

ROBUST ITERATIVE METHOD FOR SYMMETRIC QUANTUM SIGNAL PROCESSING IN ALL PARAMETER REGIMES*

YULONG DONG[†], LIN LIN[‡], HONGKANG NI[§], AND JIASU WANG[†]

Abstract. This paper addresses the problem of solving nonlinear systems in the context of symmetric quantum signal processing (QSP), a powerful technique for implementing matrix functions on quantum computers. Symmetric QSP focuses on representing target polynomials as products of matrices in $SU(2)$ that possess symmetry properties. We present a novel Newton's method tailored for efficiently solving the nonlinear system involved in determining the phase factors within the symmetric QSP framework. Our method demonstrates rapid and robust convergence in all parameter regimes, including the challenging scenario with ill-conditioned Jacobian matrices, using standard double precision arithmetic operations. For instance, solving symmetric QSP for a highly oscillatory target function $\alpha \cos(1000x)$ (polynomial degree ≈ 1433) takes 6 iterations to converge to machine precision when $\alpha = 0.9$, and the number of iterations only increases to 18 iterations when $\alpha = 1 - 10^{-9}$ with a highly ill-conditioned Jacobian matrix. Leveraging the matrix product state structure of symmetric QSP, the computation of the Jacobian matrix incurs a computational cost comparable to a single function evaluation. Moreover, we introduce a reformulation of symmetric QSP using real-number arithmetics, further enhancing the method's efficiency. Extensive numerical tests validate the effectiveness and robustness of our approach, which has been implemented in the QSPPACK software package.

Key words. numerical algorithm, quantum algorithm, matrix product state, quantum signal processing, iterative method

MSC codes. 81-08, 65Z05

DOI. 10.1137/23M1598192

1. Introduction. Many scientific computing problems can be viewed as implementing matrix functions $A \mapsto f(A)$. For simplicity, we can assume that A is a Hermitian matrix with eigenvalues in $[-1, 1]$ and that $f: \mathbb{R} \rightarrow \mathbb{R}$ is a real polynomial. Quantum signal processing (QSP) [13, 14, 10, 16] provides a systematic approach and a compact quantum circuit for implementing a broad class of matrix functions on quantum computers. This leads to efficient algorithms for various quantum applications, including linear system solving, Hamiltonian system simulation, ground-state energy estimation, and quantum benchmarking [13, 10, 12, 16, 6, 4, 8, 17, 3]. Recently, the construction of QSP has been studied and generalized using advanced theoretical tools [22, 23, 21, 18, 24].

*Submitted to the journal's Numerical Algorithms for Scientific Computing section September 1, 2023; accepted for publication (in revised form) June 10, 2024; published electronically September 13, 2024.

<https://doi.org/10.1137/23M1598192>

Funding: This material is based upon work of the fourth author supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. The first author received support from the NSF Quantum Leap Challenge Institute (QLCI) program under grant OMA-2016245. Additional support was provided by the Applied Mathematics Program of the U.S. Department of Energy Office of Advanced Scientific Computing Research under contract DE-AC02-05CH1123 and by a Google Quantum Research Award. The second author is a Simons Investigator.

[†]Department of Mathematics, University of California, Berkeley, CA 94720 USA (dongyl@berkeley.edu, jiasu@berkeley.edu).

[‡]Department of Mathematics and Challenge Institute for Quantum Computation, University of California, Berkeley, and Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA (linlin@math.berkeley.edu).

[§]Institute for Computational and Mathematical Engineering (ICME), Stanford University, Stanford, CA 94305 USA (hongkang@stanford.edu).

Since any continuous function can be approximated using polynomials, the key idea behind QSP is to represent a target polynomial as a product of matrices in the special unitary group $SU(2)$, parameterized by a set of phase factors denoted as Φ . However, due to the constraints of $SU(2)$, the target polynomial must satisfy specific conditions.

DEFINITION 1.1 (target polynomial of QSP). *A polynomial $f \in \mathbb{R}[x]$ is called a target polynomial of QSP if it satisfies (1) $\deg(f) = d$, (2) the parity of f is $(d \bmod 2)$, (3) $\|f\|_\infty := \max_{x \in [-1, 1]} |f(x)| \leq 1$. Furthermore, f is called fully coherent if $\|f\|_\infty = 1$.*

The mapping from the target polynomial of degree d (described by its Chebyshev coefficients denoted by $c \in \mathbb{R}^{d+1}$) to phase factors $\Psi \in \mathbb{R}^{d+1}$ can be abstractly written as

$$(1.1) \quad F(\Psi) = c.$$

The mapping F is highly nonlinear and is not one-to-one. For a given c , our goal is to find a solution to the nonlinear system (1.1).

QSP in the fully coherent regime (or near fully coherent regime, where $\|f\|_\infty = 1 - \gamma$ for some small $\gamma > 0$) finds applications in quantum algorithms for Hamiltonian simulation [15] and time-marching-based simulation of linear differential equations [9]. As will be discussed below, this problem is particularly challenging in the fully coherent regime where the Jacobian matrix of F is very ill-conditioned (see the numerical section for an illustration of this phenomenon).

The contribution of this work is to propose a Newton's method tailored for efficiently solving the nonlinear system (1.1). Specifically, we demonstrate the following:

1. Starting from a problem-independent initial guess proposed in [7], Newton's method can converge rapidly in all parameter regimes, using standard double precision arithmetic operations.
2. The computation of the Jacobian matrix can leverage the matrix product states structure of QSP. Notably, the computational cost associated with computing the Jacobian matrix is comparable to that of a *single function evaluation*.
3. The prefactor of the numerical method can be further enhanced by reformulating the symmetric QSP using real-number arithmetics.

We have conducted extensive numerical tests, which have consistently demonstrated the efficiency and robustness of the method. Thus far we have not encountered any instances where the method fails. We have implemented the method in the QSPPACK software package.¹

Related works. The phase factor evaluation was originally conceived to be a challenging task [13, 2]. In the past few years, significant progress has been made in developing efficient algorithms to find phase factors. These algorithms fall into two categories: factorization methods [10, 11, 1, 26] and iterative methods [7, 5].

For a given *real* polynomial $f(x)$, factorization methods construct phase factors from the roots of $1 - f^2(x)$ in the complex plane, and the roots must be obtained at high precision. As a result, as the polynomial d increases, direct implementation of factorization-based methods is not numerically stable and requires $\mathcal{O}(d \log(d/\epsilon))$

¹The examples are available on the website <https://qsppack.gitbook.io/qsppack/> and the codes are open-sourced in <https://github.com/qsppack/QSPPACK>.

bits of precision [10, 11] (ϵ is the target accuracy). There have been two recent improvements in factorization-based methods: the capitalization method [1] and the Prony method [26]. Empirical results indicate that both methods are numerically stable and are applicable to large degree polynomials. Furthermore, the performance of factorization-based methods does not deteriorate near the fully coherent regime.

Compared to the elaborate construction of factorization-based methods, iterative methods are intuitive, numerically stable, and easy to implement. The idea is to directly tackle the nonlinear system (1.1) or the equivalent optimization formulation

$$(1.2) \quad \Psi^* = \operatorname{argmin}_{\Psi} \|F(\Psi) - c\|_2^2.$$

However, due to the complex energy landscape [25], direct optimization from random initial guesses can easily get stuck at local minima and can only be used for low degree polynomials. References [7, 25] propose and study the *symmetric* QSP where the set of phase factors are subjected to a symmetry condition that reduces the degrees of freedom. Reference [7] further observes that starting from a carefully chosen but problem-independent initial guess, standard optimization methods such as the LBFGS method [19] can be robust and stable and can be applied to very high degree polynomials. Recently, we proposed the fixed point iteration (FPI) algorithm [5] that directly tackles the nonlinear system (1.1) and showed that the symmetric phase factors have a well-defined limit as the polynomial degree increases toward infinity when the polynomial approaches a smooth (nonpolynomial) function. However, it is important to note that in many examples near the fully coherent regime, the assumptions of those theoretical results are violated. Consequently, gradient-based optimization methods and the FPI method may exhibit slow convergence or fail to converge altogether. The symmetric choice of phase factors implies a real-valued complementary polynomial $Q(x)$, which is used in deriving the quantum eigenvalue transformation of unitaries with real polynomials (QETU) [6]. On the other hand, recent studies [24, 18] have extended the framework of QSP to facilitate the representation of complex polynomials, where the symmetric choice of phase factors is no longer applicable.

Organization. The paper is organized as follows. The preliminaries are given in section 2. In subsection 2.1, we review relevant concepts in QSP with symmetric phase factors. Then, in subsection 2.2, we discuss the bottleneck toward the fully coherent regime and also review the iterative methods for finding phase factors in the literature. The matrix product state and its relevance in the structure of QSP are presented in subsection 2.3. Our main algorithm and its implementation are given in section 3, where we also discuss the acceleration of the algorithm by leveraging the structure of the problem and a real-number arithmetic representation of QSP. Finally, in section 4, we demonstrate our algorithm by presenting the results of numerical experiments.

2. Preliminaries.

2.1. Quantum signal processing with symmetric phase factors. QSP represents a class of polynomials in terms of $SU(2)$ matrices, which is parameterized by phase factors [10, Theorem 4]. The phase factors $\Psi = (\psi_0, \psi_1, \dots, \psi_d) \in \mathbb{R}^{d+1}$ are symmetric if they satisfy the constraint $\psi_i = \psi_{d-i}$ for any $i = 0, \dots, d$. Reference [25] proposes a variant of QSP representation focusing on symmetric phase factors.

THEOREM 2.1 (QSP with symmetric phase factors [25, Theorem 1]). *Consider any $P \in \mathbb{C}[x]$ and $Q \in \mathbb{R}[x]$ satisfying the following conditions:*

1. $\deg(P) = d$ and $\deg(Q) = d - 1$.
2. P has parity ($d \bmod 2$) and Q has parity ($d - 1 \bmod 2$).
3. (Normalization condition) $\forall x \in [-1, 1] : |P(x)|^2 + (1 - x^2)|Q(x)|^2 = 1$.
4. If d is odd, then the leading coefficient of Q is positive.

There exists a unique set of symmetric phase factors $\Psi := (\psi_0, \psi_1, \dots, \psi_d) \in D_d$ such that

$$(2.1) \quad U(x, \Psi) = e^{i\psi_0 Z} \prod_{j=1}^d [W(x) e^{i\psi_j Z}] = \begin{pmatrix} P(x) & iQ(x)\sqrt{1-x^2} \\ iQ(x)\sqrt{1-x^2} & P^*(x) \end{pmatrix},$$

where

$$(2.2) \quad D_d = \begin{cases} \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]^{\frac{d}{2}} \times \left[-\pi, \pi\right] \times \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]^{\frac{d}{2}}, & d \text{ is even,} \\ \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]^{d+1}, & d \text{ is odd.} \end{cases}$$

In the above equations, X and Z denote Pauli matrices,

$$W(x) = e^{i\arccos(x)X} = \begin{pmatrix} x & i\sqrt{1-x^2} \\ i\sqrt{1-x^2} & x \end{pmatrix},$$

and $P^*(x)$ represents the complex conjugate of the complex polynomial $P(x)$ obtained by conjugating all its coefficients.

In most applications, only either the real or the imaginary part of the complex polynomial $P(x) = \langle 0 | U(x, \Psi) | 0 \rangle$ is relevant. It can be shown that these two parts can be exchanged by conjugating the unitary matrix product with $\pi/4$ - Z rotation, that is,

$$(2.3) \quad \text{Re}[\langle 0 | U(x, \Psi) | 0 \rangle] = \text{Im}[\langle 0 | e^{i\frac{\pi}{4}Z} U(x, \Psi) e^{i\frac{\pi}{4}Z} | 0 \rangle].$$

This is equivalent to shifting the edge phase factors ψ_0, ψ_d by $\pi/4$. For simplicity of presentation, this paper assumes that the imaginary part is relevant. Furthermore, we denote it as $g(x, \Psi)$,

$$(2.4) \quad g(x, \Psi) := \text{Im}[\langle 0 | U(x, \Psi) | 0 \rangle].$$

Given any symmetric phase factors $\Psi := (\psi_0, \psi_1, \dots, \psi_d)$ of length $d + 1$, we define the *reduced phase factors* Φ as

$$(2.5) \quad \Phi = (\phi_0, \phi_1, \dots, \phi_{\tilde{d}-1}) := \begin{cases} \left(\frac{1}{2}\psi_{\tilde{d}-1}, \psi_{\tilde{d}}, \dots, \psi_d\right), & d \text{ is even,} \\ (\psi_{\tilde{d}}, \psi_{\tilde{d}-1}, \dots, \psi_d), & d \text{ is odd,} \end{cases}$$

where $\tilde{d} := \lceil \frac{d+1}{2} \rceil$. For the sake of simplicity, we do not explicitly distinguish the full set of phase factors and the reduced phase factors when they are used as the argument of some function. For example, $U(x, \Phi)$ and $g(x, \Phi)$ are assumed to represent evaluations with respect to the full set of phase factors, $U(x, \Psi)$ and $g(x, \Psi)$. To be embedded in an SU(2) matrix, it is naturally required that $g(x, \Psi) \leq 1$ for any $x \in [-1, 1]$. Hence, the target function $f : \mathbb{R} \rightarrow \mathbb{R}$ is also normalized so that its norm is bounded:

$$(2.6) \quad \|f\|_\infty = \max_{x \in [-1, 1]} |f(x)| \leq 1.$$

Theorem 2.1 implies that if the target polynomial f of definite parity can be represented as symmetric QSP, it admits a Chebyshev polynomial expansion:

$$(2.7) \quad f(x) = \begin{cases} \sum_{j=0}^{\tilde{d}-1} c_j T_{2j}(x), & f \text{ is even,} \\ \sum_{j=0}^{\tilde{d}-1} c_j T_{2j+1}(x), & f \text{ is odd.} \end{cases}$$

Let \mathcal{F} denote the linear mapping from a target polynomial to its Chebyshev coefficient vector

$$(2.8) \quad c := (c_0, c_1, \dots, c_{\tilde{d}-1}) \in \mathbb{R}^{\tilde{d}}.$$

This induces the mapping from the set of reduced phase factors to the Chebyshev coefficient vector of $g(x, \Phi)$,

$$(2.9) \quad F : \mathbb{R}^{\tilde{d}} \rightarrow \mathbb{R}^{\tilde{d}}, \quad F(\Phi) := \mathcal{F}(g(x, \Phi)).$$

2.2. Iterative methods for finding phase factors and numerical difficulties near the fully coherent regime. The QSP problem can also be solved using numerical optimization,

$$(2.10) \quad \Phi^* = \underset{\Phi}{\operatorname{argmin}} \|F(\Phi) - c\|_2^2 = \underset{\Phi}{\operatorname{argmin}} \sum_{j=1}^{\tilde{d}} |g(x_j, \Phi) - f(x_j)|^2.$$

Here, $x_j = \cos\left(\frac{(2j-1)\pi}{4\tilde{d}+1}\right)$ is the j th node of the Chebyshev polynomial $T_{2\tilde{d}}(x)$. The equality in the optimization problem follows the discrete orthogonality on Chebyshev nodes. In [7], the optimization problem is first solved using the LBFGS method, and the running complexity is numerically studied. The authors also propose the use of an initial guess $\Phi^0 = (0, 0, \dots, 0)$, from which the convergence of the LBFGS method is numerically observed to be fast and stable. We remark that the initial guess in the original paper is not identical to this form, due to the difference in the definition. The original paper considers the real part of $\langle 0 | U(x, \Psi) | 0 \rangle$ to encode the polynomial of interest, whereas this paper considers the imaginary part, with equivalence established through (2.3). The choice of the initial guess is justified in the theoretical analysis in [25], which is credited to a class of optima called the maximal solution. In [25], the authors analyze the energy landscape of the optimization problem and conclude that when the target function is scaled as $\|f\|_\infty = \mathcal{O}(1/d)$, the optimization-based algorithm converges locally at $\mathcal{O}(d^2 \log(1/\epsilon))$ computational cost.

For the target function near the fully coherent regime, it is hard to guarantee the convergence of optimization-based methods. The ill-conditioned Hessian matrix around the fully coherent regime poses a challenge for the optimizer to be convergent. The numerical study in [7] shows that the condition number of the Hessian matrix at the optimum grows rapidly as the target function gets closer to the fully coherent regime. Furthermore, the theoretical analysis of the optimization landscape also suggests that the region of convergence shrinks as the target function approaches the fully coherent regime. Hence, the convergence guarantee of optimization-based methods is compromised near the fully coherent regime.

Using the mapping F defined in (2.9), the problem of finding phase factors can be formulated as solving a nonlinear equation given by (1.1). In [5], an FPI method is proposed for solving (1.1):

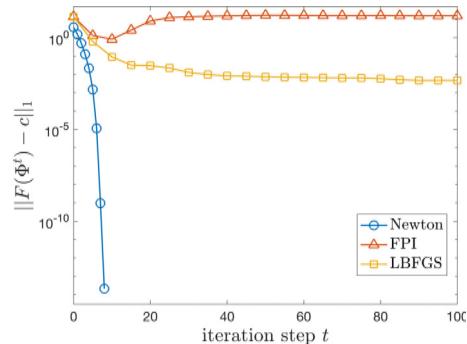


FIG. 1. The residual error after each iteration using three different methods to determine phase factors for the target function $f(x) = 0.999 \cos(500x)$ (up to the first 100 iterations). The stopping criterion is that the residual error reaches below 10^{-13} . Newton's method converges after 9 steps. The LBFGS method can eventually converge but takes over 1000 steps. The FPI method fails to converge.

$$(2.11) \quad \Phi^0 = \mathbf{0} \in \mathbb{R}^{\tilde{d}}, \quad \Phi^{t+1} = \Phi^t - \frac{1}{2} (F(\Phi^t) - c).$$

Notably, the initial guess of the FPI method coincides with that used in the LBFGS method. The analysis in [5] demonstrates that the FPI method exhibits linear convergence to the exact solution when $\|c\|_1 \leq 0.861$. This result is based on the observation that the update rule acts as a contraction mapping in a neighborhood of the initial guess $\Phi^0 = \mathbf{0}$. However, this property does not hold universally across the entire domain. The analysis in [5] reveals that the contraction property is valid when the Chebyshev coefficient vector of the target function lies within an ℓ^1 ball centered at the origin. In cases where the target function contains significant “high-frequency” components, the increasingly large Chebyshev coefficient vector may hinder the contraction of the update rule in (2.11). This situation commonly occurs in various applications; for instance, problematic convergence issues can arise when dealing with functions $\sin(\tau x)$ and $\cos(\tau x)$, where $\tau \gg 1$ is large.

To conclude the discussion on the challenges faced by iterative methods in the fully coherent regime, we present a numerical result that substantiates these difficulties. We consider the target function to be a degree-733 polynomial approximating $f(x) = 0.999 \cos(500x)$ obtained by truncating the Chebyshev expansion. This function violates the convergence analysis of iterative methods, as discussed earlier. In Figure 1, we plot the residual error at each iteration step. The FPI method does not converge at all, while the LBFGS method eventually reaches the optimum. However, the optimizer becomes trapped after the 100th step and requires over 1000 iterations to converge. In contrast, Newton's method, proposed in this paper, exhibits fast and stable convergence in the numerical results. The residual error decreases superexponentially, consistent with the expected quadratic convergence described in standard textbooks.

2.3. Matrix product state structure of quantum signal processing. The QSP problem, being a well-structured product of $SU(2)$ matrices, possesses inherent properties that allow for a special tensor structure, an MPS or tensor train (TT). These properties can be effectively leveraged to accelerate our numerical algorithm. By exploiting this tensor structure, we achieve a significant reduction in computation complexity, enabling the scalability of the solver for larger-scale applications.

In this subsection, we present a concise overview of the theory and construction of MPS/TT. Additionally, we establish its connection with our problem. To be specific,

QSP admits an MPS/TT structure with bond dimension 2 due to its $SU(2)$ -product defining (2.1).

Given a field $\mathbb{F} = \mathbb{R}$ or \mathbb{C} , an *order- r tensor* is referred to as a multidimensional array $G \in \mathbb{F}^{n_1 \times \dots \times n_r}$ where the r -tuple $(n_1, \dots, n_r) \in \mathbb{N}^r$ specifies the size of the tensor. Each entry of the tensor can be accessed with multi-indices, $G(i_1, \dots, i_r)$, where $1 \leq i_j \leq n_j \forall j = 1, \dots, r$. The contraction of two tensors yields a new tensor by summing over the specified indices. For example, if $G_1 \in \mathbb{F}^{n_1 \times n_2 \times n_3}$, $G_2 \in \mathbb{F}^{n_1 \times n_2 \times n_3}$ are two order-3 tensors, $G_3(i, j, k, l) := \sum_{s=1}^{n_3} G_1(i, j, s) G_2(k, l, s)$ elementwise defines an order-4 tensor $G_3 \in \mathbb{F}^{n_1 \times n_2 \times n_1 \times n_2}$ by contracting the index s . A graphical illustration is given in Figure 2. Specifically, the contraction of order-2 tensors (i.e., matrices) coincides with matrix multiplication.

A (parametric) tensor $G(\alpha) \in \mathbb{F}^{n_1 \times \dots \times n_r}$ is called an MPS or TT if each of its entries can be expressed as a product of matrices [20]. To be precise, an MPS/TT can be written as

$$G(i_1, \dots, i_r; \alpha) = \mathcal{G}_1(i_1; \alpha) \mathcal{G}_2(i_2; \alpha) \cdots \mathcal{G}_r(i_r; \alpha).$$

In the above expression, each $\mathcal{G}_j(\alpha)$ is an order-3 tensor. By fixing an index i_j , $\mathcal{G}_j(i_j; \alpha) = [G_j(k, l, i_j; \alpha)]_{kl} \in \mathbb{F}^{m_{j-1} \times m_j}$ becomes order-2 which is equivalent to a matrix of size (m_{j-1}, m_j) . The dangling index i_j is referred to as the mode index (or external index). The indices k, l which are dummy in contraction are referred to as rank core indices. The right-hand side of the defining equation is shorthand for contracting all nonfixed indices by matrix multiplication. The contracted tensor $G(i_1, \dots, i_r; \alpha)$ is assumed to be a scalar entrywise. Hence, the dangling components \mathcal{G}_1 and \mathcal{G}_r are set to order-2 tensors, namely, $m_0 = m_r = 1$. The bond dimension of an MPS/TT is defined to be the maximal dimension in the contraction $\max_{0 \leq j \leq r} m_j$.

Using the language of tensors, the upper-left entry of the QSP unitary matrix defined in (2.1) can be interpreted as an MPS/TT of bond dimension 2. To see this, assuming that x and a full set of phase factors $\Psi := (\psi_0, \dots, \psi_d)$ are given, the QSP matrix entry of interest is

$$(2.12) \quad \langle 0 | U(x, \Psi) | 0 \rangle = \mathcal{R}_0(\psi_0) \mathcal{W}(x) \mathcal{R}(\psi_1) \mathcal{W}(x) \cdots \mathcal{W}(x) \mathcal{R}(\psi_{d-1}) \mathcal{W}(x) \mathcal{R}_d(\psi_d) \in \mathbb{C}.$$

Here, the boundary components $\mathcal{R}_0(\psi_0) := (e^{i\psi_0}, 0)$, $\mathcal{R}_d(\psi_d) := (e^{i\psi_d}, 0)^\top$ are two-dimensional complex vectors. Furthermore, $\mathcal{W}(x) := e^{i \arccos(x) X}$ and $\mathcal{R}(\psi_j) := e^{i\psi_j Z}$ are 2-by-2 complex matrices. By identifying x, ψ as external indices, the graphical visualization of this interpretation is given in Figure 3.

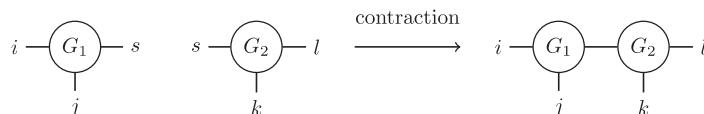


FIG. 2. A graphical visualization of two order-3 tensors and their contraction.

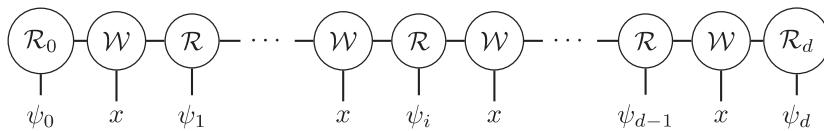


FIG. 3. A graphical visualization of $\langle 0 | U(x, \Psi) | 0 \rangle$ as a MPS/TT structure (of bond dimension 2).

3. Newton's method. When using iterative algorithms discussed in the previous section to find phase factors, the issue of convergence becomes increasingly significant when the target function is close to the fully coherent regime. To remedy these difficulties, we propose using Newton's method to solve the nonlinear equation (1.1) for phase factor evaluation. In this section, we introduce this method and discuss its implementation. The core techniques to accelerate the algorithm are fast Jacobian evaluation based on the MPS/TT structure and a real arithmetic formalism of symmetric QSP.

Newton's method can be viewed as an improvement over the FPI method by taking the local landscape into account. It can be verified that the Jacobian of the nonlinear equation in (1.1) at the origin coincides with a doubled identity matrix, that is, $DF(\mathbf{0}) = 2\mathbf{I}$. Hence, the FPI method is a variant of Newton's method, where the Jacobian is approximated along the iteration by that at the initial point, which is the origin. To be precise, the update rules $\Phi^{t+1} = T(\Phi^t)$ of both methods are written as

$$(3.1) \quad \begin{aligned} T_{\text{Newton}}(\Phi) &= \Phi - DF(\Phi)^{-1}(F(\Phi) - c), \\ \text{and } T_{\text{FPI}}(\Phi) &= \Phi - DF(\mathbf{0})^{-1}(F(\Phi) - c), \text{ where } DF(\mathbf{0}) = 2\mathbf{I}. \end{aligned}$$

The algorithm based on the first update rule is outlined in Algorithm 3.1. In the remainder of this section, we will discuss an accelerated implementation of this algorithm leveraging the structure of the symmetric QSP problem.

3.1. Jacobian of the problem. The update rule of Newton's method utilizes the Jacobian matrix of $F(\Phi)$, which is denoted as $DF(\Phi)$. According to the defining equation (2.9) of F , the i th column of $DF(\Phi)$ is

$$(3.2) \quad \frac{\partial F(\Phi)}{\partial \phi_i} = \mathcal{F} \left(\frac{\partial g(x, \Phi)}{\partial \phi_i} \right) \in \mathbb{R}^{\tilde{d}}, \quad i = 0, \dots, \tilde{d} - 1.$$

A straightforward approach for constructing the Jacobian matrix is to compute it columnwise without any optimization. This method involves performing the following procedure independently for each $0 \leq i < \tilde{d}$: evaluating $\partial g(x_k, \Phi) / \partial \phi_i$ at approximately $\mathcal{O}(d)$ distinct points and then applying a discrete Fourier transformation. Each evaluation of $\partial g(x_k, \Phi) / \partial \phi_i$ requires $\mathcal{O}(d)$ multiplications of SU(2). Consequently, the complexity for computing a column of the Jacobian is $\mathcal{O}(d^2)$. As a result, the overall complexity of this Jacobian evaluation is $\mathcal{O}(d^3)$. It is important to note that this approach does not take into account the structural characteristics of the problem, leaving room for potential optimization strategies.

Algorithm 3.1. Newton's method for finding reduced phase factors.

Input: Chebyshev coefficient vector c of a target polynomial, and stopping criteria.

```

Initiate  $t = 0$  and  $\Phi^0$  to be zero vector  $\mathbf{0}$ ;
while stopping criterion is not satisfied do
  Compute  $DF(\Phi^t)$ ;
  Update  $\Phi^{t+1} = \Phi^t - DF(\Phi^t)^{-1}(F(\Phi^t) - c)$ ;
  Set  $t = t + 1$ ;
end while
Output: Reduced phase factors  $\Phi^t$ .
```

In the subsequent subsection, we will present an accelerated evaluation method that capitalizes on the MPS/TT structure of the problem. This improved approach leads to a notable reduction in the complexity of Jacobian evaluation, from $\mathcal{O}(d^3)$ to $\mathcal{O}(d^2 \log(d))$.

In the remainder of this subsection, we delve into the structure of the Jacobian matrix columns. For the sake of simplicity, we assume that d is even. As a reminder, in the even case, the full set of phase factors can be represented as $\Psi = (\phi_{\tilde{d}-1}, \dots, \phi_1, 2\phi_0, \phi_1, \dots, \phi_{\tilde{d}-1})$. This choice does not lose generality, as a similar derivation can be obtained for the odd case.

We first observe that $\frac{\partial g(x, \Phi)}{\partial \phi_i} = \text{Im}[\langle 0 | \frac{\partial}{\partial \phi_i} U(x, \Psi) | 0 \rangle]$ and that taking a derivative on the unitary matrix $U(x, \Psi)$ is equivalent to the insertion of an additional $iZ = e^{i\pi Z/2}$ in the matrix product. Due to symmetry, the derivative leads to two matrix products with insertion. Specifically, we have

$$(3.3) \quad \frac{\partial}{\partial \phi_i} U(x, \Psi) = U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right) + U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1+i}\right) \quad \forall 0 \leq i \leq \tilde{d}-1.$$

We remark that $\Psi + \frac{\pi}{2} e_{\tilde{d}-1+i}$ is the reversed ordering of $\Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}$, which is helpful for the simplification. Consequently, $U(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1+i})$ is identical to the transpose of $U(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i})$ since the transpose reverses the order due to the symmetry of matrix Z . Hence

$$(3.4) \quad \begin{aligned} \frac{\partial g(x, \Phi)}{\partial \phi_i} &= \text{Im} \left[\left\langle 0 \left| \frac{\partial}{\partial \phi_i} U(x, \Phi) \right| 0 \right\rangle \right] \\ &= \text{Im} \left[\left\langle 0 \left| U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right) \right| 0 \right\rangle \right] + \text{Im} \left[\left\langle 0 \left| U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1+i}\right) \right| 0 \right\rangle \right] \\ &= \text{Im} \left[\left\langle 0 \left| U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right) \right| 0 \right\rangle \right] + \text{Im} \left[\left\langle 0 \left| U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right)^\top \right| 0 \right\rangle \right] \\ &= 2\text{Im} \left[\left\langle 0 \left| U\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right) \right| 0 \right\rangle \right] = 2g\left(x, \Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}\right). \end{aligned}$$

While $\Psi + \frac{\pi}{2} e_{\tilde{d}-1-i}$ is not symmetric, the evaluation of the induced polynomial is still well defined. To extract its Chebyshev coefficients, it suffices to sample this polynomial on the Chebyshev nodes $\{x_k = \cos(2\pi k/(2d+1)) : k = 0, \dots, d\}$. Subsequently, the Chebyshev coefficients can be extracted from the sample by performing fast Fourier transformation (FFT). The details are given in Algorithm 3.2.

3.2. Efficient evaluation of the Jacobian matrix. In practice, the evaluation of the Jacobian matrix $DF(\Phi)$ constitutes the most computationally demanding step in Newton's method. Although the derivative is explicitly given in (3.4), constructing the Jacobian through the straightforward application of (3.4) leads to some redundant calculations within the computational workflow. In this subsection, we propose an efficient approach to compute the Jacobian matrix, taking advantage of the MPS/TT structure of the problem. By doing so, we can significantly reduce the overall computational complexity. As we will illustrate, the computation of different columns of the Jacobian matrix exhibits substantial overlap. This indicates that we can reuse intermediate computational results and avoid redundancy, leading to increased efficiency.

Without loss of generality, we consider the case where d is even in the derivation. The odd case can be analyzed similarly. Recall that the full set of phase factors is defined as

Algorithm 3.2. Compute $\mathcal{F}(g(x, \Psi^\#))$.

Input: A full set of phase factors $\Psi^\#$ of length $d+1$ ($\Psi^\#$ is not necessarily symmetric).

Initialize $\mathbf{g} = (0, 0, \dots) \in \mathbb{R}^{2d+1}$.

Evaluate $\mathbf{g}_j \leftarrow g(x_j, \Psi^\#)$, $x_j = \cos\left(\frac{2\pi j}{2d+1}\right)$, $j = 0, \dots, d$.

Evaluate $\mathbf{g}_j \leftarrow \mathbf{g}_{2d+1-j}$, $j = d+1, \dots, 2d$.

Compute $\mathbf{v}_l \leftarrow \text{Re}(\sum_{j=0}^{2d-1} \mathbf{g}_j e^{-i\frac{2\pi}{2d+1}jl})$, $l = 0, \dots, d$, using FFT.

if $(d \bmod 2) = 0$ **then**

$\mathcal{F}(g(x, \Psi^\#)) \leftarrow \frac{2}{2d+1}(\frac{\mathbf{v}_0}{2}, \mathbf{v}_2, \mathbf{v}_4, \dots, \mathbf{v}_d)$.

else

$\mathcal{F}(g(x, \Psi^\#)) \leftarrow \frac{2}{2d+1}(\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_5, \dots, \mathbf{v}_d)$.

end if

Output: $\mathcal{F}(g(x, \Psi^\#))$.

$$(3.5) \quad \Psi := (\psi_0, \psi_1, \dots, \psi_d) = (\phi_{\tilde{d}-1}, \dots, \phi_1, 2\phi_0, \phi_1, \dots, \phi_{\tilde{d}-1}),$$

where $\Phi = (\phi_0, \dots, \phi_{\tilde{d}-1})$ are the reduced phases factors. We observe that each column of the Jacobian matrix is directly associated with taking the derivative of a QSP without symmetry, which arises from the insertion of the iZ matrix. In the absence of symmetry constraint of phase factors, each phase factor ϕ_i is independent. When calculating the derivative with respect to ϕ_i , we can separate the matrix multiplication into three parts,

$$(3.6) \quad \langle 0 | U(x, \Psi) | 0 \rangle = \mathcal{M}_{\text{left}}^{(i)} e^{i\phi_i Z} \mathcal{M}_{\text{right}}^{(i)},$$

where

$$(3.7) \quad \begin{aligned} \mathcal{M}_{\text{left}}^{(i)} &:= \mathcal{R}_0(\phi_{\tilde{d}-1}) \prod_{j=\tilde{d}-2}^{i+1} [\mathcal{W}(x) \mathcal{R}(\phi_j)] \mathcal{W}(x), \\ \mathcal{M}_{\text{right}}^{(i)} &:= \prod_{j=i-1}^0 [\mathcal{W}(x) \mathcal{R}(\phi_j)] \prod_{j=0}^{\tilde{d}-2} [\mathcal{R}(\phi_j) \mathcal{W}(x)] \mathcal{R}_d(\phi_{\tilde{d}-1}). \end{aligned}$$

The left and right components are irrelevant to taking the derivative with respect to ϕ_i because

$$(3.8) \quad \langle 0 | \partial_{\phi_i} U(x, \Phi) | 0 \rangle = 2 \left\langle 0 \left| U \left(x, \Psi + \frac{\pi}{2} e_{k-i} \right) \right| 0 \right\rangle = 2i \mathcal{M}_{\text{left}}^{(i)} Z e^{i\phi_i Z} \mathcal{M}_{\text{right}}^{(i)},$$

$$(3.9) \quad \text{and } \partial_{\phi_i} g(x, \Phi) = 2 \text{Im} \left[i \mathcal{M}_{\text{left}}^{(i)} Z e^{i\phi_i Z} \mathcal{M}_{\text{right}}^{(i)} \right].$$

Consequently, the intermediate quantities $\mathcal{M}_{\text{left}}^{(i)}$ and $\mathcal{M}_{\text{right}}^{(i)}$ can be stored and maintained in the computation process. Although the expression is similar to (3.4) as derived in the previous subsection, isolating the terms $\mathcal{M}_{\text{left}}^{(i)}$ and $\mathcal{M}_{\text{right}}^{(i)}$ allows us to remove the redundant computational steps. Subsequently, we will introduce an update mechanism for these variables that takes advantage of the MPS/TT structure in the problem. Figure 4 visually illustrates this idea.

Transiting to the next step, the intermediate quantities are updated through matrix multiplications

$$(3.10) \quad \mathcal{M}_{\text{left}}^{(i+1)} \leftarrow \mathcal{M}_{\text{left}}^{(i)} \mathcal{W}^{-1}(x) e^{-i\phi_{i+1} Z} \quad \text{and} \quad \mathcal{M}_{\text{right}}^{(i+1)} \leftarrow \mathcal{W}(x) e^{i\phi_i Z} \mathcal{M}_{\text{right}}^{(i)}.$$

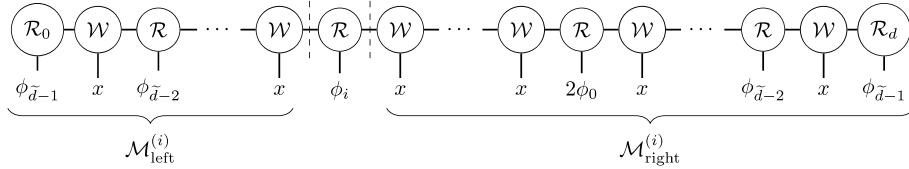


FIG. 4. A graphical visualization of the isolation and grouping when evaluating the derivative $\langle 0 | \partial_{\phi_i} U(x, \Phi) | 0 \rangle$.

Algorithm 3.3. Compute Jacobian matrix $DF(\Phi)$ using the MPS structure.

Input: Reduced phase factors Φ of length \tilde{d} and parity.
 Set $d = 2\tilde{d} - 2 +$ parity and initialize \mathbf{g} as a zero matrix of size $\tilde{d} \times (2d + 1)$.
for $j = 0, \dots, d$ **do**
 Set $x_j = \cos\left(\frac{2\pi j}{2d+1}\right)$.
 $\mathcal{M}_{\text{left}}(x_j) = (1, 0) \prod_{i=\tilde{d}-1}^1 (e^{i\phi_i Z} W(x_j))$.
 $\mathcal{M}_{\text{right}}(x_j) = e^{i\phi_0 Z} \mathcal{M}_{\text{left}}(x_j)^\top$.
if parity is odd **then**
 $\mathcal{M}_{\text{right}}(x_j) = W(x_j) \mathcal{M}_{\text{right}}(x_j)$.
end if
 $g_{0,j} \leftarrow 2\text{Im}[\mathcal{M}_{\text{left}} iZ e^{i\phi_0 Z} \mathcal{M}_{\text{right}}]$.
for $i = 1, \dots, \tilde{d} - 1$ **do**
 Update intermediate quantities according to (3.10):
 $\mathcal{M}_{\text{left}}(x_j) \leftarrow \mathcal{M}_{\text{left}}(x_j) W^{-1}(x_j) e^{-i\phi_i Z}$.
 $\mathcal{M}_{\text{right}}(x_j) \leftarrow W(x_j) e^{i\phi_{i-1} Z} \mathcal{M}_{\text{right}}(x_j)$.
 Compute $g_{i,j} \leftarrow 2\text{Im}[\mathcal{M}_{\text{left}} iZ e^{i\phi_i Z} \mathcal{M}_{\text{right}}]$ according to (3.9).
end for
end for
Evaluate $\mathbf{g}_{i,j} \leftarrow \mathbf{g}_{i,2d+1-j}$, $j = d + 1, \dots, 2d$.
Compute $\mathbf{v}_{il} \leftarrow \text{Re}(\sum_{j=0}^{2d-1} \mathbf{g}_{i,j} e^{-i\frac{2\pi}{2d+1}lj})$, $l = 0, \dots, d$, using FFT.
if parity = 0 **then**
 $\frac{\partial F(\Phi)}{\partial \phi_i} \leftarrow \frac{2}{2d+1} (\frac{\mathbf{v}_{i0}}{2}, \mathbf{v}_{i2}, \mathbf{v}_{i4}, \dots, \mathbf{v}_{id})$.
else
 $\frac{\partial F(\Phi)}{\partial \phi_i} \leftarrow \frac{2}{2d+1} (\mathbf{v}_{i1}, \mathbf{v}_{i3}, \mathbf{v}_{i5}, \dots, \mathbf{v}_{id})$.
end if
Output: $DF(\Phi)$.

By utilizing intermediate quantities, the computation of the derivatives, which are the columns of the Jacobian matrix before FFT, can be performed simultaneously, resulting in a computational cost of $\mathcal{O}(d^2)$ rather than $\mathcal{O}(d^3)$ in the straightforward method. The overall complexity of computing the Jacobian matrix is $\mathcal{O}(d^2 \log d)$ due to the use of FFT. The detailed procedure is summarized in Algorithm 3.3.

3.3. Formalism of symmetric QSP in real arithmetic operations. In the existing literature, the conventional formalism of QSP is typically presented in terms of the product of $SU(2)$ matrices, which involves complex arithmetic operations. This complex arithmetic formalism is both necessary and sufficient for general QSP, as the induced polynomials P and Q are complex without any additional symmetry

constraints. However, in the case of symmetric QSP, according to Theorem 2.1, the polynomial Q is a real polynomial. This observation raises the question of whether the formalism of QSP can be simplified to accommodate this symmetry.

In this subsection, we will introduce a formalism for symmetric QSP that utilizes real arithmetic operations. This alternative formalism not only proves to be beneficial for Newton's method proposed in this paper but also enhances the implementation of other algorithms designed to solve symmetric QSP, resulting in a constant improvement in the prefactor of the overall computational complexity.

The core idea is that $SU(2)$ is homeomorphic to $\mathbb{S}^3 \subset \mathbb{R}^4$, which arises from the parametric form of general $SU(2)$ matrices. By imposing the symmetric constraint, the upper-right entry of the consequent $SU(2)$ matrix is purely imaginary. Consequently, we can associate any symmetric QSP matrix with a vector in $\mathbb{S}^2 \subset \mathbb{R}^3$. The identification is

$$(3.11) \quad U(x, \Phi) = \begin{pmatrix} p(x, \Phi) + ig(x, \Phi) & i\sqrt{1-x^2}q(x, \Phi) \\ i\sqrt{1-x^2}q(x, \Phi) & p(x, \Phi) - ig(x, \Phi) \end{pmatrix} \in SU(2)$$

$$\leftrightarrow V(x, \Phi) := (p(x, \Phi), g(x, \Phi), \sqrt{1-x^2}q(x, \Phi))^\top \in \mathbb{S}^2.$$

Under the identification we introduced, the matrix multiplication in symmetric QSP is equivalent to interleaved rotations in $SO(3)$. This relation is quantified by the following recurrence equation:

$$(3.12) \quad V(x, (\phi_k, \phi_{k-1}, \dots)) = R_z(2\phi_k)R_x(2\arccos(x))V(x, (\phi_{k-1}, \dots)),$$

where the $SO(3)$ rotations are

$$(3.13) \quad R_z(2\phi) = \begin{pmatrix} \cos 2\phi & -\sin 2\phi & 0 \\ \sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 1 \end{pmatrix} \text{ and } R_x(2\theta) = \begin{pmatrix} \cos(2\theta) & 0 & -\sin(2\theta) \\ 0 & 1 & 0 \\ \sin(2\theta) & 0 & \cos(2\theta) \end{pmatrix}.$$

For further details on this identification, we refer readers to Appendix A. By leveraging this identification, the QSP polynomials can be derived from the product of real matrices, leading to a faster computation with a constant improvement in the prefactor, compared to evaluating them using the product of complex matrices. This real arithmetic improvement can also apply to the LBFGS method and the FPI method.

4. Experiments. In this section, we demonstrate the performance of Newton's method in solving phase factors through various numerical examples. These examples are essential for solving scientific computing problems using quantum algorithms. We begin by introducing the numerical examples, followed by the presentation and discussion of the numerical results in the rest of this section.

4.1. Setup of numerical examples. The core of a quantum algorithm design based on QSP lies in the abstraction of the original problem as a matrix function transformation. This transformation allows us to encode the desired function or its polynomial approximation by finding the appropriate phase factors. In order to illustrate this procedure, we present the following examples.

Quantum Hamiltonian simulation. The problem of quantum Hamiltonian simulation involves finding an efficient method for implementing the time evolution matrix of a Hamiltonian matrix, denoted as $H \mapsto \exp(-i\tau H)$, for a given evolution time τ . In [13], a near-optimal quantum Hamiltonian simulation algorithm based on QSP is proposed. This algorithm abstracts the problem into a function approximation task,

where the target function $f(x) = e^{-i\tau x}$ is parametrized using QSP. The Chebyshev series expansion, known as the Jacobi–Anger expansion [13], is commonly employed to approximate this target function:

$$(4.1) \quad e^{-i\tau x} = J_0(x) + 2 \sum_{k \text{ even}} (-1)^{k/2} J_k(\tau) T_k(x) + 2i \sum_{k \text{ odd}} (-1)^{(k-1)/2} J_k(\tau) T_k(x),$$

where J_k 's are the Bessel functions of the first kind. As a result, by truncating the Jacobi–Anger series, a polynomial approximation of the target function can be obtained. The real and imaginary parts of the truncated series, which approximate $\cos(\tau x)$ and $\sin(\tau x)$, respectively, serve as the target polynomials for two separate QSP phase-evaluation problems. To ensure that the truncation error is upper-bounded by ϵ_0 , it is sufficient to choose the degree of truncation as $d = e|\tau|/2 + \log(1/\epsilon_0)$.

Quantum Gaussian filter. The quantum Gaussian filter is a matrix function parameterized by μ and σ . It is proportional to $\exp(-(H - \mu I)^2/\sigma^2)$, where H is the Hamiltonian matrix. This matrix function is designed to localize around the given “energy level” μ , with the degree of localization controlled by the bandwidth parameter σ . The function suppresses eigenvalues of H that are far from μ . Ideally, one would choose μ to be close to an eigenvalue of H , allowing the matrix function to approximate the projection onto the corresponding eigenspace.

The quantum Gaussian filter serves as an intermediate subroutine for near-optimal quantum linear system solvers [12]. However, directly decomposing the defining function of the quantum Gaussian filter may result in exponentially large scaling factors due to hyperbolic functions. To address this issue and improve numerical stability, one can shift and rescale the Hamiltonian so that its eigenvalues lie in a smaller subinterval $D_\kappa = [1/\kappa, 1]$ within the positive half-axis. By employing this eigenvalue shifting technique, it is sufficient to approximate the Gaussian density function in the positive half-axis. Thus, the target function is set to $f(x) = e^{-(|x| - \mu)^2/\sigma^2}$ as an even extension.

Heaviside energy filter. The Heaviside function is widely used in classical applications such as signal processing and filter design. It also plays a crucial role as a subroutine in quantum algorithms for ground-state energy estimation and ground-state preparation [6].

Consider a Hamiltonian matrix that has been shifted and scaled so that its eigenvalues lie in the interval $[0, 1]$. The Heaviside energy filter $f(H)$ attenuates the high-energy components of the Hamiltonian. The function $f(x)$ is defined as follows:

$$(4.2) \quad f(x) = \begin{cases} 1, & |x| < 0.5, \\ \frac{1}{2}, & |x| = 0.5, \\ 0, & |x| > 0.5. \end{cases}$$

To address the singularity at 0.5, we assume that the target function only needs to be approximated within the interval $D_\delta = [0, (1 - \delta)/2] \cup [(1 + \delta)/2, 1]$. This allows us to focus on the desired energy range and mitigate the effects of the singularity.

Matrix inversion. Matrix inversion is a fundamental topic in numerical linear algebra with wide-ranging applications, including numerical optimization and least squares problems. In the context of function transformation, the equivalent problem is to implement the transformation $H \mapsto f(H) = H^{-1}$. If the matrix has a condition number of $\kappa = \text{cond}(H)$, it suffices to approximate the target function $f(x) = 1/x$ on the interval $D_\kappa = [1/\kappa, 1]$ using an odd function. This allows us to focus on the desired range of the function and effectively approximate the matrix inversion operation.

4.2. Constructing target polynomials approximating target functions.

In order to apply QSP to solve the problems discussed above, we need to find polynomial approximations to the target functions that satisfy the conditions outlined in Theorem 2.1. Various methods have been proposed in the literature for constructing these polynomial approximations in a streamlined manner.

One approach is to directly truncate the Chebyshev series expansion of the target function. This can be efficiently achieved using FFT applied to the transformed target function $f(\cos(\theta))$. However, when the target function is not defined on the entire interval $[-1, 1]$, the truncated series polynomial may not be bounded by 1 on the entire interval, making it unsuitable for representation using QSP. To address this issue, one approach is to use the Remez exchange algorithm proposed in [7] to find the best polynomial approximation for the partially defined target function. Another method involves numerically finding the best polynomial approximation using a convex-optimization-based approach as described in [6, section IV].

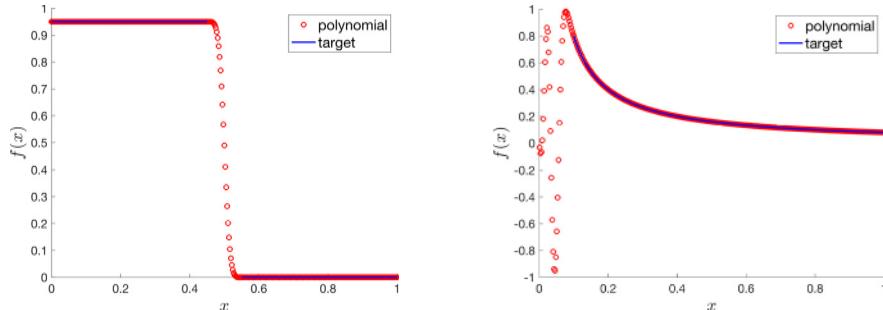
In the presented numerical examples, we use the truncated Chebyshev series for quantum Hamiltonian simulation and quantum Gaussian filter, where the target functions are defined on the interval $[-1, 1]$. For other examples where the target function is defined on a further subinterval, we employ the convex-optimization-based method to find the target polynomial approximation. The resulting target polynomials, obtained using the convex-optimization-based method, are visualized in Figure 5.

4.3. Numerical results. We evaluate the performance of Newton's method for finding phase factors in the presented numerical tests. All experiments are conducted using MATLAB R2020a on a computer with an Intel Core i5 Quad CPU running at 2.11 GHz and 8 GB of RAM.

The performance metrics used to evaluate the performance of Newton's method are the runtime and the residual error. The runtime refers to the amount of time it takes for the method to converge and find the desired phase factors. The residual error measures the discrepancy between the polynomial parametrized by the computed phase factors and the true polynomial which is defined as

$$(4.3) \quad \text{residual error} = \|F(\Phi) - c\|_1.$$

The numerical results for the error metric of Newton's method are presented in Figure 6. It is evident from the results that Newton's method exhibits significantly



(a) Heaviside energy filter function and its polynomial approximation with $\delta = 0.1$.

(b) Matrix inversion function and its polynomial approximation with $\kappa = 10$.

FIG. 5. Polynomial approximation of the target functions obtained by the convex-optimization-based method. These target functions and their related parameters are defined in subsection 4.1.

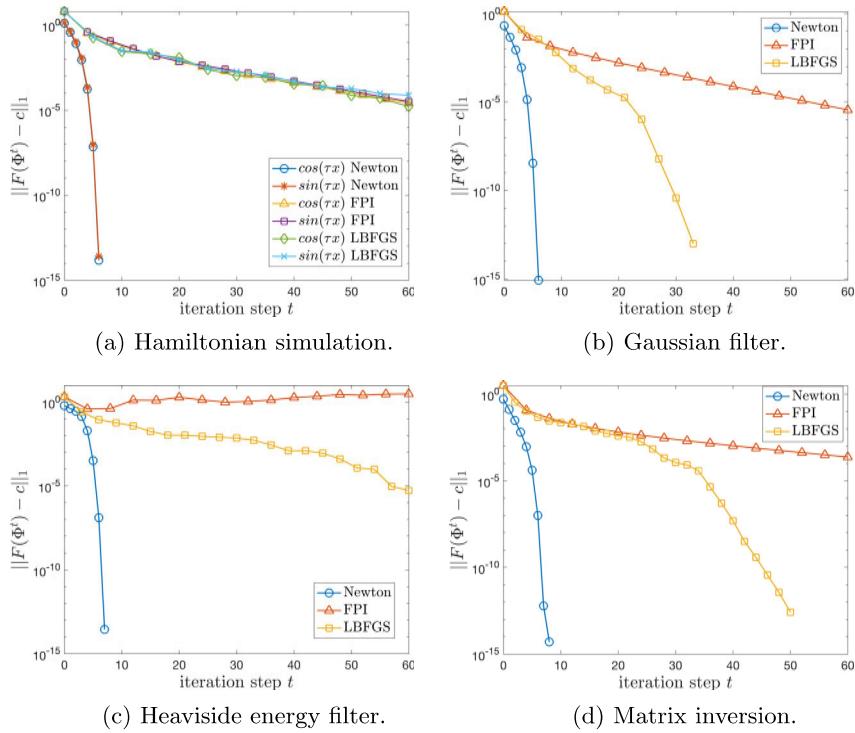


FIG. 6. Residual error as a function of iteration steps for different numerical examples. The target polynomials in (a)–(d) are chosen to be near the fully coherent regime, with the maximum absolute value set to 0.99 in (a)–(c) and 0.998 in (d). (a) Quantum Hamiltonian simulation with $\tau = 100$. The target polynomials, obtained by truncating the Jacobi–Anger series with truncation error $\epsilon_0 = 10^{-14}$, have degrees of 167 and 168. (b) Quantum Gaussian filter with $\mu = 0.5$ and $\sigma = 0.1$. The target polynomial is derived from the Chebyshev series expansion using FFT, resulting in a degree-100 polynomial. (c) Heaviside energy filter with $\delta = 0.1$. The target polynomial is a degree-250 polynomial obtained from a convex-optimization-based method. (d) Matrix inversion with $\kappa = 10$. The target polynomial is a degree-301 polynomial derived from a convex-optimization-based method.

faster convergence compared to other iterative methods for solving phase factors. The error curve aligns well with the expected quadratic convergence of Newton's method in general analysis. Besides, Newton's method exhibits greater stability in terms of runtime as the target function approaches the fully coherent regime. Figure 7 depicts the numerical results for the runtime of three iterative methods for determining phase factors. It also clearly illustrates the superior speed of Newton's method compared to the other two iterative methods.

To further analyze the performance of Newton's method near the fully coherent regime, the runtime and the number of iterations are plotted as a function of the distance to the fully coherent regime ($1 - \|f\|_\infty$) in Figure 8. It is noteworthy that even when the target function is extremely close to being fully coherent ($1 - \|f\|_\infty \leq 1 \times 10^{-9}$), Newton's method is capable of locating the optimum within a small number of iterations. This result highlights the robustness of Newton's method for finding phase factors in the nearly fully coherent regime.

Finally, we investigate the condition number of the Jacobian matrices at the phase factors obtained by Newton's method for different target functions, as presented

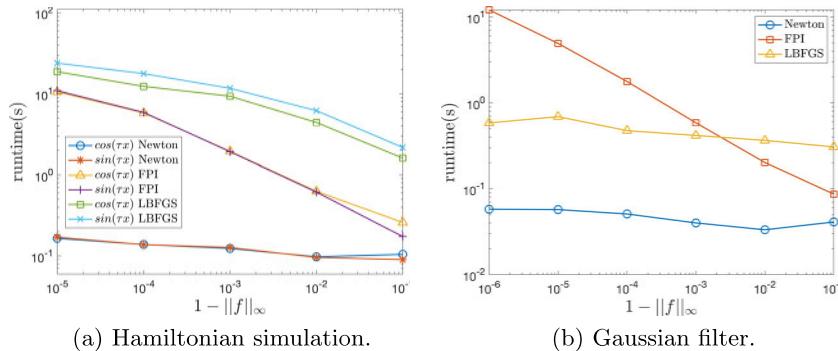


FIG. 7. Runtime analysis of numerical examples for different target functions. The target polynomials in (a) and (b) are scaled by a constant to make the problem increasingly close to the fully coherent regime $\|f\|_\infty \rightarrow 1$. (a) Quantum Hamiltonian simulation with $\tau = 100$ and truncation error $\epsilon_0 = 1 \times 10^{-14}$. The target polynomials have degrees of 167 and 168. (b) Quantum Gaussian filter with $\mu = 0.5$ and $\sigma = 0.01$. The target polynomial is of degree 100.

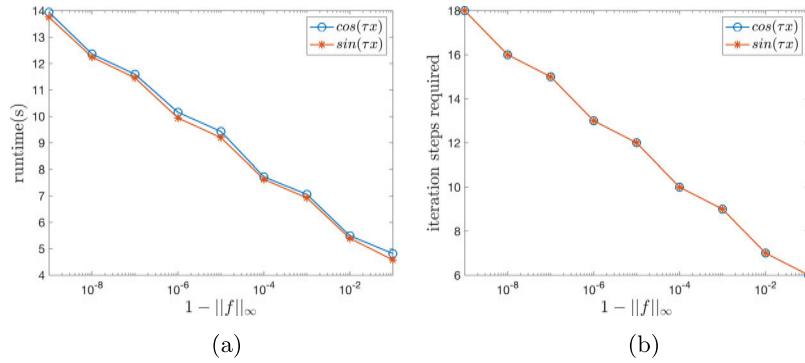


FIG. 8. Convergence analysis of Newton's method near the fully coherent regime. The target polynomial in each example is scaled by a constant to make the problem increasingly close to the fully coherent regime $\|f\|_\infty \rightarrow 1$. The problem is set to be quantum Hamiltonian simulation with $\tau = 1000$ and truncation error $\epsilon_0 = 1 \times 10^{-14}$. The degrees of the target polynomials are 1390 and 1391. (a) Runtime of Newton's method. (b) Number of iterations before convergence.

in Figure 9. The results indicate that as the target function approaches the fully coherent regime, the condition number of the Jacobian matrix becomes increasingly ill-conditioned. Despite this challenge, Newton's method continues to exhibit remarkable performance in finding phase factors. This emphasizes the effectiveness and reliability of Newton's method in phase factor determination, even in challenging scenarios near the fully coherent regime.

5. Conclusion. This paper presents a novel approach for solving the vector-valued, nonlinear system that arises in QSP using Newton's method. Numerical results indicate that the proposed method can robustly find phase factors in all parameter regimes, in particular the challenging fully coherent regime with ill-conditioned Jacobian matrices. Our method takes advantage of the matrix product states structure of QSP, enabling efficient computation of the Jacobian matrix. Additionally, the use of real-number arithmetics further enhances the prefactor of the numerical method. It is worth noting that these structure properties not only can be used to increase the efficiency of Newton's method, but also benefit the other algorithms including the

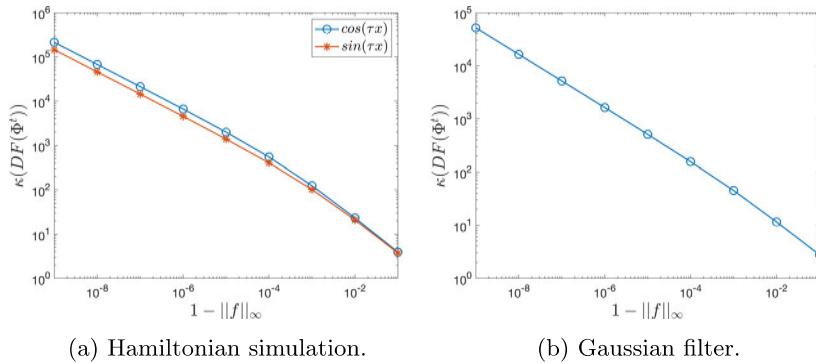


FIG. 9. Condition number of the Jacobian matrix at the numerical optimum. Each problem is solved by using Newton's method. The target polynomial in each example is scaled by a constant to make the problem increasingly close to the fully coherent regime. (a) Quantum Hamiltonian simulation with $\tau = 1000$ and truncation error $\epsilon_0 = 1 \times 10^{-14}$. The target polynomials have degrees of 1390 and 1391. (b) Quantum Gaussian filter with $\mu = 0.5$ and $\sigma = 0.01$. The target polynomial is of degree 100.

optimization-based (LBFGS) method and FPI. The Newton method, together with the option to enable real-number arithmetics computation in all aforementioned methods, have been implemented in the QSPPACK software package, providing a practical tool for solving QSP problems in scientific computing on quantum computers.

From a theoretical perspective, there are open problems regarding the impressive performance of Newton's method. While convergence in an ℓ^1 neighborhood of $\mathbf{0}$ can be understood using the same contraction mapping technique as in [5], the theoretical understanding of the effectiveness of the method in the fully coherent regime remains elusive. Additionally, extensive numerical experiments consistently converged to the maximal solution which is a special class of symmetric phase factor solutions proposed and studied in [25]. Further investigations are needed to understand the properties of the mapping F near this maximal solution.

Appendix A. Details about the formalism of symmetric QSP in real arithmetic operations. In the main text, we present a concise idea of the real-number arithmetic representation of QSP. In this section, we aim to provide a more comprehensive discussion and present additional details on this topic.

The computation of the QSP matrix boils down to that of a sequence of unitary matrix multiplications in (2.1). Furthermore, the QSP matrix admits the following decomposition as a consequence of Theorem 2.1

$$(A.1) \quad U(x, \Phi) = \begin{pmatrix} a_{\tilde{d}-1}(x) + id_{\tilde{d}-1}(x) & i\alpha_{\tilde{d}-1}(x) \\ i\alpha_{\tilde{d}-1}(x) & a_{\tilde{d}-1}(x) - id_{\tilde{d}-1}(x) \end{pmatrix}.$$

Here, $a_{\tilde{d}-1}(x)$, $d_{\tilde{d}-1}(x)$, and $\alpha_{\tilde{d}-1}(x)/\sqrt{1-x^2}$ are real polynomials in the variable x . According to the convention presented in the main text, $d_{\tilde{d}-1}$ stands for the component of interest, also known as $g(x, \Phi)$ in the main text to emphasize the dependence in phase factors Φ . As the goal in this section is to derive a simple recipe for computing the QSP matrix with a given set of phase factors Φ , we drop the Φ dependence in this section for the notational simplicity.

Let the entrywise value of the phase factors be $\Phi = (\phi_0, \phi_1, \dots, \phi_{\tilde{d}-1})$. For ease of discussion, we refer to $\Phi^{(k)} = (\phi_0, \phi_1, \dots, \phi_k)$ as the k th truncated phase factors

for each $k = 0, 1, \dots, \tilde{d} - 1$. The corresponding sequence of QSP matrices is denoted entrywise as

$$(A.2) \quad U(x, \Phi^{(k)}) = \begin{pmatrix} a_k(x) + id_k(x) & i\alpha_k(x) \\ i\alpha_k(x) & a_k(x) - id_k(x) \end{pmatrix}.$$

We remark that each truncated set of phase factors also gives a symmetric QSP. Hence, the decomposition in (A.1) applies, implying that $a_k(x)$, $\alpha_k(x)$, and $d_k(x)$ are well defined. By appending ϕ_k to the $(k - 1)$ th truncation $\Phi^{(k-1)}$, the recurrence relation follows

$$(A.3) \quad \begin{pmatrix} a_k(x) + id_k(x) & i\alpha_k(x) \\ i\alpha_k(x) & a_k(x) - id_k(x) \end{pmatrix} \\ = e^{i\phi_k Z} W(x) \begin{pmatrix} a_{k-1}(x) + id_{k-1}(x) & i\alpha_{k-1}(x) \\ i\alpha_{k-1}(x) & a_{k-1}(x) - id_{k-1}(x) \end{pmatrix} W(x) e^{i\phi_k Z}.$$

It can be verified that the following rearrangement is equivalent to the recurrence relation:

$$(A.4) \quad \begin{pmatrix} a_k(x) \\ d_k(x) \\ \alpha_k(x) \end{pmatrix} = R_z(2\phi_k) R_x(2 \arccos(x)) \begin{pmatrix} a_{k-1}(x) \\ d_{k-1}(x) \\ \alpha_{k-1}(x) \end{pmatrix},$$

where

$$(A.5) \quad R_z(2\phi) = \begin{pmatrix} \cos 2\phi & -\sin 2\phi & 0 \\ \sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

and $R_x(2 \arccos(x)) = \begin{pmatrix} 2x^2 - 1 & -2x\sqrt{1-x^2} & 0 \\ 2x\sqrt{1-x^2} & 1 & 0 \\ 0 & 0 & 2x^2 - 1 \end{pmatrix}$

are the induced $\text{SO}(3)$ rotation matrices. It can also be shown that the base cases of the recurrence are

(1) when d is even

$$(A.6) \quad \begin{pmatrix} a_0(x) \\ d_0(x) \\ \alpha_0(x) \end{pmatrix} = \begin{pmatrix} \cos(2\phi_0) \\ \sin(2\phi_0) \\ 0 \end{pmatrix},$$

and (2) when d is odd

$$(A.7) \quad \begin{pmatrix} a_0(x) \\ d_0(x) \\ \alpha_0(x) \end{pmatrix} = \begin{pmatrix} \cos(2\phi_0)x \\ \sin(2\phi_0)x \\ \sqrt{1-x^2} \end{pmatrix}.$$

Remarkably, the equivalent recurrence relation (A.4) involves only real quantities, and R_x and R_z only actively act as a rotation on two entries. In contrast to the complex recurrence relation (A.3), the real recurrence has lower time and space complexity. This improvement is due to the simplified structure of symmetric QSP compared with the original formalism without symmetry.

The MPS/TT structure still holds in the real recurrence relation. We refer \mathcal{X} and \mathcal{Z} to the parametric order-2 tensor standing for the induced $\text{SO}(3)$ rotations,

namely, $\mathcal{Z}(\phi) = R_z(2\phi)$ and $\mathcal{X}(x) = R_x(2\arccos(x))$. Let \mathcal{I} be the order-1 tensor representing the base of the recurrence in (A.6) and (A.7). Furthermore, to extract the component of the computational interest, let \mathcal{H} be the order-1 tensor representing the last operation, which is

$$(A.8) \quad \mathcal{H}(\phi_{\tilde{d}-1}) = (0 \ 1 \ 0) R_z(2\phi_{\tilde{d}-1}) = (\sin(2\phi_{\tilde{d}-1}) \ \cos(2\phi_{\tilde{d}-1}) \ 0).$$

Then, the recurrence relation in real-number arithmetic can be visualized graphically in Figure 10. In contrast to the computation in the complex-arithmetic representation, the symmetry constraint of the QSP phase factors is reflected in the doubled argument in the Z tensor of the real-number arithmetic representation. Hence, when computing the derivative, it does not need tricks to arrange the derivatives coming from two symmetric sites. Specifically, the following identity holds:

$$(A.9) \quad \frac{\partial g(x, \Phi)}{\partial \phi_i} = \mathcal{N}_{\text{left}}^{(i)} \frac{d\mathcal{Z}(\phi_i)}{d\phi_i} \mathcal{N}_{\text{right}}^{(i)} = 2\mathcal{N}_{\text{left}}^{(i)} \mathcal{Z}(\phi_i + \pi/4) \mathcal{N}_{\text{right}}^{(i)}.$$

Here, the left and right parts under the partition are given by

$$(A.10) \quad \mathcal{N}_{\text{left}}^{(i)} := \mathcal{H} \prod_{j=\tilde{d}-2}^{i+1} (\mathcal{X}(x) \mathcal{Z}(\phi_j)) \mathcal{X}(x), \quad \mathcal{N}_{\text{right}}^{(i)} := \prod_{j=i-1}^1 \mathcal{X}(x) \mathcal{Z}(\phi_j) \mathcal{I},$$

whose graphical visualizations are presented in Figure 10. The update of these quantities in the computational process is

$$(A.11) \quad \mathcal{N}_{\text{left}}^{(i+1)} \leftarrow \mathcal{N}_{\text{left}}^{(i)} \mathcal{X}^{-1}(x) \mathcal{Z}(-\phi_{i+1}), \quad \text{and} \quad \mathcal{N}_{\text{right}}^{(i+1)} \leftarrow \mathcal{Z}(\phi_i) \mathcal{X}(x) \mathcal{N}_{\text{right}}^{(i)}.$$

For completeness, we provide the algorithm for computing the Jacobian matrix using the MPS/TT structure and the real-number arithmetic representation in Algorithm A.1.

In Figure 11, we numerically demonstrate that using the real-number arithmetic formalism of QSP improves the time complexity of iterative methods by a constant prefactor. Notably, this improvement is not limited to Newton's method but also applies to other iterative methods for finding phase factors.

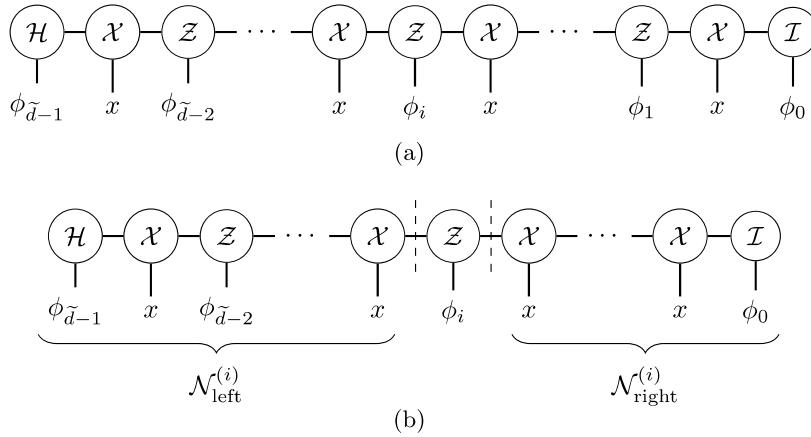


FIG. 10. A graphical visualization of the MPS/TT structure of the problem in the real-number arithmetic representation. (a) The structure of the recurrence relation. (b) The partition when computing the Jacobian.

Algorithm A.1. computing the Jacobian matrix using the MPS/TT structure and the real-number arithmetic representation.

Input: A set of reduced phase factors Φ , its length \tilde{d} , and its parity $p \in \{0, 1\}$.

Set $d = 2\tilde{d} - 2 + p$ and initialize \mathbf{g} as a zero matrix of size $\tilde{d} \times (2d + 1)$.

for $j = 0, \dots, d$ **do**

 Set $x_j = \cos\left(\frac{2\pi j}{2d+1}\right)$.

$\mathcal{N}_{\text{left}}(x_j) = \mathcal{H} \prod_{i=\tilde{d}-2}^1 (\mathcal{X}(x_j) \mathcal{Z}(\phi_i)) \mathcal{X}(x_j)$.

if parity is even ($p = 0$) **then**

$\mathcal{N}_{\text{right}}(x_j) = (1, 1, 0)^\top$.

else

$\mathcal{N}_{\text{right}}(x_j) = (x, x, \sqrt{1 - x^2})^\top$.

end if

$g_{0,j} \leftarrow 2\mathcal{N}_{\text{left}} \mathcal{Z}(\phi_0 + \pi/4) \mathcal{N}_{\text{right}}(x_j)$.

for $i = 1, \dots, \tilde{d} - 1$ **do**

$\mathcal{N}_{\text{left}}(x_j) \leftarrow \mathcal{N}_{\text{left}}(x_j) \mathcal{X}^{-1}(x_j) \mathcal{Z}(-\phi_i)$.

$\mathcal{N}_{\text{right}}(x_j) \leftarrow \mathcal{X}(x_j) \mathcal{Z}(\phi_{i-1}) \mathcal{N}_{\text{right}}(x_j)$.

$g_{i,j} \leftarrow 2\mathcal{N}_{\text{left}}(x_j) \mathcal{Z}(\phi_i + \pi/4) \mathcal{N}_{\text{right}}(x_j)$.

end for

end for

Set $\mathbf{g}_{i,j} \leftarrow \mathbf{g}_{i,2d+1-j}$, $j = d + 1, \dots, 2d$.

Compute $\mathbf{v}_{il} \leftarrow \text{Re}(\sum_{j=0}^{2d-1} \mathbf{g}_{i,j} e^{-i\frac{2\pi}{2d+1}lj})$, $l = 0, \dots, d$ by using FFT.

if parity is even ($p = 0$) **then**

$\frac{\partial F(\Phi)}{\partial \phi_i} \leftarrow \frac{2}{2d+1} (\frac{\mathbf{v}_{i0}}{2}, \mathbf{v}_{i2}, \mathbf{v}_{i4}, \dots, \mathbf{v}_{id})$.

else

$\frac{\partial F(\Phi)}{\partial \phi_i} \leftarrow \frac{2}{2d+1} (\mathbf{v}_{i1}, \mathbf{v}_{i3}, \mathbf{v}_{i5}, \dots, \mathbf{v}_{id})$.

end if

Output: $DF(\Phi)$.

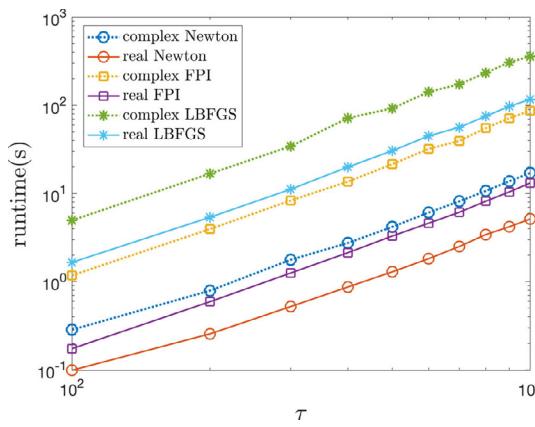


FIG. 11. Comparing the runtime of iterative methods for finding phase factors using the real-number and complex-number arithmetic. The problem is set to quantum Hamiltonian simulation with variable τ parameters. The target polynomial is derived by truncating the Jacobi–Anger series with truncation error $\epsilon_0 = 1 \times 10^{-14}$. The maximal value of the target polynomial is scaled by a constant so that $\|f\|_\infty = 0.9$.

REFERENCES

- [1] R. CHAO, D. DING, A. GILYEN, C. HUANG, AND M. SZEGEDY, *Finding Angles for Quantum Signal Processing with Machine Precision*, preprint, arXiv:2003.02831, 2020.
- [2] A. M. CHILDS, D. MASLOV, Y. NAM, N. J. ROSS, AND Y. SU, *Toward the first quantum simulation with quantum speedup*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 9456–9461.
- [3] Y. DONG, J. GROSS, AND M. Y. NIU, *Optimal Low-Depth Quantum Signal-Processing Phase Estimation*, preprint, arXiv:2407.01583, 2024.
- [4] Y. DONG AND L. LIN, *Random circuit block-encoded matrix and a proposal of quantum LINPACK benchmark*, Phys. Rev. A, 103 (2021), 062412.
- [5] Y. DONG, L. LIN, H. NI, AND J. WANG, *Infinite Quantum Signal Processing*, preprint, arXiv: 2209.10162, 2022.
- [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 (2022), 040305.
- [7] Y. DONG, X. MENG, K. B. WHALEY, AND L. LIN, *Efficient phase factor evaluation in quantum signal processing*, Phys. Rev. A, 103 (2021), 042419.
- [8] Y. DONG, K. B. WHALEY, AND L. LIN, *A quantum hamiltonian simulation benchmark*, npj Quantum Inf., 8 (2022), 131.
- [9] D. FANG, L. LIN, AND Y. TONG, *Time-marching based quantum solvers for time-dependent linear differential equations*, Quantum, 7 (2023), 955.
- [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, 2019, pp. 193–204.
- [11] J. HAAH, *Product decomposition of periodic functions in quantum signal processing*, Quantum, 3 (2019), 190.
- [12] L. LIN AND Y. TONG, *Optimal quantum eigenstate filtering with application to solving quantum linear systems*, Quantum, 4 (2020), 361.
- [13] G. H. LOW AND I. L. CHUANG, *Optimal hamiltonian simulation by quantum signal processing*, Phys. Rev. Lett., 118 (2017), 010501.
- [14] G. H. LOW AND I. L. CHUANG, *Hamiltonian simulation by qubitization*, Quantum, 3 (2019), 163.
- [15] J. M. MARTYN, Y. LIU, Z. E. CHIN, AND I. L. CHUANG, *Efficient fully-coherent quantum signal processing algorithms for real-time dynamics simulation*, J. Chem. Phys., 158 (2023), 024106.
- [16] J. M. MARTYN, Z. M. ROSSI, A. K. TAN, AND I. L. CHUANG, *Grand unification of quantum algorithms*, PRX Quantum, 2 (2021), 040203.
- [17] S. McARDLE, A. GILYÉN, AND M. BERTA, *Quantum State Preparation Without Coherent Arithmetic*, preprint, arXiv:2210.14892, 2022.
- [18] D. MOTLAGH AND N. WIEBE, *Generalized quantum signal processing*, PRX Quantum, 5 (2024), 020368.
- [19] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer-Verlag, Berlin, 1999.
- [20] I. V. OSELEDETS, *Tensor-train decomposition*, SIAM J. Sci. Comput., 33 (2011), pp. 2295–2317.
- [21] Z. M. ROSSI, V. M. BASTIDAS, W. J. MUNRO, AND I. L. CHUANG, *Quantum Signal Processing with Continuous Variables*, preprint, arXiv:2304.14383, 2023.
- [22] Z. M. ROSSI AND I. L. CHUANG, *Multivariable quantum signal processing (M-QSP): Prophecies of the two-headed oracle*, Quantum, 6 (2022), 811.
- [23] Z. M. ROSSI AND I. L. CHUANG, *Semantic embedding for quantum algorithms*, J. Math. Phys., 64 (2023), 122202.
- [24] T. DE LIMA SILVA, L. BORGES, AND L. AOLITA, *Fourier-Based Quantum Signal Processing*, preprint, arXiv:2206.02826, 2022.
- [25] J. WANG, Y. DONG, AND L. LIN, *On the energy landscape of symmetric quantum signal processing*, Quantum, 6 (2022), 850.
- [26] L. YING, *Stable factorization for phase factors of quantum signal processing*, Quantum, 6 (2022), 842.