

Fast Quantum Simulation of Electronic Structure by Spectral Amplification

Guang Hao Low^{1,*}, Robbie King^{1,2}, Dominic W. Berry³, Qiushi Han,⁴ A. Eugene DePrince III^{1,5}, Alec F. White⁶,
Ryan Babbush¹, Rolando D. Somma¹, and Nicholas C. Rubin^{1,†,‡}

¹*Google Quantum AI, Venice, California 90291, USA*

²*Department of Computing and Mathematical Sciences, California Institute of Technology (Caltech), Pasadena, California 91125, USA*

³*School of Mathematical and Physical Sciences, Macquarie University, Sydney, New South Wales 2109, Australia*

⁴*Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

⁵*Department of Chemistry and Biochemistry, Florida State University, Tallahassee, Florida, USA*

⁶*Quantum Simulation Technologies Inc., Boston, Massachusetts, USA*



(Received 6 March 2025; revised 17 July 2025; accepted 9 September 2025; published 31 October 2025)

The most advanced techniques using fault-tolerant quantum computers to estimate the ground-state energy of a chemical Hamiltonian involve compression of the Coulomb operator through tensor factorizations, enabling efficient block encodings of the Hamiltonian. A natural challenge of these methods is the degree to which block-encoding costs can be reduced. We address this challenge through the technique of spectral amplification, which magnifies the spectrum of the low-energy states of Hamiltonians that can be expressed as sums of squares. Spectral amplification enables estimating ground-state energies with significantly improved cost scaling in the block encoding normalization factor Λ to just $\sqrt{2\Lambda E_{\text{gap}}}$, where $E_{\text{gap}} \ll \Lambda$ is the lowest energy of the sum-of-squares Hamiltonian. To achieve this, we show that sum-of-squares representations of the electronic structure Hamiltonian are efficiently computable by a family of classical simulation techniques that approximate the ground-state energy from below. In order to further optimize, we also develop a novel factorization that provides a trade-off between the two leading Coulomb integral factorization schemes—namely, double factorization and tensor hypercontraction—that when combined with spectral amplification yields a factor of 4 to 195 speedup over the state of the art in ground-state energy estimation for models of iron-sulfur complexes and a CO₂-fixation catalyst.

DOI: [10.1103/pb2g-j9cw](https://doi.org/10.1103/pb2g-j9cw)

Subject Areas: Quantum Physics,
Quantum Information

I. INTRODUCTION

High-accuracy simulation of the electronic structure of molecules and materials is one of the primary motivations for building a quantum computer. This motivation originates from the difficulty of robustly simulating electrons with classical techniques at low enough cost, to accommodate large systems necessary for quantitative prediction of chemical phenomena such as reaction chemistry, molecular spectra, and ground-state bonding information. The phase estimation algorithm (PEA) [1], implemented on a fault-tolerant quantum computer, provides one such route

towards quantitative electronic structure simulation at polynomial cost in the system size [2,3]. The asymptotic algorithmic advantage of the PEA led to a large body of work examining the constant factors associated with ground-state energy estimation. While initial quantum resource estimates were exceedingly high [4], the focus on optimizing quantum algorithms has resulted in drastic reductions. As a representative example, over the past six years, the quantum resources (gate complexity) required for PEA of the iron-molybdenum co-factor FeMoco have steadily improved by factors of 2–4 with each algorithmic innovation [5–9], resulting in a 7 orders of magnitude reduction in resources. These reductions enable efficient ground-state energy estimation of classically intractable systems that are predicted to run on future, fault-tolerant quantum computers faster than lower-accuracy classical methods [10,11]. Computational prediction of chemistry usually requires many calculations and thus even further reductions will be required to enable routine large-scale simulation [12]. Herein, we introduce a new algorithmic strategy and take the first steps towards an optimized

*Contact author: guanghaolow@google.com

†Contact author: nickrubin@google.com

‡These authors contributed equally to this work.

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

implementation, resulting in a substantial reduction over prior art.

At the core of recent algorithmic improvements in second quantized electronic structure simulation is the exploitation of tensor decompositions of the Coulomb operator yielding an efficient linear combination of unitaries (LCU) representation. Using the LCU representation, a unitary (a quantum walk operator) implementing a function of the Hamiltonian can be efficiently constructed and used by the PEA [13–15]. Ultimately, the quantum cost of performing phase estimation on a walk operator constructed using the LCU representation scales linearly in the normalization constant associated with a “block encoding” of the Hamiltonian [16,17]. For electronic structure simulation, the normalization constant is intimately related to the degree of compression of the two-electron integral operator when constructing the LCU. The two primary strategies for compressing the large two-electron integrals using tensor factorizations are double factorization (DF) [9] and tensor hypercontraction (THC) [7], each providing normalization constants that scale as low-order polynomials of system size. The freedom of the LCU representation and symmetries of chemical Hamiltonians allow for modification of the Hamiltonian and the LCU normalization constant, by adding operators that annihilate any state with the desired symmetries. This strategy, called the block-invariant symmetry shift (BLISS), was recently demonstrated to reduce LCU normalizations by a factor of 2 when using particle number symmetry shifts [18,19]. Bounds identifying the limits of symmetry shifting also suggest that current integral factorizations combined with symmetry shifting are within a factor of 4 of the optimal LCU normalization [8,20]. Therefore, finding the algorithmic improvements necessary to make quantum chemistry calculations routine will not be possible solely through more sophisticated tensor factorizations.

To address the stagnating returns of symmetry-shifting and compression, we show that further quantum speedups are achievable by the quantum algorithmic idea of “spectral amplification” (SA). Fundamentally, SA allows one to relax the precision requirements in certain estimates, potentially resulting in significantly improved resource counts. The concept was initially introduced to achieve a quadratic quantum speedup within the context of adiabatic quantum computation [21], and was more recently used to obtain improved quantum algorithms for simulating dynamics of low-energy states [22,23]. SA first produces a “square root” of the Hamiltonian, so that the low-energy spectrum is amplified (i.e., $\sqrt{x} \gg x$ for $x \ll 1$). A precise estimate of a small eigenvalue of the original Hamiltonian can then be obtained by squaring a coarser estimate of an eigenvalue of its square root. The latter problem is expected to be solved more efficiently using the PEA [24]; however, to obtain a net gain, we must ensure that simulating the Hamiltonian square root is not considerably more difficult

than simulating the original Hamiltonian. Otherwise, the improvements due to requiring less precision in the estimate could be negated by the overhead due to simulating the square root.

To this end, we note that SA is directly applicable to Hamiltonians represented as a sum of squares (SOS):

$$H = \sum_{\alpha} O_{\alpha}^{\dagger} O_{\alpha} + E_{\text{SOS}} = H_{\text{sqr}}^{\dagger} H_{\text{sqr}} + E_{\text{SOS}}, \quad (1)$$

where E_{SOS} is an energy shift ensuring that $(H - E_{\text{SOS}})$ is positive semidefinite. The matrix representations of operators O_{α} are not necessarily square and we can define $H_{\text{sqr}} := \sum_{\alpha} |\chi_{\alpha}\rangle \otimes O_{\alpha}$ for any orthogonal set of states $|\chi_{\alpha}\rangle$. [Eq. (1) is also known as a “gap-amplifiable” Hamiltonian in Ref. [23].] SA focuses then on simulating H_{sqr} rather than H . To quantify the potential improvement due to SA, we can consider a dimensionless parameter $\Delta_{\text{gap}} := (E_{\text{gs}} - E_{\text{SOS}})/\lambda_{\text{sqr}}^2$, where $E_{\text{gs}} \geq E_{\text{SOS}}$ is the ground-state energy of H , and λ_{sqr} is a normalization factor that will appear in the process that block encodes H_{sqr} [22]. Ideally, we would like to satisfy $\Delta_{\text{gap}} \ll 1$, in which case $\sqrt{\Delta_{\text{gap}}} \gg \Delta_{\text{gap}}$, allowing us to achieve a large amplification of the low-energy spectrum and exploit the power of SA. However, this is not always easy to obtain; SOS representations for Hamiltonians of practical relevance with this property might not be efficiently computed [21,23]. And, even in cases where H_{sqr} can be efficiently computed, we need to ensure that it *can be* simulated efficiently, which is not obvious using known block-encoding techniques. In this work we are able to address these difficulties with three results.

First, we introduce practical quantum algorithms for spectral amplification of the SOS representation. As H_{sqr} has singular values $\sqrt{\Delta_{\text{gap}}} \lambda_{\text{sqr}} = \sqrt{E_{\text{gs}} - E_{\text{SOS}}}$, an asymptotic improvement in ground-state energy estimation can already be obtained by singular value estimation of H_{sqr} [24,25]. However, current methods for block encoding matrix square roots [23,26] and for singular value estimation have much larger constant factors than standard block encoding and phase estimation techniques [14]. We resolve this by introducing rectangular block encodings of H_{sqr} that are conceptually simpler and cheaper than previous approaches based on Hermitian dilations [27]. A quantum walk unitary can then be defined on rectangular block encodings with eigenphases $\arccos(\sqrt{\Delta_{\text{gap}}})$ [25]. A subtlety arises from this kind of quantum walk as it maps right singular values to left singular values and vice versa, which complicates a naive application of standard phase estimation techniques [14]. Nevertheless, we prove that two quantum walk steps restores the use of standard phase estimation techniques whilst maintaining the SA advantage.

Second, we elucidate how to express a second quantized representation of H as an SOS Hamiltonian and compute an

optimal Δ_{gap} in polynomial time. The protocol for achieving the equality of Eq. (1) is related to SOS relaxations of noncommutative polynomial optimization problems which have been extensively used in the quantum chemistry, condensed matter, and quantum information communities [28–30]. The optimization relaxations rely on an SOS representation of non-negative model Hamiltonians which are used to construct an optimization protocol based on semidefinite programming. The primal variables of the semidefinite program (SDP) are the pseudoexpectation values defining approximations to the quantum marginals of a state [30–34]. With minor modification, the dual of the lower bound program provides the solution to Eq. (1) with an optimal Δ_{gap} given a particular form for the SOS algebra defining the functional form of O_α . It is well understood that expanding the algebra leads to tighter lower bound estimates E_{SOS} for bosonic and fermionic Hamiltonians [30,35] and that E_{SOS} can be made arbitrarily close to E_{gs} , and thus $\Delta_{\text{gap}} \ll 1$, with exponential classical resources [36]. In the context of SA, minimizing the Δ_{gap} by using a large SOS algebra is in tension with a low-cost block encoding of the SOS Hamiltonian. We find for interesting quantum systems represented by the spin-free Hamiltonian

$$H = \sum_{pq} h_{pq}^{(1)} E_{pq} + \frac{1}{2} \sum_{pqrs} h_{pqrs}^{(2)} E_{pq} E_{rs},$$

$$E_{pq} := \sum_{\sigma \in \{0,1\}} a_{p\sigma}^\dagger a_{q\sigma}, \quad (2)$$

where $h_{pq}^{(1)}$ and $h_{pqrs}^{(2)}$ are coefficients for the one- and two-electron operators, respectively, that a spin-free variant of the level-2 sum-of-squares algebra (where O_α are of quadratic many-body order) already gives small gaps and leads to improvements in ground-state energy estimation. The spin-free variant of the level-2 SOS algebra is the simplest algebra in the SOS hierarchy for chemistry Hamiltonians

and more bespoke constructions of the SOS algebra provides a route to investigate the trade-off between gap minimization and block-encoding costs.

Third, we introduce a new tensor factorization, double-factorized-tensor-hypercontraction (DFTHC), that realizes an SOS Hamiltonian representation that further optimizes the block-encoding cost. Knowledge of the optimal Δ_{gap} found by the SDP illuminates ultimate limits achievable by SA. However, the SOS Hamiltonian found by SDP does not optimize other quantum-algorithm-relevant parameters. To overcome this limitation, DFTHC is a variational ansatz that allows for a direct minimization of the block-encoding normalization factor λ_{sqr} , and guided by prior knowledge of the optimal gap Δ_{gap} , almost achieves the optimal E_{SOS} . We optimize the DFTHC ansatz by gradient descent on a nonlinear cost function that uses Δ_{gap} as a regularization parameter but, unlike the SDP, cannot be provably solved optimally in polynomial time. Nevertheless, we reliably find suboptimal solutions that are extremely compact and chemically accurate approximations to the two-body interaction tensor. DFTHC is also readily combined with BLISS. For benchmarks on molecules with up to $N = 150$ orbitals, we use the DFTHC representation in all our following resource estimates to report chemically accurate representations with block-encoding Toffoli gate costs scaling like $\sim N^{0.96}$, $\lambda_{\text{sqr}}^2 \sim N^{1.46}$, and $E_{\text{gap}} \sim N^{0.88}$.

Combining the three aforementioned results, we report the quantum resources needed for ground-state energy estimation of FeMoco in Table I compared with prior works. The efficiency of DFTHC + BLISS + SA enables a factor of 4 reduction in costs for the larger FeMoco active space [38] using the simplest SOS algebra of the hierarchy. To verify the robustness of the method we also compute resource estimates for small iron-sulfur clusters [40], a CO₂ ruthenium catalyst [9] series, and the CPD1 active space of P450 [11]; all of which have been targets of recent resource estimation work and demonstrate the scaling of the

TABLE I. Improvements in Toffoli and qubit costs on benchmark molecules for performing phase estimation targeting ground-state energies with a standard deviation of at most chemical accuracy $\epsilon_{\text{chem}} = 1.6$ mHa. Total Toffoli counts are based on a phase estimation uncertainty of $\sigma_{\text{PEA}} = 1$ mHa.

Year	Innovation	FeMoco-54 [37]			FeMoco-76 [38]		
		Qubits	Toffolis	Reference	Qubits	Toffolis	Reference
2017	First resource estimate by Trotterization [37]	111	5.0×10^{13}	[37]	...	...	...
2019	Qubitization of single factorization [17]	3320	9.5×10^{10}	[7]	3628	1.2×10^{11}	[7]
2020	Qubitization of double factorization (DF) [9]	3600	2.3×10^{10}	[9]	6404	5.3×10^{10}	[7]
2020	Tensor hypercontraction (THC) [7]	2142	5.3×10^9	[7]	2196	3.2×10^{10}	[7]
2024	Symmetry compression of DF [39]	1994	2.6×10^9	[39]	...	...	...
2025	Symmetry compression of THC [8]	...	...	...	1512	4.3×10^9	[8]
This work	Spectral amplification and DFTHC	1132	3.42×10^8		1454	1.00×10^9	
Improvement of this work over [39,8] ^a		1.8×	7.0×		1.0×	4.3×	

^aOur approach is distinct and independent of recent results by Caesura *et al.* [8].

effective LCU normalization. For these systems, improvement factors of up to 200 are obtained for unsymmetry-shifted Hamiltonians. Like prior work [7,9,17] the reported Toffoli and qubit complexities correspond with the task of obtaining an estimate of E_{gs} with a standard deviation σ_E of at most chemical accuracy $\epsilon_{\text{chem}} = 1.6$ mHa. This is accomplished using the PEA to accuracy σ_{PEA} , which projects an input trial state $|\psi_{\text{trial}}\rangle$ onto the true ground state with probability almost $p = |\langle\psi_{\text{gs}}|\psi_{\text{trial}}\rangle|^2$ and then returns an estimate of E_{gs} with standard deviation σ_{PEA} . For many interesting electronic structure problems, we can obtain $p \approx 1$ by preparing matrix product states (MPS) trial states based on modest bond dimension DMRG calculations which are still estimated to need further energy refinement [10,41]. For consistency with previous work, we assume $p = 1$ in our resource estimates and the cost of PEA is then $[\pi/(2\sigma_{\text{PEA}})]\lambda_{\text{eff}}\text{Cost}[\text{BE}]$, where $\text{Cost}[\text{BE}]$ is the cost of block encoding a function of the H_{sqr} .

Attaining chemical accuracy requires bounding error contributions from σ_{PEA} and the DFTHC tensor approximation error. Tightly and efficiently bounding the error associated with an approximate Hamiltonian in DFTHC form presents a numerical and theoretical challenge that can ultimately impact the overall cost of phase estimation. In order to fully quantify how to deploy the DFTHC factorization we challenge the commonly used truncation threshold metric based on the high-spin CCSD(T) correlation energy ϵ_{corr} of the factorized representation. If errors are systematic, chemical accuracy is attained if these error satisfy

$$\sigma_E \leq \sigma_{\text{PEA}} + |\epsilon_{\text{corr}}| \leq \epsilon_{\text{chem}}, \quad (3)$$

enabling us to select $\sigma_{\text{PEA}} = 1.0$ mHa, and $|\epsilon_{\text{corr}}| \leq 0.6$ mHa consistent with prior resource estimates [7,9]. However, we demonstrate that it is possible to choose much larger $\sigma_{\text{PEA}} = 1.4$ mHa by randomizing bits-of-precision truncation and initial conditions in the DFTHC optimization, so that ϵ_{corr} is a random variable with standard deviation $\sigma_{\text{corr}} \leq 0.774$ mHa and zero mean. The randomization protocol has the added benefit that the variances of random errors satisfy

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{corr}}^2} \leq \epsilon_{\text{chem}}. \quad (4)$$

We present resource estimates for this approach in a later section as it is not directly comparable to previous work but constitutes an improved strategy for accounting for errors in ground-state energy estimation.

The remainder of the paper is structured as follows: Section II reviews the general SA protocol and describes circuit improvements framed by modifications to singular value estimation, Sec. III describes the SDP approach to minimizing the SOS gap for electronic structure problems

and demonstrates that $\Delta_{\text{gap}} \ll 1$ are achievable at polynomial classical cost, Sec. IV introduces DFTHC, the quantum circuits for its block encoding, and numerical result. In Sec. V we present resource estimates for a variety of benchmark systems and report improvement factors over prior work. We also highlight potential challenges in using high-spin correlation energy from CCSD(T) and improved strategies for optimizing errors for each algorithmic component to achieve chemical precision.

II. QUANTUM CIRCUITS FOR SPECTRAL AMPLIFICATION

We describe SA within the context of SOS representations for Hamiltonians, discuss block encodings and PEA, and present practical quantum circuits for their implementations. Let

$$H_{\text{SA}} := \sum_{\alpha=0}^{L-1} O_{\alpha}^{\dagger} O_{\alpha}, \quad H = H_{\text{SA}} + E_{\text{SOS}}, \quad (5)$$

where O_{α} corresponds to a linear combination of operators. Assuming an SOS representation like Eq. (5) for an arbitrary Hamiltonian acting on a finite-dimensional space is without loss of generality; in Sec. III, we show how this representation can be computed efficiently. As $O_{\alpha}^{\dagger} O_{\alpha} \geq 0$ is positive semidefinite, the eigenvalues of H_{SA} are non-negative. The special case where the lowest eigenvalue of H_{SA} is zero corresponds to H_{SA} being “frustration free,” in which case the smallest eigenvalues of all $O_{\alpha}^{\dagger} O_{\alpha}$ ’s are also zero. For a given H , it is always possible to shift it and find a frustration-free H_{SA} , but this might require exponential-cost preprocessing and thus is not feasible. Despite this, it might be possible to efficiently find an H_{SA} that can be written as an SOS and that is “near” frustration free, with lowest eigenvalue $E_{\text{gap}} = E_{\text{gs}} - E_{\text{SOS}}$ close to zero, which is the approach considered in this work. Note that the square root H_{sqr} of $H_{\text{SA}} = H_{\text{sqr}}^{\dagger} H_{\text{sqr}}$ has a singular value $\sqrt{E_{\text{gap}}}$, having an amplified sensitivity to variations in E_{gs} .

The quantum circuits we present for SA are based on the block-encoding framework [16]. We say a finite-dimensional operator O is block encoded with normalization factor $\lambda > 0$ such that $\|O\| \leq \lambda$, if there is a unitary quantum circuit

$$\text{BE}[O/\lambda] = \begin{pmatrix} O/\lambda & \cdots \\ \vdots & \ddots \end{pmatrix} \in \mathbb{C}^{n \times n} \Rightarrow \Pi_l \text{BE}[O/\lambda] \Pi_m = O/\lambda, \quad (6)$$

where $O/\lambda \in \mathbb{C}^{l \times m}$ appears as a block of $\text{BE}[O/\lambda]$, with $l \leq n$ and $m \leq n$. The operators Π_l and Π_m are orthogonal projectors onto subspaces of dimension l and m , specified by the (all-zero) strings $|0\rangle_l$, and $|0\rangle_m$, respectively.

(Note that $l \neq m$ if O is not square.) Highly efficient block encoding of H_{sqr} are possible to construct. Assume we are given access to controlled-block encodings $\text{BE}[O_\alpha/\lambda_\alpha]$, for some $\lambda_\alpha \geq \|O_\alpha\|$. Then, one may apply unitary iteration and coherent alias sampling state preparation [14] to synthesize the “SELECT” and “PREPARE” circuits

$$\text{SEL} := \sum_{\alpha=0}^{L-1} |\alpha\rangle\langle\alpha| \otimes \text{BE}\left[\frac{O_\alpha}{\lambda_\alpha}\right], \quad \text{PREP}|0\rangle_a = \frac{1}{\lambda_{\text{sqr}}} \sum_{\alpha=0}^{L-1} \lambda_\alpha |\alpha\rangle_a |\text{garb}_a\rangle_a = \frac{1}{\lambda_{\text{sqr}}} \sum_{\alpha=0}^{L-1} \lambda_\alpha |\chi_\alpha\rangle_a, \quad (7)$$

where $\lambda_{\text{sqr}} := \sqrt{\sum_{\alpha=0}^{L-1} \lambda_\alpha^2}$, “a” refers to ancillary qubits, and $|\text{garb}_a\rangle_a$ represents a “garbage” (unit) state that is irrelevant. Then, we can let $H_{\text{sqr}} = \sum_{\alpha=0}^{L-1} |\chi_\alpha\rangle \otimes O_\alpha$ and use these operations to construct the desired block encoding of H_{sqr} ; that is,

$$\Pi_l \text{SEL} \Pi_m \cdot \text{PREP}|0\rangle_a = \frac{1}{\lambda_{\text{sqr}}} \sum_{\alpha=0}^{L-1} |\chi_\alpha\rangle \otimes O_\alpha \Rightarrow \text{SEL} \cdot \text{PREP} = \text{BE}\left[\frac{H_{\text{sqr}}}{\lambda_{\text{sqr}}}\right] := U. \quad (8)$$

Note that U acts on an enlarged space and the (column) dimension of the block containing $H_{\text{sqr}}/\lambda_{\text{sqr}}$ is larger than that of the blocks containing the individual O_α/λ_α ’s. Also, observe that $H_{\text{sqr}}^\dagger H_{\text{sqr}} = H_{\text{SA}}$ as desired and, if $E_k \geq 0$ is an eigenvalue of H_{SA} , then $\sqrt{E_k}$ is a singular value of H_{sqr} . Our definition of H_{sqr} as a rectangular matrix enables a more efficient block encoding Eq. (7) over previous works based on a Hermitian dilation [21,23], which use more qubits and are at least a factor of 2 more expensive to block encode as shown in Appendix A.

Conventionally in spectral amplification the advantage is gained from the nonlinearity of the square root near zero. Here we find that a similar result is obtained from the nonlinear transformation of the eigenvalues in qubitization. For H with eigenvalues $E_1, E_2, \dots$, it is well known that qubitization [16] converts a block encoding $\text{BE}[H/\lambda]$ to a quantum walk with spectrum $e^{\pm i \arccos(E_k/\lambda)}$ [42]. In eigenvalue estimation it is more efficient to estimate the phase of this walk operator than that of Hamiltonian evolution, as has been used in many recent works [7,9,14]. The eigenphases of the quantum walk have positive and negative

values, but this has no effect on the estimated Hamiltonian eigenvalues, and will be ignored in the following.

The situation is more complicated with H_{sqr} because it is not Hermitian, but it is possible to construct two walk operators W_1, W_2 , each of which applies a phase factor of $\arccos(\sqrt{E_j}/\lambda_{\text{sqr}})$ on alternating subspaces (see Appendix A 2). When E_j is close to 0, the amplification is from the nonlinearity of the square root, but it is in the linear region of the arccos. However, the phase estimation needs to be on the product of the two walk operators $W = W_2 W_1$, which gives a net phase of

$$2 \arccos(\sqrt{E_j}/\lambda_{\text{sqr}}) = \arccos(2E_j/\lambda_{\text{sqr}}^2 - 1). \quad (9)$$

Moreover, the product of the two walk operators for H_{sqr} corresponds exactly to the walk operator for $2H_{\text{SA}}/\lambda_{\text{sqr}}^2 - \mathbb{I}$. The formula for the spectrum of walk operators gives the phase factor on the right-hand side of Eq. (9). From this perspective, the amplification for E_j close to 0 is due to the nonlinearity of arccos near -1 .

More specifically, the block encoding is

$$\text{BE}\left[\frac{H_{\text{SA}}}{\Lambda} - \mathbb{I}\right] \equiv U^\dagger \text{REF}_{\text{BE}} U = \text{PREP}^\dagger \cdot \text{SEL}^\dagger \cdot \text{REF}_{\text{BE}} \cdot \text{SEL} \cdot \text{PREP}, \quad \Lambda \equiv \frac{1}{2} \lambda_{\text{sqr}}^2, \quad (10)$$

where REF_{BE} is a reflection over the subspace spanned by the columns of $H_{\text{sqr}}/\lambda_{\text{sqr}}$. This circuit is equivalent to constructing a controlled-block encoding $\sum_{\alpha} |\alpha\rangle\langle\alpha| \otimes \text{BE}[2O_\alpha^\dagger O_\alpha/\lambda_\alpha^2 - 1] = \sum_{\alpha} |\alpha\rangle\langle\alpha| \otimes$

$\text{BE}[T_2(\text{BE}[O_\alpha^\dagger/\lambda_\alpha], \text{BE}[O_\alpha/\lambda_\alpha])]$ using a single step of qubitization as shown in Fig. 1, followed by taking a linear combination block encodings of $T_2(\text{BE}[O_\alpha/\lambda_\alpha], \text{BE}[O_\alpha^\dagger/\lambda_\alpha])$, using preparation Prep and “unpreparation”




FIG. 1. Controlled block encoding of the second Chebyshev polynomial of O_α . The control register contains the control state, $|\alpha\rangle$, of Eq. (8), which provides the appropriate weighting to each T_2 in the LCU. For simplicity, we omit the garbage register.

PREP[†] of the control state PREP|0⟩_a in Eq. (8). This generalizes similar constructions in prior work: Oblivious amplitude amplification [43] when $H_{\text{sqr}}t$ is unitary, or for obtaining a $\frac{1}{2}$ improvement in the block-encoding normalization for Hermitian $H_{\text{sqr}}t$ [9].

We now perform the PEA on W , which provides an estimate of the eigenphase with standard deviation σ_{PEA} , using $\pi/(2\sigma_{\text{PEA}})$ total queries to controlled U and $U^\dagger$ [14]. The Hamiltonian H_{SA} has lowest eigenvalue E_{gap} , and this results in eigenphases $|\phi_{\text{gap}}| = \arccos(E_{\text{gap}}/\Lambda - 1)$ for the walk operator W formed from Eq. (10). We use a first-order error propagation formula to obtain the query cost for estimating E_{gap} with variance $\sigma_{E_{\text{gap}}}^2$, being proportional to

$$\text{PEA queries to } W = \frac{\pi\lambda_{\text{eff}}}{2\sigma_{E_{\text{gap}}}}, \quad \lambda_{\text{eff}} = \sqrt{E_{\text{gap}}(2\Lambda - E_{\text{gap}})}. \quad (11)$$

Thus, the “effective lambda” λ_{eff} for performing PEA on a single step of qubitization depends on the block-encoding scaling factor for H_{SA} and $\sqrt{E_{\text{gap}}}$. There can be some inaccuracy in the error propagation formula near a divergence of the derivative when $E_{\text{gap}} \rightarrow 0$. However, we show in Appendix G that the fractional corrections to $\sigma_{E_{\text{gap}}}$ are negligible for the parameters we consider.

In summary, SA has a smaller effective scaling constant λ_{eff} when considering the uncertainty scaling of the arccos function of the Hamiltonian. The main idea in the next section is to construct an SOS representation H_{SA} that has a low value of E_{gap} . We will see that methods for computing good lower bounds to E_{gs} translate directly to maximizing E_{SOS} , which in turn minimizes E_{gap} .

III. SUM-OF-SQUARES REPRESENTATIONS OF THE CHEMISTRY HAMILTONIAN

Given a Hamiltonian H that is a $2k$ -degree polynomial in a set of operators $\{\mathfrak{o}_j\}$, which do not necessarily commute, the Hamiltonian plus a constant shift can be represented as an SOS:

$$H - E_{\text{SOS}} = \sum_{\alpha=1}^{L-1} O_\alpha^\dagger O_\alpha, \quad (12)$$

where each O_α is itself a polynomial in $\{\mathfrak{o}_j\}$ of degree at least k , and E_{SOS} is a constant. While this is not generally possible for continuous commutative variables (the Motzkin polynomial being one example [44]), operators that are constructed from polynomials of noncommutative variables can be represented as SOS of Hermitian operators [45]. More generally, it is possible to represent the Hamiltonian with O_α involving polynomials of degree $\geq k$ which can lead to E_{SOS} being closer to the lowest

eigenvalue of H . Such a representation of the Hamiltonian constitutes a proof for the lowest eigenvalue of the Hamiltonian being at least E_{SOS} . In fact, the problem of expressing the Hamiltonian in this form is the dual problem of identifying a lower bound to E_{gs} through the pseudomoment method [28,31,32]. The variational two-electron reduced-density-matrix method is one such example relevant to electronic structure, where a hierarchy of operator witnesses constrains the two-particle reduced-density-matrix (2-RDM) to have a fermionic n -electron wave function preimage while directly optimizing with respect to the 2-RDM degrees of freedom. We now describe the efficient protocol for determining an SOS representation of a Hamiltonian through semidefinite programming.

We first note that by fixing a degree k , we define SOS^k to be the subset of operators which can be written as a sum of squares of degree k polynomials in $\{\mathfrak{o}_j\}$. Given the generator set involving operators in SOS^k

$$O_\alpha \in \text{Span} \left(\mathfrak{o}_i, \mathfrak{o}_i \mathfrak{o}_j, \dots, \underbrace{\mathfrak{o}_i \mathfrak{o}_j \dots \mathfrak{o}_n}_{k/2\text{-terms}} \right) \quad (13)$$

we can express the determination of the optimal shifted Hamiltonian as a semidefinite program

$$\begin{aligned} \max \quad & E_{\text{SOS}} \\ \text{s.t.} \quad & H - E_{\text{SOS}}\mathbb{I} = (\vec{\mathfrak{o}}^\dagger)^T G \vec{\mathfrak{o}}, \\ & G \succeq 0, \end{aligned} \quad (14)$$

where $\vec{\mathfrak{o}}$ is the vector of operators in Eq. (13), and G is symmetric matrix of coefficients, known as the Gram matrix, defining the sum-of-squares polynomial. The SOS generators, $\mathfrak{o}_{j,d,\alpha}$, can be an arbitrary algebra as long as it is possible to represent H in that algebra. Considering fermionic Hamiltonians, the natural SOS generating algebra is the set of fermionic ladder operators, and thus a specific example of $\vec{\mathfrak{o}}$ is

$$\vec{\mathfrak{o}} = (\mathbb{I}, a_{i\sigma}, a_{j\tau}^\dagger, a_{i\sigma} a_{j\tau}, a_{i\sigma} a_{j\tau}^\dagger, \dots)^T, \quad (15)$$

where $a_{i\sigma}$ represents an annihilator for fermionic mode i with spin projection σ . The equality constraint on the Gram matrix, G , in Eq. (14) is shorthand for relating coefficients of normal ordered operators. We include the identity operator in this set as it can arise from the anticommutation relations. As a simple example, consider a spinless quadratic Hamiltonian

$$H = \sum_j v_j a_j^\dagger a_j, \quad (16)$$

and the SOS algebra $\{a_i, a_i^\dagger\}$. One can determine an SOS representation of Eq. (16) by solving the program in

Eq. (14) to find the lower bound. Factorizations of the Gram matrix G provide the SOS representation of the Hamiltonian via Eq. (13). In this example, it is also possible to analytically solve for the E_{SOS} and O_α . First, note that operators associated with positive v_j are already in SOS form. For negative v_j the anticommutation relations can be applied to return an SOS operator with a positive coefficient plus an identity term. More explicitly,

$$H = \sum_{j: v_j \geq 0} v_j a_j^\dagger a_j + \sum_{j: v_j < 0} (-v_j) a_j a_j^\dagger + \sum_{j: v_j < 0} v_j = \sum_{\alpha} O_\alpha^\dagger O_\alpha + E_{\text{SOS}}, \quad (17)$$

where the O_α 's are the $a_j^\dagger$'s or a_j 's depending on the sign of v_j , and the shift is $E_{\text{SOS}} = \sum_{j: v_j < 0} v_j$. This proves that for diagonal quadratic Hamiltonians, and using the SOS algebra $\{a_i, a_i^\dagger\}$, a lower bound energy equal to the ground state energy can be determined. The result can be extended to nondiagonal quadratic Hamiltonians involving terms like $a_i^\dagger a_j$, as they can be diagonalized efficiently. For Hamiltonians beyond quadratic many-body order more sophisticated SOS generating algebras are required to represent the Hamiltonian. Returning to the electronic structure Hamiltonian of Eq. (2), we can consider the entire level-2 SOS algebra spanned by generators

$$\{a_{j\sigma}, a_{j\sigma}^\dagger\} \cup \{a_{i\tau} a_{j\sigma}, a_{i\tau}^\dagger a_{j\sigma}^\dagger, a_{i\tau} a_{j\sigma}^\dagger, a_{i\tau}^\dagger a_{j\sigma}\} \quad (18)$$

with an element of the algebra defined by a particular linear combination of the generators

$$O_\alpha = \sum_{j,\sigma} c_{j\sigma}^\alpha a_{j\sigma} + \sum_{j,\sigma} d_{j\sigma}^\alpha a_{j\sigma}^\dagger + \sum_{i,j,\sigma,\tau} b_{i\sigma j\tau}^\alpha a_{i\sigma} a_{j\tau} + \sum_{i,j,\sigma,\tau} g_{i\sigma j\tau}^\alpha a_{i\sigma}^\dagger a_{j\tau}^\dagger + \sum_{i,j,\sigma,\tau} e_{i\sigma j\tau}^\alpha a_{i\sigma}^\dagger a_{j\tau} + \sum_{i,j,\sigma,\tau} f_{i\sigma j\tau}^\alpha a_{i\sigma} a_{j\tau}^\dagger, \quad (19)$$

where the coefficients $\{c, d, b, g, e, f\}$ are complex values and thus

$$H_{\text{SOS}} := \sum_{\alpha} O_\alpha^\dagger O_\alpha = \sum_{i,j,k,l,\sigma,\tau,\kappa,\lambda} D_2 G_{kkll}^{i\sigma j\tau} a_{i\sigma}^\dagger a_{j\tau}^\dagger a_{l\lambda} a_{k\kappa} + Q_2 G_{i\sigma j\tau}^{kkll} a_{kk} a_{l\lambda} a_{j\tau}^\dagger a_{i\sigma}^\dagger + \dots, \quad (20)$$

as is done commonly in the chemistry application of outer approximations—2-RDM constraints that are known as DQG [46]. While optimization over this set of SOS-generators gives a tight E_{gap} mirroring the high-quality

solutions from the variational 2-RDM method [46,47], the block-encoding would require a full spinful representation of H_{SOS} as the different spin sectors of the SOS Hamiltonian would contain different coefficients. Furthermore, the largest DQG-lower bound calculations that have been performed are on the order of 60 orbitals [48] and thus it is potentially challenging to scale the classical preprocessing needed to employ SA. Therefore, we consider the simplified algebra spanned by

$$\left\{ \mathbb{I}, a_{j\sigma}, a_{j\sigma}^\dagger, \sum_{\sigma} a_{i\sigma}^\dagger a_{j\sigma} \right\} \quad (21)$$

with elements defined by

$$O_\alpha := \sum_{j,\sigma} c_{j\sigma}^\alpha a_{j\sigma} + \sum_{j,\sigma} d_{j\sigma}^\alpha a_{j\sigma}^\dagger + \left(e^\alpha \mathbb{I} + \sum_{i,j} g_{ij}^\alpha \sum_{\sigma} a_{i\sigma}^\dagger a_{j\sigma} \right). \quad (22)$$

The number of generators that span this algebra is slightly larger than the minimum SOS algebra to represent the Hamiltonian. The addition of $\mathbb{I}$ makes this algebra equivalent to the spin-free algebra augmented with hole-rotation generators $\sum_{\sigma} a_{i\sigma} a_{j\sigma}^\dagger$.

The equality constraint of the SDP, Eq. (14), mapping operators in the algebra of Eq. (22) to the chemistry Hamiltonian Eq. (2) can be obtained by reordering the fermionic ladder operators according to the anticommutation relations. We provide a complete derivation of the linear constraints in Appendix F. In order to numerically verify that E_{gap} is small for challenging electronic structure simulations we use a combination of variational wave functions, primarily low bond dimension MPS, to upper bound the size of the gap. In Table II we report known or computed variational upper bounds for common benchmarked systems. The lower bounds are computed from the solution of the semidefinite program defined in Eq. (14) using the spin-free algebra defined in Eq. (22) modified to the isospectral algebra containing the hole rotation generators. To solve the semidefinite program we use the cuLoRADS solver [49] which implements a GPU-accelerated low-rank factorization-based solver [50]. In all SDPs we use a rank on the primal solution scaling as $\mathcal{O}[\log(m)]$, where m is the number of linear constraints. This scaling is inspired by the existence of approximate solutions with that order of rank [51]. Rank scaling as the logarithm of the number of constraints is substantially smaller than the rank necessary to certify that the local minimum corresponds to the global minimum [52]. Despite this, we find stable solutions with a tight gap. In Table II we report variational upper bounds, spin-free algebra lower bounds, and energy gaps ($E_{\text{gs}} - E_{\text{SOS}}$) that we use in the next section to inform the search for symmetry-shifted Hamiltonians. In the final column of Table II we report a

TABLE II. Variational energies from above and below computed with different methods. The lower bounding energies are obtained by solving the SDP program Eq. (14) defined over the spin-free algebra described in Eq. (22). The SDP finds a lower bound to the lowest energy in Fock space versus the variational energies are for a fixed particle sector. We note that solving the SDP with the boundary-point method [53], the iteration complexity scales cubically with matrix size and thus is $\mathcal{O}(N^6)$ for an N -spatial orbital system. The gaps in the penultimate column are the difference between the upper and lower bound. All variational energies are for the active space Hamiltonians without a core shift $H_{\text{core}} = 0$. The approximation ($\approx$) indicates the systems where, due to a runtime constraints, we did not converge the SDP for the ≥ 100 orbital CO_2 systems and instead approximated the gap up to the accuracy of the DFTHC ansatz and its suboptimal optimizer as described and validated in Appendix E. The last column reports a lower bound on the SA $\lambda_{\text{eff}} \geq \sqrt{E_{\text{gap}}(\Delta_\mu - E_{\text{gap}})}$ Eq. (11) using the LCU normalization lower bound $\Delta_\mu/2 \leq \Lambda$ computed using the protocol in Ref. [20]. The bound for FeMoco-[113e,76o] was computed using the $S = 3/2$ sector, for CPD1-P450 both the X and G Hamiltonian bounds were computed using $S = 5/2$. For all other systems the bound was computed using the singlet $S = 0$ sector.

System	Method	Variational energy (Ha)			λ_{eff} lower bound (Ha)
		Above	Below	E_{gap}	
Fe_2S_2 [30e, 20o]	DMRG $M = 1500$	-116.605 51	-117.582	0.976	3.632
Fe_4S_4 [54e, 36o]	DMRG $M = 600$	-327.212 57	-329.079	1.866	7.472
FeMoco [54e, 54o]	DMRG $M = 500$	-269.060 20	-272.524	3.464	16.025
FeMoco [113e, 76o]	DMRG $M = 6000$ [10]	-1119.059 33	-1123.300	4.240	21.691
CO_2 [XVIII] [64e, 56o]	DMRG $M = 500$	-295.244 41	-296.941	1.695	11.065
CO_2 [XVIII] [100e, 100o]	DMRG $M = 500$	-632.320 45	-635.134	≈ 2.813	21.478
CO_2 [XVIII] [150e, 150o]	CISD	-1350.651 68	-1355.571	≈ 4.919	37.365
CPD1-P450G [47e, 43o]	DMRG $M = 1500$ [11]	-232.382 95	-235.136	2.753	11.720
CPD1-P450X [63e, 58o]	DMRG $M = 500$	-419.849 03	-424.550	4.700	21.141

lower bound on the SA effective LCU normalization. The lower bound is computed using the reported gaps and the lower bound to the LCU normalization computed by single determinant wave functions for the maximum and minimum energy state in a given particle and spin sector [20].

The lower bounds in Table II are far from accuracies required for thermochemistry, but this is not the most important metric for SA. The parameter scaling walk operator queries is the square root of the ground-state energy of H_{SOS} and thus extremely tight gaps—close to frustration free representations of H —are not necessary to achieve a large speedup. The bounds reported in Table II demonstrate that SOS representations with small enough ground-state energies can be found with polynomial

classical computing effort. Tighter lower bounds could be derived by the addition of penalty parameters or BLISS [18] like operators to the SOS representation.

IV. DOUBLE FACTORIZED-TENSOR-HYPERCONTRACTION SUM OF SQUARES: DFTHC-SOS

We introduce the DFTHC of two-body electronic structure tensor, which is designed to span the spin-free generating set Eq. (22). DFTHC generalizes double factorization [9,54], its variations [39,55], and also the factorized form of tensor hypercontraction [7]. For spin-free Hamiltonians, the DFTHC representation is

$$H_{\text{DFTHC}} = \sum_{pq} h^{(1)'} E_{pq} + \frac{1}{2} \sum_{r \in [R]} \sum_{c \in [C]} \left(W^{(rc)} \mathbb{I} + \sum_{b=0}^{B-1} w_b^{(rc)} \sum_{pq} u_{b,p}^{(r)} u_{b,q}^{(r)} E_{pq} \right)^2 - \frac{1}{2} \sum_{r \in [R]} \sum_{c \in [C]} |W^{(rc)}|^2, \quad (23)$$

$$h_{pqmn}^{(2)'} = \sum_{r \in [R], c \in [C]} \left(\sum_{b \in [B]} w_b^{(rc)} u_{b,p}^{(r)} u_{b,q}^{(r)} \right) \left(\sum_{b' \in [B]} w_{b'}^{(rc)} u_{b',m}^{(r)} u_{b',n}^{(r)} \right), \quad (24)$$

where $h^{(2)'}$ is the approximated two-body tensor, rank, bases, and copies (R, B, C) parameters may be varied freely, and all parameters $h^{(1)'}, W, w, u$ are real with $\|\vec{u}_b^{(r)}\|_2 = 1$. Note the inclusion of the identity component $W_B^{(rc)} \mathbb{I}$ as a free parameter, which is necessary according

to Eq. (22). We also find it convenient to define $w_B^{(rc)} = W^{(rc)} + \sum_{b=0}^{B-1} w_b^{(rc)}$. Thus, double factorization is special case of varying R while fixing $B = N$, $w_B^{(rc)} = 0$, and $C = 1$ [9,54] or $C = 2$ [39] or $C = N$ [55], in addition to imposing a mutual orthogonality

constraint on $\vec{u}_b^{(r)} \cdot \vec{u}_{b'}^{(r)} = \delta_{b,b'}$ for each rank component. Furthermore, $(R, B, C) = (1, R_{\text{THC}}, R_{\text{THC}})$ captures the factorized THC representation of two-body terms $\sum_{c \in [R_{\text{THC}}]} (\sum_{b \in [R_{\text{THC}}]} w_b^{(c)} \sum_{pq} u_{b,p} u_{b,q} E_{pq})^2$.

Let us now express H_{DFTHC} in the SOS form and evaluate its normalization factor Λ . We represent regular fermion operators with Majorana operators

$$a_{p\sigma} = \frac{\gamma_{p\sigma 0} + i\gamma_{p\sigma 1}}{2}, \quad a_{p\sigma}^\dagger = \frac{\gamma_{p\sigma 0} - i\gamma_{p\sigma 1}}{2},$$

$$\{\gamma_{p\sigma x}, \gamma_{q\tau y}\} = 2\delta_{pq}\delta_{\sigma\tau}\delta_{xy}. \quad (25)$$

In the following, we also find it convenient to define the rotated Majorana operator $\gamma_{\vec{u}\sigma x} := \sum_j u_j \gamma_{j\sigma x}$ for any real unit vector $\vec{u}$. This representation is convenient for compiling to quantum circuits as each Majorana operator maps to a Pauli operator. Let the eigendecomposition of $h_{pq}^{(1)'} = \sum_{s \in \{-, +\}} \sum_r s w_s^{(r)} u_{s,p}^{(r)} u_{s,q}^{(r)}$ be split into semipositive and negative parts where all $w_\pm^{(r)} \geq 0$. Then $h_{pq}^{(1)'} E_{pq}$ is diagonalized and, following Eq. (17), we find its optimal E_{SOS} contribution by

$$h_{pq}^{(1)'} E_{pq} = \sum_{\sigma \in \{0,1\}} \left(\sum_j w_+^{(j)} a_{\vec{u}_+^{(j)}\sigma}^\dagger a_{\vec{u}_+^{(j)}\sigma} + \sum_j w_-^{(j)} a_{\vec{u}_-^{(j)}\sigma}^\dagger a_{\vec{u}_-^{(j)}\sigma} \right) - 2 \sum_j w_-^{(j)}. \quad (26)$$

It is possible to verify the SOS representation:

$$H \approx H_{\text{DFTHC}} = \sum_{G \in \{D_1, Q_1\}} \sum_r \sum_{\sigma \in \{0,1\}} O_{G^\sigma, r}^\dagger O_{G^\sigma, r} + \sum_{r \in [R], c \in [C]} O_{\text{SF}, rc}^\dagger O_{\text{SF}, rc} + E_{\text{SOS}}, \quad (27)$$

$$O_{D_1^\sigma, r} = \sqrt{w_+^{(r)}} \frac{\gamma_{\vec{u}_+^{(r)}\sigma 0} + i\gamma_{\vec{u}_+^{(r)}\sigma 1}}{2},$$

$$O_{Q_1^\sigma, r} = \sqrt{w_-^{(r)}} \frac{\gamma_{\vec{u}_-^{(r)}\sigma 0} - i\gamma_{\vec{u}_-^{(r)}\sigma 1}}{2}, \quad (28)$$

$$O_{\text{SF}, rc} = \frac{1}{\sqrt{2}} \left(w_B^{(rc)} + \frac{i}{2} \sum_{b \in [B]} \sum_{\sigma \in \{0,1\}} w_b^{(rc)} \gamma_{\vec{u}_b^{(r)}\sigma 0} \gamma_{\vec{u}_b^{(r)}\sigma 1} \right), \quad (29)$$

$$E_{\text{SOS}} = -2 \sum_r w_-^{(r)} - \frac{1}{2} \sum_{r \in [R]} \sum_{c \in [C]} |W^{(rc)}|^2. \quad (30)$$

As the operators $\gamma_{\vec{u}_s^{(j)}\sigma x}, \gamma_{\vec{u}_s^{(r)}\sigma 0} \gamma_{\vec{u}_b^{(r)}\sigma 1}$ are also unitary, the generators Eqs. (28) and (29) may be block encoded using the LCU technique to obtain

$$\text{BE} \left[\frac{O_{G_1^\sigma, r}}{\lambda_{G_1^\sigma, r}} \right], \quad \lambda_{G_1^\sigma, r} = \sqrt{w_s^{(j)}}, \quad \forall (s, G) \in \{(+, D), (-, Q)\}, \quad (31)$$

$$\text{BE} \left[\frac{O_{\text{SF}, rc}}{\lambda_{\text{SF}, rc}} \right], \quad \lambda_{\text{SF}, rc} = \frac{1}{\sqrt{2}} \sum_{b \in [B+1]} |w_b^{(rc)}|. \quad (32)$$

Following Eqs. (7) and (10), we obtain the block encodings

$$\text{BE} \left[\frac{H_{\text{sqr}}^\dagger H_{\text{sqr}}}{\Lambda} - \mathbb{I} \right], \quad \Lambda = \frac{1}{2} \left(\sum_{\sigma} \lambda_{D_1^\sigma}^2 + \lambda_{Q_1^\sigma}^2 + \sum_{rc} \lambda_{\text{SF}, rc}^2 \right) = \|h^{(1)'}\|_1 + \frac{1}{4} \sum_{r \in [R]} \sum_{c \in [C]} \left| \sum_{b \in [B+1]} |w_b^{(rc)}| \right|^2. \quad (33)$$

At a high level, the quantum circuit that block encodes the DFTHC representation Eq. (33) has some components similar to that of double factorization, with the main difference arising from block encoding the one-body terms in an SOS format. We provide all details of the block-encoding circuit in Fig. 2 with a step-by-step description of the compilation in Appendix B. From the summary of block-encoding costs in Table III, we see that cost is dominated by the following table-lookup oracles [56] and Sel:

- (i) Rot controlled by $N + RB$ indices to outputs $(N - 1) b_{\text{rot}}$ rotation angles, where each of the $N - 1$ is represented by b_{rot} bits, that rotate $\gamma_{0\sigma x}$ to $\gamma_{\vec{u}_b^{(r)}\sigma x}$.
 - (ii) Qroam controlled by $N + R(B + 1)C$ indices to output $b_Q = \lceil \log_2 B \rceil + b_{\text{coeff}}$ coefficients $w_s^{(r)}, w_b^{(rc)}$.
 - (iii) Sel which takes the table-lookup outputs to block-encode the generators $O_{D_1^\sigma, r}, O_{Q_1^\sigma, r}, O_{\text{SF}, rc}$.
- The Toffoli counts of these components then each scale very differently with N, R, B, C .




FIG. 2. The complete quantum circuit for the block encoding. First, “prep” is the state preparation on x_0 . The operation “eq” creates an equal superposition on the b register with succ b flagging success. Then the QROM on x_0 and b is used to output alt and keep data for the coherent alias sampling, as well as G , r (which is independent of b). The operation with “In” on alt and keep, with prep on b indicates the inequality test and controlled swap used in coherent alias sampling. Then b , r are used to output the data for the rotations as well as flipping a flag qubit to indicate if $b \neq B$. For simplicity we indicate that the rotation data for both SF and D_1 , Q_1 is output at the same stage here, though these can be separated to reduce Toffoli complexity (see Appendix B). There are two spin qubits $|\sigma_0\rangle$ and $|\sigma_1\rangle$, with one used for SF (the inner sum) and the other for D_1 and Q_1 (the outer sum). The controlled operation copies it out into $|\sigma\rangle$ depending on G , and that is used to control the swap between the spin up and spin down system registers. The operation “Maj” indicates the circuit for performing the Majorana operator as shown in Fig. 4. The operations are all inverted, except for the Hadamard on $|\sigma_1\rangle$ and prep on x_0 , which are part of the outer sum. There is a reflection on the control qubits used in the inner block encoding, then the entire procedure is performed again. Note that we apply the Hermitian conjugate of this operator using $\text{Maj}^\dagger$.

TABLE III. FeMoco-54 cost breakdown. $N = 54$, $(R, B, C) = (10, 27, 27)$, $b_{\text{rot}} = 15$, $b_{\text{coeff}} = b_{k1} = b_{k2} = 15$. Total qubit costs is 1131 because the 313 temporary qubits for inner state prep QROM can use the 795 persistent qubits for RPrep. We denote this reuse with an asterisks on the temporary ancilla for the inner Prep subroutine.

	Subroutine	Cost	Persistent ancilla	Temporary ancilla	Calls/ Walk
Outer Prep	Uniform State prep	$4\lceil \log(N + RC) \rceil = 36$	$\lceil \log(N + RC) \rceil + 2 = 11$	$\lceil \log(N + RC) \rceil - 2 = 7$	1
	AliasSampling	$\lceil \frac{N+RC}{2^{k_1}} \rceil + 2^{k_1} b_1 = 177$, $b_1 = 24$, $k_1 = 2$	$b_1(\text{QROM} - \text{out}) + b_{k_1}(\text{uniform}) = 39$	$2^{k_1} b_1 + \lceil \log_2(\frac{N+RC}{2^{k_1}}) \rceil = 103$	1
	AliasSampling †	$b_1 + \lceil \frac{N+RC}{2^{k_1}} \rceil + 2^{k_1} b_1 = 61$			1
	Uniform State prep †	$4\lceil \log(N + RC) \rceil = 36$			1
	Total	310	50	110	
Inner Prep	for x_0 do: QROAM(b)	$RC\lceil \frac{B+1}{2^{k_2}} \rceil + 2^{k_2} b_2 = 860$, $b_2 = 20$, $k_2 = 4$	$b_2 + b_{k_2} + \lceil \log(R) \rceil + \lceil \log(C) \rceil + 2 = 46$	$\lceil \log(RC\lceil \frac{B+1}{2^{k_2}} \rceil) \rceil + (2^{k_2} - 1)b_2 = 310$	2
	(for x_0 do: QROAM(b)) †	$b_2 + N + \lceil \frac{RC}{2^{k_2}} \rceil + 2^{k_2}(B + 1) = 254$, $k_4 = 2$			2
	uniform(b)	$4\lceil \log(B + 1) \rceil = 20$	$\lceil \log(B + 1) \rceil + 2 = 7$	$\lceil \log(B + 1) \rceil - 2 = 3$	2
	uniform(b) †	$4\lceil \log(B + 1) \rceil = 20$			2
	Total	1154	53	313*	
RPrep	QROM($\vec{u}$)	$N + RB = 324$	$(N - 1)b_{\text{rot}} = 795$	$\max(\lceil \log(N + RC) \rceil = 9, \lceil \log(RB) \rceil = 9)$	2
	QROM($\vec{u}$) †	$R + B = 37$			2
	Total	361	795	9	
Sel	ROT	$4b_{\text{rot}}(N - 1) = 3180$	$b_{\text{rot}} = 15$	0	2
	CSWAP	$2N = 108$	0	0	2
	Maj-control	7	0	0	2
	Total	3295	15	0	
REF	T_2	$\lceil \log(B + 1) \rceil + b_{k_2} + 1 = 21$	0	0	1
	Walk	$\lceil \log(N + RC) \rceil + \lceil \log(B + 1) \rceil + b_{k_1} + b_{k_2} + 2 = 46$	0	0	1
	Total	9997	913 + 108 (system) = 1021	max(110, 9)	

Hence, DFTHC is beneficial following the intuition that the expressive power of the factorization can be greatly improved by making the costs of Rot and Qroam comparable. By taking advantage of spacetime trade-offs [56] and measurement-based uncomputation [17] to reduce the Toffoli gate cost of Qroam using the idle qubits that later occupied by the output of Rot, the Toffoli counts of these components are

- (i) Computing and uncomputing Rot twice: Each computation costs $\approx N + RB$.
- (ii) Computing and uncomputing Qroam twice: Each computation costs $\approx [\text{RBC}(\lambda + 1)] + \lambda b_Q$ for any choice of $\lambda \leq \{(N - 1)b_{\text{Rot}}\}/b_Q$.
- (iii) Applying Sel twice: Each application costs $\approx 4(N - 1)b_{\text{Rot}}$.

For instance, the cost Rot strongly dominates in DF, whereas Qroam strongly dominates in THC.

DFTHC achieves then chemical accuracy with a most compact representation for molecules such as FeMoco [113e,76o] when (R, B, C) is roughly in-between the equivalent DF and THC parameters as shown in Table X. The selected DFTHC values lead to lookup-table sizes that are already substantially smaller than previously reported DF or THC methods, and an exhaustive search for optimal (R, B, C) would lead to even further improvement. For the typical DFTHC solutions that we found, the three contributions Rot, Qroam, Sel capture $\sim 90\%$ of the Toffoli cost, with Sel being the dominant cost accounting for $\sim 60\%$. For example, Table III gives a full cost breakdown.

We use the following heuristic for choosing (R, B, C) .

- (i) Amdahl's law: $C = \Theta(N)$ so that the Toffoli counts of Rot and Qroam are roughly similarly with $\Theta(RB)$ scaling.
- (ii) Linear cost scaling: THC is notable for exhibiting an almost linear $\tilde{\Theta}(N)$ block-encoding Toffoli with $R_{\text{THC}} = \tilde{\Theta}(N)$. Hence we choose $B = \Theta(N)$ and $R = \tilde{\Theta}(1)$.
- (iii) Extensivity: Suppose a given N_0 orbital system is well approximated by some $(R, B, C) = [R_0, \Theta(N_0), \Theta(N_0)]$. Then given K noninteracting copies of the same system, one expects a good approximation with $(R, B, C) = [R_0 K, \Theta(N_0), \Theta(N_0)]$. In other words, $B, C \lesssim N$ suffices.

We also carefully study the bits of precision needed in Appendix D. Roughly speaking, chemical precision is attained with $b_{\text{rot}} \in [11, 18]$, $b_{\text{coeff}} \in [6, 14]$, depending on the system. This is consistent with previous work [7,57] and all resource estimates use bit counts reported in Appendix D.

A. Numerical optimization

DFTHC allows us to simultaneously find electronic structure SOS with good gaps, block-encoding costs, and normalization factors, without directly solving the SDP in Eq. (14). We find good DFTHC approximations

$h^{(2)'}_{pq}$ of the exact two-body integrals $h^{(2)}$ using a cost function that minimizes the Frobenius error $\epsilon_{\text{fro}} = \|h^{(2)} - h^{(2)'}\|_{\text{fro}}$ together with regularization on Λ and E_{gap} .

It is only necessary to approximate the two-body terms in DFTHC as we have freedom to choose the one-body terms. By expanding Eq. (23) and then comparing to the original exact Hamiltonian Eq. (2),

$$H_{\text{DFTHC}} = \sum_{pq} \left(h^{(1)'}_{pq} + \sum_{rc} W^{(rc)} \sum_{b \in [B]} w_b^{(rc)} u_{b,p}^{(r)} u_{b,q}^{(r)} \right) E_{pq} + \frac{1}{2} \sum_{rc} \left(\sum_{b \in [B]} w_b^{(rc)} \sum_{pq} u_{b,p}^{(r)} u_{b,q}^{(r)} E_{pq} \right)^2, \quad (34)$$

$$H = H_{\text{DFTHC}} + \frac{1}{2} \sum_{pqrs} (h^{(2)}_{pqrs} - h^{(2)'}_{pqrs}) E_{pq} E_{rs}, \quad (35)$$

we see for any computed w, u tensors, the choice $h^{(1)'}_{pq} = h^{(1)}_{pq} - \sum_{rc} W^{(rc)} \sum_{b \in [B]} w_b^{(rc)} u_{b,p}^{(r)} u_{b,q}^{(r)}$ leads to a one-body contribution equal to the exact Hamiltonian, with some error in the two-body DFTHC tensor $h^{(2)'}$.

We reduce the contribution of error terms by using the identity $W^{(rc)} = w_B^{(rc)} - \sum_{b \in [B]} w_b^{(rc)}$ to rewrite

$$h^{(1)'}_{pq} = h^{(1)}_{pq} + \sum_r h^{(2)'}_{pqrr} - \sum_{rc} w_B^{(rc)} \sum_{b \in [B]} w_b^{(rc)} u_{b,p}^{(r)} u_{b,q}^{(r)}, \quad (36)$$

$$\begin{aligned} 2 \sum_j w_j^{(j)} &= \|h^{(1)'}\|_1 - \text{Tr}[h^{(1)'}] \\ &= \|h^{(1)'}\|_1 - \text{Tr}[h^{(1)}] - \sum_{pq} h^{(2)'}_{ppqq} \\ &\quad + \sum_{rc} w_B^{(rc)} \sum_{b \in [B]} w_b^{(rc)}, \end{aligned} \quad (37)$$

$$\begin{aligned} \frac{1}{2} \sum_{rc} |W^{(rc)}|^2 &= \frac{1}{2} \sum_{pq} h^{(2)'}_{ppqq} + \frac{1}{2} \sum_{rc} |w_B^{(rc)}|^2 \\ &\quad - \sum_{rc} w_B^{(rc)} \sum_{b \in [B]} w_b^{(rc)}, \end{aligned} \quad (38)$$

where $\|\cdot\|_1$ is the Schatten one-norm. Hence,

$$\begin{aligned} E_{\text{gap}} &= E_{\text{gs}} + \|h^{(1)'}\|_1 - \text{Tr}[h^{(1)}] - \frac{1}{2} \sum_{pq} h^{(2)'}_{ppqq} \\ &\quad + \frac{1}{2} \sum_{rc} |w_B^{(rc)}|^2. \end{aligned} \quad (39)$$

TABLE IV. Runtime for evaluating a single gradient descent Adam step of Eq. (41) for typical examples in a Jax implementation running on a single Nvidia L4 GPU, and typical hyperparameters. In all cases, the learning rate is geometrically decreasing from 10^{-1} to a final value that is typically $10^{-4, -5, -6}$.

Molecule	(N, R, B, C)	λ_{reg} (Ha)	ϵ_{reg} (mHa)	E_{reg} (Ha)	Time per step (ms)	Steps to solution
Fe ₂ S ₂ [30e, 20o]	(20, 14, 15, 5)	20	10	[1, 4]	0.8	$10^{[5,6]}$
Fe ₄ S ₄ [54e, 36o]	(36, 9, 18, 18)	20	10	$1.87 \times [1, 2]$	1.2	$10^{[5,6]}$
FeMoco [54e, 54o]	(54, 10, 27, 27)	50	10	$3.46 \times [1, 1.6]$	1.3	$10^{[5,6]}$
FeMoco [113e, 76o]	(76, 15, 57, 19)	100	10	[4, 8]	4.4	$10^{[5,6]}$
CO ₂ [XVIII] [64e, 56o]	(56, 5, 28, 28)	50	10	[1, 4]	1.8	$10^{[5,6]}$
CO ₂ [XVIII] [100e, 100o]	(100, 9, 75, 25)	50	10	[2, 8]	8.2	$10^{[5,6]}$
CO ₂ [XVIII] [150e, 150o]	(150, 13, 112, 37)	200	10	[4, 16]	49.0	$10^{[5,6]}$
CPD1-P450X [63e, 58o]	(58, 9, 29, 14)	50	10	$4.70 \times [1, 1.5]$	1.3	$10^{[5,6]}$

Then using prior knowledge of the limit where DFTHC is exact, we make the replacements $\sum_r h_{pqrr}^{(2)'} \leftarrow \sum_r h_{pqrr}^{(2)}$. A similar replacement was made previous work [17] when going from the normal-ordered representation to Eq. (2) or to the Majorana representation [9].

We expand the space of good solutions using symmetry shifting (BLISS) [18,19], which is based on the observation that the spectrum of the symmetry sector with fixed particle

number η is invariant under any operator that multiplies the annihilator $(\hat{N} - \eta)$. In this work, we consider first and second order annihilators constructed from particle number conservation, which correspond to a continuous set of spectrally equivalent Hamiltonian where original electronic structure tensor $(h^{(2)}, h^{(1)}, h^{(0)})$ may be replaced by any $(h_S^{(2)}, h_S^{(1)}, h_S^{(0)})$ in

$$\mathcal{H}_\eta = \left\{ \left(h^{(2)} + \frac{h^{(\text{sym})} \otimes \mathbb{I}_N + \mathbb{I}_N \otimes h^{(\text{sym})}}{2}, h^{(1)} - \frac{\eta}{2} h^{(\text{sym})} + \beta_1 \mathbb{I}_N, h^{(0)} - \frac{\eta}{2} \beta_1 \right) \right. \\ \left. \left| \forall \beta_1 \in \mathbb{R}, h^{(\text{sym})} = h^{(\text{sym})\dagger} \in \mathbb{R}^{N \times N} \right\}, \quad (40)$$

where the tensor product of any two matrices $A \otimes B$ has entry $(A \otimes B)_{pqrs} = A_{pq} B_{rs}$ and $\mathbb{I}_N$ is the $N \times N$ identity matrix.

We optimize the DFTHC ansatz for good block-encoding parameters by defining the nonlinear program

$$\underset{\substack{u \in \mathbb{R}^{R \times B \times N}, \\ w \in \mathbb{R}^{R \times (B+1) \times C}, \\ (h_S^{(2)}, h_S^{(1)}, h_S^{(0)}) \in \mathcal{H}_\eta}}{\text{argmin}} \left[\frac{\|h_S^{(2)} - h^{(2)'}\|_{\text{fro}}}{\epsilon_{\text{reg}}} + \frac{\Lambda}{\lambda_{\text{reg}}} + \text{Relu} \left(\frac{E_{\text{gap}} - E_{\text{reg}}}{E_{\text{reg}}} \right) \right], \quad \text{Relu}(x) := \begin{cases} x, & x \geq 0, \\ 0, & x < 0. \end{cases} \quad (41)$$

where $\cdot_{\text{reg}}$ are regularization hyperparameters. Equation (41) minimizes the error between the DFTHC approximation $h^{(2)'}$ which is a function of u and w , and any exact two-body tensor $h_S^{(2)}$ in the η -particle symmetry sector. At the same time, it places some priority on minimizing the block-encoding normalization factor Λ , which is implicitly a function of $u, w, h_S^{(2)}, h_S^{(1)}$ as seen in Eq. (36). This program also attempts to get solutions close to the optimal gap through E_{gap} , which implicitly depends on all parameters. Note the Relu function, which avoids optimizing for unrealistic gaps less than the optimal. Whereas prior approaches to solving similar nonlinear programs for DF or THC required chemically motivated initial conditions (u_0, w_0) for (u, w) from which gradient descent begins, we

find that it suffices to make the much simpler choice of initial parameter values $\mathcal{P} = (u_0, w_0, \beta_{1,0}, h_0^{(\text{sym})})$, where $u_b^{(r)}$ is a random unit vector, $w_0, h_0^{(\text{sym})} \sim \mathcal{N}(0.01)$ and $\beta_1 \sim \mathcal{N}(0.05)$ to be element-wise sampled from a distribution where $\mathcal{N}(\sigma)$ is the normal distribution with mean zero and standard deviation σ .

We approximately solve Eq. (41) by first-order gradient descent using the Adam optimizer implemented in the Optax and Jax libraries for the Python programming language. The most expensive operation in each step is the tensor contraction in Eq. (24), which costs $\mathcal{O}(RCN^4)$ multiplications. A major benefit is this approach is autodifferentiation, which allows for easily computing gradients even through complicated functions, such as the Schatten norm




FIG. 3. FeMoco-54 CCSD(T) correlation energy error vs rank of DFTHC truncation (small dots). The dashed line at 1.6 mHa is the chemical accuracy threshold and the shaded region is the threshold 0.3 mHa used to select the rank R^* (large dots) of good solutions.

in Eq. (39). Our approach is also GPU accelerated, which allows us to collect robust statistics and optimize over a large number of (R, B, C) parameters for systems up to $N = 150$ orbitals, as shown in Table IV together with typical hyperparameters, with a discussion of hyperparameter selection in Appendix C.

B. Selection of good solutions

Although finding the optimal solution to Eq. (41) could take exponential time in N, R, B, C , we emphasize that it suffices to find approximate solutions more rapidly, so long as they are chemically accurate. Chemical accuracy for each DFTHC solution is validated by calculating the difference in CCSD(T) correlation energy ϵ_{corr} compared to the original representation H , as commonly recommended [7,17], as a proxy to the difference in exact correlation energy ϵ_{corr}^* . Prior work [12,39] computed good DF or THC solutions by scanning over ranks, searching for good hyperparameters, and finally picking the lowest quantum cost solution with $|\epsilon_{\text{corr}}|$ below some threshold ϵ_{th} . In other words, using $\cdot'$ to refer to variables computed from a DFTHC solution,

$$\begin{aligned} E_{\text{gs}} &= E_{\text{HF}} + E_{\text{corr}}, & E'_{\text{gs}} &= E'_{\text{HF}} + E'_{\text{corr}}, \\ \epsilon_{\text{corr}} &= E_{\text{corr}} - E'_{\text{corr}}, \end{aligned} \quad (42)$$

where E_{HF} is the Hartree-Fock energy, which can be computed in $\mathcal{O}(N^3)$ time. Hence, using a quantum computer to estimate $\hat{E}'_{\text{gs}}$ with standard deviation σ_{PEA} from the PEA, we may obtain the true ground-state energy $\hat{E}_{\text{gs}} = \hat{E}'_{\text{gs}} + (E_{\text{HF}} - E'_{\text{HF}}) + \epsilon_{\text{corr}}^*$ with standard deviation


$$\sigma_E \leq \sigma_{\text{PEA}} + |\epsilon_{\text{corr}}^*| \approx \sigma_{\text{PEA}} + |\epsilon_{\text{corr}}| \leq \sigma_{\text{PEA}} + \epsilon_{\text{th}}. \quad (43)$$

We use the same approach to generate candidate DFTHC solutions below the threshold $\epsilon_{\text{th}} \leq 0.3$ mHa as illustrated for FeMoco-54 in Fig. 3. Following the heuristics in Sec. IV, we choose (R, B, C) values with $B \in \{N, \lfloor 3N/4 \rfloor, \lfloor N/4 \rfloor\}$ and $C \in \{\lfloor N/2 \rfloor, \lfloor N/4 \rfloor\}$, and scan over R . To make a comparison with prior art, we also study DF-like $(R, N, 1)$ and THC-like $(1, R_{\text{THC}}, R_{\text{THC}})$ factorizations. The best Toffoli gate costs for points below the ϵ_{th} threshold is reported in Table V together with relevant parameter like E_{gap} —see Tables X and XI for best solutions for each choice of (B, C) . As the absolute energy of each DFTHC factorization may shift by a large amount $E'_{\text{HF}} - E_{\text{HF}}$, we also recompute E_{gap} using the CCSD(T) energy to approximate the ground-state energy. CCSD(T) is not an upper bound on E'_{gs} , but we expect the correction to be small compared E_{gap} , which is on the order of Hartrees. CCSD(T) in the spin-0 sector is run on the systems FeMoco-54 and the XVIII complexes. For other molecules, convergence requires targeting the high-spin sector. We choose $n_{\uparrow} - n_{\downarrow}$ for Fe_2S_2 , Fe_4S_4 , FeMoco-76, and CPD1-P450X to be 8, 16, 35, and 5, respectively, which are choices identical to prior art [7,10], and made for the purposes of comparison.

Drawing general conclusions about the best hyperparameters is not straightforward, as ϵ_{corr} fluctuates significantly in Fig. 3 and does not improve monotonically with rank. Similar to DF and THC, these fluctuations arise as the DFTHC program of Eq. (41) is nonlinear and the final solution is sensitive to initial conditions. This is easily resolved by collecting statistics on ϵ_{corr} , by restarting the DFTHC program for multiple different initial conditions $\mathcal{P}$, in a randomized fashion as described in Sec. IV A. This allows us to make statistically robust conclusions in Appendix C about the scaling of various parameters of interest, such as the standard deviation σ_{corr} , Λ , and E_{gap} .

TABLE V. (top) Improvements in simulating electronic structure to chemical accuracy with DFTHC and spectral amplification compared to best prior double-factorization methods and tensor hypercontraction methods. Improvement factors are computed as the ratio of the product the LCU 1-norm λ , and the block-encoding cost $\text{Cost}[\text{BE}]$. The source for each cost is provided as a citation reference in the λ row. The number of calls to the block encoding is computed as $\lceil \frac{\pi\lambda}{2\sigma_{\text{PEA}}} \rceil$ [14]. In line with previous work we take $\sigma_{\text{PEA}} = 1.0$ mHa. (bottom) Scaling of parameters Toffoli cost (green), block-encoding normalization Λ , spectrum amplified normalization λ_{eff} , and SOS gap E_{gap} .

	Molecule	FeMoco				CO ₂ [XVIII]			
		Fe ₂ S ₂ -20 [30e, 20o]	Fe ₄ S ₄ -36 [54e, 36o]	[54e, 54o]	[113e, 76o]	CPD1-P450X [63e, 58o]	[64e, 56o]	[100e, 100o]	[150e, 150o]
DF	$\lambda_{\text{DF}}/\text{Ha}$	41.13 [10]	154.83 [10]	78.00 [39]	584.54 [10]	111.3 [39]	293.50 [9]	782.1 [9]	1923.8 [9]
	$C_{\text{B.E.}}$	7177	16 505	19 589	40 442	21 736	54 227	97 679	182 005
	Qubits	989	3115	3722	6404	2596	3700	18 017	34 218
THC	$\lambda_{\text{THC}}/\text{Ha}$	63.81 [10]	168.34 [10]	306.30 [7]	198.90 [8]	130.9 [8]	...	...	...
	$C_{\text{B.E.}}$	4626	8573	11 016	13 763	8334	...	...	...
	Qubits	579	1081	2142	1512	999	...	...	...
DFTHC + BLISS + SA	(R, B, C)	(14, 15, 5)	(9, 18, 18)	(10, 27, 27)	(15, 57, 19)	(9, 29, 14)	(5, 28, 28)	(8, 75, 25)	(9, 112, 37)
	$\epsilon_{\text{corr}}/\text{mHa}$	0.0131	0.0041	0.1254	0.1362	0.2766	0.2697	0.042	0.062
	$\epsilon_{\text{fno}}/\text{Ha}$	0.0221	0.1669	0.2383	0.0825	0.3083	0.4182	0.2187	0.3510
	Λ/Ha	17.5299	49.8149	58.3440	179.7296	97.4395	55.4773	155.5	336.1
	E_{gap}/Ha	1.2381	2.3070	4.0535	5.3820	5.6837	2.6918	4.565	6.454
	$\lambda_{\text{eff}}/\text{Ha}$	6.4690	14.9842	21.3674	43.6538	32.7923	17.0712	37.68	65.87
	$\text{Cost}[\text{BE}]$	3924	7356	10203	14629	9535	7685	18041	27367
	Qubits	463	868	1132	1454	1150	919	1955	2864
	Total Toffoli	3.99×10^7	1.73×10^8	3.42×10^8	1.00×10^9	4.91×10^8	2.06×10^8	1.07×10^9	2.83×10^9
	Improvement (DF)	11.6	23.2	7.0	36.9	7.7	121.3	112.4	194.2
	Improvement (THC)	11.6	13.1	15.5	4.3	3.5	...	...	...



with DFTHC rank R , and enables the use of a different bound on error contributions like

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{corr}}^2} + |\bar{\epsilon}_{\text{corr}}| \leq \epsilon_{\text{chem}}. \quad (44)$$

This bound also tolerates larger errors, say $\sigma_{\text{PEA}} = 1.4$ mHa and $\sigma_{\text{corr}} = 0.7$ mHa. In Table XI of Appendix C, we present an alternate set of resource estimates based on Eq. (44), as they represent average-case solutions rather than best-case solutions and are of independent interest.

V. RESOURCE ESTIMATES

In Table V, we report the quantum resource estimates of the DFTHC + BLISS + SA protocol for estimating the

ground-state energy, and compare these with prior art using the double-factorization, tensor hypercontraction, and BLISS where applicable. We consider a variety of chemical systems: Iron-sulfur complexes (Fe₂S₂ [40], Fe₄S₄ [40], 54 orbital FeMoco active space [37], 76 orbital FeMoco active space [38]), a representation of the cytochrome P450 enzyme active site (CPD1) [11], and the CO₂-fixation Ruthenium catalyst [9]. Beyond being chemically distinct and a robust test of the method, the varying size and degree of correlation in each of the systems allows us to extrapolate scalings in Table V, based on the eight system considered spanning an order of magnitude difference in number of orbitals and electrons.

The DFTHC resources were computed by scanning over (R, B, C) values to find lowest cost solutions balancing gap, rotation, and coefficient oracle costs along with the

difference in correlation energy with respect to the untruncated integrals ϵ_{corr} . We report the LCU rescaling value for the SOS Hamiltonian Λ , differences between DFTHC two-electron integral tensors, the effective DFTHC + SA LCU rescaling value λ_{eff} , block-encoding space and time resources, and improvement factors over prior work. The improvement factors are computed as the ratio of the products of block-encoding and λ . The bits for LCU coefficients and rotations b_{coeff} and b_{rot} , respectively, were estimated according to Appendix D, where we design a randomized rounding scheme that converts bit truncation errors to a variable ϵ_{trunc} contribution to ϵ_{corr} . We choose the maximum number of bits that satisfies the error bound

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{trunc}}^2} + |\epsilon_{\text{corr}}| + |\bar{\epsilon}_{\text{trunc}}| \leq \epsilon_{\text{chem}}. \quad (45)$$

We obtain a Λ that is smaller than those computed by prior techniques, which combined along with a small E_{gap} , indicates that the DFTHC factorization of spin-free algebra is sufficient for finding spectrum amplifiable representations where speedups in phase estimation are possible.

Our results are at least a factor of 3 times cheaper than the best costs reported in the literature for this set of systems. For the FeMoco-76 orbital system from Li *et al.* [38] we find a factor of 4 improvement over recent work integrating tensor hypercontraction and BLISS [8] resulting in a cost of 9.99×10^8 Toffoli gates for ground-state energy estimation. For smaller Iron-Sulfur systems (Fe_2S_2 , Fe_4S_4 , FeMoco-54) and the CO_2 systems where symmetry shifting is used to a lesser extent in prior art, DFTHC + BLISS + SA constitutes a substantial improvement by a factor of 10 to 195. The SA effective LCU normalization, λ_{eff} determined by DFTHC are within a factor of 2 of the lower bounds to the optimal values provided in Table II. This does not indicate a limit of the SA technique—only that for a

given SOS algebra we have found a tight gap and a highly compressed form of the two-electron integral tensor.

Utilizing the same quantum error correction model as Ref. [7] we can estimate the physical resources required for ground-state energy estimation using DFTHC + BLISS + SA. Using four CCZ factories, a physical error rate of 0.001, a surface code cycle time of $1 \mu\text{s}$ and reaction time of $10 \mu\text{s}$ we estimate that the FeMoco-76 example will require 4.5×10^6 physical qubits running for 8.6 h on a set of $d = 27$ surface code patches. FeMoco-54 is estimated to require 3.2×10^6 physical qubits running for 2.7 h on a set of $d = 27$ surface code patches. For reference, prior THC based runtime estimates from Ref. [7], not considering the additional space-time trade-offs made in that work, result in a FeMoco-54 runtime of 2 days and FeMoco-76 runtime of 12 days [7].

The DFTHC protocol requires selecting a large enough (R, B, C) such that a faithful representation of the original Hamiltonian is preserved while optimizing for the desired phase estimation cost. Prior works have used the difference in correlation energy, ϵ_{corr} , with respect to untruncated representation of the two-electron integrals in a high-spin electronic configuration. We follow this trend and report factorized solutions with ϵ_{corr} of ≤ 0.3 mHa in a similar fashion to Ref. [8]. Recovering correlations from a quantum computer is a viable way to compute high accuracy energy differences and thus correlation energy difference is a reasonable target when using an approximate wave function with a faithful representation of the correlation energy and multireference nature of the target eigenstate. High-spin electronic configuration are selected such that CCSD(T) converges without issue. In this regime the wave function is primarily single reference and thus correlation energy based metrics can differ substantially from low-spin, potentially, multireference targets for phase estimation. In Table VI we

TABLE VI. Various energies and energy difference between a DFTHC representation of the two-electron integrals and the true two-electron integrals. All simulations used the exact one-electron integrals. The DMRG calculations for Fe_2S_2 used a bond dimension of $M = 2000$ while Fe_4S_4 used $M = 1000$. The key difference in correlation energies is shown in columns labeled $\epsilon_{\text{CCSD(T)}}(\text{corr})$ and $\epsilon_{\text{DMRG}}(\text{corr})$. The variance and orders-of-magnitude difference suggest high spin CCSD(T) correlation energy may not be an optimal metric when selecting tensor factorization thresholds even with a randomization protocol proposed in Sec. IV B. The last two columns show absolute energy differences of substantial magnitude suggesting that low threshold truncations may be representing a Hamiltonian with fundamentally different physics.

(R, B, C)	In units of (Ha)			In units of (mHa)				
	E_{scf}	$E_{\text{CCSD(T)}}(\text{corr})$	$E_{\text{DMRG}}(\text{corr})$	ϵ_{scf}	$\epsilon_{\text{CCSD(T)}}(\text{corr})$	$\epsilon_{\text{DMRG}}(\text{corr})$	$\epsilon_{\text{CCSD(T)}}$	ϵ_{DMRG}
Fe_2S_2 (12, 20, 10)	−116.381 358	−0.159 411	−0.217 612	6.516 516	0.077 150	0.028 436	6.593 666	6.544 952
	(14, 10, 10)	−116.387 588	−0.159 194	−0.219 803	0.285 888	0.294 471	2.162 635	0.580 359
	(14, 15, 5)	−116.344 198	−0.159 501	−0.218 300	43.675 643	0.013 070	0.659 010	43.662 573
	(11, 15, 10)	−116.372 589	−0.159 387	−0.215 839	15.284 690	0.101 229	1.801 982	15.385 919
Fe_4S_4 (16, 27, 9)	−326.903 132	−0.178 192	−0.219 675	101.275 269	0.405 329	2.645 671	101.680 599	103.920 941
	(9, 18, 18)	−326.772 089	−0.178 601	−0.213 938	232.317 519	0.004 119	8.383 212	232.313 399
	(12, 18, 18)	−326.865 711	−0.173 126	−0.210 854	138.695 868	0.242 423	11.467 001	138.938 291
	(12, 27, 18)	−326.944 912	−0.178 400	−0.218 820	59.494 843	0.197 676	3.500 753	59.692 519

demonstrate a substantial difference in correlation energy differences computed from high-spin CCSD(T) energies and low-spin DMRG wave functions. This suggests that higher truncation thresholds may be needed for high accuracy electronic structure calculations on quantum computers. In order to assess this, a numerical technique benchmarking the quality of the integral truncation needs to be developed. Clearly, a facile CCSD(T) calculation at high-spin insufficiently trades off model complexity of correlation energy with computational complexity. We raise this question for further research.

VI. DISCUSSION

In this work we have reduced the cost of ground-state energy estimation of second quantized electronic structure Hamiltonians by utilizing ideas from SA. For SA to provide a net gain, an SOS Hamiltonian with a sufficiently small gap and a sufficiently simple H_{sqt} is required. Our work demonstrates that an SOS algebra that is similar to the algebra needed to express the spin-free chemistry Hamiltonian is sufficient for tight energy gaps and low block-encoding costs yielding reduced quantum resources. To ensure the block encoding of H_{sqt} is low cost we developed a new tensor factorization that provides an optimal trade-off between the best prior art (DF and THC factorizations). The new factorization efficiently compresses the two-electron integral tensor, results in a tight SOS gap $\Delta_{\text{gap}} \ll 1$, and balances the trade-off between the two dominant table lookup oracles. The efficiency of this new integral compression scheme is demonstrated by the fact that the effective LCU normalization is within a factor of 2 of the optimal value for all systems studied. These results establish that the major obstacles of combining SA and ground-state energy estimation can be surmounted.

Application of DFTHC and SA to FeMoco continues the trend of cost reductions. In this case, a factor of 4 over prior art employing THC combined with block-invariant symmetry shifting [8]. With this reduction in costs the phase estimation component (9.99×10^8 Toffoli gates) of ground-state energy estimation is now similar cost as state preparation (between 7.6×10^8 and 1.3×10^9) for MPS with estimated ground-state overlaps between 0.95 and 0.99 [10].

Even more refinements to DFTHC and SA are possible. For DFTHC, improved optimal parameter search that minimizes space-time instead of simply the Toffoli complexity can lead to improved circuits for execution on real systems. Furthermore, it may be possible to convert the large nonlinear optimization to a low-rank semidefinite program which also includes various hyperparameter optimizations [58]. Such a conversion would provide a provable polynomial-time preprocessing that optimizes

quantum algorithm parameters and not just E_{gap} . SA can be further optimized by exploring minimal expansions to the SOS algebra to decrease the ground-state energy of H_{SA} . For example, it is well known from the variational 2-RDM literature that the full level 2-algebra can achieve lower bounds to the ground state energy as tight as $1 \times 10^{-2} - 1 \times 10^{-3}$ Hartree [59]. It is interesting to explore how pieces of the full level 2-algebra (or even higher-order algebra terms like in Ref. [60]) to improve E_{gap} without drastically increasing the block encoding cost. In terms of physical resources, the cost models of Ref. [7] are now out of date with the invention of magic state cultivation in Ref. [61]. Considering the small space-time volume of magic state cultivation, it is possible that the predicted 8 h of runtime is a substantial overestimate. At the logical level, with the reduction in space-time costs for T gates a reanalysis of space-time trade-offs at the QROM level, which will then completely dominate the spacetime volume of Sel, will be required.

Applying SA beyond second quantized chemistry will require detailed examination of the trade-off between the growing SOS complexity needed to express a Hamiltonian close to a frustration free form and the commensurate increased block-encoding cost. Possible avenues forward are improving the preprocessing cost of using the full spinful level-2 algebra or, based on recent work on perturbative analysis of SOS representations [60,62,63], including a subset of level-3 algebra components. Application of SA to lattice models will require a careful trade-off between SOS algebra size and block-encoding cost due to the existing methods already providing linear scaling in system size [14]. We expect SA to be integral to further reducing ground-state energy estimation costs. We further note that it may be possible to apply SOS representations to classical quantum simulation algorithms. For example, in Ref. [63] it is argued that optimized SOS representations provide sign problem mitigation within free-projection auxiliary field quantum Monte Carlo. Application of SOS forms to perturbation theory also advocates that preprocessing a Hamiltonian into an SOS form can provide utility beyond quantum algorithms [60].

ACKNOWLEDGMENTS

D. W. B. worked on this project under a sponsored research agreement with Google Quantum AI. D. W. B. is also supported by Australian Research Council Discovery Projects No. DP210101367 and No. DP220101602.

DATA AVAILABILITY

The data that support the findings of this article are openly available [65], embargo periods may apply.

APPENDIX A: EQUIVALENT QUANTUM WALK REPRESENTATIONS OF SPECTRAL AMPLIFIABLE HAMILTONIANS

In this section, we present alternate methods to generate quantum walks specialized to “spectrum-amplifiable” Hamiltonians H_{SA} as in Eq. (5), given block encodings of its constituent parts $\text{BE}[O_\alpha/\lambda_\alpha]$. Section II presents a specific block encoding

$$\text{BE}\left[\frac{2H}{\sum_\alpha \lambda_\alpha^2} - \mathbb{I}\right], \quad (\text{A1})$$

which can be converted into a quantum walk with eigenphases $\pm \arccos\{[2H/(\sum_\alpha \lambda_\alpha^2)] - 1\}$ though qubitization [16]. The quadratic amplification of gradient for phases near π then provides a quadratic advantage in phase estimation for estimating the ground-state energy. One may then interpret the quantum walk as the source of advantage stemming from the nonlinear dependence of phase on the eigenvalue of the original Hamiltonian. Similarly, it is possible to consider the source of advantage by defining an appropriate square root of Hamiltonian.

1. Quantum walk on Hermitian square-rooted Hamiltonians

One possible definition for the square-root of Eq. (5) is the Hermitian operator

$$H_{\text{sqrt}} = \sum_{\alpha=0}^{L-1} (|\chi_\alpha\rangle\langle 0| \otimes O_\alpha + |0\rangle\langle \chi_\alpha| \otimes O_\alpha^\dagger). \quad (\text{A2})$$

Note that we are free to replace $|\chi_\alpha\rangle$ with any set of mutually orthogonal quantum states that are also orthogonal to the multiqubit computational basis state $|0\rangle$. This squares to

$$H_{\text{sqrt}}^2 = \begin{pmatrix} H_{\text{SA}} & \mathbf{0} \\ \mathbf{0} & \ddots \end{pmatrix}, \quad (\text{A3})$$

where the upper diagonal block is the original gap-amplifiable Hamiltonian and $\mathbf{0}$ is an all-zero matrix of corresponding dimension. We can obtain H_{SA} by applying a projection over the ancilla space specified by the ancillary state $|0\rangle$. As such, if H_{SA} has eigenvalues $E_j \geq 0$ and eigenstates $|\psi_j\rangle$, then H_{sqrt} has eigenvalues $\pm\sqrt{E_j}$ with closely related eigenstates $|\phi_j^\pm\rangle$. To prove this, first observe that the subspace spanned by $\{|0\rangle|\psi_j\rangle, H_{\text{sqrt}}|0\rangle|\psi_j\rangle\}$ is invariant under H_{sqrt} :

$$H_{\text{sqrt}}|0\rangle|\psi_j\rangle = \sum_\alpha |\chi_\alpha\rangle \otimes (O_\alpha|\psi_j\rangle), \quad (\text{A4})$$

$$\begin{aligned} H_{\text{sqrt}}^2|0\rangle|\psi_j\rangle &= H_{\text{sqrt}} \sum_\alpha |\chi_\alpha\rangle \otimes (O_\alpha|\psi_j\rangle) \\ &= \sum_\alpha (\mathbb{I} \otimes O_\alpha^\dagger O_\alpha)|0\rangle \otimes |\psi_j\rangle = E_j|0\rangle \otimes |\psi_j\rangle. \end{aligned} \quad (\text{A5})$$

Second, we diagonalize H_{sqrt} on the subspace spanned by this set of states, to find that it has eigenvalues and eigenvectors

$$|\phi_j^\pm\rangle = \frac{1}{\sqrt{2}} \left(|\psi_j\rangle \pm \frac{1}{\sqrt{E_j}} H_{\text{sqrt}}|\psi_j\rangle \right), \quad (\text{A6})$$

$$H_{\text{sqrt}}|\phi_j^\pm\rangle = \pm\sqrt{E_j}|\phi_j^\pm\rangle. \quad (\text{A7})$$

We now discuss a block encoding for Eq. (A2). Let $\mathcal{B}$ also denote an ancillary system. Given access to controlled block encodings $\text{BE}[O_\alpha/\lambda_\alpha]$, where $\langle 0|_{\mathcal{B}} \text{BE}[O_\alpha/\lambda_\alpha] |0\rangle_{\mathcal{B}} = O_\alpha/\lambda_\alpha$ and $\lambda_\alpha \geq \|O_\alpha\|$, unitary iteration and coherent-alias-sampling state preparation [14] allows us to synthesize the Select and PREPARE circuits:

$$\text{SEL} = |0\rangle\langle 0| \otimes \mathbb{I} + \sum_\alpha |\alpha\rangle\langle \alpha| \otimes \text{BE}[O_\alpha/\lambda_\alpha], \quad (\text{A8})$$

$$\text{PREP}: |0\rangle_{\text{a}} \mapsto \sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqrt}}} |\alpha\rangle_{\text{a}} |\text{garb}_\alpha\rangle_{\text{a}} = \sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqrt}}} |\chi_\alpha\rangle_{\text{a}} = |\Psi\rangle_{\text{a}}, \quad \lambda_{\text{sqrt}} = \sqrt{\sum_\alpha \lambda_\alpha^2}. \quad (\text{A9})$$

Let $X_c = |0\rangle\langle 1|_c + |1\rangle\langle 0|_c$ be the “Pauli X” and let $\text{CPREP} = |0\rangle\langle 0|_c \otimes \mathbb{I} + |1\rangle\langle 1|_c \otimes \text{PREP}$, where c is an ancillary qubit. Observe that

$$\begin{aligned} \text{CPREP}(X_c \otimes |0\rangle\langle 0|_{\text{a}}) \text{CPREP}^\dagger &= |0\rangle\langle 1|_c \otimes |0\rangle\langle 0|_{\text{a}} \text{PREP}^\dagger + |1\rangle\langle 0|_c \otimes \text{PREP}|0\rangle\langle 0|_{\text{a}} \\ &= |0\rangle\langle 1|_c \otimes |0\rangle\langle \Psi|_{\text{a}} + |1\rangle\langle 0|_c \otimes |\Psi\rangle\langle 0|_{\text{a}} \\ &= |00\rangle\langle 1\Psi|_{\text{ac}} + |1\Psi\rangle\langle 00|_{\text{ac}}. \end{aligned} \quad (\text{A10})$$

Let $\text{AND} := \mathbb{I}_b \otimes |0\rangle\langle 0|_a + X_b \otimes (\mathbb{I}_a - |0\rangle\langle 0|_a)$ be a two-qubit operator, where b is another ancilla qubit, with $\mathbb{I}_{a,b}$ denoting the 2×2 identity. Hence

$$\langle b|_b \text{CPREP}(X_c \otimes \text{AND}) \text{CPREP}^\dagger |0\rangle_b = \text{CPREP}(X_c \otimes |0\rangle\langle 0|_a) \text{CPREP}^\dagger, \quad (\text{A11})$$

so $\text{CPREP}(X_c \otimes \text{AND}) \text{CPREP}^\dagger$ that block encodes Eq. (A10).

Hence, the circuit

$$U = \text{SEL} \cdot \text{CPREP} \cdot (X_c \otimes \text{AND}) \cdot \text{CPREP}^\dagger \cdot \text{SEL}^\dagger, \quad (\text{A12})$$

is a block encoding of $H_{\text{sqr}}/\lambda_{\text{sqr}}$:

$$\begin{aligned} \langle 0|_b \langle 0|_B U |0\rangle_B |0\rangle_b &= \langle 0|_B \text{SEL}(|1\Psi\rangle\langle 00|_{ac} + |00\rangle\langle 1\Psi|_{ac}) \text{SEL}^\dagger |0\rangle_B \\ &= \langle 0|_B \text{SEL} \left(\sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqr}}} |\alpha\rangle |\text{garb}_\alpha\rangle |1\rangle\langle 0| + |0\rangle \sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqr}}} \langle \alpha| \langle \text{garb}_\alpha| \langle 1| \right) \text{SEL}^\dagger |0\rangle_B \\ &= \langle 0|_B \left(\sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqr}}} |\alpha\rangle |\text{garb}_\alpha 1\rangle\langle 0| \otimes \text{BE} \left[\frac{O_\alpha}{\lambda_\alpha} \right] + |0\rangle \sum_\alpha \frac{\lambda_\alpha}{\lambda_{\text{sqr}}} \langle \alpha| \langle \text{garb}_\alpha 1| \otimes \text{BE}^\dagger \left[\frac{O_\alpha}{\lambda_\alpha} \right] \right) |0\rangle_B \\ &= \frac{1}{\lambda_{\text{sqr}}} \sum_\alpha (|\chi_\alpha\rangle\langle 0| \otimes O_\alpha + |0\rangle\langle \chi_\alpha| \otimes O_\alpha^\dagger) \\ &= \frac{H_{\text{sqr}}}{\lambda_{\text{sqr}}}. \end{aligned} \quad (\text{A13})$$

Observe that $U^2 = \mathbb{I}$ is self-inverse. Hence, following qubitization [16], the quantum walk $W = \text{REF}U$ of H_{sqr} is obtained by applying a reflection $\text{REF} = 2|0\rangle\langle 0|_{bB} - I$. This quantum walk has eigenphases $\pm \arccos(\sqrt{E_j}) \approx \pm[(\pi/2) - \sqrt{E_j}]$. Hence, applying the PEA with precision δ using c/δ queries to W learns $\sqrt{E_j}/\lambda_{\text{sqr}}$, for some constant c , which also translates to a quadratic advantage in the error of E_j . In the limit of small δ , the precision ϵ of learning E_j is

$$\epsilon = \left| \frac{d\sqrt{E_j}/\lambda_{\text{sqr}}}{dE_j} \right|^{-1} \delta = 2\sqrt{E_j \lambda_{\text{sqr}}^2} \delta. \quad (\text{A14})$$

Overall, each quantum walk step has cost dominated by two queries to Sel and two queries to Prep . Hence, the overall query complexity is $4c\sqrt{E_j \sum_\alpha \lambda_\alpha^2}/\epsilon$ to Sel and to Prep .

2. Quantum walk on rectangular matrices

One may more directly define a non-Hermitian matrix square root by choosing

$$H'_{\text{sqr}} = \sum_{\alpha=0}^{L-1} |\chi_\alpha\rangle \otimes O_\alpha, \quad (\text{A15})$$

which has a singular value decomposition

$$H'_{\text{sqr}} = \sum_j \sqrt{E_j} |\tilde{\psi}_j\rangle \langle \psi_j|, \quad |\tilde{\psi}_j\rangle = \frac{H'_{\text{sqr}} |\psi_j\rangle}{|H'_{\text{sqr}} |\psi_j\rangle|}. \quad (\text{A16})$$

Just like Eq. (A2), we are free to replace $|\chi_\alpha\rangle$ with any set of mutually orthogonal quantum states like basis states. As expected, the singular values (SVs) of H'_{sqr} are the square-root of those of H_{SA} following

$$\begin{aligned} \text{SV}(H'_{\text{sqr}}) &= \sqrt{\text{SV}(H_{\text{sqr}}'^\dagger H'_{\text{sqr}})} \\ &= \sqrt{\text{SV} \left(\sum_\alpha |\chi_\alpha\rangle \otimes O_\alpha^\dagger \sum_{\alpha'} |\chi_{\alpha'}\rangle \otimes O_{\alpha'} \right)} \\ &= \sqrt{\text{SV} \left(\sum_\alpha O_\alpha^\dagger O_\alpha \right)} \\ &= \sqrt{\text{SV}(H_{\text{SA}})}. \end{aligned} \quad (\text{A17})$$

By forming a quantum walk over rectangular matrices [25], we can again use the PEA to estimate the energies $\sqrt{E_j}$. Compared to Appendix A 1, the block encoding of the rectangular matrix H'_{sqr} is more straightforward.

To see this, given access to controlled block encodings $\text{BE}[O_\alpha/\lambda_\alpha]$, where $\langle 0|_B \text{BE}[O_\alpha/\lambda_\alpha] |0\rangle_B = O_\alpha/\lambda_\alpha$, unitary iteration and coherent alias sampling state preparation [14] allows us to synthesize the Select and PREPARE circuits

$$\text{SEL} = \sum_{\alpha} |\alpha\rangle\langle\alpha| \otimes \text{BE}[O_{\alpha}/\lambda_{\alpha}], \quad (\text{A18})$$

$$\text{PREP}|0\rangle_{\text{a}} = \sum_{\alpha} \frac{\lambda_{\alpha}}{\lambda_{\text{sqrt}}} |\alpha\rangle_{\text{a}} |\text{garb}_{\alpha}\rangle_{\text{a}} = \sum_{\alpha} \frac{\lambda_{\alpha}}{\lambda_{\text{sqrt}}} |\chi_{\alpha}\rangle_{\text{a}},$$

$$\lambda_{\text{sqrt}} := \sqrt{\sum_{\alpha} \lambda_{\alpha}^2}. \quad (\text{A19})$$

Then we have the block encoding

$$U := \text{SEL} \cdot \text{PREP}^{\dagger} = \text{BE} \left[\frac{H'_{\text{sqrt}}}{\lambda_{\text{sqrt}}} \right], \quad (\text{A20})$$

following

$$\langle 0|_{\text{B}} \text{SEL} \cdot \text{PREP}^{\dagger} |0\rangle_{\text{a}} |0\rangle_{\text{B}} = \frac{\sum_{\alpha} |\chi_{\alpha}\rangle \otimes O_{\alpha}}{\lambda_{\text{sqrt}}} = \frac{H'_{\text{sqrt}}}{\lambda_{\text{sqrt}}}, \quad (\text{A21})$$

where we used the freedom to choose any set of mutually orthogonal states $|\chi_{\alpha}\rangle$.

Let us now define the quantum walks $W_1 := \text{REF}_{\text{B}} U$ and $W_2 := \text{REF}_{\text{aB}} U^{\dagger}$ over non-Hermitian matrices following [25]. We find it useful to introduce the following states

$$\begin{aligned} W_1 |\perp_j\rangle &= \text{REF}_{\text{B}} U \frac{(I - |0\rangle\langle 0|_{\text{aB}}) U^{\dagger} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} \\ &= \frac{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} - \text{REF}_{\text{B}} U \frac{(|0\rangle\langle 0|_{\text{aB}}) U^{\dagger} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} \\ &= \frac{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} - (2|0\rangle\langle 0|_{\text{B}} - I) U |0\rangle_{\text{aB}} \frac{\sqrt{E_j}/\lambda_{\text{sqrt}} |\psi_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} \\ &= \frac{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} - 2|0\rangle\langle 0|_{\text{B}} U |0\rangle_{\text{aB}} \frac{\sqrt{E_j}/\lambda_{\text{sqrt}} |\psi_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} + U |0\rangle_{\text{aB}} |\psi_j\rangle \frac{\sqrt{E_j}/\lambda_{\text{sqrt}}}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}} \\ &= -\sqrt{1 - \frac{E_j}{\lambda_{\text{sqrt}}^2}} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle + \frac{\sqrt{E_j}}{\lambda_{\text{sqrt}}} |\tilde{\perp}_j\rangle, \end{aligned} \quad (\text{A25})$$

and similarly for W_2 , which maps the subspace spanned by $\{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle, |\tilde{\perp}_j\rangle\}$ to $\{|0\rangle_{\text{aB}} |\psi_j\rangle, |\perp_j\rangle\}$.

We can diagonalize W_1 and W_2 on these subspaces by defining

$$|\phi_j^{\pm}\rangle := \frac{|0\rangle_{\text{aB}} |\psi_j\rangle \pm i |\perp_j\rangle}{\sqrt{2}}, \quad |\tilde{\phi}_j^{\pm}\rangle := \frac{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle \pm i |\tilde{\perp}_j\rangle}{\sqrt{2}}. \quad (\text{A26})$$

Then by substituting from Eq. (A25),

$$\begin{aligned} |\perp_j\rangle &:= \frac{(\mathbb{I} - |0\rangle\langle 0|_{\text{aB}}) U^{\dagger} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{|(\mathbb{I} - |0\rangle\langle 0|_{\text{aB}}) U^{\dagger} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle|} \\ &= \frac{(\mathbb{I} - |0\rangle\langle 0|_{\text{aB}}) U^{\dagger} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}}. \end{aligned} \quad (\text{A22})$$

$$\begin{aligned} |\tilde{\perp}_j\rangle &:= \frac{(\mathbb{I} - |0\rangle\langle 0|_{\text{B}}) U |0\rangle_{\text{a}} |\psi_j\rangle}{|(\mathbb{I} - |0\rangle\langle 0|_{\text{B}}) U |0\rangle_{\text{a}} |\psi_j\rangle|} \\ &= \frac{(\mathbb{I} - |0\rangle\langle 0|_{\text{B}}) U |0\rangle_{\text{a}} |\psi_j\rangle}{\sqrt{1 - E_j/\lambda_{\text{sqrt}}^2}}. \end{aligned} \quad (\text{A23})$$

Consider the eigenstates $|\psi_j\rangle$ of H_{SA} of eigenvalue $E_j \geq 0$. Then W_1 maps the subspace spanned by $\{|0\rangle_{\text{aB}} |\psi_j\rangle, |\perp_j\rangle\}$ to $\{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle, |\tilde{\perp}_j\rangle\}$. We can show this by direct computation:

$$\begin{aligned} W_1 |0\rangle_{\text{aB}} |\psi_j\rangle &= \text{REF}_{\text{B}} \left(\frac{\sqrt{E_j}}{\lambda_{\text{sqrt}}} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle + \sqrt{1 - \frac{E_j}{\lambda_{\text{sqrt}}^2}} |\tilde{\perp}_j\rangle \right) \\ &= \frac{\sqrt{E_j}}{\lambda_{\text{sqrt}}} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle - \sqrt{1 - \frac{E_j}{\lambda_{\text{sqrt}}^2}} |\tilde{\perp}_j\rangle, \end{aligned} \quad (\text{A24})$$

$$\begin{aligned} W_1 |\phi_j^{\pm}\rangle &= \frac{\sqrt{E_j} |0\rangle_{\text{B}} |\tilde{\psi}_j\rangle \pm i |\tilde{\perp}_j\rangle}{\lambda_{\text{sqrt}} \sqrt{2}} \\ &\quad \pm i \sqrt{1 - \frac{E_j}{\lambda_{\text{sqrt}}^2}} \frac{|0\rangle_{\text{B}} |\tilde{\psi}_j\rangle \pm i |\tilde{\perp}_j\rangle}{\sqrt{2}} \\ &= e^{\pm i \arccos(\sqrt{E_j}/\lambda_{\text{sqrt}})} |\tilde{\phi}_j^{\pm}\rangle, \end{aligned} \quad (\text{A27})$$

$$W_2 |\tilde{\phi}_j^{\pm}\rangle = e^{\pm i \arccos(\sqrt{E_j}/\lambda_{\text{sqrt}})} |\phi_j^{\pm}\rangle. \quad (\text{A28})$$

Hence, the quantum walks W_1 and W_2 apply eigenphases $\pm \arccos(\sqrt{E_j}/\lambda_{\text{sqr}})$ on the alternating subspaces $|\phi_j^\pm\rangle$ and $|\tilde{\phi}_j^\pm\rangle$.

We may thus perform phase estimation on the quantum walks, so long as we alternate between W_1 and W_2 . For instance, if the PEA calls for the k th power of a unitary to be applied, we would replace this with the length k sequence $\cdots W_2 W_1 W_2 W_1$, ending with W_1 if k is odd and with W_2 if k is even. Hence, we may estimate E_j with the same sensitivity as Eq. (A14), but with half the number of queries to Sel and Prep due to the simpler block encoding of the rectangular H'_{sqr} compared to the Hermitian H_{sqr} in Eq. (A2).

Moreover, we may also choose to perform phase estimation on the product $W_2 W_1$. From Eq. (A27),

$$\begin{aligned} W_2 W_1 |\phi_j^\pm\rangle &= e^{\pm i 2 \arccos(\sqrt{E_j}/\lambda_{\text{sqr}})} |\phi_j^\pm\rangle \\ &= (\cos(\pm 2 \arccos(\sqrt{E_j}/\lambda_{\text{sqr}})) \pm i \cdots) |\phi_j^\pm\rangle \\ &= (T_2[\sqrt{E_j}/\lambda_{\text{sqr}}] \pm i \cdots) |\phi_j^\pm\rangle \\ &= (2E_j/\lambda_{\text{sqr}}^2 - 1) \pm i \cdots |\phi_j^\pm\rangle \\ &= e^{\pm i \arccos(2E_j/\lambda_{\text{sqr}}^2 - 1)} |\phi_j^\pm\rangle, \end{aligned} \quad (\text{A29})$$

where we have used the definition of Chebyshev polynomials $T_k[\cos(\theta)] = \cos(k\theta)$ with $T_2[x] = 2x^2 - 1$. If we expand the quantum circuit

$$W_2 W_1 = \text{REF}_{\text{aB}}(U^\dagger \text{REF}_B U), \quad (\text{A30})$$

the attentive reader will notice that the term in round brackets is a block-encoding $\text{BE}[2H_{\text{SA}}/\lambda_{\text{sqr}}^2 - \mathbb{I}]$ following

$$\begin{aligned} \frac{2H_{\text{SA}}}{\lambda_{\text{sqr}}^2} - \mathbb{I} &= \langle 0 |_{\text{aB}} W_2 W_1 | 0 \rangle_{\text{aB}} \\ &= \langle 0 |_{\text{aB}} \text{REF}_{\text{aB}} U^\dagger \text{REF}_B U | 0 \rangle_{\text{aB}} \\ &= \langle 0 |_{\text{aB}} U^\dagger \text{REF}_B U | 0 \rangle_{\text{aB}}. \end{aligned} \quad (\text{A31})$$

APPENDIX B: DFTHC BLOCK-ENCODING COST

For block encoding we will leverage the representation of the spin-free operator in Majorana representation. There two high-level components to the block encoding: (i) the linear combination of $T_2(\text{BE}[O_\alpha^\dagger], \text{BE}[O_\alpha])$ and (ii) implementation of each qubitization step on each O_α . The

reflection on the ancilla to construct the walk operator is controlled on all ancilla associated with the linear combination of T_2 and qubitization signal states. In the Majorana representation a DFTHC ansatz with $r \in [R], c \in [C], b \in [B]$ has SOS generators

$$\begin{aligned} O_{\text{SF},(rc)} &= w_B^{(rc)} I + \frac{i}{2} \sum_{\sigma \in \{0,1\}} \sum_{b \in [B]} w_b^{(rc)} \tilde{\gamma}_{\vec{u}_b^{(r)} \sigma 1} \tilde{\gamma}_{\vec{u}_b^{(r)} \sigma 0}, \\ \tilde{\gamma}_{\vec{u} \sigma x} &= \sum_p u_p \gamma_{p \sigma x}, \end{aligned} \quad (\text{B1})$$

$$O_{\text{D}_1, r}^e = \frac{\sqrt{w_+^{(r)}}}{2} (\tilde{\gamma}_{\vec{u}_+^{(r)} \sigma 0} + i \tilde{\gamma}_{\vec{u}_+^{(r)} \sigma 1}), \quad (\text{B2})$$

$$O_{\text{Q}_1, r}^e = \frac{\sqrt{w_-^{(r)}}}{2} (\tilde{\gamma}_{\vec{u}_-^{(r)} \sigma 0} - i \tilde{\gamma}_{\vec{u}_-^{(r)} \sigma 1}). \quad (\text{B3})$$

As previously mentioned, the optimal D_1 and Q_1 generators are found by an eigendecomposition of the one-body $N \times N$ coefficient matrix. Hence, we can identify $w_+^{(r)}$ with the positive semidefinite N_{D_1} eigenvalues of the coefficient matrix, and $w_-^{(r)}$ with the absolute value of the $N_{\text{Q}_1} = N - N_{\text{D}_1}$ remaining eigenvalues, which are negative. In total, there are $N + RC(B + 1)$ different coefficients w and $N + RB$ real unit vectors $\vec{u}$ across the different generators.

In the following, we will use qubit registers $|r\rangle_{\text{r}}, |G\rangle_{\text{G}}, |b\rangle_{\text{b}}, |c\rangle_{\text{c}}$ storing integer the integers r, b, c . It is useful to index the different operators with a single integer $x_0 \in [X_0]$ that is a function of G, r, c as follows:

$$x_0(G, r, c) = \begin{cases} r, & G = G_{\text{D}_1}, \\ N_{\text{D}_1} + r, & G = G_{\text{Q}_1}, \\ N + rC + c, & G = G_{\text{SF}}, \end{cases} \quad X_0 = N + RC. \quad (\text{B4})$$

We adopt the convention that $b = B$ indexes the identity term.

The sum-of-square Hamiltonian is then

$$H_{\text{SOS}} = \sum_g \sum_{rc} O_{g,rc}^\dagger O_{g,rc} = \sum_{x_0} O_{x_0}^\dagger O_{x_0}. \quad (\text{B5})$$

For each term, we associate a certain block-encoding normalization factor that will be later relevant.

$$\lambda_{\text{SF}} = \sum_{rc} \lambda_{\text{SF},rc}^2, \quad \lambda_{\text{SF},rc} = \frac{1}{\sqrt{2}} \left(|w_B^{(rc)}| + \sum_{b=0}^{B-1} |w_b^{(rc)}| \right), \quad (\text{B6})$$

$$\lambda_{D_1} = \sum_r \lambda_{D_1^0, r}^2 + \sum_r \lambda_{D_1^1, r}^2, \quad \lambda_{D_1^c, r} = \sqrt{w_+^{(r)}}, \quad (B7)$$

$$\lambda_{Q_1} = \sum_r \lambda_{Q_1^0, r}^2 + \sum_r \lambda_{Q_1^1, r}^2, \quad \lambda_{Q_1^c, r} = \sqrt{w_-^{(r)}}. \quad (B8)$$

The overall block-encoding normalization of the sum-of-squares Hamiltonian will be

$$\Lambda = \frac{1}{2}(\lambda_{Q_1} + \lambda_{D_1} + \lambda_{SF}). \quad (B9)$$

The key idea is that using an efficiently computable indexing scheme to each distinct coefficient and unit vector, we may form table-lookups of the smallest possible size to this data, and hence block encode with the smallest constant factors.

The method for block encoding the Hamiltonian uses similar principles to those used for double factorization in prior work [7,9], with some key differences. The main parts that are in common are as follows.

- (1) A register is prepared with the appropriate amplitudes over x_o , which is used to control the block encoding of O_{x_o} .
- (2) The single step of oblivious amplitude amplification is used to provide the factor of 2 improvement, as explained above. This means that we block encode O_{x_o} , reflect on the ancilla qubits of the block encoding, then apply the inverse of the block encoding.
- (3) A control register is prepared with an equal superposition to represent σ for the spin, and used to control a swap between the system registers representing spin up and spin down.
- (4) A superposition over b needs to be prepared.
- (5) Each O_{x_o} involves application of Majorana operators in a transformed basis, with the choice of basis controlled by x_o . This is achieved by applying QROM on x_o to output the rotation angles to change the basis, then applying a sequence of Givens rotations. These rotations are only applied on one half the qubits representing spin up (for example), because of the previous step.
- (6) After the basis transformation, the Majorana operators need only be applied to a single qubit representing a single (transformed) spin-orbital.

Next we detail the subtleties involved in each of these steps.

1. Preparing the superposition over x_o

The new subtlety is that the contiguous register containing x_o is used to index multiple variables G, r, c . Recall that G is used to select between the operators $O_{SF, (rc)}$,

$O_{D_1, r}$, and $O_{Q_1, r}$. The variable r is used for further indexing over all three of these operators, and c is used to index only $O_{SF, (rc)}$. At this stage we do not yet output these variables, and just prepare the state with the appropriate weightings

$$\text{PREP}|0\rangle = \sum_{x_o=0}^{X_o-1} \frac{\lambda^{(G_{x_o} r_{x_o} c_{x_o})}}{\sqrt{2\Lambda}} |x_o\rangle |\text{garbage}_{x_o}\rangle_{\text{garb}}, \quad (B10)$$

where the composite index $(G_{x_o} r_{x_o} c_{x_o})$ from Eq. (B4) selects one of the coefficients $\lambda_{SF, rc}, 2\lambda_{D_1^c, r}, 2\lambda_{Q_1^c, r}$ from Eq. (B6). The dominant cost of this step is the alias sampling step which invokes Qroam [17,56]. This is the control state which we will use to index controlled block encoding.

2. The step of oblivious-amplitude amplification or qubitization

The use of oblivious-amplitude amplification is the same as in previous work [9], with the only difference being the number of qubits reflected upon due to the different operators, and the non-Hermitian nature of the block encoded operators giving $T_2(O_a^\dagger, O_a)$. The qubits we need to reflect on are as follows.

- (i) The $\lceil \log(B+1) \rceil$ qubits that the state for b is prepared on.
- (ii) The b_{k2} qubits used for the equal superposition state in the coherent alias sampling.
- (iii) One qubit that is rotated as part of the preparation of the equal superposition over b .
- (iv) One qubit for the sum over spins in $O_{SF, rc}$ (see below).

These registers are similar to those used in prior work.

The preparation of equal superposition states (as part of coherent alias sampling) has a small amplitude for failure, and operations should be controlled on success of these preparations. That results in a very small (less than about 0.1%) increase in λ , and has a smaller impact than contamination from unwanted operations if the control was not used. The reflection in this part needs to be controlled on the success of preparation of the register for x_o , so has a Toffoli cost of

$$\lceil \log(B+1) \rceil + b_{k2} + 1. \quad (B11)$$

In prior work there was one extra control which yielded an additional Toffoli of complexity (on $\ell = 0$ in Ref. [7]).

3. Selecting between the spins

A new subtlety in the selection of the spins is that for $O_{\text{SF},(rc)}$ the summation over the spin σ is required for the block encoding of that operator, whereas $O_{D_1^{\sigma,r}}$ and $O_{Q_1^{\sigma,r}}$ are dependent on σ , and the summation of spin is part of the outer sum (with x_0). This means that we need separate control variables for the spin for SF versus D_1 and Q_1 .

That is, we have one control qubit in a $|+\rangle$ state, that is used to control the swap of the spin up and spin down parts of the system register for D_1 or Q_1 . A second control qubit is used to control the swap for the case of SF. These controlled swaps are performed in the block encoding of O_{x_0} . The second qubit is part of the block encoding as well, so is reflected upon in the oblivious amplitude amplification. That ensures that the summation over σ for SF is obtained. The first qubit is not reflected upon in the step of oblivious amplitude amplification, ensuring that it is part of the outer summation. When we perform the block encoding of O_{x_0} , we extract the variables G, r , so we have the variable G selecting between these operators available to control upon.

To be more specific, the Majorana operators of different spin can be related by

$$\gamma_{p1x} = \left(\prod_{j=0}^{N-1} \text{SWAP}_{j,j+N} \right) \gamma_{p0x} \left(\prod_{j=0}^{N-1} \text{SWAP}_{j,j+N} \right) \cdot \vec{Z}_{N-1}, \quad (\text{B12})$$

where $\vec{Z}_{N-1}$ accounts for the extra string of operators needed for the Jordan-Wigner representation. In the Jordan-Wigner representation, $\gamma_{p\sigma 0} = X_{\sigma N+p} \vec{Z}_{\sigma N+p-1}$ and $\gamma_{p\sigma 1} = Y_{\sigma N+p} \vec{Z}_{\sigma N+p-1}$, where $\vec{Z}_{\sigma N+p-1} = Z_{\sigma N+p-1} \cdots Z_0$.

Because it can be performed with Clifford gates, it does not add to the Toffoli complexity considered in our costings. Let us define the controlled operations

$$U_{\text{Swap}} := \sum_{\sigma \in \{0,1\}} |\sigma\rangle\langle\sigma| \otimes \left(\prod_{j=0}^{N-1} \text{SWAP}_{j,j+N} \right)^{\sigma}, \quad (\text{B13})$$

$$U_1 := \sum_{\sigma \in \{0,1\}} |\sigma\rangle\langle\sigma| \otimes \vec{Z}_{N-1}^{\sigma}. \quad (\text{B14})$$

Then we can apply the appropriate operators using U_{Swap} and U_1 , for example,

$$\sum_{\sigma} |\sigma\rangle\langle\sigma| \otimes \gamma_{p\sigma x} = U_{\text{Swap}} \gamma_{p0x} U_{\text{Swap}} U_1. \quad (\text{B15})$$

To select the appropriate σ register to use, we will encode G using two qubits, with one selecting between SF and the pair D_1 and Q_1 . Calling that G_0 , we will have registers with $|G_0\rangle|\sigma_0\rangle|\sigma_1\rangle$. By applying two Toffolis (and X gates),

we obtain an output qubit $|\sigma\rangle$, where $\sigma = 1$ if $\sigma_{G_0} = 1$. After the controlled swaps, $|\sigma\rangle$ may be erased using an X measurement and Clifford phase fixups in the usual way, so there is no extra Toffoli cost for erasure.

Therefore, producing the qubit with $|\sigma\rangle$ adds an extra cost of two Toffolis to the block encoding of each O_{x_0} , for a total of 4 Toffolis for the overall block encoding of the Hamiltonian. The four controlled swaps add a cost of $4N$ Toffolis.

4. Preparing the superposition over b

This preparation of the inner superposition is similar to prior work with two key differences. The first is that we use this step to also output G, r , which are used in later steps. Prior to this we have just prepared a superposition over the contiguous register x_0 , and the unary iteration over this variable in this step is used to output G, r as well as the information needed for the coherent alias sampling for preparing the appropriate superposition over b . Note that we do not need to output c , because all operations which use its value are performed directly on x_0 .

The second difference is that no superposition over b is needed for the operators $O_{D_1^{\sigma,r}}, O_{Q_1^{\sigma,r}}$. When iterating over the values of x_0 corresponding to these operators, we can simply not output values used for coherent alias sampling for b . However, it is convenient to output the set of rotations needed for D_1, Q_1 in the next step when iterating over these values of x_0 . This is an additional optimization over what is shown in Fig. 2, where all rotations for both SF and D_1, Q_1 are output in the same step.

The signs needed in the superposition over b (i.e., the signs of $w_b^{(rc)}$) can be obtained by using a single qubit output in the coherent alias sampling preparation to flag whether the sign should be changed when the controlled swap is performed. That is, this qubit is $|1\rangle$ if the sign for the alt value is different. The remaining sign can be applied in the QROM erasure. This is because the QROM erasure is performed by measuring the data qubits in the X basis and performing a sign correction. Any signs needed can be combined with the sign correction and performed in the same step. This is distinct from prior work which used two output qubits for the sign.

To be more specific, in coherent alias sampling one initially prepares an equal superposition state over ℓ , and can output a sign qubit in addition to the alt and keep values to give

$$\sqrt{\frac{1}{L}} \sum_{\ell=0}^{L-1} |\ell\rangle |\text{alt}_{\ell}\rangle |\Delta\epsilon_{\ell}\rangle |\text{keep}_{\ell}\rangle, \quad (\text{B16})$$

where $\Delta\epsilon_{\ell}$ is 1 if the signs needed for ℓ and alt_{ℓ} are different. The inequality test with an equal superposition state and controlled swap give

$$\sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} |\ell\rangle |\text{alt}_\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=0}^{\text{keep}_\ell-1} |m\rangle |0\rangle + \sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} |\text{alt}_\ell\rangle |\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=\text{keep}_\ell-1}^{M-1} |m\rangle |1\rangle. \quad (\text{B17})$$

A controlled phase gate between the qubit flagging the result of the inequality test and the sign qubit then gives

$$\sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} |\ell\rangle |\text{alt}_\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=0}^{\text{keep}_\ell-1} |m\rangle |0\rangle + \sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} (-1)^{\Delta\varsigma_\ell} |\text{alt}_\ell\rangle |\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=\text{keep}_\ell-1}^{M-1} |m\rangle |1\rangle. \quad (\text{B18})$$

Then applying the sign for ℓ gives

$$\sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} (-1)^{\varsigma_\ell} |\ell\rangle |\text{alt}_\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=0}^{\text{keep}_\ell-1} |m\rangle |0\rangle + \sqrt{\frac{1}{ML}} \sum_{\ell=0}^{L-1} (-1)^{\varsigma_\ell + \Delta\varsigma_\ell} |\text{alt}_\ell\rangle |\ell\rangle |\Delta\varsigma_\ell\rangle |\text{keep}_\ell\rangle \sum_{m=\text{keep}_\ell-1}^{M-1} |m\rangle |1\rangle. \quad (\text{B19})$$

Since $\Delta\varsigma_\ell$ is 1 if the signs for ℓ and alt_ℓ are different, $(-1)^{\varsigma_\ell + \Delta\varsigma_\ell}$ yields the correct sign for alt_ℓ .

5. Rotating the basis

Before rotating the basis, we need to first output the set of rotations needed in an ancilla register to be used as control. The distinction from prior work is what registers are used to control the output of these rotations. For the operator $O_{\text{SF}(rc)}$, the rotations are selected using r and b , but they are *not* selected using c . This reduces the cost of the iteration needed to output the set of rotations.

Another difference is that for D_1 , Q_1 , the rotations are selected only using r and b is not used. A further subtlety is that we are using $b = B$ to select the identity term. The way this is accounted for is that we prepare the superposition over b including $b = B$, but then in the iteration to output the rotations, we set 1 for $b \neq B$ or 0 for $b = B$. The qubit is used as part of the control in the application of the Majorana operators (see below).

Otherwise the rotation of the basis is applied in a similar way as in prior work. Note that for any unit vector $\vec{u}$, there exists a unitary rotation $U(\vec{u})$ comprised of $N - 1$ different arbitrary single-qubit rotations with rotation angle θ_j and $\mathcal{O}(N)$ two-qubit Clifford gates such that $U^\dagger(\vec{u})\gamma_{00x}U(\vec{u}) = \tilde{\gamma}_{\vec{u}0x}$. In particular any sequence of rotation angles $\vec{\theta} = (\theta_0, \dots, \theta_{N-2})$ corresponds to some $u = f(\vec{\theta})$. By having each of the $j \in [N - 1]$ rotation angles θ_j be specified by a b_{rot} -bit binary numbers, one may further define the controlled rotation [9]

$$U_{\text{Rot}}|\vec{\theta}\rangle|\psi\rangle = |\vec{\theta}\rangle U(u(\vec{\theta}))|\psi\rangle, \quad (\text{B20})$$

$$U_{\text{Rot}}^\dagger \gamma_{00x} U_{\text{Rot}} = \sum_{\vec{\theta}} |\vec{\theta}\rangle \langle \vec{\theta}| \otimes \tilde{\gamma}_{f(\vec{\theta})0x}, \quad (\text{B21})$$

and each U_{Rot} can be implemented with $2(N - 1)b_{\text{rot}}$ Toffoli gates. To simplify notation, we use $\vec{u}$ and its representation through $\vec{\theta}$ interchangeably. For instance, $|\vec{u}\rangle = |\vec{\theta} = f^{-1}(\vec{u})\rangle$. As there are multiple solutions $\vec{\theta}$ representing the same u , we arbitrarily select any one solution from $f^{-1}(u)$.

Note also that the transformation can be applied in common between the Majorana operators $\tilde{\gamma}_{\vec{u}0x}$ with $x = 0$ and 1. we can transform between the two operators using the phase gate S , because $X = S^\dagger Y S$. One could apply S and $S^\dagger$ before and after (respectively) on all system qubits, but because the tensor product of S commutes through the Givens rotations, they need only be applied to the single qubit representing the transformed spin orbital.

6. Applying the Majorana operators on a single qubit

This last step is significantly changed here as compared to prior work. In Ref. [9] the number operator was written as a sum of the identity and Z , and the identity was removed from the sum so that only the Z gate needed to be applied. Because of the transformation of basis, it need only be applied to a single qubit.

In our operators, we have a similar term in $O_{\text{SF}(rc)}$, with the product of Majorana operators $\tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_1} \tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_0}$. The basis transformation means that these Majorana operators are only applied to a single qubit, and the sequence of Z gates on the other qubits is not required. This means that $i\tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_1} \tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_0}$ simplifies to just a Z gate on a single qubit, similar to prior work.

The difference is that we also have D_1 and Q_1 , where the operators $\tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_1}$ and $\tilde{\gamma}_{\vec{u}_b^{(r)}\sigma_0}$ need to be applied individually. However, these are simplified to Y and X operators on a single qubit, so are still simple to apply. A single control




FIG. 4. The quantum circuit used to apply the Majorana operators on the single qubit, with the target qubit at the bottom. The four control qubits at the top correspond to the numbered list (i) to (iv) in the text.

qubit is used to provide the sum $\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 0} \pm i\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 1} = X \pm iY$. A qubit flagging the difference between D_1 and Q_1 can be used to control the sign in this implementation. A second qubit flagging whether SF is performed can be used to control whether the Z gate is performed, or whether it is X or Y.

To be more specific, we have the following control qubits.

- (i) The first qubit for G , which distinguishes between SF and the pair D_1, Q_1 .
- (ii) The second qubit for G , which distinguishes between D_1 and Q_1 .
- (iii) A qubit which flags whether we have the identity component of SF.
- (iv) A qubit in superposition used to distinguish between $\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 1}$ and $\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 0}$ for D_1, Q_1 .

The quantum circuit used to apply the Majorana operators with these control qubits is then shown in Fig. 4. The first operation is to apply the Z gate for SF. The second controlled operation is giving the required X operation for $\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 0}$. The third controlled operation applies a Z gate, provided we have D_1, Q_1 , and the fourth qubit is selecting $\tilde{\gamma}_{\tilde{u}^{(r)}\sigma 1}$ (so $iY = ZX$ is needed). The fourth controlled operation applies the correct sign factor in $X \pm iY$ needed to select between D_1 and Q_1 . This application of the Majorana operators uses only two Toffolis. We also need controls by successful preparation of the equal

superposition states for x_0 and $|b\rangle$. A single Toffoli is used to flag success for both, and control by this flag will add another 4 Toffolis for a total of 7.

7. Total costs

a. Summary of costs

Now we enumerate the total costs for the block encoding, with the complete circuit illustrated in Fig. 2. A more detailed explanation of the block-encoding costs is given in Appendix B 7 b. The total costs for constructing the walk operator are as follows.

- (1) Preparation of the control register for the outer sum over x_0 and erasure, corresponding to everything in Table VII.
- (2) The reflection in the block encoding of the Hamiltonian for $T_2(O_a^\dagger, O_a)$, with cost

$$\lceil \log(B+1) \rceil + b_{k2} + 1. \quad (\text{B22})$$

- (3) The four controlled swaps for the implementation of the spin, with cost $4(N+1)$.
- (4) Preparation of the control register for b in the inner sum, given by the costs in Table VIII, with a factor of 2 for the two calls.
- (5) The basis rotation, including the output and erasure of the rotation angles, as well as the rotations themselves, as given in Table IX (again with a factor of 2).
- (6) The Majorana operators on the single qubit with cost of 7 Toffolis. This is called twice per T_2 .
- (7) A reflection on all the ancilla for the block encodings to construct the walk operator of Toffoli cost

$$\lceil \log(N+RC) \rceil + \lceil \log(B+1) \rceil + b_{k1} + b_{k2} + 2. \quad (\text{B23})$$

b. More details of block-encoding costs

Here we give further details of the block-encoding costs, excluding the controlled swaps for the spin, reflection for

TABLE VII. Preparation costs for the state preparation over unitaries $T_2(O_a^\dagger, O_a)$. The right column gives the corresponding step number from Appendix B 7 b, and the “Calls” column gives the number of calls for each step of the walk (so the multiplying factor for the Toffoli cost). The selection register is x_0 , which takes on $N+RC$ values. The quantity $b_1 = \lceil \log(N+RC) \rceil + b_{k1}$, where b_{k1} is the number of bits used for the “keep” register in coherent alias sampling. The AliasSampling involves an inequality test at the cost of b_{k1} and controlled swap costing $\lceil \log(N+RC) \rceil$. The b_1 term in the fourth line is for inverting the inequality tests and controlled swaps. The integers k_1 and k_5 may be chosen to minimize costs. The ancillas are not listed for the inverse operations because they are considerably smaller than used for the forward operations, so do not affect the maximum qubit usage of the algorithm.

Operation	Cost	Persistent ancilla	Temporary ancilla	Calls	Step
Uniform(x_0)	$4\lceil \log(N+RC) \rceil$	$\lceil \log(N+RC) \rceil + 2$	$\lceil \log(N+RC) \rceil - 2$	1	1
AliasSampling(x_0)	$\lceil \frac{N+RC}{2^{k_1}} \rceil + 2^{k_1} b_1$	$b_1 + b_{k1}$	$2^{k_1} b_1 + \lceil \log_2(\frac{N+RC}{2^{k_1}}) \rceil$	1	2
AliasSampling(x_0) [†]	$b_1 + \lceil \frac{N+RC}{2^{k_5}} \rceil + 2^{k_5}$	...	...	1	10
Uniform(x_0) [†]	$4\lceil \log(N+RC) \rceil$	...	...	1	11

TABLE VIII. Preparation costs for the b register. Unary iteration is performed over the first RC values of x_o , combined with Qroam over b , to output data needed for the coherent alias sampling state preparation on b . The quantity $b_2 = \lceil \log(B+1) \rceil + b_{k_2} + 1$, where b_{k_2} is the bit size of the keep register. While performing the unary iteration over x_o , we also output G, r with no further Toffoli cost. The right column gives the corresponding step from Appendix B 7 b.

Operation	Toffoli	Persistent ancilla	Temporary ancilla	Calls	Step
Uniform(b)	$4\lceil \log(B+1) \rceil$	$\lceil \log(B+1) \rceil + 2$	$\lceil \log(B+1) \rceil - 2$	2	3
AliasSampling(b)	$RC\lceil \frac{B+1}{2^{k_2}} \rceil + 2^{k_2}b_2$	$b_2 + b_{k_2} + \lceil \log R \rceil + 2$	$\lceil \log(RC\lceil \frac{B+1}{2^{k_2}} \rceil) \rceil + (2^{k_2} - 1)b_2$	2	4
AliasSampling(b) [†]	$b_2 + N + \lceil \frac{RC}{2^{k_4}} \rceil + 2^{k_4}(B+1)$	...	...	2	8
Uniform(b) [†]	$4\lceil \log(B+1) \rceil$	...	...	2	9

TABLE IX. QROM to output rotations, implement them, and erase the data. The integer b_{rot} is the bits of precision used for each rotation angle.

Operation	Toffoli	Persistent ancilla	Temporary ancilla	Calls	Step
QROM($\vec{u}$)	$N + RB$	$(N-1)b_{\text{rot}}$	$\max(\lceil \log(N+RC) \rceil, \lceil \log(RB) \rceil)$	2	5
basis rotations	$2(N-1)b_{\text{rot}}$	...	b_{rot}	2	6
QROM($\vec{u}$) [†]	$R + B$	...	...	2	7

the step of oblivious amplitude amplification, and implementation of the Majorana operator on a single qubit, as these are detailed above.

- (1) First we prepare an equal superposition for x_o with cost

$$4\lceil \log(N+RC) \rceil. \quad (\text{B24})$$

This is omitting small $O(1)$ costs, and also ignoring savings when $N+RC$ has factors of 2. For this costing see Sec. III E 2 of Ref. [64] (pages 42 and 43), where the cost is $4k+2s-13$ Toffolis for preparation on k qubits and using s bits of accuracy for a rotation. That rotation can be used with $s=7$ to give very accurate results, for a cost of $4k+1$. This is why we give the cost above as 4 times the number of qubits, ignoring $O(1)$ costs.

Qubit cost. The number of qubits used in this step is $\lceil \log(N+RC) \rceil$ to store x_o , the extra qubit that is rotated in the equal superposition, and a qubit which is used to flag success of the preparation. Extra temporary qubits are used in the coherent arithmetic (inequality test), which may be ignored as they are fewer than those used in later steps. We also use a phase gradient register for rotations, which is used in the entire algorithm so the qubits for this register are quantified separately. The total qubit cost is therefore

$$\lceil \log(N+RC) \rceil + 2 \quad (\text{B25})$$

persistent qubits, and

$$\lceil \log(N+RC) \rceil - 2 \quad (\text{B26})$$

temporary qubits for the arithmetic. It is also possible that the number of temporary qubits used for the rotation on the ancilla qubit can be larger for small values of $N+RC$.

- (2) Perform Qroam on x_o with cost

$$\left\lceil \frac{N+RC}{2^{k_1}} \right\rceil + (2^{k_1} - 1)b_1, \quad (\text{B27})$$

for some choice of integer k_1 , where b_1 is the output size

$$b_1 = \lceil \log(N+RC) \rceil + b_{k_1}, \quad (\text{B28})$$

with b_{k_1} the number of bits for the keep register in coherent alias sampling. Then the inequality test is cost b_{k_1} and the controlled swap is cost $\lceil \log(N+RC) \rceil$, for total cost in step 2 of

$$\left\lceil \frac{N+RC}{2^{k_1}} \right\rceil + 2^{k_1}b_1. \quad (\text{B29})$$

Qubit cost. The persistent qubits used in this step are

$$b_1 + b_{k_1}, \quad (\text{B30})$$

with b_1 for the output of the QROM, and b_{k_1} qubits used for the equal superposition register. (Those must be retained because they are not reset.) There are also

$$\left\lceil \log \left(\frac{N + RC}{2^{k_1}} \right) \right\rceil + (2^{k_1} - 1)b_1 \quad (\text{B31})$$

temporary qubits for the output of the QROM as well as the temporary qubits used in the unary iteration. Recall that there are $2^{k_1}b_1$ qubits output, then we swap the b_1 output qubits into the correct register, and $(2^{k_1} - 1)b_1$ qubits are erased (via an X measurement and a later phase fixup). There are also temporary qubits used in the inequality test for coherent alias sampling, but these are much less than those used in the QROM so may be ignored.

- (3) For the block encoding of O_{x_0} , we first prepare an equal superposition over b with cost

$$4\lceil \log(B + 1) \rceil. \quad (\text{B32})$$

This preparation does not need to be controlled.

Qubit cost. Similar to step 1, there is a persistent ancilla cost of

$$\lceil \log(B + 1) \rceil + 2 \quad (\text{B33})$$

qubits for b and the rotated qubit, and

$$\lceil \log(B + 1) \rceil - 2 \quad (\text{B34})$$

temporary qubits in the arithmetic.

- (4) Next we can perform unary iteration over x_0 , combined with Q_{roam} on b . This has Toffoli cost

$$RC \left\lceil \frac{B + 1}{2^{k_2}} \right\rceil + (2^{k_2} - 1)b_2. \quad (\text{B35})$$

Here k_2 is a new choice of integer, and b_2 is the output size for coherent alias sampling on b of

$$b_2 = \lceil \log(B + 1) \rceil + b_{k_2} + 1, \quad (\text{B36})$$

with b_{k_2} being another choice of the number of bits for the keep register. In b_2 we include a single qubit for the sign. Again there will be a b_2 cost for the inequality test as well as a controlled swap, to give total cost

$$RC \left\lceil \frac{B + 1}{2^{k_2}} \right\rceil + 2^{k_2}b_2. \quad (\text{B37})$$

The idea of performing the unary iteration in this step is furthermore that we can output other useful

information. We can output G, r with no further Toffoli cost, since we are already iterating over x_0 . We separate out the QROM for the final N values of x_0 into a separate step.

Depending on the values of the parameters, it may also be suitable to output data for all values of b at once, then use the controlled swaps to swap the desired data into the correct location. That is, use the entire b register like the less significant bits in Q_{roam} . Then the Toffoli cost is instead

$$RC + (B + 1)b_2. \quad (\text{B38})$$

Qubit cost. The persistent ancilla cost is

$$b_2 + b_{k_2} + \lceil \log R \rceil + 2. \quad (\text{B39})$$

This includes b_2 for the QROM output, b_{k_2} for the equal superposition state, and the qubits for storing the output values of G, r . The temporary qubits needed are

$$\left\lceil \log \left(RC \left\lceil \frac{B + 1}{2^{k_2}} \right\rceil \right) \right\rceil + (2^{k_2} - 1)b_2. \quad (\text{B40})$$

Here $(2^{k_2} - 1)b_2$ qubits are used as temporary qubits in the QROM output, and the first term is temporary qubits in the unary iteration.

- (5) In this step we consider the output of the set of rotations. In Fig. 2 we indicate all rotations being output in the same step. It is possible to improve on that approach by taking advantage of the QROM that was used in the preceding step which used iteration over x_0 for the coherent alias sampling. Continuing the unary iteration, we can iterate over the final N values corresponding to D_1 and Q_1 , and output the rotations (as well as the values of G, r).

Next, to output the rotations for SF, perform QROM on the r, b registers controlled by $G_0 = 0$. In previous work there was no advantage to Q_{roam} since the size of the output set of rotations is quite large, so here we cost this as regular QROM (unary iteration) with cost RB . This gives a total Toffoli cost for this step of

$$N + RB. \quad (\text{B41})$$

We also will flip an ancilla qubit to 0 for the case that the identity is needed, corresponding to $b = B$. This only needs one additional Toffoli, because we may iterate over b first then iterate over r . That is, for each value of b in the unary iteration from 0 to $B - 1$, we use it to control an iteration over r . The cost of the controlled iteration over r is

$R - 1$, which is multiplied by B for the number of values of b where we iterate over r . The total number of values of b that we iterate over is $B + 1$, with no iteration over r needed for $b = B$. The cost of the iteration over $B + 1$ values of b is B (including one Toffoli for the control on $G_0 = 0$ to select the SF part of the Hamiltonian), for a total of

$$B(R - 1) + (B - 1) = RB. \quad (\text{B42})$$

Together with the N cost for D_1 and Q_1 , the cost is $N + RB$ as above.

Qubit cost. The persistent ancilla cost is, for the output rotations

$$(N - 1)b_{\text{rot}}. \quad (\text{B43})$$

We do not include another ancilla cost for G , r output here as that was accounted for the preceding step. There is a temporary ancilla cost of

$$\lceil \log(N + RC) \rceil, \quad (\text{B44})$$

for the first unary iteration, and

$$\lceil \log R \rceil + \lceil \log(B + 1) \rceil, \quad (\text{B45})$$

for the second unary iteration. The total temporary ancilla cost here is the maximum of these. (Part of the temporary ancillas are used from the preceding step to save Toffolis.)

- (6) Perform the rotations for the select part of the block-encoding. For N orbitals with b_{rot} bits of accuracy for the rotations, the Toffoli cost of the forward and reverse Givens rotations is

$$2(N - 1)b_{\text{rot}}. \quad (\text{B46})$$

Qubit cost: There is just a temporary ancilla cost of b_{rot} for the additions into the phase gradient register.

- (7) The second QROM from Item 5 can be erased with phase fixup cost

$$R + B. \quad (\text{B47})$$

This can be achieved by performing binary-to-unary conversion on one register, then unary iteration on the other register with controlled phases on the unary qubits. There is a temporary ancilla cost of $\min(R, B)$, with the minimum corresponding to which register we perform the binary to unary conversion on. The steps from here on have negligible temporary ancilla costs as compared to the earlier steps, so will be omitted from this discussion.

- (8) Invert the coherent alias sampling from Item 4. Inverting the controlled swaps and inequality tests has cost b_2 . For the phase fixup, we note that typically RC is larger than B , and also that the method of phase fixup from Ref. [17] is to perform binary-to-unary conversion on one subset of qubits, and unary iteration on the other set of qubits (see Fig. 6 of Ref. [17]). Here, we can perform binary-to-unary conversion on the register containing B , and unary iteration on the x_0 register. That gives the phase fixup cost as

$$RC + B + 1. \quad (\text{B48})$$

We can also combine some of the qubits from x_0 with B in this procedure to give a cost

$$\left\lceil \frac{RC}{2^{k_4}} \right\rceil + 2^{k_4}(B + 1). \quad (\text{B49})$$

Similar to the procedure for the QROM, we can perform iteration over the final N values, but here to perform the phase fixup. This gives a total cost (with the inversion of swaps and inequality test)

$$b_2 + N + \left\lceil \frac{RC}{2^{k_4}} \right\rceil + 2^{k_4}(B + 1). \quad (\text{B50})$$

In the case that RC is divisible by some power of 2, we can use a similar phase fixup approach for N to give cost

$$b_2 + \left\lceil \frac{N}{2^{k'_4}} \right\rceil + 2^{k'_4} + \left\lceil \frac{RC}{2^{k_4}} \right\rceil + 2^{k_4}(B + 1), \quad (\text{B51})$$

where RC is divisible by $2^{k'_4}$.

- (9) Invert the equal superposition from item 3 with cost

$$4\lceil \log(B + 1) \rceil \quad (\text{B52})$$

Note that all these costs from item 3 to item 9 are multiplied by two for the squared operator.

- (10) Invert the coherent alias sampling from item 2 with cost

$$b_1 + \left\lceil \frac{N + RC}{2^{k_5}} \right\rceil + 2^{k_5}, \quad (\text{B53})$$

where b_1 is for inverting the controlled swaps and inequality test.

- (11) Invert the equal superposition preparation from item 1 with cost again

$$4\lceil \log(N + RC) \rceil. \quad (\text{B54})$$

Further minor savings can be made by retaining ancillas to eliminate the non-Clifford cost of inverting the inequality test and controlled swaps for the coherent alias sampling. That would remove the b_2 cost from item 8 and b_1 cost from item 10, at the expense of retaining the same number of ancillas.

Note that we order the values so that the alternatives for G_{SF} start at $x_0 = 0$. This ensures that we do not use the temporary qubits for QROM at the same time as the output angles. In important cases we can also obtain r, c from the appropriate subset of bits. If C is a power of 2 (like 2 or 4), then using $x_0 = rC + c$ this works, and similarly, if r is a power of 2 and we use $x_0 = cR + r$.

In addition to the block-encoding costs above, a reflection is used to provide the complete qubitised operator. The registers that need be reflected upon are as follows.

- (i) Registers of size $\lceil \log(N + RC) \rceil$ and $\lceil \log(B + 1) \rceil$ used for x_0 and b , respectively.
- (ii) The b_{k1} and b_{k2} qubits used for the equal superposition states (that we preform the inequality test with) in the two steps of coherent alias sampling.
- (iii) Two qubits that are rotated in the preparation in the equal superposition states for x_0 and b .
- (iv) Two qubits for the spin (with one for SF in the inner sum, and one for D_1 and Q_1 in the outer sum).

The reflection on these qubits needs to be controlled by a single qubit for the phase estimation, and each step there is another Toffoli used for the unary iteration on the control register for phase estimation. That gives a Toffoli cost in each step of

$$\lceil \log(N + RC) \rceil + \lceil \log(B + 1) \rceil + b_{k1} + b_{k2} + 2. \quad (\text{B55})$$

APPENDIX C: ROBUST SELECTION OF DFTHC SOLUTIONS

We find that resource estimates for DFTHC solutions to the benchmark molecules are relatively insensitive to a broad range hyperparameters. Based on the heuristics described in Sec. IV, we explore the choice $(R, B, C) \in \mathbb{Z}^+ \times \{\lfloor N/4 \rfloor, \lfloor N/2 \rfloor\} \times \{\lfloor N/2 \rfloor, \lfloor 3N/4 \rfloor, N\}$ and scan over the rank $R \in \mathbb{Z}^+$ to find the solutions that are chemically accurate and minimize the product of λ_{eff} and Toffoli cost of each block encoding. We also explore the “DF-like” and “THC-like” selection $(R, N, 1)$ and $(1, B, B)$, respectively, that also optimizes for E_{gap} . The results tabulated in Table X of Appendix C indicate that except for the noticeably worse DF-like factorization, resources estimates vary by at most 20% across these choices for FeMoco-54 and FeMoco-76.

The E_{reg} hyperparameter can be chosen to be equal to, or slightly larger than the *a priori* known optimal gap E_{gap}^* computed in Sec. III. Moreover, we find that the nonlinear

program Eq. (41) is robust to making poor choices of E_{reg} . To accommodate cases where E_{gap}^* is unavailable, we explore sweeping over the ranges of E_{reg} in Table IV. Guided by the insensitivity of DFTHC to the specific choice of B, C , we chose $(B, C) = (\lfloor 3N/4 \rfloor, \lfloor N/4 \rfloor)$. We find in Table XII of Appendix E that E_{gap}^* can be heuristically estimated from the smallest E_{gap} found by the nonlinear program, even if an unphysical $E_{\text{reg}} < E_{\text{gap}}^*$ is chosen.

Generally speaking, increasing the number of gradient descent steps improved the quality of the solution, manifesting in a smaller Frobenius error and a smaller Λ . However, we found that the improvement between 10^5 and 10^6 steps was noticeable, but not significant. Nevertheless, all the DFTHC results presented in this work we obtained with 10^6 steps. As for the other hyperparameters, λ_{reg} was selected to be smaller than the block-encoding normalization Λ found from poorly converged DFTHC solutions. As a matter of practice, we find that λ_{reg} does not need to be particularly precise. One could, for instance, choose λ_{reg} based on efficiently computable lower bounds [20], such as the $\Delta_\mu/2$ parameter described in Table II. We did not significantly explore varying $\epsilon_{\text{reg}} = 10$ mHa, which was chosen to be significantly smaller than the Frobenius errors of previously reported chemically accurate solutions for DF and THC [7,9].

However, finding the rank R^* of the best DFTHC approximation is not straightforward as the classically computed ϵ_{corr} fluctuates significantly and does not improve monotonically with rank. For instance, best solutions selected by thresholding on the lowest cost with $|\epsilon_{\text{corr}}| \leq 0.3$ mHa in Table X show no clear pattern of what choice of hyperparameters B, C are best. Similar to DF and THC, these fluctuations arise as the DFTHC program Eq. (41) is nonlinear and the final solution is sensitive to hyperparameters and the initial values of, say, u, w . However, we can obtain robust estimates of rank and mitigate the effect of fluctuations by averaging the classically computed ϵ_{corr} over many samples for each (R, B, C) , rather than taking the best value. This is accomplished by converting ϵ_{corr} to a random variable by initializing the DFTHC solver with randomly selected parameters values $\mathcal{P}$ for each sample as described in Sec. IV A. Hence, each classically computed ϵ_{corr} is drawn from a distribution with classically estimated mean $\bar{\epsilon}_{\text{corr}}$ and variance σ_{corr}^2 , such as in Fig. 5, where we classically sample enough ϵ_{corr} to interpolate $R_{\leq \epsilon_{\text{th}}}$ according to the criteria $\sigma_{\text{corr}} \leq \epsilon_{\text{th}}$ and $\bar{\epsilon}_{\text{corr}} \ll \epsilon_{\text{th}}$. Although this leads to a significantly larger rank than R^* , another major advantage is that we may instead bound error contributions to σ_E differently like

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{corr}}^2} + |\bar{\epsilon}_{\text{corr}}| \leq \epsilon_{\text{chem}}. \quad (\text{C1})$$

TABLE X. Best solutions found satisfying $|\epsilon_{\text{corr}}| \leq 0.3$ mHa. We provide the lowest ϵ_{corr} solution found for cases where the target ϵ_{corr} was not attained. O_{Rot} and O_{Coeff} Toffolis count the cost of computing and uncomputing twice. Total Toffolis is the dominant costs in block encoding $\text{BE}[H_{\text{SA}}/\Lambda - 1]$, corresponding to the sum of Rot Toffolis, Coeff Toffolis and $8(N-1)b_{\text{rot}}$ from the dominant cost of two select unitaries.


	(R, B, C)	DF-like $(M, N, 1)$	$C = \lfloor N/4 \rfloor$			$C = \lfloor N/2 \rfloor$			THC-like $(1, M, M)$
			$(M, \lfloor \frac{N}{2} \rfloor, \cdot)$	$(M, \lfloor \frac{3N}{4} \rfloor, \cdot)$	$(M, N, \cdot)$	$(M, \lfloor \frac{N}{2} \rfloor, \cdot)$	$(M, \lfloor \frac{3N}{4} \rfloor, \cdot)$	$(M, N, \cdot)$	
Fe ₂ S ₂ -20o30e	$M_{ \epsilon_{\text{corr}} \leq 0.3 \text{ mHa}}$	85	15	14	13	14	11	12	80
	$\epsilon_{\text{corr}}/\text{mHa}$	0.016	0.938	-0.019	0.253	0.288	0.094	0.187	0.073
	$\epsilon_{\text{in}}/\text{mHa}$	3.0	43.9	22.1	18.7	37.4	20.7	8.5	15.6
	λ/Ha	15.9	19.4	17.5	17.5	19.8	19.5	16.3	16.1
	E_{gap}/Ha	0.988	1.03	1.02	1.028	0.977	0.992	0.984	1.002
	$\lambda_{\text{eff}}/\text{Ha}$	5.52	6.24	5.89	5.9	6.15	6.14	5.58	5.58
	$(b_{\text{coeff}}, b_{\text{rot}})$	(11, 15)	(11, 15)	(11, 15)	(11, 15)	(11, 15)	(11, 15)	(11, 15)	(11, 15)
	Rot Toffolis	6800	600	840	1040	560	660	960	320
	Coeff Toffolis	814	530	622	706	734	786	972	1670
	Select Toffolis	2280	2280	2280	2280	2280	2280	2280	2280
	Total Toffolis	9894	3410	3742	4026	3574	3726	4212	4270
	$\lambda \times$ total Toffolis	1.57×10^5	6.62×10^4	6.55×10^4	7.05×10^4	7.08×10^4	7.27×10^4	6.87×10^4	6.87×10^4
	$\lambda_{\text{eff}} \times$ total Toffolis	5.46×10^4	2.13×10^4	2.20×10^4	2.38×10^4	2.20×10^4	2.29×10^4	2.35×10^4	2.39×10^4
Fe ₄ S ₄ -36o54e	$M_{ \epsilon_{\text{corr}} \leq 0.3 \text{ mHa}}$	105	23	17	15	9	12	12	175
	$\epsilon_{\text{corr}}/\text{mHa}$	-0.289	0.104	-0.297	-0.276	-0.004	0.198	-0.092	-0.188
	$\epsilon_{\text{in}}/\text{mHa}$	38.8	38.3	33.1	29.4	166.9	43.7	23.9	20.5
	λ/Ha	39.6	45.1	49.6	47.2	49.8	46.2	44.3	40.4
	E_{gap}/Ha	2.119	2.082	1.967	1.977	2.099	1.926	1.902	1.946
	$\lambda_{\text{eff}}/\text{Ha}$	12.78	13.55	13.84	13.52	14.31	13.2	12.85	12.39
	$(b_{\text{coeff}}, b_{\text{rot}})$	(13, 17)	(13, 17)	(13, 17)	(13, 17)	(13, 17)	(13, 17)	(13, 17)	(13, 17)
	Rot Toffolis	15 120	1656	1836	2160	648	1296	1728	700
	Coeff Toffolis	1298	1280	1336	1478	1126	1594	1878	4002
	Select Toffolis	4760	4760	4760	4760	4760	4760	4760	4760
	Total Toffolis	21 178	7696	7932	8398	6534	7650	8366	9462
	$\lambda \times$ total Toffolis	8.39×10^5	3.47×10^5	3.93×10^5	3.96×10^5	3.25×10^5	3.53×10^5	3.71×10^5	3.82×10^5
	$\lambda_{\text{eff}} \times$ total Toffolis	2.71×10^5	1.04×10^5	1.10×10^5	1.14×10^5	9.35×10^4	1.01×10^5	1.07×10^5	1.17×10^5
FeMoco-54o54e	$M_{ \epsilon_{\text{corr}} \leq 0.3 \text{ mHa}}$	190	10	7	13	10	8	6	200
	$\epsilon_{\text{corr}}/\text{mHa}$	0.192	0.217	0.11	-0.214	-0.014	-0.166	0.219	-0.094
	$\epsilon_{\text{in}}/\text{mHa}$	36.2	279.4	243.7	80.1	238.3	142.3	136.1	121.5
	λ/Ha	57.9	60.7	61.9	65.2	58.3	61.2	65.3	63.6
	E_{gap}/Ha	3.567	3.772	4.789	3.609	3.691	3.617	3.626	3.607
	$\lambda_{\text{eff}}/\text{Ha}$	20.01	21.06	23.88	21.39	20.42	20.73	21.46	21.11
	$(b_{\text{coeff}}, b_{\text{rot}})$	(9, 16)	(9, 16)	(9, 16)	(9, 16)	(9, 16)	(9, 16)	(9, 16)	(9, 16)
	Rot Toffolis	41 040	1080	1120	2808	1080	1280	1296	800
	Coeff Toffolis	1962	1116	1162	1848	1620	1804	1810	4074
	Select Toffolis	6784	6784	6784	6784	6784	6784	6784	6784
	Total Toffolis	49 786	8980	9066	11 440	9484	9868	9890	11 658
	$\lambda \times$ total Toffolis	2.88×10^6	5.45×10^5	5.61×10^5	7.46×10^5	5.53×10^5	6.04×10^5	6.46×10^5	7.41×10^5
	$\lambda_{\text{eff}} \times$ total Toffolis	9.96×10^5	1.89×10^5	2.16×10^5	2.45×10^5	1.94×10^5	2.05×10^5	2.12×10^5	2.46×10^5

(Table continued)


TABLE X. (Continued)

		$C = \lfloor N/4 \rfloor$				$C = \lfloor N/2 \rfloor$			THC-like (1, M, M)
		DF-like (M, N, 1)	$(M, \lfloor \frac{N}{2} \rfloor, \cdot)$	$(M, \lfloor \frac{3N}{4} \rfloor, \cdot)$	$(M, N, \cdot)$	$(M, \lfloor \frac{N}{2} \rfloor, \cdot)$	$(M, \lfloor \frac{3N}{4} \rfloor, \cdot)$	$(M, N, \cdot)$	
FeMoco-76o113e	$M_{ e_{\text{corr}} \leq 0.3 \text{ mHa}}$	130	27	15	18	16	18	7	220
	$\epsilon_{\text{corr}}/\text{mHa}$	0.257	0.063	0.134	0.359	0.152	0.059	-0.132	-0.038
	$\epsilon_{\text{in}}/\text{mHa}$	287.2	65.7	82.5	43.5	131.0	45.3	120.0	234.1
	λ/Ha	140.6	185.6	179.7	181.7	206.0	162.0	197.2	144.7
	E_{gap}/Ha	6.149	4.529	5.037	4.53	4.571	4.526	4.982	5.335
	$\lambda_{\text{eff}}/\text{Ha}$	41.13	40.75	42.25	40.32	43.15	38.03	44.05	38.93
	$(b_{\text{coeff}}, b_{\text{rot}})$	(9, 15)	(9, 15)	(9, 15)	(9, 15)	(9, 15)	(9, 15)	(9, 15)	(9, 15)
	Rot Toffolis	39 520	4104	3420	5472	2432	4104	2128	880
	Coeff Toffolis	1970	2726	2476	3214	2970	3852	2830	4484
	Select Toffolis	9000	9000	9000	9000	9000	9000	9000	9000
	Total Toffolis	50 490	15 830	14 896	17 686	14 402	16 956	13 958	14 364
	$\lambda \times \text{total Toffolis}$	7.10×10^6	2.94×10^6	2.68×10^6	3.21×10^6	2.97×10^6	2.75×10^6	2.75×10^6	2.08×10^6
	$\lambda_{\text{eff}} \times \text{total Toffolis}$	2.08×10^6	6.45×10^5	6.29×10^5	7.13×10^5	6.22×10^5	6.45×10^5	6.15×10^5	5.59×10^5
XVIII-56o64e	$M_{ e_{\text{corr}} \leq 0.3 \text{ mHa}}$	110	8	6	5	5	5	4	120
	$\epsilon_{\text{corr}}/\text{mHa}$	-0.253	0.013	0.121	0.222	0.27	-0.083	-0.109	-0.181
	$\epsilon_{\text{in}}/\text{mHa}$	62.0	284.2	221.0	191.7	418.2	243.3	184.7	237.3
	λ/Ha	58.1	59.2	61.0	65.1	55.5	59.9	62.8	59.7
	E_{gap}/Ha	2.34	2.425	2.58	2.681	2.569	2.299	2.294	2.554
	$\lambda_{\text{eff}}/\text{Ha}$	16.33	16.77	17.55	18.48	16.69	16.44	16.82	17.28
	$(b_{\text{coeff}}, b_{\text{rot}})$	(7, 12)	(7, 12)	(7, 12)	(7, 12)	(7, 12)	(7, 12)	(7, 12)	(7, 12)
	Rot Toffolis	24 640	896	1008	1120	560	840	896	480
	Coeff Toffolis	1432	992	1082	1138	1114	1404	1446	2256
	Select Toffolis	5280	5280	5280	5280	5280	5280	5280	5280
	Total Toffolis	31 352	7168	7370	7538	6954	7524	7622	8016
	$\lambda \times \text{total Toffolis}$	1.82×10^6	4.24×10^5	4.50×10^5	4.91×10^5	3.86×10^5	4.51×10^5	4.79×10^5	4.79×10^5
	$\lambda_{\text{eff}} \times \text{total Toffolis}$	5.12×10^5	1.20×10^5	1.29×10^5	1.39×10^5	1.16×10^5	1.24×10^5	1.28×10^5	1.38×10^5
Cpd1X-58o63e	$M_{ e_{\text{corr}} \leq 0.3 \text{ mHa}}$	200	9	6	5	8	6	4	120
	$\epsilon_{\text{corr}}/\text{mHa}$	-0.214	0.277	-0.19	-0.019	0.098	-0.206	-0.223	-0.183
	$\epsilon_{\text{in}}/\text{mHa}$	16.0	308.3	359.6	412.0	341.6	195.0	218.9	288.1
	λ/Ha	96.2	97.4	101.4	99.5	90.6	99.4	106.0	93.1
	E_{gap}/Ha	4.815	5.243	5.677	6.324	5.467	5.039	5.207	5.559
	$\lambda_{\text{eff}}/\text{Ha}$	30.05	31.53	33.44	34.9	30.99	31.25	32.82	31.69
	$(b_{\text{coeff}}, b_{\text{rot}})$	(10, 15)	(10, 15)	(10, 15)	(10, 15)	(10, 15)	(10, 15)	(10, 15)	(10, 15)
	Rot Toffolis	46 400	1044	1032	1160	928	1032	928	480
	Coeff Toffolis	2142	1168	1184	1254	1596	1716	1622	2434
	Select Toffolis	6840	6840	6840	6840	6840	6840	6840	6840
	Total Toffolis	55 382	9052	9056	9254	9364	9588	9390	9754
	$\lambda \times \text{total Toffolis}$	5.33×10^6	8.82×10^5	9.18×10^5	9.21×10^5	8.48×10^5	9.53×10^5	9.95×10^5	9.08×10^5
	$\lambda_{\text{eff}} \times \text{total Toffolis}$	1.66×10^6	2.85×10^5	3.03×10^5	3.23×10^5	2.90×10^5	3.00×10^5	3.08×10^5	3.09×10^5

(a)



(b)



(c)




Fig. 5. (Continued).




FIG. 5. For molecule (a) Fe_2S_2 , (b) Fe_4S_4 , (c) FeMoco-54, (d) FeMoco-76, (e) Cpd1X-58, and (f) XVIII-56: (left) CCSD(T) correlation energy error vs rank of DFTHC truncation (small dots). Error bars are the standard deviation of the mean $\bar{\epsilon}_{\text{corr}}$ estimated with 64 samples. Dashed lines at ± 1.6 mHa are at the chemical accuracy threshold. (right) $\bar{\epsilon}_{\text{corr}}$ and estimated standard deviation σ_{corr} of ϵ_{corr} . The shaded region is the threshold $\epsilon_{\text{th}} = 0.7$ mHa used to select the rank of good solution. (below) Additional fits of other parameters relevant to estimation quantum simulation costs.

TABLE XI. Cost for simulating electronic structure to chemical accuracy with DFTHC and spectral amplification with the DFTHC rank R chosen by robust selection with $\sigma_{\text{corr}} \leq 0.7$ mHa. The number of calls to the block encoding is computed as $\lceil \frac{\pi \lambda_{\text{eff}}}{2\sigma_{\text{PEA}}} \rceil$ [14]. We take $\sigma_{\text{PEA}} = 1.4$ mHa and $\sigma_{\text{trunc}} = 0.331$ mHa.

	Fe ₂ S ₂ -20	Fe ₄ S ₄ -36	FeMoco-54	FeMoco-76	p450CPD1X-58	CO ₂ [XVIII]-56
(R, B, C)	(26, 15, 5)	(49, 27, 9)	(24, 40, 13)	(15, 64, 32)	(18, 43, 14)	(7, 42, 14)
$\epsilon_{\text{corr}}/\text{mHa}$	0.6809	0.6893	0.6865	0.6966	0.6597	0.6783
$\epsilon_{\text{fro}}/\text{mHa}$	5.7744	4.3616	42.8963	48.9071	49.6692	181.6920
$(b_{\text{coeff}}, b_{\text{rot}})$	(12, 16)	(14, 18)	(10, 17)	(11, 16)	(11, 17)	(8, 13)
E_{gap}/Ha	0.9924	2.1867	3.5156	6.1981	5.6782	2.5696
$\lambda_{\text{eff}}/\text{Ha}$	6.1604	14.0722	21.4486	50.7167	33.3112	17.9276
C_{walk}	4814	11 493	12 930	17 588	12 897	8554
Qubits	499	926	1201	1560	1285	989
Total Toffoli	3.33×10^7	1.81×10^8	3.11×10^8	1.00×10^9	4.82×10^8	1.72×10^8

We find empirically that $|\bar{\epsilon}_{\text{corr}}|$ concentrates to zero, up to statistical accuracy. Hence, this enables significantly larger choice of $\sigma_{\text{PEA}} = 1.4$ mHa, and $\sigma_{\text{corr}} \leq 0.7$ mHa. The costs of simulation in Table XI under this criteria is not directly comparable to prior work, but its statistical robustness might prevent an unintentional bias of solutions towards good ϵ_{corr} but bad ϵ_{corr}^* .

APPENDIX D: TRUNCATION TO FINITE BITS OF PRECISION

A final source of error σ_{trunc} arises from choosing a finite number of bits of precision b_{rot} and b_{coeff} to represent the u and w tensors in the block encoding. Although previous approaches used a deterministic rounding strategy [7], exploration over yet more hyperparameters ($b_{\text{rot}}, b_{\text{coeff}}, R, B, C$) for each sample would introduce prohibitive cost. We instead use an unbiased randomized rounding strategy which allows us to explore hyperparameters independently. This randomization induces a distribution on ϵ_{trunc} which allows us to use either of the bound

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{trunc}}^2} + |\epsilon_{\text{corr}}| + |\bar{\epsilon}_{\text{trunc}}| \leq \epsilon_{\text{chem}}, \quad (\text{D1})$$

$$\sigma_E \leq \sqrt{\sigma_{\text{PEA}}^2 + \sigma_{\text{corr}}^2 + \sigma_{\text{trunc}}^2} + |\bar{\epsilon}_{\text{corr}}| + |\bar{\epsilon}_{\text{trunc}}| \leq \epsilon_{\text{chem}}. \quad (\text{D2})$$

The change in correlation energy between the truncated and original DFTHC solutions by our randomized rounding strategy is presented in Fig. 6 for some FeMoco-54 solution. See Fig. 7 for data on the other molecules. We see that σ_{trunc} scales like $O(2^{-b_{\text{rot}}, -b_{\text{coeff}}})$ with bits of precision as presented, and $|\bar{\epsilon}_{\text{trunc}}|$ concentrates to zero rapidly, and so we assume it is exactly zero in the following for simplicity. We assume that these solutions are typical for FeMoco-54 to determine bits of precision required. The linear error bound Eq. (43) leaves a budget of $\sigma_{\text{trunc}} \leq 0.830$ mHa, which is achieved with $(b_{\text{coeff}}, b_{\text{rot}}) = (9, 15)$, and the statistical bound Eq. (C1) leave a budget of $\sigma_{\text{trunc}} \leq 0.331$ mHa which is achieved with $(b_{\text{coeff}}, b_{\text{rot}}) = (10, 17)$.




FIG. 6. (top) Change in correlation energy due to randomized rounding of rotation and coefficients to b_{rot} and b_{coeff} bits, respectively, for selected DFTHC solutions with different (R, B, C) . The estimated mean change $\bar{\Delta\epsilon}_{\text{corr}}$ and the standard deviation of the mean are computed with 64 samples for each bit setting. (bottom) Estimated standard deviation of $\Delta\epsilon_{\text{corr}}$ and best fit with respect to $b \in \{b_{\text{rot}}, b_{\text{coeff}}\}$. Chemical accuracy at 1.6 mHa is plotted for reference (dashed line).




Fig. 7. (Continued).




FIG. 7. For molecule (a) Fe_2S_2 , (b) Fe_4S_4 , (c) FeMoco-76, (d) Cpd1X-58, (e) XVIII-56, (f) XVIII-100 and XVIII-150: (top) change in correlation energy due to randomized rounding of rotation and coefficients to b_{rot} and b_{coeff} bits, respectively, for selected DFTHC solutions with different (R, B, C) . The estimated mean change $\overline{\Delta\epsilon_{\text{corr}}}$ and the standard deviation of the mean are computed with 64 samples for each bit setting. (bottom) Estimated standard deviation of $\Delta\epsilon_{\text{corr}}$ and best fit with respect to $b \in \{b_{\text{rot}}, b_{\text{coeff}}\}$. Chemical accuracy at 1.6 mHa is plotted for reference (dashed line).

As discussed in the DFTHC block-encoding Fig. 2 each $\vec{u}_b^{(r)}$ is generated by a sequence of $N - 1$ rotation angles $\vec{\theta}_b^{(r)} \in [0, 2\pi)^{N-1}$ that realize the Majorana rotation from γ_{000} to $\gamma_{\vec{u}_b^{(r)}}$. Discretizing these rotations angles to b_{rot} bits of precision leads to a Toffoli cost of $(N - 1)b_{\text{bits}}$. For every angle $\theta \in [0, 2\pi)$, we use the randomized rounding strategy in [57] by choosing the first $b_{\text{rot}} - 1$ bits to be the integer $x = \lfloor (2^{b_{\text{rot}}-1})\theta/(2\pi) \rfloor$. The last bit is sampled randomly from the Bernoulli distribution $\mathbb{B}(p)$ with probability $p = (2^{b_{\text{rot}}-1})\theta/(2\pi) \bmod 1$. Then the rounded $\theta^{(b_{\text{rot}})}$ is sampled from the distribution $2\pi(x + \mathbb{B}(p))/2^{b_{\text{rot}}-1}$. By a triangle inequality, each sampled vector generated by these truncated rotations deviates by at most $|\vec{u}_b^{(r)b_{\text{rot}}} - \vec{u}_b^{(r)}| = \mathcal{O}(N2^{-b_{\text{rot}}})$, but the expectation $\mathbb{E}[|\vec{u}_b^{(r)b_{\text{rot}}} - \vec{u}_b^{(r)}|] = \mathcal{O}(N4^{-b_{\text{rot}}})$ has effectively double the bits of precision. A key difference from [57] is that we sample each DFTHC sample by rounding once in the entire phase estimation algorithm instead of for every application of the block-encoding circuit.

The coefficients $w_b^{(rc)}$ are implemented by alias sampling [14] to b_{coeff} bits of precision. As seen in the block encoding of $O_{\text{SF},rc}$ in Eq. (29), for each component (r, c) , alias sampling prepares an approximation to the quantum state $[1/|w^{(rc)}|_1] \sum_b \sqrt{|w_b^{(rc)}|} |b\rangle |\text{garb}_b\rangle$. To simplify notation, let us drop the (r, c) superscript. Then in essence, alias sampling on a dimension $B + 1$ state discretizes the distribution w into $M = (B + 1)2^{b_{\text{coeff}}-1}$ bins, where the subtracted bit stores the sign of w . We now describe our randomized rounding strategy that assigns indices b to each bins so that the fraction of bins marked with b is $|w_b|/|w|_1$ in expectation. First, for each b , assign $M_b = \lfloor M|w_b|/|w|_1 \rfloor$ bins. This can leave up to $K \leq B$ unassigned bins. Second, assign the remaining bins by

sampling a Hamming weight K bit-string according to the probability distribution $p_b \propto M|w_b|/|w|_1 \bmod 1$ on $(B + 1)$ bins.

We note that the block encoding of one-body terms $H'_1 = \sum_{pq} h_{pq}^{(1)'} E_{pq}$ can essentially be treated as exact with minimal additional cost, which is not included in our later resources estimates. First, observe that the weights $\sqrt{|w_{\pm}^{(r)}|}$ of one-body terms $O_{D_{1,r}^c}, O_{Q_{1,r}^c}$ in Eq. (28) are prepared by outer Prep alias sampling on only $N + RC$ elements, compared to inner Prep alias sampling for $w_b^{(rc)}$ on $RC(B + 1)$ elements. Second, the bits of precision in outer Prep can be chosen separately from that of inner Prep, and even doubling the bits of precision comes at a very small additive cost. Hence, we may treat the coefficients $w_{\pm}^{(r)}$ and also $|w^{(rc)}|_1$ to be prepared with machine precision. Third, the finite bits of precision b_{rot} means that we block encode an approximate one-body term

$$\begin{aligned} H_1^{(b_{\text{rot}})} &= \frac{1}{2} \sum_{s \in \{+,1\},r} s w_s^{(r)} \sum_{\sigma} \left(i \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 0} \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 1} + \mathbb{I} \right) \\ &= \frac{1}{2} \sum_{pq} \sum_{\sigma} h_{pq}^{(1)'} h_{pq}^{(b_{\text{rot}})} (i \gamma_{p\sigma 0} \gamma_{q\sigma 1} + \mathbb{I}), \end{aligned} \quad (\text{D3})$$

with error $H_1'' = H'_1 - H_1^{(b_{\text{rot}})} = \frac{1}{2} \sum_{pq} \sum_{\sigma} h^{(1)''} \times (i \gamma_{p\sigma 0} \gamma_{q\sigma 1} - 1)$. Fourth, the key idea is that we get a higher accuracy block-encoding of $H^{(1)'}$ by taking a linear combination of $H_1^{(b_{\text{rot}})}$ and H_1'' . This comes at a small additive cost of N additional addresses to inner Prep and RPrep. Using a triangle inequality and the fact $\|\gamma_{\vec{u}\sigma 0}\| = 1$, we see that this residual term is an exponentially small correction in the bits of rotation like

$$\begin{aligned} \|H_1''\| &= \frac{1}{2} \left\| \sum_{s \in \{+,1\},r} s w_s^{(r)} \sum_{\sigma} \left[\left(\gamma_{\vec{u}_{\pm}^{(r)}\sigma 0} - \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 0} \right) \gamma_{\vec{u}_{\pm}^{(r)}\sigma 1} + \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 0} \left(\gamma_{\vec{u}_{\pm}^{(r)}\sigma 1} - \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 1} \right) \right] \right\| \\ &\leq 2 \sum_{s \in \{+,1\},r} |w_s^{(r)}| \|\gamma_{\vec{u}_{\pm}^{(r)}\sigma 0} - \gamma_{\vec{u}_{\pm}^{(r)b_{\text{rot}}}\sigma 0}\| \leq 2 \sum_{s \in \{+,1\},r} |w_s^{(r)}| \|\vec{u}_{\pm}^{(r)} - \vec{u}_{\pm}^{(r)b_{\text{rot}}}\| = \mathcal{O}(\|h^{(1)'}\|_1 N 2^{-b_{\text{rot}}}). \end{aligned} \quad (\text{D4})$$

Hence, the change $\|h^{(1)''}\|_1$ to block-encoding normalization is also small. This process may be nested multiple times if the desired accuracy is still not achieved.


APPENDIX E: ESTIMATING E_{gap}^* WITH DFTHC

For the large XVIII complexes on 100 and 150 orbitals, we obtain an estimate on the upper bound for the SOS gap E_{gap}^* by using the DFTHC ansatz. Instead of choosing the regularization parameter E_{reg} in the DFTHC cost function Eq. (41) to be E_{gap}^* , which is unavailable *a priori* for these

systems, we scan over the range of E_{reg} presented in Table IV. We compute the CCSD(T) energies E_{gs} for DFTHC solutions obtained in this manner to estimate the actual gaps E_{gap} as described in Sec. IV B, and find that all solutions have $E_{\text{gap}} > E_{\text{gap}}^*$, including those where E_{reg} was chosen to be smaller than E_{gap}^* , such as for XVIII-56 and FeMoco-76. As E_{gap} appears to be robust under DFTHC optimization and does not attain unphysical values, we estimate E_{gap}^* as the smallest among all the found E_{gap} . For this dataset, we chose the DFTHC bases

TABLE XII. (top) Best DFTHC solutions found by scanning over E_{reg} and R and selecting those satisfying $|\epsilon_{\text{corr}}| \leq 0.3$ mHa. Three samples are generated for each choice (E_{reg} , R). (bottom) Plot of correlation energy error and computed E_{gap} for all data in this scan. The dashed line is chemical accuracy 1.6 mHa and the solid line is E_{gap}^* from Table II of the main text. For XVIII-100o and XVIII-150o, E_{gap}^* is taken from the solution with the smallest E_{gap} .

	XVIII 56o64e	XVIII 100o100e	XVIII 150o150e	Fe ₂ S ₂ 20o30e	Fe ₄ S ₄ 36o54e	FeMoco 54o54e	FeMoco 76o113e	Cpd1X 58o63e
$(R, B, C) = (\cdot, \lfloor \frac{3N}{4} \rfloor, \lfloor \frac{N}{4} \rfloor)$	(·, 42, 14)	(·, 75, 25)	(·, 112, 37)	(·, 15, 5)	(·, 27, 9)	(·, 40, 13)	(·, 57, 19)	(·, 43, 14)
$R_{ \epsilon_{\text{corr}} \leq 0.3 \text{ mHa}}$	9	8	9	17	16	7	15	6
$\epsilon_{\text{corr}}/\text{mHa}$	-0.259	-0.042	-0.062	-0.236	0.291	0.11	0.134	-0.19
$\epsilon_{\text{in}}/\text{mHa}$	137.1	218.7	351.0	12.4	35.8	243.7	82.5	359.6
Λ/Ha	61.0	155.5	336.1	18.2	46.7	61.9	179.7	101.4
E_{gap}/Ha	2.222	4.565	6.454	1.066	2.19	4.789	5.037	5.677
$\lambda_{\text{eff}}/\text{Ha}$	16.47	37.68	65.87	6.23	14.3	24.36	42.55	33.92
$(b_{\text{coeff}}, b_{\text{rot}})$	(7,12)	(8,16)	(9,16)	(11,15)	(13,17)	(9,16)	(9,15)	(10,15)
Rot Toffolis	1512	3300	4032	1020	1728	1120	3420	1032
Coeff Toffolis	1330	2788	3846	688	1294	1162	2476	1184
Select Toffolis	5280	12 672	19 072	2280	4760	6784	9000	6840
Total Toffolis	8122	18 760	26 950	3988	7782	9066	14 896	9056
$\lambda \times \text{total Toffolis}$	4.95×10^5	2.93×10^6	9.06×10^6	7.26×10^4	3.63×10^5	5.61×10^5	2.68×10^6	9.18×10^5
$\lambda_{\text{eff}} \times \text{total Toffolis}$	9.46×10^4	4.63×10^5	1.26×10^6	1.76×10^4	7.87×10^4	1.56×10^5	4.48×10^5	2.17×10^5



and copies hyperparameters to be $B = \lfloor (3N)/4 \rfloor$ and $C = \lfloor N/4 \rfloor$, respectively, and tabulate in Table XII the solutions found to have the best quantum simulation costs.

APPENDIX F: SPIN-FREE FORMALISM

Here we derive the equality constraints of the semi-definite program to determine the optimal gap and H_{SOS} representation using a spin-free algebra for the sum-of-squares generators. Consider a rank-2 spin-free (SF) particle-hole and hole-particle generators of the form

$$O_{\text{SF}} = \sum_{ij} \left(g_{ij} \sum_{\sigma} a_{j\sigma}^{\dagger} a_{i\sigma} + \bar{g}_{ij} \sum_{\sigma} a_{j\sigma} a_{i\sigma}^{\dagger} \right). \quad (\text{F1})$$

The resulting SOS is not flexible enough to describe the electronic Hamiltonian well, so we also introduce the following rank-1 generators

$$O_{D_{\sigma}} = \sum_i d_{i\sigma} a_{i\sigma}, \quad (\text{F2})$$

$$O_{Q_{\sigma}} = \sum_i q_{i\sigma} a_{i\sigma}^{\dagger}, \quad (\text{F3})$$

and the SOS Hamiltonian is

$$\hat{H}_{\text{SOS}} = O_{\text{SF}}^\dagger O_{\text{SF}} + \sum_{\sigma} (O_{D_{\sigma}}^\dagger O_{D_{\sigma}} + O_{Q_{\sigma}}^\dagger O_{Q_{\sigma}}), \quad (\text{F4})$$

with

$$\begin{aligned} O_{\text{SF}}^\dagger O_{\text{SF}} = & \sum_{i,j,k,l} \left(G_{ij}^{ik} \sum_{\sigma,\tau} a_{i\sigma}^\dagger a_{k\sigma} a_{j\tau}^\dagger a_{l\tau} + G_{jl}^{ik} \sum_{\sigma,\tau} a_{i\sigma}^\dagger a_{k\sigma} a_{l\tau}^\dagger a_{j\tau} \right. \\ & \left. + G_{ij}^{''ki} \sum_{\sigma,\tau} a_{k\sigma} a_{i\sigma}^\dagger a_{j\tau}^\dagger a_{l\tau} + G_{jl}^{''ki} \sum_{\sigma,\tau} a_{k\sigma} a_{i\sigma}^\dagger a_{l\tau}^\dagger a_{j\tau} \right), \end{aligned} \quad (\text{F5})$$

$$O_{D_{\sigma}}^\dagger O_{D_{\sigma}} = \sum_{i,j} D_{j\sigma}^{i\sigma} a_{i\sigma}^\dagger a_{j\sigma}, \quad (\text{F6})$$

$$O_{Q_{\sigma}}^\dagger O_{Q_{\sigma}} = \sum_{i,j} Q_{i\sigma}^{j\sigma} a_{j\sigma}^\dagger a_{i\sigma}. \quad (\text{F7})$$

The matrix representation of $O_{\text{SF}}^\dagger O_{\text{SF}}$ has four blocks, arranged as

$$\mathbf{G}_{\text{SF}} = \begin{pmatrix} \mathbf{G} & \mathbf{G}' \\ \mathbf{G}'' & \mathbf{G}''' \end{pmatrix}. \quad (\text{F8})$$

After bringing the creation and annihilation operators in Eq. (6) to a common order, we can see that

$$\hat{H}_{\text{SOS}}^{(2)} = \sum_{i,j,k,l} (G_{ij}^{ik} - G_{jl}^{ik} - G_{ij}^{''ki} + G_{jl}^{''ki}) \sum_{\sigma,\tau} a_{i\sigma}^\dagger a_{k\sigma} a_{j\tau}^\dagger a_{l\tau}, \quad (\text{F9})$$

$$\begin{aligned} \hat{H}_{\text{SOS}}^{(1)} = & 2 \sum_{i,j} \sum_{\sigma} \left(\sum_p (G_{pp}^{ij} + G_{ji}^{pp} - G_{ij}^{pp} - G_{pp}^{ji}) \right. \\ & \left. + \frac{1}{2} (D_{j\sigma}^{i\sigma} - Q_{i\sigma}^{j\sigma}) \right) a_{i\sigma}^\dagger a_{j\sigma}, \end{aligned} \quad (\text{F10})$$

$$E_{\text{SOS}} = 4 \sum_{pq} G_{qq}^{pp} + \sum_p \sum_{\sigma} Q_{p\sigma}^{p\sigma}. \quad (\text{F11})$$

Note that two-electron terms for a given i, j, k , and l have the same coefficient regardless of the spins associated with the labels. The coefficients of $\hat{H}_{\text{SOS}}^{(2)}$ and $\hat{H}_{\text{SOS}}^{(1)}$ are then equated to the one- and two-electron integrals of the spin-free chemistry Hamiltonian.

APPENDIX G: CORRECTION FOR DIVERGENT DERIVATIVE

In this Appendix, we address the issue that the error propagation formula would predict an unreasonably large improvement arising from the divergent derivative of a

function like $\sqrt{x}$ near zero. It is possible to give a more accurate estimate of the uncertainty by deriving corrections to the error propagation formula. To illustrate this method, we initially consider the estimation of $y = f(x) = \sqrt{x}$, with a normal distribution for the error in measuring y . Then the estimate of x can be taken to be the square of the estimated value of y . If the estimation of y has mean y_0 and standard error σ_y , then the mean value of y^2 will be $y_0 + \sigma_y^2$, assuming the normal distribution. The standard error in x will then be

$$\sigma_x = 2\sigma_y \sqrt{y_0^2 + \sigma_y^2/2}. \quad (\text{G1})$$

In contrast, if we were to assume the usual formula for propagation of uncertainty, we would obtain

$$\sigma_x = \sigma_y \left| \left(\frac{df}{dx} \right) \Big|_{y=y_0} \right|^{-1} \Rightarrow \sigma_x = 2\sigma_y y_0. \quad (\text{G2})$$

That is, the correction replaces y_0 with $\sqrt{y_0^2 + \sigma_y^2/2}$, which means that the uncertainty in x does not approach zero for small y_0 .

More generally, an inverse function $f^{-1}(y)$ can be expanded in a quadratic about y_0 , so that

$$\begin{aligned} x \approx & f^{-1}(y_0) + (y - y_0) \left(\frac{df^{-1}}{dy} \Big|_{y=y_0} \right) \\ & + \frac{1}{2} (y - y_0)^2 \left(\frac{d^2 f^{-1}}{dy^2} \Big|_{y=y_0} \right). \end{aligned} \quad (\text{G3})$$

Then we have the mean value of x

$$\bar{x} \approx f^{-1}(y_0) + \frac{1}{2} \sigma_y^2 \left(\frac{d^2 f^{-1}}{dy^2} \Big|_{y=y_0} \right), \quad (\text{G4})$$

and the variance is

$$\begin{aligned} \langle (x - \bar{x})^2 \rangle \approx & \left\langle \left[(y - y_0) \left(\frac{df^{-1}}{dy} \Big|_{y=y_0} \right) \right. \right. \\ & \left. \left. + \frac{1}{2} [(y - y_0)^2 - \sigma_y^2] \left(\frac{d^2 f^{-1}}{dy^2} \Big|_{y=y_0} \right) \right]^2 \right\rangle \\ \approx & \sigma_y^2 \left(\frac{df^{-1}}{dy} \Big|_{y=y_0} \right)^2 + \frac{1}{4} (\kappa - 1) \sigma_y^4 \left(\frac{d^2 f^{-1}}{dy^2} \Big|_{y=y_0} \right)^2, \end{aligned} \quad (\text{G5})$$

where κ is the kurtosis of the distribution (the fourth standardized moment of the distribution).

In the case where $f(x) = \sqrt{x}$ this gives $4\sigma_y^2 y_0^2 + 2\sigma_y^2$ as above, with the 2 coming from $\kappa - 1$, so will be larger for distributions with excess kurtosis. In Sec. II the function is $\pm \arccos(h/\Lambda - 1)$, which yields

$$\begin{aligned}\left.\frac{df^{-1}}{dy}\right|_{y=y_0} &= -\frac{1}{2}\lambda \sin(y_0) = -\sqrt{h(2\Lambda - h)}, \\ \left.\frac{d^2f^{-1}}{dy^2}\right|_{y=y_0} &= -\frac{1}{2}\lambda \cos(y_0) = \Lambda - h.\end{aligned}\quad (\text{G6})$$

That then gives

$$\sigma_{E_k}^2 = \sigma_{\text{QPE}}^2 E_{\min}(2\Lambda - E_{\min}) + \frac{1}{4}(\kappa - 1)\sigma_{\text{QPE}}^4(2\Lambda - E_{\min})^2. \quad (\text{G7})$$

Using $\kappa = 3$ and approximating small $E_{\min}$ in the last term gives

$$\sigma_{E_k} = \sigma_{\text{QPE}} \sqrt{E_{\min}(2\Lambda - E_{\min}) + \Lambda^2 \sigma_{\text{QPE}}^2 / 2}. \quad (\text{G8})$$

To give an estimate of the parameter range where this correction will be significant, let us solve for σ_{QPE} , and expand for small σ_{E_k} , to give

$$\sigma_{\text{QPE}} \approx \frac{\sigma_{E_k}}{\sqrt{E_{\min}(2\Lambda - E_{\min})}} \left(1 - \frac{\sigma_{E_k}^2 \Lambda^2}{4E_{\min}^2(2\Lambda - E_{\min})^2} \right). \quad (\text{G9})$$

In the cases we are interested in, $\Lambda \gg E_{\min}$, and the correction will be negligible provided

$$E_{\min} \gg \sigma_{E_k}. \quad (\text{G10})$$

If we substitute an initial choice for $\sigma_{\text{QPE}} = \sigma'_{E_k} / \lambda_{\text{eff}}$, where $\lambda_{\text{eff}} = \sqrt{E_{\min}(2\Lambda - E_{\min})}$, then the correction

$$\sigma_{E_k} = \sigma'_{E_k} \sqrt{1 + \frac{\Lambda^2 \sigma_{E_k}^2}{2\lambda_{\text{eff}}^4}} \approx \sigma'_{E_k} \left(1 + \frac{\Lambda^2 \sigma_{E_k}^2}{4\lambda_{\text{eff}}^4} \right). \quad (\text{G11})$$

For typical parameters we consider like $\Lambda \sim 100$ Ha, $\lambda_{\text{eff}} \sim 10$ Ha, $\sigma'_{E_k} \sim 1$ mHa, the fractional correction to σ_{E_k} is negligible at $\sim 2.5 \times 10^{-7}$.

-
- [1] A. Y. Kitaev, *Quantum measurements and the Abelian stabilizer problem*, [arXiv:quant-ph/9511026](#).
[2] D. S. Abrams and S. Lloyd, *Simulation of many-body Fermi systems on a universal quantum computer*, *Phys. Rev. Lett.* **79**, 2586 (1997).
[3] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Simulated quantum computation of molecular energies*, *Science* **309**, 1704 (2005).
[4] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, *Gate-count estimates for performing quantum chemistry on small quantum computers*, *Phys. Rev. A* **90**, 022305 (2014).

- [5] E. Campbell, *Random compiler for fast Hamiltonian simulation*, *Phys. Rev. Lett.* **123**, 070503 (2019).
[6] I. D. Kivlichan, C. E. Granade, and N. Wiebe, *Phase estimation with randomized Hamiltonians*, [arXiv:1907.10070](#).
[7] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, *Even more efficient quantum computations of chemistry through tensor hypercontraction*, *PRX Quantum* **2**, 030305 (2021).
[8] A. Caesura, C. L. Cortes, W. Poll, S. Sim, M. Steudtner, G.-L. R. Anselmetti, M. Degroote, N. Moll, R. Santagati, M. Streif *et al.*, *Faster quantum chemistry simulations on a quantum computer with improved tensor factorization and active volume compilation*, [arXiv:2501.06165](#).
[9] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, *Quantum computing enhanced computational catalysis*, *Phys. Rev. Res.* **3**, 033055 (2021).
[10] D. W. Berry, Y. Tong, T. Khattar, A. White, T. I. Kim, G. H. Low, S. Boixo, Z. Ding, L. Lin, S. Lee, G. Kin-Lic Chan, R. Babbush, and N. C. Rubin, *Rapid initial-state preparation for the quantum simulation of strongly correlated molecules*, *PRX Quantum* **6**, 020327 (2025).
[11] J. J. Goings, A. White, J. Lee, C. S. Tautermann, M. Degroote, C. Gidney, T. Shiozaki, R. Babbush, and N. C. Rubin, *Reliably assessing the electronic structure of cytochrome P450 on today's classical computers and tomorrow's quantum computers*, *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2203533119 (2022).
[12] N. C. Rubin, D. W. Berry, F. D. Malone, A. F. White, T. Khattar, A. Eugene DePrince, S. Siculo, M. Kühn, M. Kaicher, J. Lee, and R. Babbush, *Fault-tolerant quantum simulation of materials using Bloch orbitals*, *PRX Quantum* **4**, 040303 (2023).
[13] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, *Quantum algorithm for spectral measurement with a lower gate count*, *Phys. Rev. Lett.* **121**, 010501 (2018).
[14] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, *Encoding electronic spectra in quantum circuits with linear T complexity*, *Phys. Rev. X* **8**, 041015 (2018).
[15] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, and R. Babbush, *Improved techniques for preparing eigenstates of fermionic Hamiltonians*, *npj Quantum Inf.* **4**, 22 (2018).
[16] G. H. Low and I. L. Chuang, *Hamiltonian simulation by qubitization*, *Quantum* **3**, 163 (2019).
[17] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, *Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization*, *Quantum* **3**, 208 (2019).
[18] I. Loaiza and A. F. Izmaylov, *Block-invariant symmetry shift: Preprocessing technique for second-quantized Hamiltonians to improve their decompositions to linear combination of unitaries*, *J. Chem. Theory Comput.* **19**, 8201 (2023).
[19] S. Patel, A. S. Brahmachari, J. T. Cantin, L. Wang, and A. F. Izmaylov, *Global minimization of electronic Hamiltonian l -norm via linear programming in the block invariant symmetry Shift (BLISS) method*, *J. Chem. Theory Comput.* **21**, 703 (2025).

- [20] C. L. Cortes, D. Rocca, J. F. Gonthier, P. J. Ollitrault, R. M. Parrish, Gian-Luca R. Anselmetti, M. Degroote, N. Moll, R. Santagati, and M. Streif, *Assessing the query complexity limits of quantum phase estimation using symmetry-aware spectral bounds*, *Phys. Rev. A* **110**, 022420 (2024).
- [21] R. D. Somma and S. Boixo, *Spectral gap amplification*, *SIAM J. Comput.* **42**, 593 (2013).
- [22] G. H. Low and I. L. Chuang, *Hamiltonian simulation by uniform spectral amplification*, [arXiv:1707.05391](https://arxiv.org/abs/1707.05391).
- [23] A. Zlokapa and R. D. Somma, *Hamiltonian simulation for low-energy states with optimal time dependence*, *Quantum* **8**, 1449 (2024).
- [24] S. Gu, R. D. Somma, and B. Şahinoğlu, *Fast-forwarding quantum evolution*, *Quantum* **5**, 577 (2021).
- [25] 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*, STOC 2019 (Association for Computing Machinery, New York, NY, USA, 2019), p. 193–204.
- [26] A. N. Chowdhury and R. D. Somma, *Quantum algorithms for Gibbs sampling and hitting-time estimation*, *Quantum Inf. Comput.* **17**, 41 (2017).
- [27] A. W. Harrow, A. Hassidim, and S. Lloyd, *Quantum algorithm for linear systems of equations*, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [28] A. J. Coleman, *The convex structure of electrons*, *Int. J. Quantum Chem.* **11**, 907 (1977).
- [29] D. A. Mazziotti, *Quantum chemistry without wave functions: Two-electron reduced density matrices*, *Acc. Chem. Res.* **39**, 207 (2006).
- [30] S. Pironio, M. Navascués, and A. Acín, *Convergent relaxations of polynomial optimization problems with noncommuting variables*, *SIAM J. Optim.* **20**, 2157 (2010).
- [31] H. Kummer, *n-representability problem for reduced density matrices*, *J. Math. Phys. (N.Y.)* **8**, 2063 (1967).
- [32] R. M. Erdahl, *Representability*, *Int. J. Quantum Chem.* **13**, 697 (1978).
- [33] J. Eisert, *A note on lower bounds to variational problems with guarantees*, [arXiv:2301.06142](https://arxiv.org/abs/2301.06142).
- [34] W. Hall, *Compatibility of subsystem states and convex geometry*, *Phys. Rev. A* **75**, 032102 (2007).
- [35] D. A. Mazziotti, *Structure of fermionic density matrices: Complete n-representability conditions*, *Phys. Rev. Lett.* **108**, 263002 (2012).
- [36] Y.-K. Liu, *Consistency of local density matrices is qma-complete*, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2006 and 10th International Workshop on Randomization and Computation, RANDOM 2006, Barcelona, Spain, 2006. Proceedings* (Springer, New York, 2006), pp. 438–449.
- [37] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, *Elucidating reaction mechanisms on quantum computers*, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7555 (2017).
- [38] Z. Li, J. Li, N. S. Dattani, C. Umrigar, and G. K. Chan, *The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations*, *J. Chem. Phys.* **150**, 024302 (2019).
- [39] D. Rocca, C. L. Cortes, J. F. Gonthier, P. J. Ollitrault, R. M. Parrish, G.-L. Anselmetti, M. Degroote, N. Moll, R. Santagati, and M. Streif, *Reducing the runtime of fault-tolerant quantum simulations in chemistry through symmetry-compressed double factorization*, *J. Chem. Theory Comput.* **20**, 4639 (2024).
- [40] Z. Li and G. K.-L. Chan, *Spin-projected matrix product states: Versatile tool for strongly correlated systems*, *J. Chem. Theory Comput.* **13**, 2681 (2017).
- [41] M. Mörchén, G. H. Low, T. Weymuth, H. Liu, M. Troyer, and M. Reiher, *Classification of electronic structures and state preparation for quantum computation of reaction chemistry*, [arXiv:2409.08910](https://arxiv.org/abs/2409.08910).
- [42] A. M. Childs, *On the relationship between continuous- and discrete-time quantum walk*, *Commun. Math. Phys.* **294**, 581 (2010).
- [43] D. W. Berry, A. M. Childs, and R. Kothari, *Hamiltonian simulation with nearly optimal dependence on all parameters*, in *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on* (IEEE, New York, 2015), pp. 792–809, [10.1109/FOCS.2015.54](https://doi.org/10.1109/FOCS.2015.54).
- [44] T. S. Motzkin, *The arithmetic-geometric inequality*, *Inequalities* (Proc. Sympos. Wright-Patterson Air Force Base, Ohio, 1965); 205, 54 (1967).
- [45] J. W. Helton, *“Positive” noncommutative polynomials are sums of squares*, *Ann. Math.* **156**, 675 (2002).
- [46] D. A. Mazziotti, *Variational minimization of atomic and molecular ground-state energies via the two-particle reduced density matrix*, *Phys. Rev. A* **65**, 062511 (2002).
- [47] N. C. Rubin and D. A. Mazziotti, *Comparison of one-dimensional and quasi-one-dimensional Hubbard models from the variational two-electron reduced-density-matrix method*, *Theor. Chem. Acc.* **133**, 1492 (2014).
- [48] J. W. Mullinax, E. Maradzike, L. N. Koulias, M. Mostafanejad, E. Epifanovsky, G. Gidofalvi, and A. E. DePrince III, *Heterogeneous CPU + GPU algorithm for variational two-electron reduced-density matrix-driven complete active-space self-consistent field theory*, *J. Chem. Theory Comput.* **15**, 6164 (2019).
- [49] Q. Han, Z. Lin, H. Liu, C. Chen, Q. Deng, D. Ge, and Y. Ye, *Accelerating low-rank factorization-based semidefinite programming algorithms on GPU*, [arXiv:2407.15049](https://arxiv.org/abs/2407.15049).
- [50] Q. Han, C. Li, Z. Lin, C. Chen, Q. Deng, D. Ge, H. Liu, and Y. Ye, *A low-rank ADMM splitting approach for semidefinite programming*, [arXiv:2403.09133](https://arxiv.org/abs/2403.09133).
- [51] A. M.-C. So, Y. Ye, and J. Zhang, *A unified theorem on SDP rank reduction*, *Math. Oper. Res.* **33**, 910 (2008).
- [52] I. Waldspurger and A. Waters, *Rank optimality for the Burer-Monteiro factorization*, *SIAM J. Optim.* **30**, 2577 (2020).
- [53] J. Povh, F. Rendl, and A. Wiegele, *A boundary point method to solve semidefinite programs*, *Computing* **78**, 277 (2006).
- [54] M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, and G. K.-L. Chan, *Low rank representations for quantum simulation of electronic structure*, *npj Quantum Inf.* **7**, 83 (2021).
- [55] O. Oumarou, M. Scheurer, R. M. Parrish, E. G. Hohenstein, and C. Gogolin, *Accelerating quantum computations of chemistry through regularized compressed double factorization*, *Quantum* **8**, 1371 (2024).

- [56] G. H. Low, V. Kliuchnikov, and L. Schaeffer, *Trading T gates for dirty qubits in state preparation and unitary synthesis*, [Quantum](#) **8**, 1375 (2024).
- [57] G. H. Low, *Halving the cost of quantum multiplexed rotations*, [arXiv:2110.13439](#).
- [58] S. Burer and R. D. Monteiro, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, [Math. Program., Ser. B](#) **95**, 329 (2003).
- [59] M. Nakata, H. Nakatsuji, M. Ehara, M. Fukuda, K. Nakata, and K. Fujisawa, *Variational calculations of fermion second-order reduced density matrices by semidefinite programming algorithm*, [J. Chem. Phys.](#) **114**, 8282 (2001).
- [60] M. B. Hastings, *Perturbation theory and the sum of squares*, [arXiv:2205.12325](#).
- [61] C. Gidney, N. Shutty, and C. Jones, *Magic state cultivation: Growing T states as cheap as CNOT gates*, [arXiv:2409.17595](#).
- [62] M. B. Hastings, *Improving perturbation theory with the sum-of-squares: Third order*, [arXiv:2412.03564](#).
- [63] M. B. Hastings, *Field theory and the sum-of-squares for quantum systems*, [arXiv:2302.14006](#).
- [64] Y. R. Sanders, D. W. Berry, P. C. S. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, *Compilation of fault-tolerant quantum heuristics for combinatorial optimization*, [PRX Quantum](#) **1**, 020312 (2020).
- [65] N. Rubin, *Fast quantum simulation of electronic structure by spectral amplification*, Zenodo (2025), [10.5281/zenodo.17066718](#).