle^{\otimes a} \otimes A|\psi\rangle + \sum_{j\neq 0} |j\rangle \otimes |\psi_j^\perp\rangle, \tag{5}$$

for some garbage states $|\psi_j^\perp\rangle$. This implies that, within the "successful" subspace (i.e., the subspace on which A is applied), the associated unitary U leaves $|0\rangle^{\otimes a}$ invariant, whereas all other terms map the ancillary register to nonzero configurations. Consequently, if U is applied a second time, some of these nonzero configurations may be mapped back to $|0\rangle^{\otimes a}$, thereby contaminating the successful branch and resulting in the top-left block encoding A^2 plus an additional garbage term. With this construction, however, the incrementer causes the new b-qubit register to end up in a nonzero configuration; thus, even if subsequent calls to the block encoding map the original ancillas back to $|0\rangle^{\otimes a}$, the overall main branch associated with $|0\rangle^{\otimes b+a}$ remains clean. This remains true until Q_n is applied n times, after which some branch (in particular, the one that has failed n times) is mapped from $|1\rangle^{\otimes b}$ back to $|0\rangle^{\otimes b}$. This is still acceptable, since the original ancillas are not in the state $|0\rangle^{\otimes a}$; in other words, this yields an n-regular block encoding, but not an $(n + 1)$-regular one, because a branch of the form $|0\rangle^{\otimes b}|k\rangle^{\otimes a}$ can contaminate the main branch. A formal proof is provided in the Supplemental Material [32].

Note that this technique is also known as the compression gadget (Lemma 13 of Ref. [33]). The novelty here is the observation that the gadget-augmented circuit is an n-regular unitary.

Note that, within this approach, standard techniques for multiplying block encodings [34,35] can be directly applied. The incrementer requires only $\mathcal{O}(b) = \mathcal{O}(\log n)$ ancillary qubits and can be diagonalized by an approximate quantum Fourier transform [36], one at the beginning and another one at the end of the circuit, with a cost of $\mathcal{O}(b \log b) = \tilde{\mathcal{O}}(\log n)$ two-qubit gates. Taking into account each of the n controlled rotations, from which $\log(n/\epsilon)$ are retained, the overall two-qubit gates would be $\mathcal{O}(n \log(n/\epsilon))$.

Quantum eigenvalue transformation of arbitrary matrices. We can thus apply the construction from the previous section to obtain an eigenvalue transformation of A.

Theorem 3 (Quantum eigenvalue transformation). Let $P(z)$ be a polynomial of degree n and let A be an arbitrary square matrix. Then, a block encoding of $P(A)$ can be obtained using $\mathcal{O}(\log n)$ ancillary qubits, $\mathcal{O}(n \log n)$ elementary operations, and n calls to a block encoding of A.

A full circuit implementing a polynomial transformation of a given matrix A is depicted in Fig. 1. Notice that, although we use the term *eigenvalue transformation*, we do not even need to assume A is diagonalizable, since we do not rely on the spectral mapping theorem, unlike previous constructions [3,10]. The polynomial transformation achieved by this construction can be better understood via the Jordan decomposition of the matrix, $A = SJS^{-1}$, where $J = \bigoplus_i J_i$. This decomposition remains valid even for nondiagonalizable matrices. Since any polynomial of A can be written as $P(A) = SP(J)S^{-1}$, with $P(J) = \mathrm{diag}(P(J_i))$, the transforma-

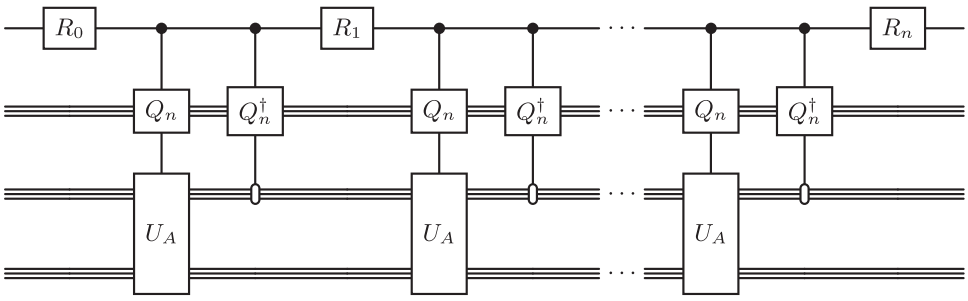


FIG. 1. Circuit for quantum eigenvalue transformation of general matrices based on (G)QSP. If $R_0, \ldots, R_n$ are the processing operators realizing some polynomial $P(z)$, then the above circuit implements a block encoding of $P(A)$, where A is block encoded by U_A.

tion amounts to a polynomial transformation of the Jordan blocks. In the case when A is diagonalizable, each J_i reduces to the eigenvalue λ_i, and the transformation coincides with the polynomial transformation of the eigenvalues. On the other hand, in the case of a nontrivial Jordan block, $P(J_i)$ is the upper-triangular Toeplitz matrix carrying $P^{(k)}(\lambda_i)/k!$ on its kth superdiagonal [37]; throughout this Letter, "eigenvalue transformation" is to be understood in this Jordan sense, i.e., as the implementation of the full matrix function $P(A)$, and not merely of the map $\lambda_i \mapsto P(\lambda_i)$.

Examples. As a first example, we consider the matrix exponential e^A. We simply implement the function $f(z) = e^z/3$ via GQSP, which is analytic and satisfies $|f(z)| \leqslant e/3 < 1$ on the unit disk $\mathbb{D}$. The Taylor coefficients $c_k = z^k/k!$ of the exponential decay factorially on $\mathbb{D}$, so that

$$\left| e^z - \sum_{k=0}^{N} \frac{z^k}{k!} \right| \leqslant \epsilon \tag{6}$$

holds on $\mathbb{D}$ already for $N = \mathcal{O}\left(\frac{\log 1/\epsilon}{\log\log 1/\epsilon}\right)$. Note that the subnormalization is the *constant* 3, independent of the degree and of the precision, so the success amplitude is $\Theta(1)$; in LCU-based implementations, the subnormalization is instead set by the 1-norm of the coefficients of the linear combination.

In the short-time simulation regime (i.e., implementing e^{tA} for $t = \mathcal{O}(1)$), this method requires fewer copies of A than the LCHS method for simulating time-independent homogeneous linear differential equations presented in Ref. [20], and it can be used to address unstable evolutions as well. The number of ancillary qubits is also significantly lower (see Table I).

As a second example, providing a physically relevant instance of the matrix exponential, consider the damped oscillator $\ddot{\boldsymbol{x}} + C\dot{\boldsymbol{x}} + K\boldsymbol{x} = 0$, written in state-space form $\dot{\mathbf{v}} = A\,\mathbf{v}$ where, for a one-dimensional system with $C = 2\gamma$ and $K = \omega_0^2$,

$$A(\gamma, \omega_0) = \begin{pmatrix} 0 & 1 \\ -\omega_0^2 & -2\gamma \end{pmatrix}. \tag{7}$$

The solution is given by e^{At}. At *critical damping* $\gamma = \omega_0$, the two eigenvalues collide at $\lambda = -\gamma$ and A becomes a single Jordan block: It is *defective*, $A + A^\dagger \not\geqslant 0$, and the eigenvector matrix S satisfies $\kappa(S) \to \infty$. This is the prototypical *exceptional point* of a non-Hermitian generator, ubiquitous in open and *PT*-symmetric systems, photonic lattices, and dissipative many-body models. Since our construction never references the eigenvector basis, it implements e^{tA} in short-time regime at the exceptional point with exactly the same cost $\mathcal{O}\left(\frac{\log 1/\epsilon}{\log\log 1/\epsilon}\right)$ as in the diagonalizable case, whereas Lap-LCHS is inapplicable ($A + A^\dagger \not\geqslant 0$) and the Chebyshev/Faber guarantees degrade with $\kappa(S) \to \infty$. Besides having no condition on the transformed matrix, for short-time exponentiation the method presented here has lower query complexity, extra two-qubit gates, and extra ancillas compared with the rest of the methods (see Table D.1 in the Supplemental Material [38] for a direct comparison).

Finally, consider the wave equation $\ddot{\boldsymbol{u}} + A\boldsymbol{u} = 0$, whose solution can be written as $u(t) = \cos(t\sqrt{A})u(0) + A^{-1/2}\sin(t\sqrt{A})\dot{u}(0)$. These functions are holomorphic in the complex unit disk and do not have a limit for $t \to \infty$: Thus, there is no locally integrable inverse Laplace transform. These functions have recently been implemented via contour-based matrix decomposition [39], with a complexity that scales with the condition number and $\sqrt{1/\epsilon}$. We can use Theorem 3 to directly implement these functions with complexity $\mathcal{O}\left(\frac{\log 1/\epsilon}{\log\log 1/\epsilon}\right)$ for short times [40].

An LCHS-type extension that lifts the stability requirement has recently been proposed [22]; however, its complexity $\mathcal{O}(S\log(S/\epsilon))$ is governed by $S \geqslant \max_{z\in\overline{\mathbb{D}}}\{|f|, |f'|, |f''|\}$, and thus in the short-time regime the complexity is $\mathcal{O}(\log 1/\epsilon)$, so the direct implementation via Theorem 3 remains favorable.

In general, we do not have a simple closed-form expression for the Taylor series of an arbitrary function we wish to implement, so it is not straightforward to bound its convergence rate directly. For convenience, we state the following well-known result from complex analysis (with a proof provided in the Supplemental Material [41]), which turns out to be particularly useful in the application of the technique presented here.

Theorem 4 (Taylor series convergence). Let $f(z)$ be a function analytic and bounded in $\mathbb{D}$, and suppose f admits an analytic continuation to the disk $\{z \in \mathbb{C} : |z| < R\}$ for some radius $R = 1 + r > 1$, on which $|f(z)| \leqslant M$. Then, the Taylor expansion of f centered at $z = 0$,

$$f(z) = \sum_{k=0}^{\infty} a_k z^k,$$

satisfies the error estimate

$$\left| \sum_{k=0}^{N} a_k z^k - f(z) \right| \leqslant \epsilon \tag{8}$$

TABLE I. Comparison of spectral-processing methods. "Def." indicates whether a nondiagonalizable input is admissible with cost *independent of the eigenvector conditioning* $\kappa(S)$: ✓, yes; •, formally representable but with hypotheses/bounds that degrade as $\kappa(S) \to \infty$ (Low and Su) or that are violated for indefinite Hermitian part/unstable spectra (Lap-LCHS); ×, not applicable. Here, f is the polynomial function, n is the polynomial degree, K is the number of LCHS quadrature nodes and T is a truncation parameter determined by the decay of the inverse Laplace g, $B \geqslant \max_{\|z\| \leqslant 1+\rho}\{|f(z)|, |f'(z)|, |f''(z)|\}$, $W(A)$ is the numerical range of A, and ρ is the margin between $W(A)$ and the nearest point where f ceases to be analytical. The success amplitude is the inverse subnormalization (probability = amplitude^2).

Method	Spectral object	Matrix class	Def.	Queries to U_A	Extra ancillas	Success amplitude	Function class
QSVT [10]	Singular values	General (block encoded)	×	n	$\mathcal{O}(1)$	$\mathcal{O}(1)$ (set by $\|p\|_{\mathbb{D}}$)	Real definite parity
QET-U [6,11]	Eigenvalues	Unitary U only	×	n	1	$\mathcal{O}(1)$	Any $\|P\| \leqslant 1$ on $\mathbb{T}$
Faber history states [17]	Eigenvalues	Spectrum in region Ω	•	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$ (history reg.)	Faber capacity of Ω; bounds $\propto \kappa(S)$	Faber series on Ω
Lap-LCHS [21]	Eigenvalues	$L = \frac{A+A^\dagger}{2} \succeq 0$ (stable)	•	$\widetilde{\mathcal{O}}\Big(\|g\|T \log^{1+o(1)} \frac{1}{\varepsilon}\Big)$	$\mathcal{O}\Big(\log T + \log\log^{1+1/\beta} \frac{1}{\varepsilon}\Big)$	$\propto 1/\|f\|_1 \|g\|_1$; state-prep optimal	Inverse Laplace transforms g
Weyl-LCHS [22]	Jordan blocks	Arbitrary square	✓	$\mathcal{O}\Big(\frac{B}{\rho} \log \frac{B}{\rho\epsilon}\Big)$,	$\mathcal{O}(\log\log 1/\epsilon)$	$\propto 1/\|f\|_1$	f analytic on $W(A)$ + margin ρ
This work	Jordan blocks	Arbitrary square	✓	n	$\mathcal{O}(\log n)$	$\mathcal{O}(1)$ if $\|f\| \leqslant$ const on $\mathbb{D}$	Any f bounded analytic on $\mathbb{D}$

on $\mathbb{D}$ whenever $N = \Theta(\frac{1}{r} \log \frac{M}{r\epsilon})$.

By Theorem 4, this method ensures $\mathcal{O}(\log M/\epsilon)$ copies of A are sufficient to block encode $f(A)$ up to precision ϵ, whenever f extends analytically across $\mathbb{T}$.

Conclusions. In this Letter, we propose a method to apply a polynomial to the eigenvalues (or, more generally, to the Jordan blocks) of a block-encoded square matrix A. We have shown that, if the kth power U^k block encodes A^k, a property we call *regularity*, then $P(A)$ can be obtained by simply applying P as a GQSP polynomial to U, and that any block encoding can be turned into an n-regular one by equipping it with a quantum incrementer.

Recently proposed methods cannot be generally applied to any arbitrary square matrix (see Table I for a comparison). Lap-LCHS can only be applied to stable matrices, with positive semidefinite Hermitian part [42]. Besides, it relies on the existence of the inverse Laplace transform for the function class, which might not generally exist (see the example of the wave equation). Chebyshev/Faber history states have a complexity and estimation guarantees that scale with $\kappa(S)$, which diverges as the matrix approaches defectiveness. Weyl-LCHS is able to handle unstable matrices, but its complexity scales with $|f(z)|, |f'(z)|, |f''(z)|$, making it inefficient for rapidly varying functions, such as the mentioned wave-equation solution or step functions.

Moreover, the aforementioned methods suffer from the high gate and ancilla cost of LCU, in contrast with the comparatively fewer requirements of the construction in Theorem 2, of $\mathcal{O}(\log n)$ ancilla qubits and $\mathcal{O}(n \log(n/\epsilon))$ two-qubit gates, where n is the maximum degree of the polynomial. Note that the ancilla count for LCHS in Table I refers only to the implementation of e^{-At}; the discretized Laplace transformation increases this cost drastically. We remark that GQSP itself requires one singly controlled application of U_A per degree—a constant-factor overhead over the uncontrolled block encoding—whereas the select oracles underlying LCU-based methods require multiply controlled applications of U_A or of its powers; the comparison in Table I therefore also holds at the level of controlled queries.

We also note that unitaries not only block encode themselves, but are also ∞-regular, which provides an alternative viewpoint on why applying GQSP to unitaries effectively yields an eigenvalue transformation. At the same time, if we consider a block encoding of a matrix A that is close to a unitary, then the block encoding will be approximately regular, since the failed branches multiply by blocks of small operator norm.

A limitation of this method lies in the rigidity of the domain: GQSP requires the construction of functions that are analytic on the *entire* unit disk $\mathbb{D}$, even if we can assume that the eigenvalues of A lie in a proper subset of it (e.g., the stability region for the simulation of nonunitary dynamics). We can circumvent this problem via a conformal map $\mathcal{M}(z)$ from the region in which the eigenvalues lie to the unit circle: We first construct $\mathcal{M}^{-1}(A)$ and apply the function $f(\mathcal{M}(z))$ to it, with $\mathcal{M}(z)$ chosen such that the composite function is easier to handle. We provide an example for the exponential in the Supplemental Material [43], with a total complexity $\mathcal{O}(t^2 + \log 1/\epsilon)$ in the long-time regime.

Acknowledgments. The authors would like to thank Dong An for important clarifications on Ref. [23] and to acknowledge insightful discussion with Mario Berta. L.L. was supported by the Swiss National Science Foundation (SNSF), Project No. 200020-214808. X.G. and M.S. thank support from the Basque Government BasQ initiative under the Q-STREAM project. They also acknowledge support from OpenSuperQ + 100 (Grant No. 101113946) of the EU Flagship on Quantum Technologies, from Project Grant No. PID2024-156808NB-I00 and Spanish Ramón y Cajal Grant No. RYC-2020-030503-I funded by MICIU/AEI/10.13039/501100011033 and by "ERDF A way of making Europe" and "ERDF Invest in your Future," and from the Spanish Ministry for Digital Transformation and of Civil Service of the Spanish Government through the QUANTUM ENIA project call Quantum Spain, by the EU through the Recovery, Transformation and Resilience

Plan—Next Generation EU within the framework of the Digital Spain 2026 Agenda, and from Basque Government through Grant No. IT1887-26.

Data availability. There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] G. H. Low and I. L. Chuang, Optimal Hamiltonian simulation by quantum signal processing, Phys. Rev. Lett. **118**, 010501 (2017).

[2] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, Grand unification of quantum algorithms, PRX Quantum **2**, 040203 (2021).

[3] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, Quantum **3**, 163 (2019).

[4] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, SIAM J. Comput. **46**, 1920 (2017).

[5] L. Lin and Y. Tong, Near-optimal ground state preparation, Quantum **4**, 372 (2020).

[6] Y. Dong, L. Lin, and Y. Tong, Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices, PRX Quantum **3**, 040305 (2022).

[7] S. McArdle, A. Gilyén, and M. Berta, Quantum state preparation without coherent arithmetic, Phys. Rev. Lett. **136**, 240603 (2026).

[8] L. Laneve, Robust black-box quantum-state preparation via quantum signal processing, arXiv:2305.04705.

[9] G. H. Low, T. J. Yoder, and I. L. Chuang, Methodology of resonant equiangular composite quantum gates, Phys. Rev. X **6**, 041067 (2016).

[10] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, 2019), pp. 193–204.

[11] D. Motlagh and N. Wiebe, Generalized quantum signal processing, PRX Quantum **5**, 020368 (2024).

[12] M. Alexis, G. Mnatsakanyan, and C. Thiele, Quantum signal processing and nonlinear Fourier analysis, Rev. Mat. Complut. **37**, 655 (2024).

[13] M. Alexis, L. Lin, G. Mnatsakanyan, C. Thiele, and J. Wang, Infinite quantum signal processing for arbitrary Szegö functions, Commun. Pure Appl. Math. **79**, 123 (2026).

[14] L. Laneve, Generalized quantum signal processing and nonlinear Fourier transform are equivalent, arXiv:2503.03026.

[15] H. Ni, R. Sarkar, L. Ying, and L. Lin, Inverse nonlinear fast Fourier transform on SU(2) with applications to quantum signal processing, arXiv:2505.12615.

[16] C. Sünderhauf, Generalized quantum singular value transformation, arXiv:2312.00723.

[17] G. H. Low and Y. Su, Quantum eigenvalue processing, SIAM J. Comput. **55**, 135 (2026).

[18] D. An, J.-P. Liu, and L. Lin, Linear combination of Hamiltonian simulation for nonunitary dynamics with optimal state preparation cost, Phys. Rev. Lett. **131**, 150603 (2023).

[19] X. Huang and D. An, Fourier transform-based linear combination of Hamiltonian simulation, arXiv:2508.19596.

[20] G. H. Low and R. D. Somma, Optimal quantum simulation of linear non-unitary dynamics, arXiv:2508.19238.

[21] D. An, A. M. Childs, L. Lin, and L. Ying, Laplace transform–based quantum eigenvalue transformation via linear combination of Hamiltonian simulation, SIAM J. Comput. **55**, 376 (2026).

[22] H. Ni and L. Ying, Quantum eigenvalue transformation via linear combination of Hamiltonian simulation: A Weyl calculus approach, arXiv:2606.29848.

[23] S. Jiang and D. An, Contour-integral based quantum eigenvalue transformation: Analysis and applications, arXiv:2601.11959.

[24] Y. Wang, Y. Liang, H. Chen, and J.-P. Liu, Optimal quantum eigenvalue transformation via linear combinations of Hermitian matrices, arXiv:2607.25812.

[25] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, Quantum Inf. Comput. **12**, 901 (2012).

[26] J.-P. Liu, H. O. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs, Efficient quantum algorithm for dissipative nonlinear differential equations, Proc. Natl. Acad. Sci. USA **118**, e2026805118 (2021).

[27] H. Krovi, Improved quantum algorithms for linear and nonlinear differential equations, Quantum **7**, 913 (2023).

[28] A. M. Childs, J.-P. Liu, and A. Ostrander, High-precision quantum algorithms for partial differential equations, Quantum **5**, 574 (2021).

[29] H. Weimer, A. Kshetrimayum, and R. Orús, Simulation methods for open quantum many-body systems, Rev. Mod. Phys. **93**, 015008 (2021).

[30] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for a proof of Lemma 1 (Appendix A).

[31] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for the example (Appendix B).

[32] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for the proof of Theorem 2 (Appendix C).

[33] G. H. Low and N. Wiebe, Hamiltonian simulation in the interaction picture, arXiv:1805.00675.

[34] D. Fang, L. Lin, and Y. Tong, Time-marching based quantum solvers for time-dependent linear differential equations, Quantum **7**, 955 (2023).

[35] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, G. Salton, S. Wang, and F. G. S. L. Brandão, *Quantum Algorithms: A Survey of Applications and End-to-End Complexities* (Cambridge University Press, Cambridge, 2025).

[36] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).

[37] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for an example of a polynomial transformation of a Jordan block (Appendix B).

[38] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for a comparison with other methods in the short-time regime (Appendix D), which includes Ref. [44].

[39] C. Wang, H.-Y. Liu, C. Xue, X.-N. Zhuang, M. Dou, Z.-Y. Chen, and G.-P. Guo, Quantum simulation of non-Hermitian special functions and dynamics via contour-based matrix decomposition, Quantum Sci. Technol. **11**, 035027 (2026).

[40] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for an application to second-order wave dynamics (Appendix E).

[41] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for the proof of Theorem 4 (Appendix F).

[42] Lap-LCHS cannot in general be applied to nilpotent matrices. As a simple example, consider the Jordan block $A + A^\dagger = X \not\succeq 0$ such that all eigenvalues are 0 but $A + A^\dagger = X \not\succeq 0$.

[43] See Supplemental Material at http://link.aps.org/supplemental/10.1103/sm6w-5ttl for the domain manipulation through conformal maps (Appendix G).

[44] D. W. Berry, D. Motlagh, G. Pantaleoni, and N. Wiebe, Doubling the efficiency of Hamiltonian simulation via generalized quantum signal processing, Phys. Rev. A **110**, 012612 (2024).