

Random circuit block-encoded matrix and a proposal of quantum LINPACK benchmarkYulong Dong^{1,2} and Lin Lin^{3,4,*}¹*Berkeley Center for Quantum Information and Computation, Berkeley, California 94720, USA*²*Department of Chemistry, University of California, Berkeley, California 94720, USA*³*Department of Mathematics, University of California, Berkeley, California 94720, USA*⁴*Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

(Received 5 February 2021; accepted 6 April 2021; published 14 June 2021)

The LINPACK benchmark reports the performance of a computer for solving a system of linear equations with dense random matrices. Although this task was not designed with a real application directly in mind, the LINPACK benchmark has been used to define the list of TOP500 supercomputers since the debut of the list in 1993. We propose that a similar benchmark, called the quantum LINPACK benchmark, could be used to measure the whole machine performance of quantum computers. The success of the quantum LINPACK benchmark should be viewed as the minimal requirement for a quantum computer to perform a useful task of solving linear algebra problems, such as linear systems of equations. We propose an input model called the Random Circuit Block-Encoded Matrix (RACBEM), which is a proper generalization of a dense random matrix in the quantum setting. The RACBEM model is efficient to be implemented on a quantum computer and can be designed to optimally adapt to any given quantum architecture, with relying on a black-box quantum compiler. Besides solving linear systems, the RACBEM model can be used to perform a variety of linear algebra tasks relevant to many physical applications, such as computing spectral measures, time series generated by a Hamiltonian simulation, and thermal averages of the energy. We implement these linear algebra operations on IBM Q quantum devices as well as quantum virtual machines, and demonstrate their performance in solving scientific computing problems.

DOI: [10.1103/PhysRevA.103.062412](https://doi.org/10.1103/PhysRevA.103.062412)**I. INTRODUCTION**

Quantum computers hold the promise of dramatically accelerating calculations in a wide range of fields, and quantum supremacy was achieved in 2019 via sampling random quantum circuits [1]. Assume that there are 10 000 quantum computers (or many more) available now, how should we select the top 500 best performing computers for scientific computing applications? The answer in the context of classical supercomputers is given by the LINPACK benchmark [2], which measures the floating point computing power of a classical computer via its performance for solving linear systems of equations $Ax = b$. The input matrix A is a dense pseudorandom matrix, and there is no immediate application associated with such a matrix (similar to a quantum supremacy experiment in this sense). There has been much controversy over its effectiveness in measuring the capability of classical computers in scientific computing applications since the very beginning. However, LINPACK is widely used and performance numbers are available for almost all relevant systems. The LINPACK benchmark has been used as the defining criterion of TOP500 supercomputers since the debut of the list in 1993 [4]. One important reason is that dense matrices and dense matrix operations are relatively easy to implement and to optimize on classical computers. These operations have been tuned to be highly scalable, which

enabled the performance benchmark of systems that cover a performance range of 12 orders of magnitude in the past 20 years.

In order to mimic the success of the LINPACK benchmark on quantum computers, we consider the problem of solving the quantum linear system problem (QLSP). Many challenging high-dimensional problems in physics, such as computing Green's functions for a quantum many-body system, can be formulated in terms of QLSP. This field has witnessed significant progress in the past few years [5–17]. Shortly speaking, given $A \in \mathbb{C}^{2^n \times 2^n}$ and $|b\rangle \in \mathbb{C}^{2^n}$, QLSP aims at obtaining an n -qubit solution vector $|x\rangle \propto A^{-1}|b\rangle$. More precisely and using the language of block encoding [11], QLSP is the problem of finding an $(m+n)$ -qubit unitary matrix U , such that

$$|x\rangle = (\langle 0^m| \otimes I_n)U(|0^m\rangle \otimes |b\rangle) = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}. \quad (1)$$

In other words, the solution is obtained upon measuring 0 for all m ancilla qubits (with a success probability p).

In this paper, we propose the quantum LINPACK benchmark, which can be concisely stated as the problem of *using the quantum computer to evaluate the success probability p for a certain random matrix A* . While the rationale of such a task will be discussed in detail, we first emphasize that the quantum LINPACK benchmark examines the *whole machine performance* of quantum computers, rather than the performance of a few qubits as often measured by methods such as randomized benchmarking [18] and gateset tomography [19]. There have been a number of whole machine

*linlin@math.berkeley.edu




FIG. 1. A cartoon illustration of the RACBEM model.

quantum benchmarks proposed in the literature, such as quantum volume [20], cycle benchmarking [21], and the linear cross-entropy benchmarking as in the supremacy experiment [1,22]. However, those benchmark methods are proposed for generic settings, and therefore they are less representative for the performance on the user-specified applications. In particular, the error of a structured circuit can be significantly larger than that of a fully randomized one [23]. The quantum LINPACK benchmark targets directly at the performance of quantum computers for scientific computing applications, as in the case of the LINPACK benchmark for classical supercomputers. We emphasize that the success of the quantum LINPACK benchmark does not guarantee that the quantum computer has solved $|x\rangle$ accurately, but should be viewed as the minimal requirement for a quantum computer to solve QLSP. Given the wide range of potential applications of QLSP from quantum many-body problems to quantum machine learning, it is important for future quantum computers to first meet the criterion of the quantum LINPACK benchmark, in order to achieve quantum advantage via the path of solving linear algebra problems.

In order to perform the quantum LINPACK benchmark, note that it would be highly inefficient if we first generate a dense pseudorandom matrix A classically and then feed it into the quantum computer using, e.g., QRAM [24]. In fact, such a strong assumption on the input model often lead to dequantized classical algorithms [25]. Instead we focus on matrices that are *inherently easy* to generate on quantum computers. In particular, the supremacy experiment inspires us to generate a random matrix directly using a random quantum circuit.

We propose an input model called the Random Circuit Block-Encoded Matrix (RACBEM). We argue that the RACBEM model is a proper generalization of dense random matrices in the quantum setting, suitable for linear algebra tasks. The RACBEM model, and its Hermitian version called H-RACBEM model, are simple to construct and allow us to get access to in principle any n -qubit matrix and n -qubit Hermitian matrix, respectively (up to a scaling factor) by adding only one ancilla qubit. Together with the recently developed technique of quantum singular value transformation (QSVT) [11], we yield a practical algorithm for performing

the quantum LINPACK benchmark on near-term devices with a shallow circuit depth.

With QSVT and the H-RACBEM model, the circuit used in the quantum LINPACK benchmark can be designed to adapt to the coupling map of almost any given gate-based quantum architecture. All operations can be carried out with straightforward usage of basic one-qubit gates and CNOT gates, and there is no complex controlled unitaries involved. Due to the use of the basic gate set and the adaptivity to the quantum architecture, the quantum LINPACK benchmark does not require the explicit use of the compiler, while the randomized benchmarking requires gate compilation. Furthermore, the number of ancilla qubits needed is minimal (usually two). By using the H-RACBEM model, the condition number of the random matrices is fully controllable which is crucial for reducing the circuit depth. While the development of quantum hardware has been very rapid in the past few years, quantum resources are expected to remain costly and limited for some time, with or without the fault-tolerant capability. The H-RACBEM model can also significantly reduce the efforts needed to optimize and to compile the QSVT algorithm. The RACBEM and its Hermitian version provides a simple and reliable way to generate random matrices on a quantum computer with a minimal number of ancilla qubits and a fully controllable circuit depth. Therefore, we expect that the quantum LINPACK benchmark uses the minimal circuit to gauge the performance of a quantum computer for solving linear algebra problems.

Furthermore, we demonstrate that using *the same quantum circuit but with different parameters*, the combination of QSVT and the H-RACBEM model can be used to perform many other linear algebra tasks, such as computing spectral measures and performing time series analysis (without Trotter splitting). Using the minimally entangled typical thermal state (METTS) algorithm [26–28], we also show how H-RACBEM simplifies the computation of the thermal average of the energy. These linear algebra tasks can also be used to construct benchmarks similar to the quantum LINPACK benchmark, whose performance reflects the minimal requirement for a quantum computer in solving corresponding scientific computing problems. We implement all algorithms on the IBM Q quantum architecture. Due to the high noise level of the currently available quantum architecture, we also demonstrate the numerical performance using quantum virtual machines (QVM) with tunable, approximate error models derived from quantum devices.

II. RANDOM CIRCUIT BASED BLOCK-ENCODING MATRIX

A. Block encoding

Inherently, quantum computers can handle only unitary operators. Hence any nonunitary operators must be encoded in terms of unitary operators. Let $A \in \mathbb{C}^{N \times N}$ be an n -qubit Hermitian matrix ($N = 2^n$). If we can find an $(n+1)$ -qubit unitary matrix U_A such that

$$U_A = \begin{pmatrix} A & \cdot \\ \cdot & \cdot \end{pmatrix} \quad (2)$$

holds, i.e., A is the upper-left matrix block of U_A , then we may get access to A via the unitary matrix U_A with $A = (\langle 0| \otimes I_n) U_A (|0\rangle \otimes I_n)$. Clearly when the operator norm $\|A\|_2$

is larger than 1, A cannot be encoded by any unitary U_A as in (2). Generally if we can find $\alpha, \epsilon \in \mathbb{R}_+$, and an $(m+n)$ -qubit matrix U_A such that

$$\|A - \alpha(|0^m\rangle \otimes I_n)U_A(|0^m\rangle \otimes I_n)\| \leq \epsilon, \quad (3)$$

then U_A is called an (α, m, ϵ) block encoding of A . Here m is called the number of ancilla qubits for block encoding. The block encoding is a powerful and versatile model, which can be used to efficiently encode density operators, Gram matrices, positive-operator valued measure (POVM), sparse-access matrices, as well as addition and multiplication of block-encoded matrices (we refer to [11] for a detailed illustration of such constructions).

The block-encoding model also has its limitation. Take an n -qubit, d -sparse matrix A (i.e., the number of nonzero entries in each row or column does not exceed d) for example, assuming access to its row and column entries via certain sparse-access oracles, then we can construct a block encoding using $n+3$ ancilla qubits [11]. Any further manipulation of U_A , such as quantum signal processing, would require using $(n+4)$ -qubit Toffoli gates, which are relatively expensive to implement [29,30]. Another example is that A is given by a linear combination of K terms, each term being a tensor product of Pauli matrices. In the setting of the Hamiltonian simulation, e^{iAt} can be simply approximated by exponentiating each term individually following a certain order via the Trotter-Suzuki formula. However, the block-encoding model would essentially require using linear combination of unitaries [31], which not only requires $\lceil \log_2 K \rceil$ ancilla qubits, but also usage of $(\lceil \log_2 K \rceil + 1)$ -qubit Toffoli gates to implement the prepare and select oracles needed to obtain the linear combination. Such operations are essentially forbidden for near-term applications due to high error rates and are still challenging when the number of qubits and the gate depth remain a limitation in reaching the desired accuracy on fault-tolerant devices.

B. RACBEM

To harness the power of the block-encoding model and to avoid its pitfalls, we propose the Random Circuit Block-Encoded Matrix (RACBEM) model as follows. Instead of first identifying A and then finding its block encoding U_A , we reverse this thought process: we first identify a unitary U_A that is easy to implement on a quantum computer, and then ask which matrix can be block encoded by U_A .

It turns out that any matrix A with $\|A\|_2 \leq 1$ can be given by a $(1,1,0)$ block encoding. Consider any n -qubit matrix with its singular value decomposition (SVD) $A = W\Sigma V^\dagger$, where all singular values in Σ belong to $[0,1]$. Then we can construct an $(n+1)$ -qubit unitary matrix

$$\begin{aligned} U_A &:= \begin{pmatrix} A & W\sqrt{I_n - \Sigma^2} \\ \sqrt{I_n - \Sigma^2}V^\dagger & -\Sigma \end{pmatrix} \\ &= \begin{pmatrix} W & 0 \\ 0 & I_n \end{pmatrix} \begin{pmatrix} \Sigma & \sqrt{I_n - \Sigma^2} \\ \sqrt{I_n - \Sigma^2} & -\Sigma \end{pmatrix} \begin{pmatrix} V^\dagger & 0 \\ 0 & I_n \end{pmatrix}, \end{aligned}$$

which is a $(1,1,0)$ -block encoding of A . Since a random circuit with $\text{poly}(n)$ depth can approximate an n -qubit Haar measure at least according to the criterion of the 2-design [32], a




FIG. 2. Coupling map of quantum computing backend `ibmq_burlington` provided on IBM Q website [3]. The encircled number stands for the label of qubit where the color represents the noise rate of one-qubit gates on it. If two qubits are related by an arrow, the CNOT gate is directly available between these two qubits.

sufficiently general $(n+1)$ -qubit unitary U_A can give access to in principle any n -qubit nonunitary matrices A (up to a scaling factor). Furthermore, such a random circuit U_A can be constructed using only basic one-qubit unitaries and CNOT gates. The matrix A obtained by measuring the first qubit (or in fact, any qubit used as the ancilla) is called a RACBEM. Since the Haar measure is the uniform distribution of unitary matrices, we conclude that RACBEM is a proper generalization of dense matrices on quantum computers suitable for performing linear algebra tasks. The layout of the two-qubit operations can be designed to be compatible with the coupling map of the hardware. We give a cartoon illustration of the RACBEM model in Fig. 1.

For instance, for the `ibmq_burlington` device with its coupling map shown in Fig. 2, we can choose qubit 1 as the RACBEM ancilla qubit, which results in a RACBEM model with three system qubits 2,3,4. The qubit 0 is not used here and is reserved as the signal qubit for quantum singular value transformation to be discussed later. An example of the RACBEM circuit is shown in Fig. 3, which is constructed with respect to the coupling map in Fig. 2.

C. H-RACBEM

In many applications, such as Hamiltonian simulation, thermal state preparation, etc., we are interested only in Hermitian matrices. It is possible to find a general circuit U_A that coincidentally block encodes a Hermitian matrix, but this can become increasingly difficult as n increases. A useful fact is that once a random circuit U_A is given, its Hermitian conjugate U_A^\dagger is easily accessible by conjugating the each gate and reversing the gate sequence. We will show below that this allows us to get access to in principle any n -qubit Hermitian matrix.

Consider the quantum circuit in Fig. 4 denoted by $U_{\mathfrak{H}}$, where $\varphi_0, \varphi_1 \in [-\pi, \pi)$. Direct calculation (see Appendix A) shows that

$$\begin{aligned} \mathfrak{H} &= (|0^2\rangle \otimes I_n)U_{\mathfrak{H}}(|0^2\rangle \otimes I_n) \\ &= [-2 \sin(2\varphi_0) \sin \varphi_1]A^\dagger A + \cos(2\varphi_0 - \varphi_1). \end{aligned} \quad (4)$$

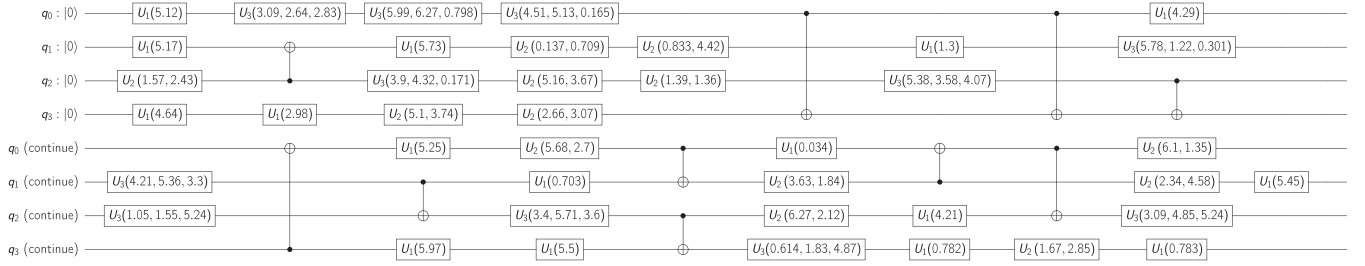


FIG. 3. A RACBEM circuit constructed using the basic gate set $\{U_1, U_2, U_3, \text{CNOT}\}$. The CNOT gates are directly implementable according to the coupling map in (a). q_0, q_1, q_2, q_3 refer to qubits 1,2,3,4 in (a), where qubit 0 is excluded as a signal qubit. The circuit at the bottom is a continuation of the top circuit. The block-encoded matrix and its singular values are given in Fig. 14.

Here \mathfrak{H} is a Hermitian matrix. We refer to it as a Hermitian RACBEM (H-RACBEM), and $U_{\mathfrak{H}}$ is its $(1,2,0)$ -block encoding. In particular, choosing $\varphi_0 = \pi/8$, $\varphi_1 = -\pi/4$, then $\mathfrak{H} = A^\dagger A$ is Hermitian positive semidefinite. This will be referred to as a canonical H-RACBEM. In other words, a canonical H-RACBEM is constructed from its (nonunique) matrix square root A .

In Fig. 4, the CNOT gate controlling on 0 instead of 1 is mainly for notational convenience, and in fact not all CNOT gates are necessary here. For example, in order to implement a canonical H-RACBEM, we need only one application of U_A , one application of U_A^\dagger , two H gates, two standard CNOT gates, one S^\dagger gate, and two T gates (see Appendix A). Since any matrix with singular values bounded by 1 can be represented as a RACBEM, we immediately have that any Hermitian positive semidefinite matrix with eigenvalues bounded by 1 can be represented as a canonical H-RACBEM, with a sufficiently flexible U_A .

D. Quantum singular value transformation

Given the SVD $A = W\Sigma V^\dagger$, and a smooth function $f(x)$ of even parity, we define the quantum singular value transformation (QSVT) as

$$f^\triangleright(A) := V f(\Sigma) V^\dagger. \quad (5)$$

Here the right-pointing triangle reflects that only the right singular vectors V are kept. Clearly $f^\triangleright(A) = f(\sqrt{A^\dagger A})$, where the right-hand side is the standard matrix function. Now let f be a real even polynomial of degree $2d$ that satisfies $|f(x)| \leq 1$ for any $x \in [-1, 1]$. Let U_A be a $(1, m, 0)$ block encoding of A . Then following [[11], Corollary 11], there exists a $(1, m+1, 0)$ block encoding of $f^\triangleright(A)$, denoted by $U_{f^\triangleright(A)}$. The circuit to implement $U_{f^\triangleright(A)}$ is given in Fig. 5, which can be constructed using d queries of U_A and d queries of U_A^\dagger . Here $\Phi := (\varphi_0, \dots, \varphi_{2d})$ are called the phase factors. One

challenge in QSVT is to find the phase factors Φ . Besides the methods for obtaining Φ by polynomial factorization [11,33], recently an optimization-based method is proposed to find Φ up to very high degrees [34]. A brief description of the method is given in Appendix B.

Therefore Fig. 4 implements a QSVT for a second-order polynomial $\mathfrak{H} = h^\triangleright(A)$ with a symmetric choice of phase factors $\Phi = (\varphi_0, \varphi_1, \varphi_0)$, where

$$h(x) = [-2 \sin(2\varphi_0) \sin \varphi_1] x^2 + \cos(2\varphi_0 - \varphi_1). \quad (6)$$

A canonical H-RACBEM is given by $h(x) = x^2$.

Consider any real polynomial $g(x)$ of degree d without parity constraint, satisfying $|g(x)| \leq 1$ for any $x \in [-1, 1]$. Then using the identity

$$g(\mathfrak{H}) = (g \circ h)^\triangleright(A) := f^\triangleright(A), \quad (7)$$

any matrix function $g(\mathfrak{H})$ can be expressed as a QSVT with respect to an even polynomial $f = g \circ h$ of degree $2d$. We remark that when g does not have a definite parity, the associated QSVT of A is much more involved. It generally requires using a linear combination of block encoding of the even and odd parts, which in turn requires implementing controlled U_A [11]. Normally such controlled operations have significant overhead. For instance, if we would like to implement a controlled-RACBEM, generally we need to convert all quantum gates in the circuit of U_A to the controlled version.

By expressing $g(\mathfrak{H})$ as a QSVT associated with an even polynomial, we have not only eliminated the need of controlled unitary operations, but also saved one additional qubit. This is because if we first construct a H-RACBEM \mathfrak{H} and then construct $g(\mathfrak{H})$, we need three ancilla qubits in total, and possibly a LCU circuit when g does not have a definite parity. On the other hand, by considering the composite function f , the parity constraint on g is completely removed, and we need only two ancilla qubits (the same as that needed for a H-RACBEM). Furthermore, each controlled gate in Fig. 5 is




FIG. 4. Quantum circuit for generating a $(1,2,0)$ -block encoding of a H-RACBEM from a $(1,1,0)$ block encoding U_A and its Hermitian conjugate. H is the Hadamard gate, and Z is the Pauli- Z gate.




FIG. 5. Quantum circuit for quantum singular value transformation of a real matrix polynomial of degree $2d$. For QSVT of a H-RACBEM, the number of ancilla qubits m can be set to 1.

a standard CNOT gate rather than a Toffoli gate. Therefore the H-RACBEM model has a salient advantage and significantly simplifies the implementation.

In many applications, such as a LCU circuit and the Hadamard test, we do need to get access to controlled $U_{\mathfrak{H}}$ or $U_{g(\mathfrak{H})}$. In such a case, the fact that \mathfrak{H} is a H-RACBEM is also helpful: Note that if we remove all Z rotations in the first line of Fig. 5, the circuit implements an identity operator since U_A, U_A^\dagger always appear in pairs due to that $f = g \circ h$ is an even polynomial. Therefore a controlled version $U_{g(\mathfrak{H})}$ can be simply implemented by controlling all the Z -rotation gates [11].

III. PROPOSAL OF QUANTUM LINPACK BENCHMARK

The RACBEM, as well as the H-RACBEM model provides a solution to the read-in problem using only basic quantum gates, and we can design them to be optimally adapted to the hardware architecture without resorting to complex quantum compilers. Hence they can be regarded as the proper generalization of “test dense matrices” in the quantum setting.

In this section, we demonstrate that the usage of the H-RACBEM model for solving QLSP. We assume $A = \mathfrak{H}$ is a H-RACBEM, and without loss of generality we may take $|b\rangle = |0^n\rangle$. Let the condition number of \mathfrak{H} be denoted by κ , which is the ratio between the maximum and the minimum of the singular values of \mathfrak{H} . It is believed that the computational complexity for solving QLSP cannot generally be better than $O(\kappa^{1-\delta})$ for any $\delta > 0$ [12], and the cost of using QSVT to solve general linear systems scales as $\tilde{O}(\kappa^2 \log(1/\epsilon))$ [11]. So the treatment of ill-conditioned matrices is very difficult especially on near-term devices. To reduce the circuit depth, in the near future it may be more productive to focus on solving well-conditioned linear systems.

Note that if A has at least one singular value that is zero (or near zero), a canonical H-RACBEM \mathfrak{H} is not invertible (or very ill-conditioned). Such events can occur more frequently as n becomes large. It can be difficult to diagnose such a problem without first obtaining some spectral information of A , which is perhaps a more difficult task than solving the linear system problem itself.

The H-RACBEM model offers a new and natural way to solve this problem. From Eq. (4), assuming $\cos(2\varphi_0 - \varphi_1) > 0$ and $-2\sin(2\varphi_0)\sin\varphi_1 > 0$, and use that $0 \leq A^\dagger A \leq 1$, the condition number of \mathfrak{H} can be bounded from above:

$$\kappa(\mathfrak{H}) \leq \frac{\cos(2\varphi_0 + \varphi_1)}{\cos(2\varphi_0 - \varphi_1)}.$$

Therefore the condition number of a H-RACBEM is fully tunable by changing the phase factors φ_0, φ_1 . According to Eq. (6), this changes the second-order polynomial function

$h(x)$, so that $\mathfrak{H} := h^\triangleright(A) \succ 0$ has a tunable, bounded condition number.

In order to solve QLSP, we are interested in finding a polynomial $g(x)$ of degree d so that

$$|g(x) - x^{-1}| \leq \epsilon, \quad x \in [\kappa^{-1}, 1].$$

Following [[35], Corollary 69], there can be satisfied by an odd polynomial $g(x)$ with degree $d \sim O(\kappa \log(1/\epsilon))$, which gives an upper bound on d . Numerical results shows that a better approximation to $g(x)$ can be obtained by solving a minimax problem using the Remez algorithm [34], and the polynomial can be chosen to be either even or odd. Figure 6 shows the shape of the optimal polynomial even or odd approximation to x^{-1} on the interval $[\kappa^{-1}, 1]$ found by the Remez algorithm to reach a target accuracy of 10^{-3} with $\kappa = 10$. In particular, the usage of an even polynomial can further reduce the polynomial degree, which will be used in the numerical tests below.

We may then construct a degree $2d$ polynomial $f = g \circ h$ as in Sec. II. Once we find the associated phase factors, the circuit in Fig. 5 implements $g(\mathfrak{H})|b\rangle$, which satisfies the error bound $\|g(\mathfrak{H})|b\rangle - \kappa^{-1}\mathfrak{H}^{-1}|b\rangle\|_2 \leq \epsilon$ for any normalized vector $|b\rangle$. The success probability of measuring the ancilla




FIG. 6. Approximate polynomials obtained by the Remez algorithm. The even and odd polynomials are generated such that the maximal approximation error to $1/(\kappa x)$ on $[\kappa^{-1}, 1]$ is 10^{-3} . We find that the polynomial degree of the even approximation is lower than that of the odd approximation.

qubits and obtaining an all 0 output (will be referred to as the success probability for short) is

$$p = \|g(\mathfrak{H})|b\rangle\|_2^2 \geq (\kappa^{-1} - \epsilon)^2,$$

which can be of modest size given κ is not too large.

Since measuring the accuracy of all entries of the solution is not a practical goal, our proposal of the quantum LINPACK benchmark is to *measure the success probability* p , i.e., the probability of measuring the two ancilla qubits and obtaining $|00\rangle$, and to compare it with the numerically exact probability (denoted by p_{exact}) computed using a classical computer. When κ, ϵ , and H-RACBEM are given, the quantum LINPACK benchmark reports the relative error $|p - p_{\text{exact}}|/p_{\text{exact}}$ to describe the accuracy of a quantum computer for solving QLSP. In the future we may also take into account the wall clock time, or a quantity analogous to the floating point operations per second (FLOPS) for classical computers to measure the efficiency of a quantum computer. The quantum LINPACK benchmark uses only basic single-qubit gates and CNOT gates, and hence is easy to implement even on near-term devices.

Clearly, the success of the quantum LINPACK benchmark is only a *necessary condition* for the accurate solution of QLSP. However, as will be shown in numerical experiments, this task can already be challenging for near-term devices due to the presence of noise, and therefore the benchmark provides a meaningful and easily implementable criterion for measuring the accuracy of quantum computers. Furthermore, due to the very flexible construction of U_A and the near-minimal usage ancilla qubits and other primitive gates, we also expect that the success of quantum LINPACK benchmark is the *minimal requirement* in order to achieve quantum advantages by solving QLSP or similar linear algebra tasks.

IV. OTHER QUANTUM LINEAR ALGEBRA APPLICATIONS

Since QSVT implements a general matrix function, its application is therefore not limited to solving QLSP. Here we demonstrate a few more applications. Throughout the section \mathfrak{H} is a H-RACBEM defined via a unitary matrix U_A and $h(x)$ is a second-order polynomial.

A. Time series analysis

Given an n -qubit state $|\psi\rangle$, consider the computation of the following time series:

$$s(t) = \langle\psi|e^{i\mathfrak{H}t}|\psi\rangle. \quad (8)$$

When \mathfrak{H} is a H-RACBEM, we can evaluate $s(t)$ by measuring the its real and imaginary components separately:

$$\begin{aligned} s(t) &= \langle\psi|\cos(\mathfrak{H}t)|\psi\rangle + i\langle\psi|\sin(\mathfrak{H}t)|\psi\rangle \\ &= 2\langle\psi|\frac{1}{2}(\cos(\mathfrak{H}t) + I_n)|\psi\rangle \\ &\quad + 2i\langle\psi|\frac{1}{2}(\sin(\mathfrak{H}t) + I_n)|\psi\rangle - (1 + i) \\ &=: 2\|g_{c,t}(\mathfrak{H})|\psi\rangle\|_2^2 + 2i\|g_{s,t}(\mathfrak{H})|\psi\rangle\|_2^2 - (1 + i). \end{aligned} \quad (9)$$

Here we introduce the functions for the cosine and sine part, respectively,

$$g_{c,t}(x) = \sqrt{\frac{\cos(xt) + 1}{2}}, \quad g_{s,t}(x) = \sqrt{\frac{\sin(xt) + 1}{2}}.$$

Now the quantities in Eq. (9) can be directly obtained via the success probability of the QSVT circuit with $f_{c,t} := g_{c,t} \circ h$ and $f_{s,t} := g_{s,t} \circ h$, respectively (with a suitable polynomial approximation of $g_{c,t}, g_{s,t}$). Note that the access to the matrix square root of \mathfrak{H} here is crucial: though $g_{c,t}$ is an even function itself, $g_{s,t}$ does not have a definite parity, and therefore the direct implementation of $g_{s,t}(\mathfrak{H})$ would require a LCU circuit and hence controlled unitary operations. In the setting of H-RACBEM, the treatment of $g_{c,t}$ and $g_{s,t}$ can be put on the same footing due to the composition with the even polynomial h .

In order to accelerate the convergence of the polynomial approximation, in practice, we may introduce another tunable parameter $\eta \geq 1$ in the formulation so that

$$g_{c,t,\eta}(x) = \sqrt{\frac{\cos(xt) + \eta}{2}}, \quad g_{s,t,\eta}(x) = \sqrt{\frac{\sin(xt) + \eta}{2}}$$

and then

$$s(t) = 2\|g_{c,t,\eta}(\mathfrak{H})|\psi\rangle\|_2^2 + 2i\|g_{s,t,\eta}(\mathfrak{H})|\psi\rangle\|_2^2 - \eta(1 + i).$$

For instance, η can be set to 1.0 or 1.5, and the details are given in Table II below.

In [36] the time series analysis is used for eigenvalue estimation, which is implemented via the Hadamard test and requires a controlled Hamiltonian evolution procedure, followed by a Fourier transform. The procedure above removes the need of performing controlled Hamiltonian evolution. We also remark that in terms of quantum estimation of eigenvalues, the method using the spectral measures to be discussed below can be more advantageous, which does not require a subsequent Fourier transform, and the resulting spectral measure is guaranteed to be non-negative.

B. Spectral measure

Given an n -qubit state $|\psi\rangle$, in order to approximately evaluate the spectral measure, we can use the Plemelj formula

$$\begin{aligned} s(E) &= \langle\psi|\delta(\mathfrak{H} - E)|\psi\rangle \\ &= \lim_{\eta \rightarrow 0^+} \frac{1}{\pi} \text{Im} \langle\psi|(\mathfrak{H} - E - i\eta)^{-1}|\psi\rangle, \quad E \in [-1, 1]. \end{aligned} \quad (10)$$

By choosing a finite η as the broadening parameter, we need to evaluate

$$s_\eta(E) = \frac{\eta}{\pi} \langle\psi|((\mathfrak{H} - E)^2 + \eta^2)^{-1}|\psi\rangle.$$

Now define

$$g_{\eta,E}(x) = \left(\frac{\eta^2}{(x - E)^2 + \eta^2} \right)^{\frac{1}{2}},$$

which satisfies $|g_{\eta,E}| \leq 1$ for $x \in [-1, 1]$, then $g_{\eta,E} \circ h$ can be approximated by an even-order polynomial. Therefore

$$s_\eta(E) = \frac{1}{\eta\pi} \|g_{\eta,E}(\mathfrak{H})|\psi\rangle\|_2^2,$$

and it can be obtained via the success probability of the QSVT circuit with $f = g_{\eta,E} \circ h$ (with a suitable polynomial approximation of $g_{\eta,E}$).

C. Thermal averages

In order to evaluate the thermal average of the energy (β is the inverse temperature)

$$E(\beta) = \frac{\text{Tr}[\mathfrak{H}e^{-\beta\mathfrak{H}}]}{\text{Tr}[e^{-\beta\mathfrak{H}}]},$$

we can use the minimally entangled typical thermal state (METTS) algorithm [26–28]. Let $Z = \text{Tr}[e^{-\beta\mathfrak{H}}]$, and

$$E(\beta) = \frac{1}{Z} \sum_{i \in [N]} \langle i | e^{-\beta\mathfrak{H}/2} \mathfrak{H} e^{-\beta\mathfrak{H}/2} | i \rangle = \sum_{i \in [N]} p_i \langle \phi_i | \mathfrak{H} | \phi_i \rangle.$$

Here i loops over all states of the computational basis, each represented by an n -bit string. From the unnormalized states $|\tilde{\phi}_i\rangle = e^{-\beta\mathfrak{H}/2} |i\rangle$, we define a probability distribution $p_i = \langle \tilde{\phi}_i | \tilde{\phi}_i \rangle / Z$, and corresponding normalized states $|\phi_i\rangle = |\tilde{\phi}_i\rangle / \sqrt{\langle \tilde{\phi}_i | \tilde{\phi}_i \rangle}$. For each i , the contribution to the energy can be expressed as

$$\langle \phi_i | \mathfrak{H} | \phi_i \rangle = \frac{\langle i | e^{-\beta\mathfrak{H}} \mathfrak{H} | i \rangle}{\langle i | e^{-\beta\mathfrak{H}} | i \rangle} = \frac{\|g_{n,\beta}(\mathfrak{H}) | i \rangle\|_2^2}{\|g_{d,\beta}(\mathfrak{H}) | i \rangle\|_2^2}. \quad (11)$$

Here we define

$$g_{n,\beta}(x) = e^{-\beta x/2} \sqrt{x}, \quad g_{d,\beta}(x) = e^{-\beta x/2},$$

for the numerator and the denominator, respectively, and without loss of generality we assume $\mathfrak{H} \geq 0$. The numerator and the denominator can be obtained via the success probability of the QSVT circuit with $f_{n,\beta} := g_{n,\beta} \circ h$ and $f_{d,\beta} := g_{d,\beta} \circ h$, respectively (with a suitable polynomial approximation of $g_{n,\beta}, g_{d,\beta}$).

Unlike Metropolis-type algorithms which follows an acceptance and rejection procedure, the METTS algorithm samples the states $\{|\phi_i\rangle\}$ as follows. We start from a computational basis state $|i\rangle$ (e.g., $|i\rangle = |0^n\rangle$). Then

(1) Evaluate the contribution to the energy from the state $|\phi_i\rangle$ via Eq. (11).

(2) Collapse $|\phi_i\rangle$ to a new state in the computational basis $|i'\rangle$ with probability $|\langle i' | \phi_i \rangle|^2$, and repeat step 1.

Note that step 2 of the METTS algorithm is very simple to implement: we only need to construct a QSVT circuit for preparing the unnormalized state $|\tilde{\phi}_i\rangle$, which is readily available when computing the denominator in Eq. (11). Then collapsing to $|i'\rangle$ can be implemented by measuring all n system qubits in the computational basis.

The evaluation of $f_{n,\beta} := g_{n,\beta} \circ h$ requires approximating the square root function. When \mathfrak{H} is a canonical H-RACBEM with $h(x) = x^2$, an alternative method is to approximate an odd function $\tilde{f}_{n,\beta}(x) = x e^{-\beta x^2/2}$ instead of $f_{n,\beta}$. Then the complexity for thermal average calculation can be only $\tilde{O}(\sqrt{\beta} \log(1/\epsilon))$ [[35], Corollary 64].

V. NUMERICAL RESULTS

The source code of RACBEM is available in the Github repository [37]. We use the optimization-based algorithm

which is described in Appendix B to generate phase factors, and the source code is available in the Github repository [38]. We demonstrate the numerical performance of the RACBEM model for solving various numerical quantum linear algebra tasks. The algorithms are performed on the IBM Q quantum device, as well as QVM with approximate noisy running environment retrieved from quantum devices. All numerical tests are implemented in python3.7 and Qiskit [39]. In order to construct an adjustable noise model, we retrieve the noise model from real quantum devices provided by IBM Q backends. The magnitude of the noise level is then made to be fully tunable through a single parameter σ . When $\sigma = 0$, the only noise contribution comes from the Monte Carlo error in measurements, which can be systematically reduced by increasing the number of samples. When $\sigma = 1$, the noise model contains all the readout errors and quantum errors associated with a probability distribution given by the retrieved noise model (see Appendix C for details). Unless otherwise noted, the number of measurements (shots) is fixed to be 8192 throughout this section. As the quantum linear algebra tasks are solved via QSVT, the overall error consists of contributions from the polynomial approximation, the noise from the device, and the Monte Carlo sampling error. We remark that in numerical results we distinguish the setup of “QSVT without error” (only taking into account the polynomial approximation error) from that of “ $\sigma = 0$ ” (including both the Monte Carlo sampling error and the polynomial approximation error).

We use Algorithm 1 (in Appendix C) to generate custom random quantum circuits with respect to a given coupling map. In each layer of the quantum circuit, we apply a one-qubit (or two-qubit) gate to each qubit (or a pair of qubits) randomly selected from the basic gate set. Though CNOT gates can increase the entanglement among the qubits, their error rate is much higher than that of one qubit gates on IBM Q backends. So we control the number of CNOT gates via a parameter to determine the probability that a CNOT gate is drawn. Unless otherwise noted, this probability is set to be $p = 0.5$. The circuit depth is the same as the number of layers in the generating circuit. We would like to emphasize that the circuit depth must not be too small. Otherwise the block-encoded matrix can sometimes become degenerate (such as a scaled identity matrix). For an n -qubit system, we empirically set the depth ℓ to be $\ell = 3$ when $n = 1$, $\ell = 7$ when $n = 2$, and $\ell = 15 + 2(n - 3)$ when $n \geq 3$. We report the statistics of singular-value distributions of the block-encoded matrix in Fig. 17 (Appendix E) to justify this adaptive choice of the circuit depth.

According to the random quantum circuit generation algorithm, we measure the total gate count of a quantum circuit by its logical gate count with respect to the basic gate set, which is upper bounded by its depth ℓ times the number of system qubits n . The available basic gates provided by IBM Q backends are the CNOT gate and U1, U2, U3 gates, which are parameterized families of generic one-qubit gates (see Appendix C for details). In order to reduce the noise due to U3 gates, we restrict the basic gate set in the custom random circuit generator to be {CNOT, U1, U2}, which is still a universal gate set. Each controlled rotation in QSVT circuit costs seven logical gates (four X gates implemented by U2 gates, two CNOT gates, and one U1 gate). Therefore, given a RACBEM whose depth is ℓ and with n system qubits, to implement




FIG. 7. Relative error in success probability of obtaining $A|0^n\rangle$ for different number of system qubits. The box ranges from the first to the third quartiles (25% and 75% percentiles, respectively) of the data set. The horizontal line inside the box stands for the median of the data set. The whisker extending from the box indicates the rest of the data set. The RACBEM's are generated using our custom random circuit generator in which the coupling map, basic gates and logical-to-physical layout are specified. When the number of system qubits is less than or equal to three, the results are obtained on the five-qubit backend `ibmq_burlington`. Other results are obtained using 15-qubit backend `ibmq_16_melbourne`. The data set consists of results from 200 RACBEM's when $n = 1, 2, 3$, 300 RACBEM's when $n = 4, 5, 6$, 400 RACBEM's when $n = 7$ and 500 RACBEM's when $n = 8$.

the QSVT of a real polynomial of degree d , the total gate count for the circuit is upper bounded by $2 + 7(d + 1) + d\ell n$. The details about QSVT phase factors used in numerical experiments can be found in Appendix E. Unless otherwise noted, the input quantum state of the quantum circuit is set to $|0^{n+2}\rangle$.

A. RACBEM

Before presenting results of various numerical linear algebra problems, we measure the effect of noise on RACBEM directly on quantum computing backends provided by IBM Q. The numerical results are displayed in Fig. 7. Each data point is obtained by generating a RACBEM using Algorithm 1, and we measure the success probability of obtaining the block-encoded matrix applied to $|0^n\rangle$, which is equal to $\|A|0^n\rangle\|_2^2$. The number of repeated measurements (shots) is 8192. Fig-

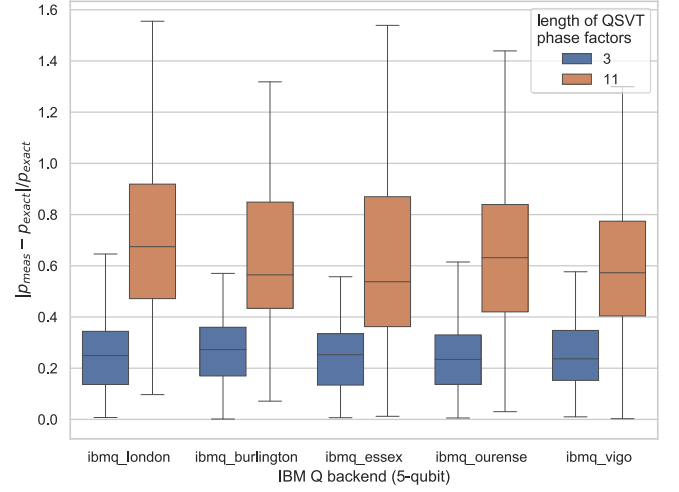


FIG. 8. Relative error of solving QLSP using the H-RACBEM model on IBM Q five-qubit backends with identical quantum architectures. The condition number of each H-RACBEM is upper bounded by two and the number of system qubits is three. The size of the box indicates the first and the third quartiles and the extending whisker indicates the rest of the data set. The phase factor sequences of length three and 11 are used to carry out QSVT to solve QLSP. Each data set is obtained using 100 samples at random. The details about phase factors are given in Table I.

ure 7 shows that the relative error of the quantum device can be considerable, ranging from 10% to 30% as n increases. Since the RACBEM is the building block of all subsequent quantum linear algebra tasks, we expect that the relative error of such tasks on quantum computing backends provided by IBM Q should be at least around the same level.

B. QLSP

According to the discussion in Sec. III, it is possible to measure the accuracy of the solution to QLSP by sampling the output distribution of $g(\mathfrak{H})|0^n\rangle$ using, e.g., a cross-entropy test similar to that in [1]. Here for simplicity we only measure $\|g(\mathfrak{H})|0^n\rangle\|_2^2$ as a success probability (of measuring all ancilla qubits and obtain zeros), and evaluate the relative error compared to the exact value $\|\alpha^{-1}\mathfrak{H}^{-1}|0^n\rangle\|_2^2$ evaluated on a classical computer where α is a scaling factor (see Table I for details). The small relative error in success probability is a necessary condition to ensure the solution is correct which yields the benchmark in a weak sense. The condition number of H-RACBEM is controlled by a second-order polynomial $h_\kappa(x) := (1 - \kappa^{-1})x^2 + \kappa^{-1}$ as in Sec. III. Another poly-

TABLE I. Parameters of the QSVT circuit in solving QLSP, i.e., in Figs. 8 and 9. Numbers in square brackets indicate the power of 10.

	κ	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
IBM Q	2	3	2.79722 [-2]	$23 + 2\ell n$	3.59306
	2	11	2.44481 [-5]	$79 + 10\ell n$	3.59306
QVM	2	5	6.18245 [-3]	$37 + 4\ell n$	2.38234
	5	7	1.90152 [-2]	$51 + 6\ell n$	5.86631
	10	13	7.45462 [-3]	$93 + 12\ell n$	11.89390
	20	19	6.65999 [-3]	$135 + 18\ell n$	23.81003

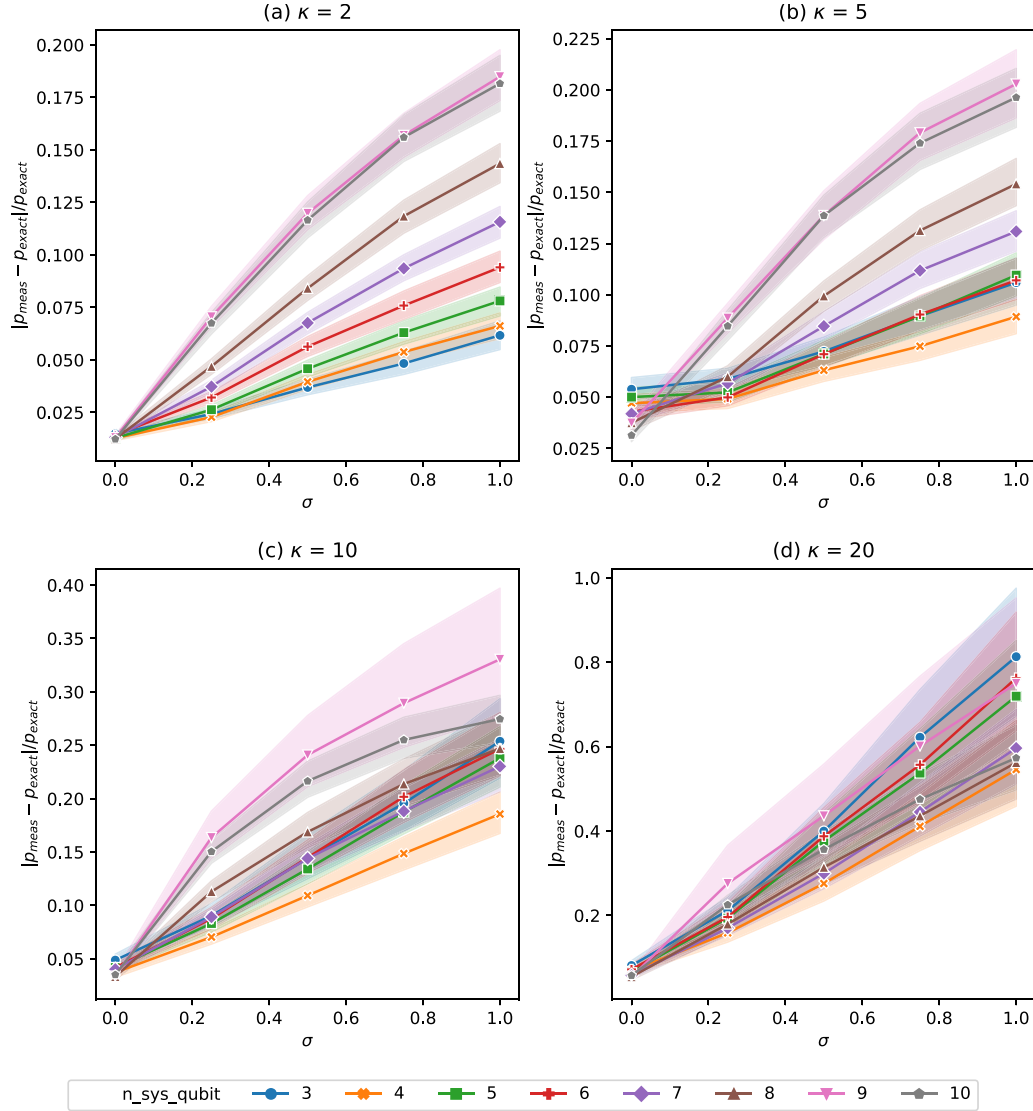


FIG. 9. Relative error of solving QLSP for different number of system qubits, condition number, and noise magnitude. The numerical results are obtained by running on QVM with architecture and noise model retrieved from IBM Q backend `ibmq_16_melbourne`. Each kind of parameter is computed by solving QLSP of 300 RACBEM's at random. The marker is the average of each data set, and the shaded area is the 95% confidence interval inferred from the data set. The details about phase factors are given in Table I.

mial $g_{\kappa}(x)$ which approximates x^{-1} is chosen to perform matrix inversion. The composite polynomial $g_{\kappa} \circ h_{\kappa}$ can be implemented by an even order QSVT circuit to carry out the overall procedure with only two ancilla qubits.

We report the performance of solving linear systems on the IBM Q device using the H-RACBEM model in Fig. 8. The architecture of the five backends are identical, and therefore we may draw a H-RACBEM at random and test it on all five backends. By tuning $h(x)$, the condition number of \mathcal{H} is upper bounded by 2. Therefore we may even use a very short QSVT circuit with $d = 2$, and the corresponding number of phase factors is three. This is essentially a linear approximation to the inverse, and the accuracy measured by the L^{∞} error is less than 0.03. In such a case, the total logical gate count is upper bounded by 113. We can refine the polynomial approximation by using a more accurate, and deeper QSVT

circuit with $d = 10$ (the number of phase factors is 11, and the L^{∞} error is less than 3×10^{-5}). In a case when $d = 10$, the total logical gate count is upper bounded by 529. The results in Fig. 8 indicate that for the shallow circuit, the relative error of the success probability is similar to that observed in Fig. 7, which measures only the success probability of the block encoding. However, the relative error is significantly larger using the deeper circuit, despite that the QSVT circuit implements a more accurate polynomial approximation to the matrix inverse. Thus, we conclude that when designing the quantum circuit, a proper choice of QSVT phase factors is needed which reflects the tradeoff between the polynomial approximation error and the error caused by the noisy running environment.

In order to further demonstrate how various parameters can affect the accuracy of the QLSP solver, we vary the number of

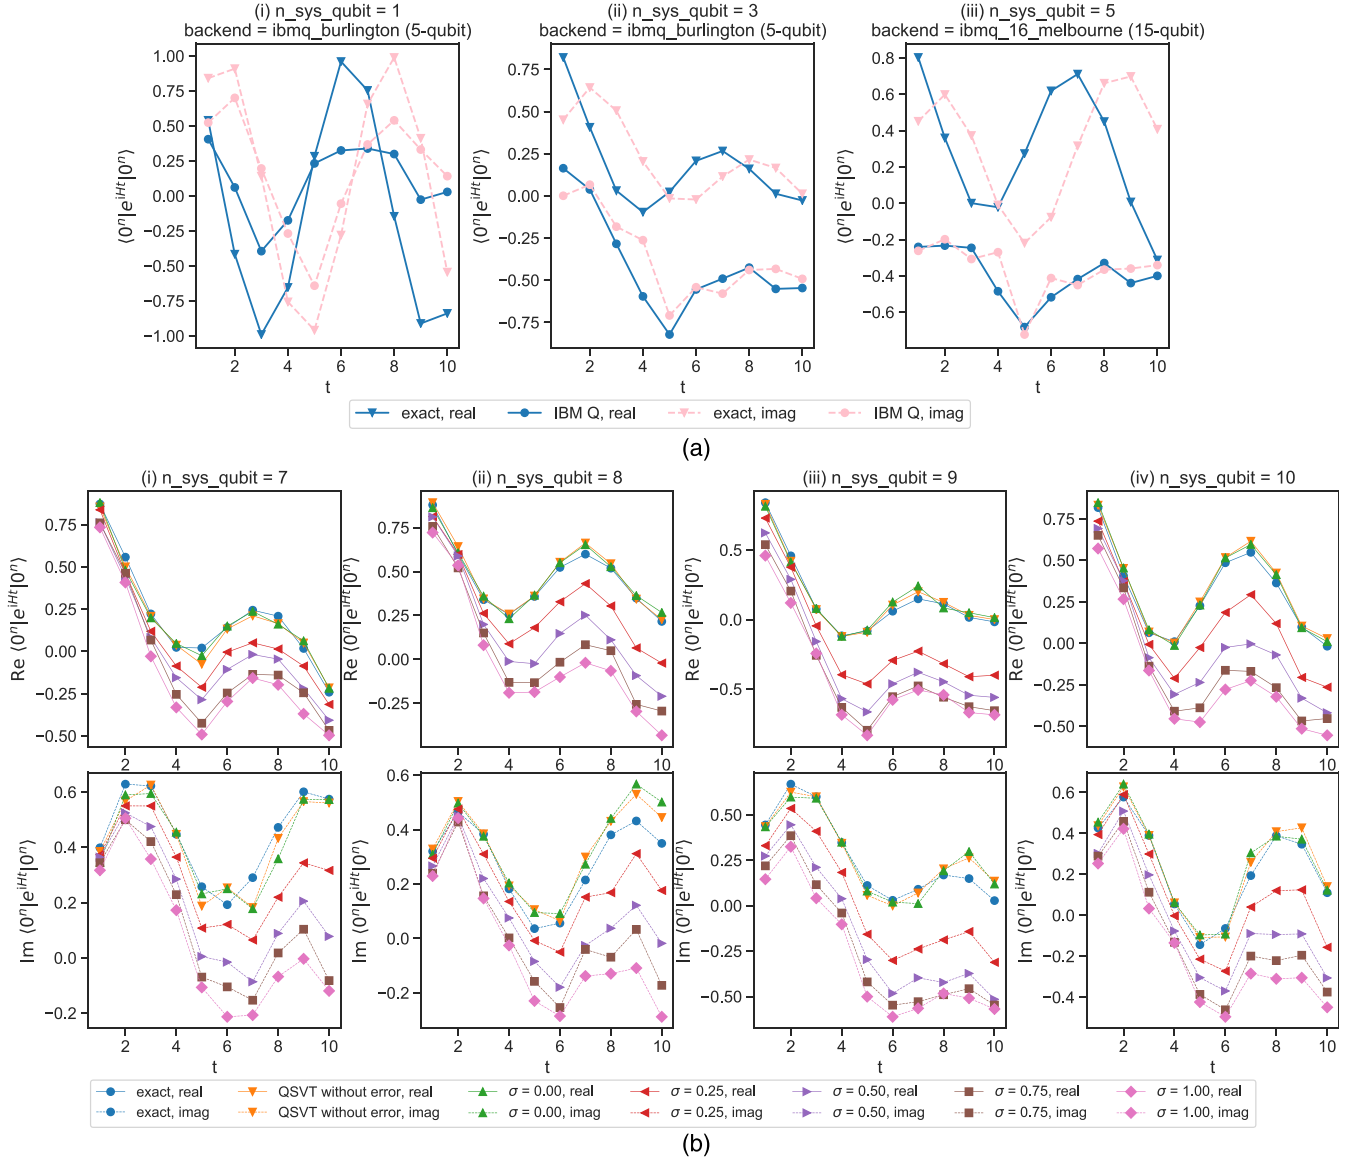


FIG. 10. Time series computed by using the canonical H-RACBEM model and QSVT. The details about phase factors are given in Table II and the details about block-encoded matrices are given in Fig. 18. (a) Results obtained on IBM Q backends. The probability of CNOT is 0.1 when generating random circuits. (b) Results obtained on QVM with architecture and noise model retrieved from the IBM Q backend `ibmq_16_melbourne`. The probability of CNOT is set to 0.5.

system qubits, the condition number, and the noise magnitude, and compute the relative error of success probability under these different settings on QVM. The numerical results are presented in Fig. 9. In all cases, we find that the QLSP solver can perform very well when σ is small. The error is due only to the polynomial approximation and the Monte Carlo sampling error. This corresponds to the setting of fault-tolerant quantum devices. However, the accuracy rapidly deteriorates as σ increases. In the noisy setup, the error also increases nearly proportionally to the condition number κ . This is because the polynomial degree d should increase as $O(\kappa)$ in order to achieve constant accuracy, and so is the circuit depth. The error also increases with respect to the number of system qubits, but the effect is less significant compared to that due to κ , which increases the circuit depth.

C. Time series analysis

The numerical results regarding the time series analysis are shown in Fig. 10. The results in Fig. 10(a) are obtained on the IBM Q backends. When the number of system qubit is 1 (the circuit uses three qubits in total), the features of the time series obtained from the quantum device can agree qualitatively with the exact solution. However, as the number of system qubits increases to five, the result is dominated by the quantum noise. In order to investigate the performance of larger systems and the effects of noise, we use the tunable QVM instead in Fig. 10(b). The length of phase factors is chosen adaptively as t increase in order to reduce the error of the polynomial approximation (details are given in Table II). When the noise level σ is tuned to 0, the results from the QSVT circuit are uniformly accurate for all t in $[1, 10]$. How-

TABLE II. Parameters of QSVT circuits in the time series analysis, i.e., in Fig. 10. Numbers in square brackets indicate the power of 10.

	t	η	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
Real part	1	1.0	3	1.23670 [-2]	$23 + 2\ell n$	1.21807
	2	1.0	3	4.25711 [-2]	$23 + 2\ell n$	1.26458
	3	1.0	5	9.64293 [-3]	$37 + 4\ell n$	1.21400
	4	1.5	7	8.48770 [-3]	$51 + 6\ell n$	1.34011
	5	2.0	7	2.66925 [-2]	$51 + 6\ell n$	1.51827
	6	1.5	9	2.10169 [-2]	$65 + 8\ell n$	1.35177
	7	1.5	9	3.47455 [-2]	$65 + 8\ell n$	1.39998
	8	1.5	9	5.78363 [-2]	$65 + 8\ell n$	1.44148
	9	1.5	11	2.84139 [-2]	$79 + 10\ell n$	1.37467
	10	1.5	11	3.26549 [-2]	$79 + 10\ell n$	1.38139
Imaginary part	1	1.0	3	1.14646 [-2]	$23 + 2\ell n$	1.16750
	2	1.0	3	4.73088 [-2]	$23 + 2\ell n$	1.24295
	3	1.0	5	1.60676 [-3]	$37 + 4\ell n$	1.19769
	4	1.0	5	8.83397 [-3]	$37 + 4\ell n$	1.18742
	5	1.5	5	7.76900 [-2]	$37 + 4\ell n$	1.23512
	6	1.5	7	3.15931 [-2]	$51 + 6\ell n$	1.36811
	7	1.5	7	6.47625 [-2]	$51 + 6\ell n$	1.35109
	8	1.5	9	5.50680 [-2]	$65 + 8\ell n$	1.43342
	9	1.5	9	5.73218 [-2]	$65 + 8\ell n$	1.40154
	10	1.5	11	6.46945 [-2]	$79 + 10\ell n$	1.41058

ever, since the circuit depth needs to increase proportionally to t , when the noise level is not 0, the error is significantly larger as t increases. We also remark that our custom noise model with $\sigma = 1$ contains only partial kinds of noise derived from the real IBM Q device (details in Appendix C). Therefore, in Fig. 10 the time series computed on the real device behaves much worse than that computed from the simulation using the noise model with $\sigma = 1$ on the QVM.

D. Spectral measure

In the limit when $\eta \rightarrow 0$, the spectral measure is given by the summation of Dirac δ functions. Hence when n is small, the spectral measure has sharp features even when η takes a finite value, which in turn requires a polynomial of higher degree to resolve. So we consider only spectral measures for $n = 7, 8, 9, 10$ on QVM, and the size of \mathfrak{H} ranges from 128 to 1024. The numerical results of spectral measures are presented in Fig. 11. The spectral measures exhibit rather different features. This is not so much related to the number of system qubits and is mostly due to the specific instance of the H-RACBEM. In all cases, the quantum algorithm can capture at least the qualitative features of the spectral measures, though the noise plays an important role particularly for the instance $n = 9$. We also remark that the functionality of the QSVT circuit depends only on the set of phase factors as E sweeps across the spectrum. The length of corresponding QSVT phase factors is set to 11 for each point. In Tables III and IV, for points closer to the middle of the spectrum ($E = 0.5$), we observe that a polynomial of larger degree is needed to achieve the same accuracy in approximation. Therefore, it can be seen from Fig. 11 that the deviation between the value given by noiseless QSVT and the exact one increases as the parameter moves towards the middle of the spectrum. The choice of the circuit depth reflects the tradeoff between approximation error and the effects of noise. To illustrate this,

we also compute the spectral measures in Appendix D by using a deeper QSVT circuits. Although the deeper circuit can

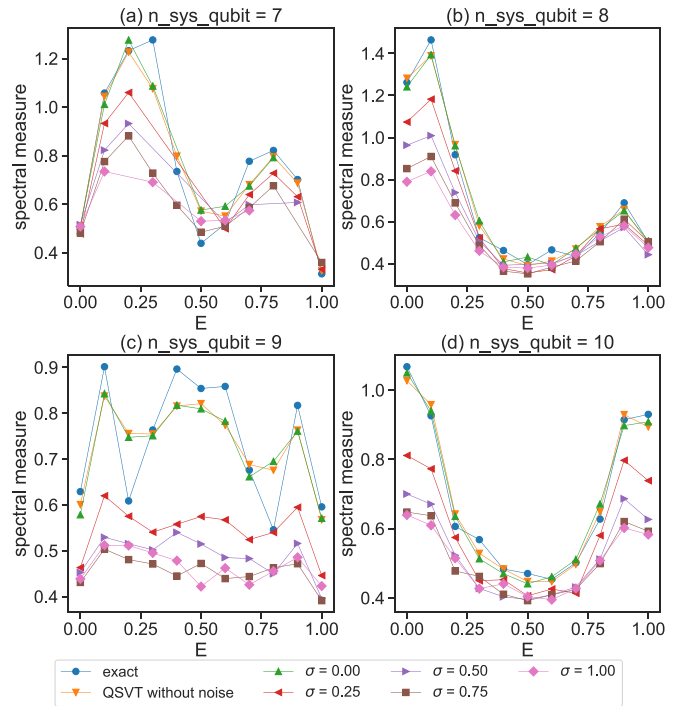


FIG. 11. Spectral measure computed by using the canonical H-RACBEM model and QSVT. The numerical results are obtained by running on QVM with architecture and noise model retrieved from IBM Q backend ibmq_16_melbourne. For each number of system qubits, we draw a canonical H-RACBEM randomly according to the backend architecture. The probability of CNOT is 0.5 when generating random circuits. The details about phase factors are given in Table III and the details about block-encoded matrices are given in Fig. 18.

TABLE III. Parameters of the QSVT circuit in computing spectral measure, i.e., in Fig. 11. Numbers in square brackets indicate the power of 10.

E	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
0.0	11	3.20242 [-2]	$79 + 10\ell n$	2.14095
0.1	11	5.59268 [-2]	$79 + 10\ell n$	2.14095
0.2	11	1.23926 [-1]	$79 + 10\ell n$	2.14095
0.3	11	1.32550 [-1]	$79 + 10\ell n$	2.14095
0.4	11	1.36737 [-1]	$79 + 10\ell n$	2.14095
0.5	11	1.71503 [-1]	$79 + 10\ell n$	2.14095
0.6	11	1.36727 [-1]	$79 + 10\ell n$	2.14095
0.7	11	1.32529 [-1]	$79 + 10\ell n$	2.14095
0.8	11	1.23922 [-1]	$79 + 10\ell n$	2.14095
0.9	11	5.59258 [-2]	$79 + 10\ell n$	2.14095
1.0	11	3.20242 [-2]	$79 + 10\ell n$	2.14095

produce better results in the noiseless setting, the quality of the spectral measure can significantly deteriorate in the presence of the quantum noise.

E. Thermal average of the energy

In Fig. 12(a) we compute the thermal average of the energy of H-RACBEM's on IBM Q backends for n from 1 to 5. We reduce the length of QSVT phase factors as much as possible while balancing the approximation error and the quantum error, and the details can be found in Table V. Since METTS is a Monte Carlo algorithm, we perform a sufficiently large number of steps to reduce the error due to the METTS algorithm, by monitoring the cumulative moving average of the thermal average of the energy in Fig. 16 in Appendix D. Compared to the results above obtained on IBM Q, the results for the thermal average of the energy are somewhat surprisingly accurate in all cases. We also compute the thermal average of the energy on QVM to further investigate the effect of the noise for $n = 3, 5$, and 7 . In Fig. 12(b), (i) and (ii), the

TABLE IV. Parameters of the QSVT circuit in computing spectral measure with higher approximation precision, i.e., in Fig. 15(b). Numbers in square brackets indicate the power of 10.

E	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
0.0	21	9.59006 [-4]	$149 + 20\ell n$	2.14095
0.1	27	5.02000 [-3]	$191 + 26\ell n$	2.14095
0.2	33	5.48205 [-3]	$233 + 32\ell n$	2.14095
0.3	37	5.90586 [-3]	$261 + 36\ell n$	2.14095
0.4	41	4.66722 [-3]	$289 + 40\ell n$	2.14095
0.5	43	4.38167 [-3]	$303 + 42\ell n$	2.14095
0.6	41	4.66727 [-3]	$289 + 40\ell n$	2.14095
0.7	39	3.70179 [-3]	$275 + 38\ell n$	2.14095
0.8	35	3.29638 [-3]	$247 + 34\ell n$	2.14095
0.9	27	5.01941 [-3]	$191 + 26\ell n$	2.14095
1.0	19	2.53240 [-3]	$135 + 18\ell n$	2.14095

identical quantum computing task is emulated on different backends, and the only difference is due to the noise model. It is evident that the behavior of the solution depends sensitive on the noise. On the other hand, we find that the solution remains remarkably accurate as n increases, even when β is relatively large. For most data points, the thermal average of the energy decreases monotonically with respect to β , and we observe that as σ increases, the energy curve shifts downwards monotonically.

VI. CONCLUSION

Analogous to the LINPACK benchmark for measuring the performance of classical supercomputers, we have proposed a quantum LINPACK benchmark to measure the performance of current and future quantum computers for scientific computing applications. The quantum LINPACK benchmark solves a well conditioned quantum linear system problem, which is enabled by the Hermitian RAndom Circuit Block-Encoded Matrix (H-RACBEM) input model and the quantum singular value transformation (QSVT). The flexibility provided by the H-RACBEM model also allows us to perform other linear algebra tasks already on near-term devices, such as computing spectral measures, and time series generated by a Hamiltonian simulation. We can also compute the thermal average of the energy, using a quantum version of the minimally entangled typical thermal state (METTS) algorithm.

We perform these linear algebra tasks on IBM Q quantum devices, and quantum virtual machines with a tunable noise model. Although present quantum devices still suffer from high noise levels, it is nonetheless encouraging to observe that the solutions can already be qualitatively obtained in the noisy environment. Among all numerical tests, the thermal average of the energy computed via the METTS algorithm appears to be particularly robust with respect to noise. When designing the quantum circuit, we need to carefully choose the length of QSVT phase factors, by balancing the polynomial approximation error and the additional quantum error caused by the increase of the circuit depth. Our numerical tests are currently limited to matrices up to 10 qubits (the corresponding matrix size is 1024) in order to compare with numerically exact results obtained on classical computers. However, the number of qubits can be directly increased to be much larger, especially on future devices with reduced noise level.

Note that Google's quantum supremacy experiment is a benchmark problem (the linear cross-entropy test) and is also hard for any classical computer. The classical hardness is justified by that sampling a certain random quantum circuit (e.g., one that implements a chaotic quantum evolution) and generating heavy outputs is a difficult task on any classical computer [32,40–47]. Although the quantum LINPACK benchmark also involves a supremacy-type random circuits, the success of the benchmark is defined by measuring only two ancilla qubits, and the task could therefore *possibly* be “classically spoofed” when the block-encoding matrix U_A is drawn from a known distribution, such as the Haar measure. We would like to therefore clarify that a quantum benchmark problem does not need to be classically hard in order to

TABLE V. Parameters of the QSVT circuit in computing the thermal average of the energy, i.e., in Fig. 12. Numbers in square brackets indicate the power of 10.

	β	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
Numerator	1	4	8.10003 [−3]	$30 + 3\ell n$	0.72602
	2	4	3.35247 [−2]	$30 + 3\ell n$	0.53351
	3	6	9.27231 [−3]	$44 + 5\ell n$	0.42483
	4	6	1.96455 [−2]	$44 + 5\ell n$	0.37029
	5	6	3.36597 [−2]	$44 + 5\ell n$	0.33182
	6	8	9.57788 [−3]	$58 + 7\ell n$	0.29986
	7	8	1.51955 [−2]	$58 + 7\ell n$	0.27823
	8	8	2.21488 [−2]	$58 + 7\ell n$	0.26052
Denominator	1	3	1.03401 [−2]	$23 + 2\ell n$	1.18530
	2	3	3.37925 [−2]	$23 + 2\ell n$	1.15324
	3	5	7.30555 [−3]	$37 + 4\ell n$	1.18960
	4	5	1.40523 [−2]	$37 + 4\ell n$	1.18014
	5	5	2.25225 [−2]	$37 + 4\ell n$	1.16846
	6	5	3.22698 [−2]	$37 + 4\ell n$	1.15529
	7	7	8.56996 [−3]	$51 + 6\ell n$	1.18781
	8	7	1.20692 [−2]	$51 + 6\ell n$	1.18290

provide useful information on the performance of quantum computers. Indeed, we expect that the success of quantum LINPACK benchmark is the *minimal requirement* in order to achieve quantum advantages via linear algebra tasks such as QLSP. On the other hand, it is possible to go beyond the quantum LINPACK benchmark, and to formulate a more elaborate cross-entropy test related to QLSP that is also classically hard. This requires detailed studied of the statistical properties of truncated unitaries in the block-encoding model and QSVT problem, which will be our future work.

ACKNOWLEDGMENTS

This work was partially supported by a Google Quantum Research Award (Y.D., L.L.), by the Department of Energy under Grants No. DE-AC02-05CH11231 and No. DE-SC0017867 and the Quantum System Accelerator project (L.L.). This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract No. DE-AC05-00OR22725. We thank Yunchao Liu, Yu Tong, and K. Birgitta Whaley for useful discussions.

APPENDIX A: DERIVATION OF H-RACBEM

To derive Eq. (4), we start from a state $|0\rangle|0\rangle|\psi\rangle$, and follow the circuit in Fig. 4. After applying U_A , the state becomes

$$\frac{1}{\sqrt{2}}(e^{i\varphi_0}|0\rangle + e^{-i\varphi_0}|1\rangle)(|0\rangle A|\psi\rangle + |1\rangle|\perp\rangle). \quad (\text{A1})$$

Here $|\perp\rangle$ is an n -qubit state defined through the relation

$$U_A|0\rangle|\psi\rangle = |0\rangle A|\psi\rangle + |1\rangle|\perp\rangle.$$

Therefore

$$\begin{aligned} |0\rangle|\psi\rangle &= U_A^\dagger U_A|0\rangle|\psi\rangle \\ &= |0\rangle(A^\dagger A)|\psi\rangle + |1\rangle|\perp'\rangle + U_A^\dagger|1\rangle|\perp\rangle \end{aligned}$$

or

$$U_A^\dagger|1\rangle|\perp\rangle = |0\rangle(I_n - A^\dagger A)|\psi\rangle - |1\rangle|\perp'\rangle. \quad (\text{A2})$$

Here $|\perp'\rangle$ is another n -qubit state. Via the relation (A2), after applying U_A^\dagger , the state (A1) is transformed to

$$\begin{aligned} &\frac{1}{\sqrt{2}}(e^{i(\varphi_0+\varphi_1)}|0\rangle + e^{-i(\varphi_0+\varphi_1)}|1\rangle)(|0\rangle A^\dagger A|\psi\rangle + |1\rangle|\perp'\rangle) \\ &+ \frac{1}{\sqrt{2}}(e^{i(\varphi_0-\varphi_1)}|0\rangle + e^{-i(\varphi_0-\varphi_1)}|1\rangle)(|0\rangle(I_n - A^\dagger A)|\psi\rangle \\ &- |1\rangle|\perp'\rangle). \end{aligned} \quad (\text{A3})$$

Finally, carrying out the remaining operations of the circuit in Fig. 4, and applying $|0^2\rangle \otimes I_n$, we obtain the form in Eq. (4).

The quantum circuit for representing a canonical H-RACBEM can be simplified using Fig. 13, still denoted by U_S . Here the two phase shift gates are

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{pmatrix}.$$

Furthermore, we removed the controlled-NOT gates near the Hadamard gates, and replaced the controlled-NOT gates by the standard CNOT gate controlled on 1.

Following the same line of calculation above, starting from $|0\rangle|0\rangle|\psi\rangle$, the state is transformed to

$$\frac{1}{\sqrt{2}}(|0\rangle + e^{i\frac{\pi}{4}}|1\rangle)(|0\rangle A|\psi\rangle + |1\rangle|\perp\rangle)$$

after applying U_A , and then to

$$\begin{aligned} &\frac{1}{\sqrt{2}}(|0\rangle + e^{-i\frac{\pi}{4}}|1\rangle)(|0\rangle A^\dagger A|\psi\rangle + |1\rangle|\perp'\rangle) \\ &+ \frac{1}{\sqrt{2}}(e^{-i\frac{\pi}{2}}|0\rangle + e^{i\frac{\pi}{4}}|1\rangle)(|0\rangle(I_n - A^\dagger A)|\psi\rangle - |1\rangle|\perp'\rangle), \end{aligned}$$

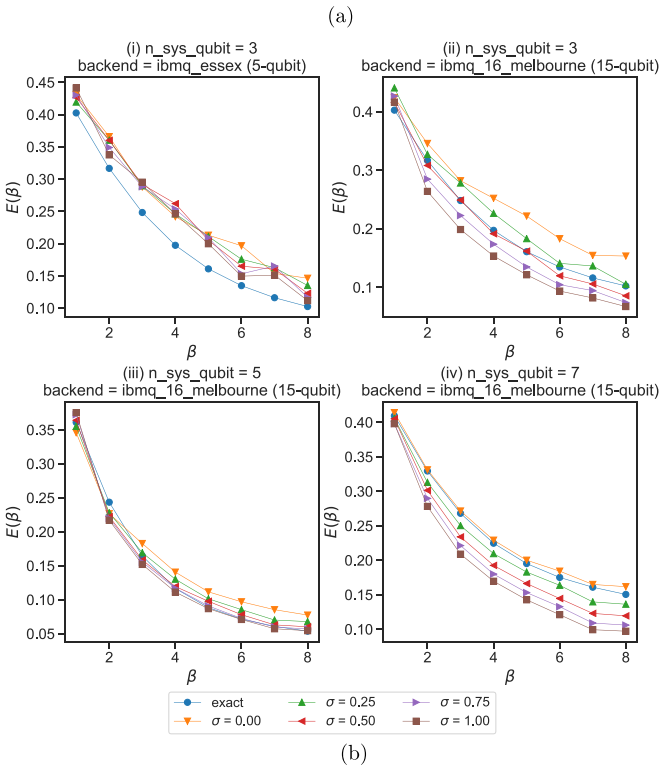



FIG. 12. Thermal average of the energy computed by using the canonical H-RACBEM model and QSVT. The details about phase factors are given in Table VI and the details about block-encoded matrices are given in Fig. 18. (a) Results obtained on IBM Q backends. The probability of CNOT is 0.1 when generating random circuits. (b) Results obtained on QVM. The probability of CNOT is set to 0.5.

after applying U_A^\dagger . Finally, applying the remaining T and H gates, as well as $\langle 0^2 | \otimes I_n$, we obtain the form

$$\mathfrak{H} |\psi\rangle = (\langle 0^2 | \otimes I_n) U_{\mathfrak{H}} |0^2\rangle |\psi\rangle = A^\dagger A |\psi\rangle.$$




FIG. 13. Quantum circuit for generating a (1,2,0) block encoding of a canonical H-RACBEM from a (1,1,0) block encoding U_A and its Hermitian conjugate. H is the Hadamard gate, and Z is the Pauli-Z gate.

APPENDIX B: OPTIMIZATION-BASED METHOD FOR FINDING PHASE FACTORS

In a recent work [34], the authors proposed an optimization-based algorithm for finding phase factors. This is an optimization problem involving matrices only in $SU(2)$, and is independent of the number of system qubits n . The goal is to find a parameterized matrix $U_\Phi(x) = e^{i\phi_0 Z} \prod_{j=1}^d [e^{i \arccos(x) X} e^{i\phi_j Z}]$, where phase factors $\Phi := (\phi_0, \dots, \phi_d)$ are related to those defined in the main text by the following relation:

$$\varphi_i = \begin{cases} \phi_0 + \frac{\pi}{4}, & i = 0, \\ \phi_i + \frac{\pi}{2}, & 1 \leq i \leq d-1, \\ \phi_n + \frac{\pi}{4}, & i = d. \end{cases} \quad (B1)$$

The objective function of the optimization problem is the distance between a real polynomial f and $\text{Re}[\langle 0|U_\Phi(\cdot)|0\rangle]$. The polynomial f is of degree d and its parity is $(d \bmod 2)$. Hence the polynomial can be determined via $\tilde{d} = \lceil \frac{d+1}{2} \rceil$ degrees of freedom. Then according to [34], the objective function can be defined as

$$L(\Phi) = \frac{1}{\tilde{d}} \sum_{j=1}^{\tilde{d}} |\text{Re}[\langle 0|U_\Phi(x_j)|0\rangle] - f(x_j)|^2, \quad (B2)$$

where $x_j = \cos(\frac{(2j-1)\pi}{4\tilde{d}})$, $j = 1, \dots, \tilde{d}$ are the positive roots of the Chebyshev polynomial $T_{2\tilde{d}}(x)$. We may further restrict the set of phase factors to be centrally symmetric. The optimization can be implemented using the standard L-BFGS algorithm, and the running time of the algorithm scales only quadratically with respect to d .

To approximate a generic real smooth function of definite parity, we can use either orthogonal projection onto the set of Chebyshev polynomials on $[-1, 1]$, or use the Remez algorithm to compute the best polynomial approximation on the given subinterval of $[-1, 1]$. This streamlines the procedure of finding the phase factors, and such a procedure can be used to identify phase factors for a very large polynomial degree ($d > 10\,000$) with high precision.

APPENDIX C: SIMULATION MODELS

1. Architecture of quantum devices

The architecture of near-term quantum devices is characterized by several features including the coupling map, basic gates, and noise error rates. The coupling map is given by a directed graph $G = (V, E)$, and the coupling maps of quantum devices provided by IBM Q are always symmetric. The nodes (vertices) V are the set of available qubits on the quantum

$$A = \begin{pmatrix} 0.096 + 0.256i & -0.041 + 0.058i & 0.096 - 0.224i & 0.120 + 0.061i & -0.138 - 0.054i & 0.013 + 0.052i & 0.189 - 0.099i & 0.152 - 0.166i \\ 0.143 - 0.023i & 0.001 - 0.335i & 0.046 - 0.237i & -0.056 + 0.007i & 0.063 + 0.016i & 0.079 - 0.063i & 0.017 + 0.276i & -0.046 + 0.007i \\ 0.054 - 0.017i & -0.073 + 0.149i & -0.002 - 0.063i & -0.128 + 0.128i & -0.371 + 0.048i & -0.163 - 0.102i & -0.069 - 0.069i & 0.126 + 0.037i \\ -0.043 - 0.208i & -0.156 - 0.170i & 0.189 - 0.080i & -0.090 + 0.142i & -0.057 + 0.075i & 0.252 + 0.080i & 0.150 + 0.057i & 0.098 - 0.043i \\ 0.145 + 0.178i & -0.325 + 0.125i & 0.114 + 0.242i & -0.136 - 0.316i & 0.145 + 0.255i & -0.120 - 0.335i & -0.046 + 0.295i & -0.142 - 0.184i \\ -0.117 + 0.149i & -0.101 + 0.338i & -0.213 - 0.018i & -0.474 + 0.081i & -0.036 - 0.121i & 0.444 + 0.147i & -0.198 + 0.035i & -0.091 - 0.054i \\ -0.063 + 0.305i & 0.001 - 0.145i & -0.177 + 0.045i & -0.209 - 0.150i & -0.041 + 0.296i & 0.046 + 0.082i & 0.387 - 0.051i & -0.430 + 0.233i \\ -0.093 - 0.127i & 0.254 + 0.307i & -0.144 - 0.265i & -0.048 - 0.353i & 0.023 + 0.060i & 0.085 - 0.156i & 0.011 + 0.225i & 0.249 + 0.420i \end{pmatrix}$$

$$\Sigma = (0.96439 \ 0.88817 \ 0.86418 \ 0.74016 \ 0.67243 \ 0.50319 \ 0.45951 \ 0.26449)$$

FIG. 14. The relevant information of RACBEM circuit in Fig. 3. The upper matrix A is the three-qubit matrix block-encoded as the upper-left block, namely, identifying q_0 as the block-encoding qubit. The elements in the lower array are the singular values of the block-encoded matrix A .

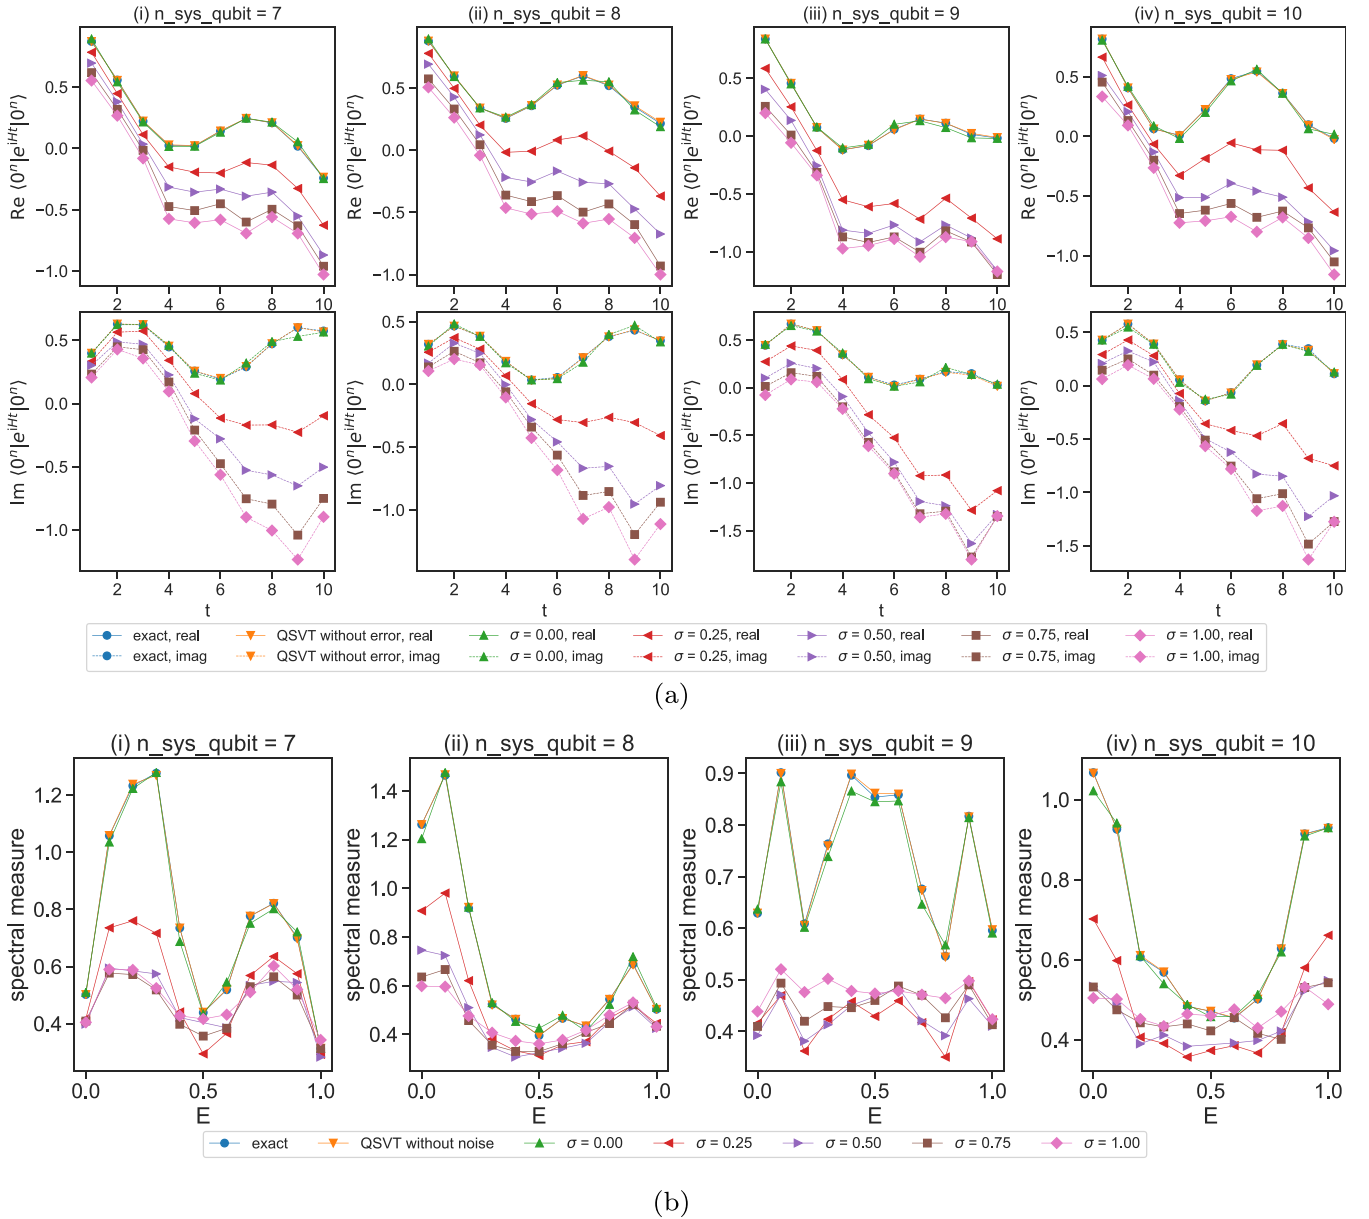


FIG. 15. Numerical results on QVM by using QSVT phase factors with larger length compared to those in the main text. Other parameters are set to be the same as those in the main text. (a) Time series. The details about phase factors are given in Table VI. (b) Spectral measure. The details about phase factors are given in Table IV.

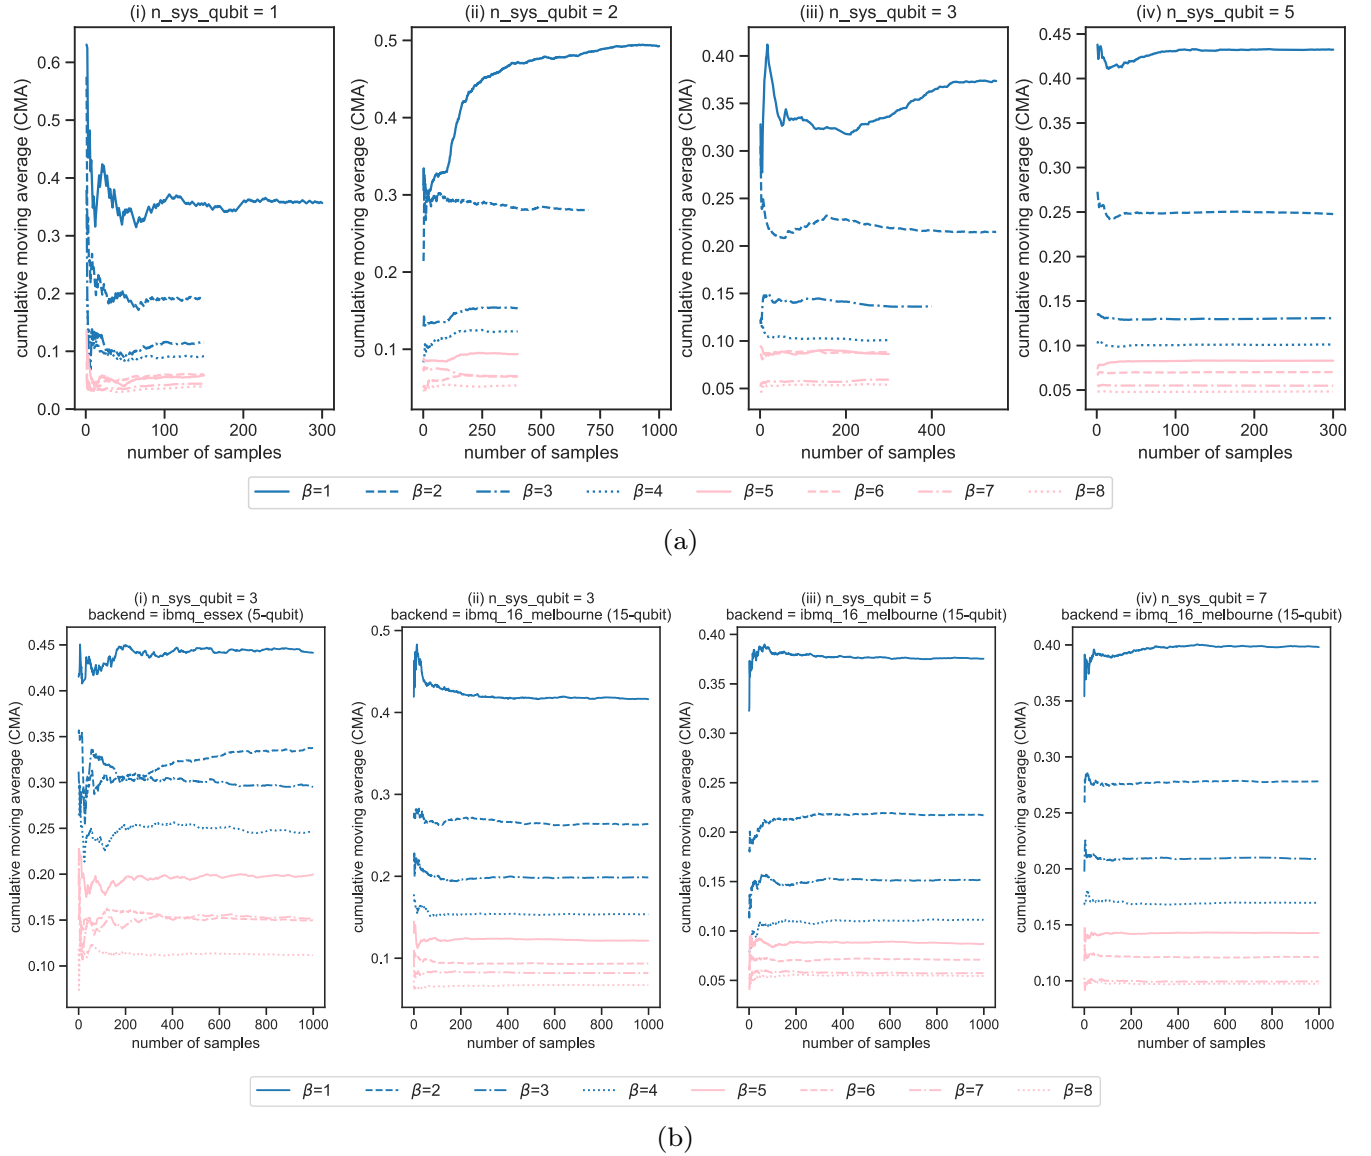


FIG. 16. Cumulative moving average (CMA) of the thermal average of the energy in the METTS algorithm. (a) CMA of the results computed on IBM Q backends. Enough samples are used for each parameter to guarantee the convergence of the result. (b) CMA of the results computed on QVM with noise magnitude $\sigma = 1$.

device, and the edges E specify the set of CNOT gates between qubit pairs that can be directly implemented on the device. The two nodes associated with an edge are assigned to be the control qubit and the target qubit of the CNOT gate, respectively. Basic gates are the building units of quantum circuits that can be directly implemented on the quantum device. If a quantum circuit involves more complicated quantum gates, these gates must first be decomposed into the composition of basic gates by a quantum transpiler before implementation on the quantum hardware. The noise error rate is a measure of the strength of noise on basic gates acting on permitted each qubit or qubit pair.

On quantum computing backends provided by IBM Q, the basic gates are identity gate, CNOT gate, U1, U2, and U3 gate. Up to a global phase factor, U3 gate is

$$U_3(\theta, \phi, \lambda) = R_z(\phi + 3\pi)R_x(\pi/2)R_z(\theta + \pi)R_x(\pi/2)R_z(\lambda),$$

which is a generic single-qubit operation parameterized by three Euler angles. The U1 and U2 gates are defined by restricting to one or two Z-rotation angles respectively:

$$\begin{aligned} U_1(\lambda) &= R_z(\lambda), \quad U_2(\phi, \lambda) \\ &= R_z(\phi + \pi/2)R_x(\pi/2)R_z(\lambda - \pi/2). \end{aligned}$$

The U3 gate can be used to generate arbitrary single-qubit operation [48]. Moreover, in the absence of error correction, the Z-rotation gates can be implemented virtually with in principle negligible error rate. Hence, the error rate of U3 gate is mainly contributed by two X rotations. Although the U1 and U2 gate are specific cases of the U3 gate, they are provided individually for the consideration of reducing error rate, since they involve only zero or one X-rotation operation, respectively. Therefore we exclude U3 from the basic gate set and draw random circuits accordingly with respect to the

Algorithm 1: Random quantum circuit generation for a given quantum device.

Input: Coupling map $G = (V, E)$ where V is the set of qubits, E is the set of qubit pairs on which CNOT is available, basic gates $\Gamma \subseteq \{U1, U2, U3, \text{CNOT}\}$, the probability of choosing CNOT gate $p \in (0, 1)$, the circuit depth ℓ .

Set $t = 0$, initialize an empty quantum circuit \mathcal{C} .

while $t < \ell$ **do**

Reset $(V, E) := G$.

while V is not empty **do**

Draw a random number r uniformly on $[0, 1]$.

if $r \leq p$ and E is not empty **then**

$P = \text{CNOT}$.

Pick a pair of qubits uniformly at random from E as operands.

else

Pick an operation P uniformly at random from $\Gamma \setminus \{\text{CNOT}\}$.

Pick n_p angles uniformly at random on $[0, 2\pi)$ to form P . $n_p = 1, 2, 3$ for $U1, U2, U3$, respectively.

Pick a qubit uniformly at random from V as the operand.

end if

Apply P to its operand(s) in the circuit \mathcal{C} .

Remove the operand(s) from V ; remove the qubit pairs in which the operand(s) involves from E .

end while

Set $t = t + 1$.

end while

Return: A random quantum circuit \mathcal{C} whose depth is ℓ and each layer is fully filled by one- or two-qubit gates.

coupling map and our restricted basic gates. Note that phase gate S and $\pi/8$ -gate T can be implemented as $U1$ gates, and Hadamard gate H is a $U2$ gate. Our restricted basic gate set still has universal representability. The control gates involved in the QSVT circuit can also be directly implemented using the restricted basic gate set, and hence our implementation reduces the usage of noisy $U3$ gates as much as possible.

2. Custom random circuit generator

The Qiskit package provides a utility routine to generate random circuits. However, it does not take into account the coupling map and basic gates available to the target backend, which can be highly inefficient and error prone. In particular, a CNOT gate cannot be directly implemented unless the two qubits are connected according to the coupling map. Therefore the random circuit generated completely randomly cannot be executed directly on the quantum hardware before invoking a quantum transpiler. As a result, we designed a custom random circuit generator in Algorithm 1. The resulting random circuit can be directly implemented on the quantum hardware without the need of using a quantum transpiler. Note that if two $U1$ gates appear consecutively, they can be combined together if needed. The matrix block-encoded in the RACBEM circuit in Fig. 3 is shown in Fig. 14.

3. Construction of the error model

In order to elucidate the impact of noise magnitude, we construct an adjustable noise model as follows. We first retrieve noise models from IBM Q backends which are calibrated by the provider. The retrieved noise model is a python dictionary consisting of the information of each error type. There are some error modes which are quantum errors and can be characterized by a discrete probability distribution, for example, the bit flip error and the phase flip error. Meanwhile,

readout errors in the measurement are also associated with a discrete probability distribution. We introduce a parameter $\sigma \in [0, 1]$ to smoothly adjust the magnitude of such errors. The correct operation corresponds to the entry with the largest magnitude in the probability distribution. Then all other entries are scaled by a factor σ , and the correct entry is adjusted accordingly to satisfy the normalization condition of the distribution. For instance, suppose a quantum error mode is given by the distribution vector $p = [0.90, 0.06, 0.04]$. Then we identify the quantum operation associated with first entry as the correct operation. The scaled distribution vector is then $p_\sigma = [1 - (1 - 0.90)\sigma, 0.06\sigma, 0.04\sigma]$. Therefore, when $\sigma = 1$, the scaled error mode is identical to that retrieved from the backends. When $\sigma = 0$, only the correct operation is applied with probability 1, and the corresponding error symptoms are eliminated.

However, in Qiskit there is another type of quantum error modes, referred to the “Kraus error.” These error modes are not characterized by a discrete probability distribution, and the corresponding Kraus operator is always applied with probability 1. Hence such “Kraus errors” cannot be adjusted using the same method illustrated above. For simplicity of implementation, we discarded such error modes in our noise dictionary. Hence the noise level of our model can be lower than that retrieved directly from the quantum device, even when $\sigma = 1$. Nonetheless, our numerical results demonstrate that the quantum noise in this “diluted” noise model can already significantly impact the output of the quantum algorithms.

APPENDIX D: ADDITIONAL NUMERICAL RESULTS

As discussed in the main text, the circuit depth has a crucial impact on the accuracy of the output. For a given RACBEM, the circuit depth can be most effectively reduced by reducing the length of QSVT phase factors. Hence the choice of a proper length of QSVT phase factors needs to

TABLE VI. Parameters of the QSVT circuit in time series analysis with higher approximation precision, i.e., in Fig. 15(a). Numbers in square brackets indicate the power of 10.

	t	η	Length of QSVT phase factors $d + 1$	QSVT approximation error	Logical gate count	Scaling factor
Real part	1	1.0	5	1.35882 [-4]	$37 + 4\ell n$	1.20020
	2	1.0	5	2.07363 [-3]	$37 + 4\ell n$	1.20298
	3	1.0	7	9.83304 [-4]	$51 + 6\ell n$	1.19894
	4	2.0	9	4.56577 [-3]	$65 + 8\ell n$	1.46804
	5	2.0	11	3.37127 [-3]	$79 + 10\ell n$	1.46510
	6	2.0	13	3.17200 [-3]	$93 + 12\ell n$	1.47531
	7	2.5	13	4.32981 [-3]	$93 + 12\ell n$	1.58693
	8	2.0	15	4.40561 [-3]	$107 + 14\ell n$	1.47748
	9	2.0	17	5.14176 [-3]	$121 + 16\ell n$	1.47880
	10	2.5	17	5.05998 [-3]	$121 + 16\ell n$	1.59715
Imaginary part	1	1.0	5	2.88576 [-4]	$37 + 4\ell n$	1.15186
	2	1.0	5	1.25689 [-3]	$37 + 4\ell n$	1.19857
	3	1.0	5	1.60676 [-3]	$37 + 4\ell n$	1.19769
	4	1.0	7	4.05038 [-3]	$51 + 6\ell n$	1.19621
	5	1.5	9	3.23071 [-3]	$65 + 8\ell n$	1.34575
	6	2.0	11	6.19908 [-3]	$79 + 10\ell n$	1.48017
	7	3.0	13	3.50032 [-3]	$93 + 12\ell n$	1.70412
	8	3.0	15	2.35584 [-3]	$107 + 14\ell n$	1.70038
	9	4.0	15	3.59205 [-3]	$107 + 14\ell n$	1.90551
	10	3.0	17	3.53947 [-3]	$121 + 16\ell n$	1.70218

balance the error due to the polynomial approximation, and that caused by the noisy quantum device. To verify this, we increase the length of QSVT phase factors for computing the time series and spectral measures (details in Tables VI and IV). From Fig. 15 it is evident that the numerical results are more accurate in the absence of noise when $\sigma = 0$. However, as the noise magnitude increases, the results with a deeper quantum circuit become significantly worse compared to those in Figs. 10 and 11.

The METTS algorithm can be used to compute thermal averages via a Markov chain. The convergence behavior, plotted using the cumulative moving average (CMA) of the thermal average of the energy, is given in Fig. 16.

APPENDIX E: SIMULATION DETAILS

In this section we provide the details of QSVT circuits used in our numerical experiments. The choice of the depth ℓ of the random quantum circuit for different number of system qubits n is discussed in the main text. The logical gate count is given by the number of basic gates used in implementing the circuit, where we set the basic gate set to $\Gamma = \{U1, U2, CNOT\}$ in numerical experiments. We display the statistics of singular-value distributions of distinct numbers of system qubits in Fig. 17. From the mean and standard deviation, it is evident that under our choice of the parameters of random circuit, the singular values of the block-encoded matrices varies largely. The singular values of relevant matrices which are used in the numerical tests in the main text are displayed in Fig. 18. Those matrices are block encoded in RACBEM circuits generated according to our setup and Algorithm 1 at random. Such results indicate that the RACBEM model can generate at least nontrivial block-encoded matrices useful for testing the performance of quantum algorithms for numerical linear algebra tasks.

The details of the parameters used in the QSVT circuits are given below. We introduce a scaling factor, which is




FIG. 17. Singular-value distributions of different numbers of system qubits. The depth of the random circuit is adaptively chosen according to the number of system qubits. When the number of system qubits is less than or equal to 3, the coupling map is retrieved from the five-qubit backend `ibmq_burlington`, otherwise, the coupling map is retrieved from the 15-qubit backend `ibmq_16_melbourne`. The probability of CNOT is 0.5, and the basic gate is restricted to $\Gamma = \{U1, U2, CNOT\}$. We generate 500 RACBEM s for each number of system qubits and compute the statistics of the difference between the maximal and minimal singular value of the block-encoded matrix.

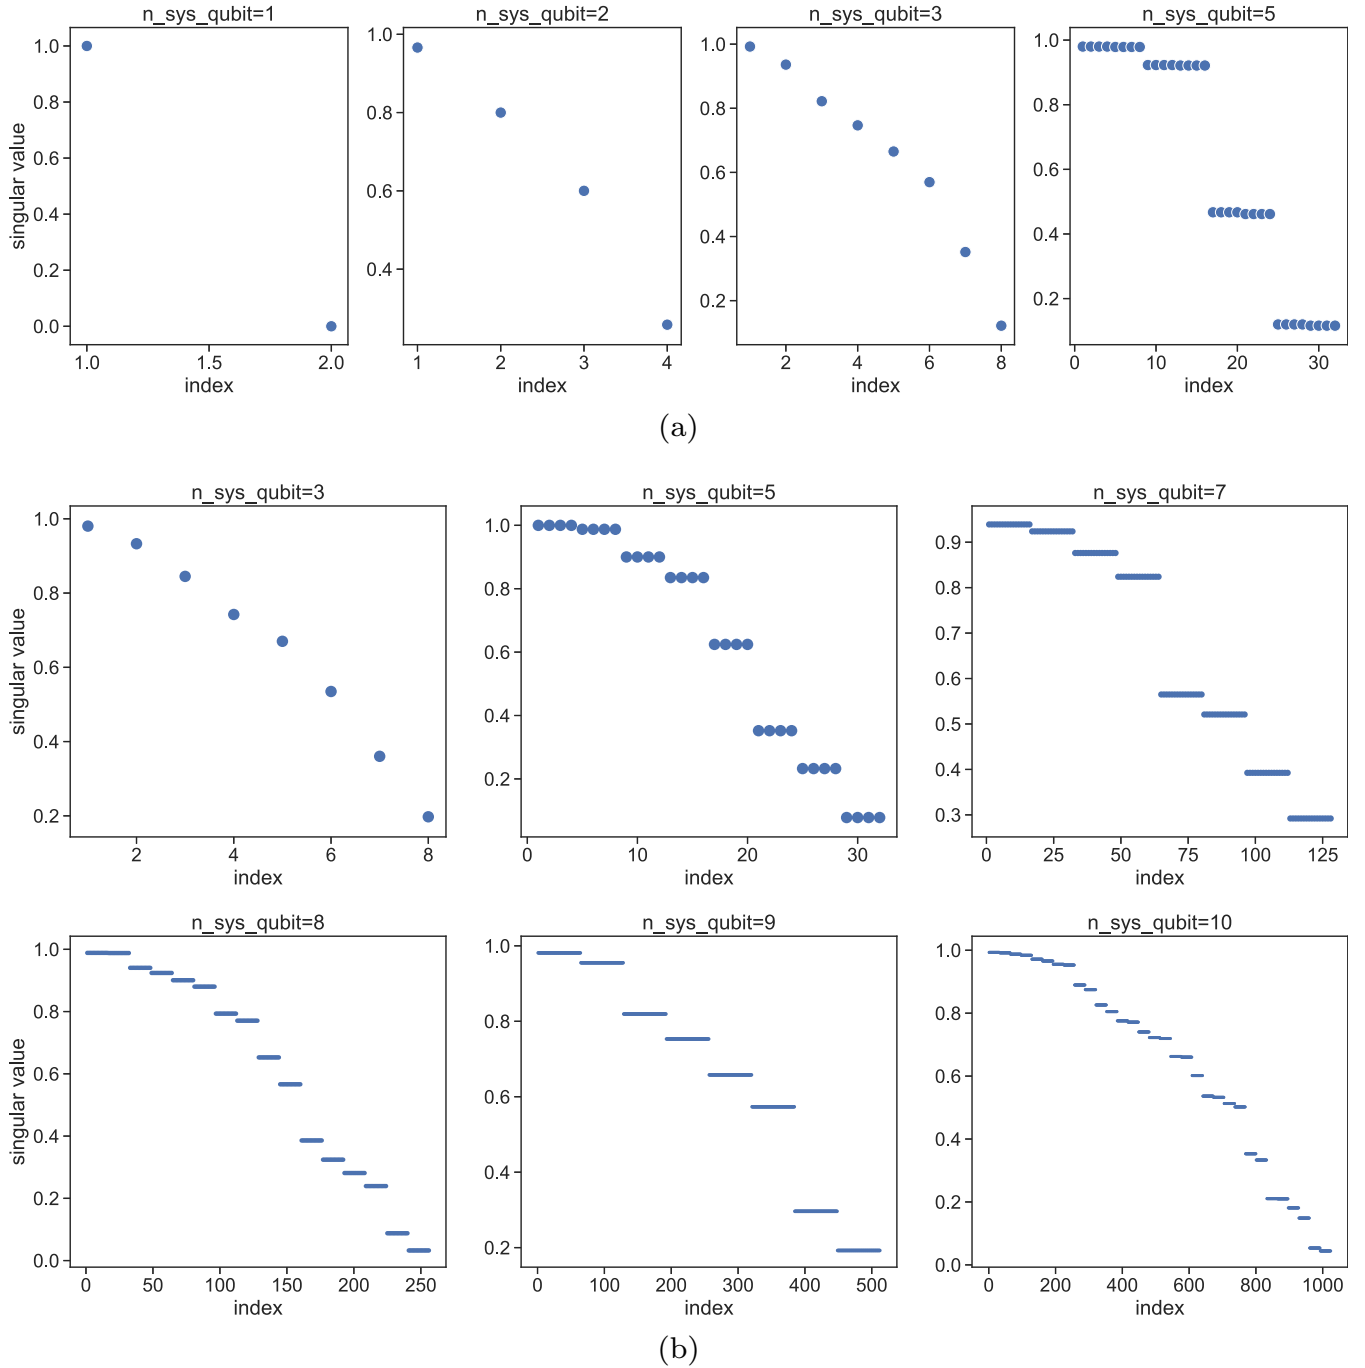


FIG. 18. Singular values of the matrices used in the numerical tests in the main text. The blue dots mark the singular values of the block-encoded matrix of a RACBEM circuit. When the number of system qubits is less than or equal to 3, the coupling map is identified as that on the five-qubit backend `ibmq_burlington`. Other results are obtained using the coupling map retrieved from the 15-qubit backend `ibmq_16_melbourne`. (a) Singular values of matrices used in the numerical experiments run on quantum computing backends provided by IBM Q. The probability of CNOT is 0.1 when generating random circuits. (b) Singular values of matrices used in the numerical experiments run on QVM. The probability of CNOT is 0.5 when generating random circuits.

slightly greater than the maximal absolute value of the target polynomial, to scale the target polynomial so that it can be

properly parameterized as QSVT phase factors. Numbers in square brackets indicate the power of 10.

[1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*,

Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).

- [2] J. J. Dongarra, P. Luszczek, and A. Petitet, The LINPACK benchmark: Past, present and future, *Concurr. Comput.-Pract. Exper.* **15**, 803 (2003).
- [3] <https://quantum-computing.ibm.com>.
- [4] <https://www.top500.org/project/>.
- [5] D. An and L. Lin, Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm, *arXiv:1909.05500* (2019).
- [6] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Variational quantum linear solver: A hybrid algorithm for linear systems, *arXiv:1909.05820* (2019).
- [7] Y. Cao, A. Papageorgiou, I. Petras, J. Traub, and S. Kais, Quantum algorithm and circuit design solving the Poisson equation, *New J. Phys.* **15**, 013021 (2013).
- [8] P. A. M. Casares and M. A. Martin-Delgado, A quantum interior-point predictor-corrector algorithm for linear programming, *J. Phys. A* **53**, 445305 (2020).
- [9] S. Chakraborty, A. Gilyén, and S. Jeffery, The power of block-encoded matrix powers: Improved regression techniques via faster hamiltonian simulation, in *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2019), pp. 33:1-33:14.
- [10] 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).
- [11] 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* (Association for Computing Machinery, New York, NY, 2019), pp. 193–204.
- [12] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [13] L. Lin and Y. Tong, Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems, *Quantum* **4**, 361 (2020).
- [14] Y. Subaşı, R. D. Somma, and D. Orsucci, Quantum Algorithms for Systems of Linear Equations Inspired by Adiabatic Quantum Computing, *Phys. Rev. Lett.* **122**, 060504 (2019).
- [15] Y. Tong, D. An, N. Wiebe, and L. Lin, Fast inversion, preconditioned quantum linear system solvers, and fast evaluation of matrix functions, *arXiv:2008.13295* (2020).
- [16] L. Wossnig, Z. Zhao, and A. Prakash, Quantum Linear System Algorithm for Dense Matrices, *Phys. Rev. Lett.* **120**, 050502 (2018).
- [17] X. Xu, J. Sun, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational algorithms for linear algebra, *arXiv:1909.03898* (2019).
- [18] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and Robust Randomized Benchmarking of Quantum Processes, *Phys. Rev. Lett.* **106**, 180504 (2011).
- [19] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography, *Nat. Commun.* **8**, 14485 (2017).
- [20] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Validating quantum computers using randomized model circuits, *Phys. Rev. A* **100**, 032328 (2019).
- [21] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, *Nat. Commun.* **10**, 5347 (2019).
- [22] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [23] T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, Measuring the capabilities of quantum computers, *arXiv:2008.11294* (2020).
- [24] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, *Phys. Rev. Lett.* **100**, 160501 (2008).
- [25] E. Tang, A quantum-inspired classical algorithm for recommendation systems, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, 2019), pp. 217–228.
- [26] M. Motta, C. Sun, A. T. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [27] E. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [28] S. R. White, Minimally Entangled Typical Quantum States at Finite Temperature, *Phys. Rev. Lett.* **102**, 190601 (2009).
- [29] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [30] M. Saeedi and M. Pedram, Linear-depth quantum circuits for n -qubit Toffoli gates with no ancilla, *Phys. Rev. A* **87**, 062318 (2013).
- [31] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, in *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science* (Institute of Electrical and Electronics Engineers, New York, NY, 2015), pp. 792–809.
- [32] A. W. Harrow and R. A. Low, Random quantum circuits are approximate 2-designs, *Commun. Math. Phys.* **291**, 257 (2009).
- [33] J. Haah, Product decomposition of periodic functions in quantum signal processing, *Quantum* **3**, 190 (2019).
- [34] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Efficient phase-factor evaluation in quantum signal processing, *Phys. Rev. A* **103**, 042419 (2021).
- [35] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, *arXiv:1806.01838*. An abbreviated version of this paper was published in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019) (Association for Computing Machinery, New York, NY, 2019), pp. 193–204.

- [36] R. D. Somma, Quantum eigenvalue estimation via time series analysis, *New J. Phys.* **21**, 123025 (2019).
- [37] <https://github.com/qsppack/RACBEM>.
- [38] <https://github.com/qsppack/QSPACK>.
- [39] H. Abraham *et al.*, Qiskit: An open-source framework for quantum computing. zenodo, <https://doi.org/10.5281/zenodo.2562111>, 2019 (2019).
- [40] S. Aaronson and A. Arkhipov, The computational complexity of linear optics. in *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, 2011), pp. 333–342.
- [41] S. Aaronson and L. Chen, Complexity theoretic foundations of quantum supremacy experiments, [arXiv:1612.05903](https://arxiv.org/abs/1612.05903) (2016).
- [42] S. Aaronson and S. Gunn, On the classical hardness of spoofing linear cross-entropy benchmarking, *Theory of Computing* **16**, 1 (2020).
- [43] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, *Nat. Phys.* **15**, 159 (2019).
- [44] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations, *Phys. Rev. Lett.* **117**, 080501 (2016).
- [45] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory, Pseudo-random unitary operators for quantum information processing, *Science* **302**, 2098 (2003).
- [46] K. Fujii, Noise threshold of quantum supremacy, [arXiv:1610.03632](https://arxiv.org/abs/1610.03632) (2016).
- [47] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, Quantum Entanglement Growth Under Random Unitary Dynamics, *Phys. Rev. X* **7**, 031016 (2017).
- [48] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient Z gates for quantum computing, *Phys. Rev. A* **96**, 022330 (2017).