Quantum Imaginary Time Evolution, Quantum Lanczos, and Quantum Thermal Averaging

Mario Motta,[1, *] Chong Sun,[1] Adrian T. K. Tan,[2] Matthew J. O'Rourke,[1] Erika Ye,[2]
Austin J. Minnich,[2] Fernando G. S. L. Brandão,[3,4] and Garnet Kin-Lic Chan[1, †]

[1]*Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125, USA*
[2]*Division of Engineering and Applied Sciences, California Institute of Technology, Pasadena, CA 91125, USA*
[3]*Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA 91125, USA*
[4]*Google LLC, Venice, CA 90291, USA*

An efficient way to compute Hamiltonian ground-states on a quantum computer stands to impact many problems in the physical and computer sciences, from quantum simulation to machine learning. Existing techniques, such as phase estimation and variational algorithms, display potential disadvantages, including requirements for deep circuits with ancillae and high-dimensional optimization. Here we describe the quantum imaginary time evolution and quantum Lanczos algorithms, analogs of classical algorithms for ground (and excited) states, but with exponentially reduced space and time requirements per iteration, and avoiding deep circuits with ancillae and high-dimensional optimization. We discuss quantum imaginary time evolution as a natural subroutine to generate Gibbs averages through an analog of minimally entangled typical thermal states. We implement these algorithms with exact classical emulation and prototype circuits on the Rigetti quantum virtual machine and Aspen-1 quantum processing unit, demonstrating the power of quantum elevations of classical algorithms.

An important application for a quantum computer is to compute the ground-state Ψ of a Hamiltonian $\hat{H}$ [1, 2]. This arises in simulations, for example, of the electronic structure of molecules and materials, [3–6] as well as in more general optimization problems. While efficient ground-state determination cannot be guaranteed for all Hamiltonians, as this is a QMA-hard problem [7], several heuristic quantum algorithms have been proposed, including adiabatic state preparation with quantum phase estimation [8, 9] (QPE) and quantum-classical variational algorithms, such as the quantum approximate optimization algorithm [10–12] and variational quantum eigensolver [13–15]. Despite many advances, these algorithms also have potential disadvantages, especially in the context of near-term quantum computing architectures with limited quantum resources. For example, phase estimation produces a nearly exact eigenstate, but appears impractical without error correction, while variational algorithms, although somewhat robust to coherent errors, are limited in accuracy for a fixed Ansatz, and involve a high-dimensional noisy classical optimization [16].

In classical simulations, different strategies are employed to numerically determine nearly exact ground-states. One popular approach is imaginary-time evolution, which expresses the ground-state as the long-time limit of the imaginary-time Schrödinger equation $-\partial_\beta |\Phi(\beta)\rangle = \hat{H}|\Phi(\beta)\rangle$, $|\Psi\rangle = \lim_{\beta \to \infty} \frac{|\Phi(\beta)\rangle}{\||\Phi(\beta)\|}$ (for $\langle \Phi(0)|\Psi\rangle \neq 0$). Unlike variational algorithms with a fixed Ansatz, imaginary-time evolution always converges to the ground-state, as distinguished from imaginary-time Ansatz optimization [17]. Another common algorithm is the iterative Lanczos algorithm [18] and its variants. The Lanczos iteration constructs the Hamiltonian matrix $\mathbf{H}$ in a Krylov subspace $\{|\Phi\rangle, \hat{H}|\Phi\rangle, \hat{H}^2|\Phi\rangle \dots\}$; diagonalizing $\mathbf{H}$ yields a variational estimate of the ground-state which tends to $|\Psi\rangle$ for a large number of iterations. For an N-qubit Hamiltonian, the classical complexity of imaginary time evolution and Lanczos algorithm scales as $\sim \exp(\mathcal{O}(N))$ in space and time. Exponential space comes from storing $\Phi(\beta)$

or the Lanczos vector, while exponential time comes from the cost of Hamiltonian multiplication $\hat{H}|\Phi\rangle$, as well as, in principle, though not in practice, the N-dependence of the number of propagation steps or Lanczos iterations. Thus it is natural to consider quantum versions of these algorithms that can overcome the exponential bottlenecks.

Here we describe the quantum imaginary time evolution (QITE) and the quantum Lanczos (QLanczos) algorithms to determine ground-states (and excited states in the case of QLanczos) on a quantum computer. As we show, under well defined assumptions, these use exponentially reduced space and time per propagation step or iteration compared to their direct classical counterparts. They also offer advantages over existing ground-state quantum algorithms as they do not use deep circuits and are guaranteed to converge to the ground-state without non-linear optimization. We further describe inexact QITE and QLanczos algorithms that present a hierarchy of approximations to apply within a limited computational budget. A crucial common component is the efficient implementation of the non-Hermitian operation of an imaginary-time step $e^{-\Delta\tau \hat{H}}$ (for small $\Delta\tau$) assuming a finite correlation length in the state. Non-Hermitian operations are not natural on a quantum computer and are usually achieved using ancillae and postselection, but we describe how to implement imaginary time evolution on a given state without these resources. The lack of ancillae and complex circuits make QITE and QLanczos potentially suitable for near-term quantum architectures. Using the QITE algorithm, we show how to sample from thermal (Gibbs) states, also without deep circuits or ancillae as is usually the case, via a quantum analog of the minimally entangled typical thermal states (QMETTS) algorithm [19, 20]. We demonstrate the algorithms on spin and fermionic Hamiltonians (short- and long-range spin and Hubbard models, MAXCUT optimization, and dihydrogen minimal molecular model) using exact classical emulation, and demonstrate proof-of-concept implementations on the Rigetti quantum virtual machine (QVM) and Aspen-1 quantum pro-

cessing units (QPUs).

Quantum Imaginary-Time Evolution. Define a geometric k-local Hamiltonian $\hat{H} = \sum_m \hat{h}_m$ (where each term $\hat{h}_m$ acts on at most k neighbouring qubits on an underlying graph) and a Trotter decomposition of the corresponding imaginary-time evolution,

$$e^{-\beta \hat{H}} = (e^{-\Delta\tau \hat{h}_1} e^{-\Delta\tau \hat{h}_2} \ldots)^n + \mathcal{O}(\Delta\tau) \, ; \; n = \frac{\beta}{\Delta\tau} \quad (1)$$

applied to a state $|\Psi\rangle$. After a single Trotter step, we have

$$|\Psi'\rangle = e^{-\Delta\tau \hat{h}_m} |\Psi\rangle. \quad (2)$$

The basic idea is that the normalized state $|\bar{\Psi}'\rangle = |\Psi'\rangle / \|\Psi'\|$ is generated from $|\Psi\rangle$ by a unitary operator $e^{-i\Delta\tau \hat{A}[m]}$ acting on a neighbourhood of the qubits acted on by $\hat{h}_m$, where $\hat{A}[m]$ can be determined from tomography of $|\Psi\rangle$ in this neighbourhood up to controllable errors. This is illustrated by the simple example where $|\Psi\rangle$ is a product state. The squared norm $c = \|\Psi'\|^2$ can be calculated from the expectation value of $\hat{h}_m$, requiring measurements over k qubits,

$$c = \langle\Psi| e^{-2\Delta\tau \hat{h}[m]} |\Psi\rangle = 1 - 2\Delta\tau \langle\Psi|\hat{h}_m|\Psi\rangle + \mathcal{O}(\Delta\tau^2) \quad (3)$$

Because $|\Psi\rangle$ is a product state, $|\Psi'\rangle$ is obtained applying the unitary operator $e^{-i\Delta\tau \hat{A}[m]}$ also on k qubits. $\hat{A}[m]$ can be expanded in terms of an operator basis, e.g. the Pauli basis $\{\sigma_i\}$ on k qubits,

$$\hat{A}[m] = \sum_{i_1 i_2 \ldots i_k} a[m]_{i_1 i_2 \ldots i_k} \sigma_{i_1} \sigma_{i_2} \ldots \sigma_{i_k}. \quad (4)$$

Up to $\mathcal{O}(\Delta\tau)$, the coefficients $a[m]_{i_1 i_2 \ldots i_k}$ are defined by the linear system $\mathbf{S}\mathbf{a}[m] = \mathbf{b}$ where the elements of $\mathbf{S}$ and $\mathbf{b}$ are expectation values over k qubits,

$$S_{i_1 i_2 \ldots i_k, i'_1 i'_2 \ldots i'_k} = \langle\Psi| \sigma^\dagger_{i_1} \sigma^\dagger_{i_2} \ldots \sigma^\dagger_{i_k} \sigma_{i'_1} \sigma_{i'_2} \ldots \sigma_{i'_k} |\Psi\rangle$$
$$b_{i_1 i_2 \ldots i_k} = -i \, c^{-\frac{1}{2}} \langle\Psi| \sigma^\dagger_{i_1} \sigma^\dagger_{i_2} \ldots \sigma^\dagger_{i_k} \hat{h}[m] |\Psi\rangle \quad (5)$$

In general, $\mathbf{S}$ has a null space; to ensure $\mathbf{a}[m]$ is real, we minimize $\|c^{-1/2}\Psi' - (1 - i\Delta\tau \hat{A}[m])\Psi\|^2$ w.r.t. real variations in $\mathbf{a}[m]$ (see SI). Because the solution is determined from a linear problem, there are no local minima.

In this simple case, the normalized result of the imaginary time evolution step could be represented by a unitary update over k qubits, because $|\Psi\rangle$ had correlation length zero. After the initial step, this is no longer the case. However, for a more general $|\Psi\rangle$ with finite correlations over at most C qubits (i.e. correlations between observables separated by distance L are bounded by $\exp(-L/C)$), $|\Psi'\rangle$ can be generated by a unitary acting on a domain of width at most $O(C)$ qubits surrounding the qubits acted on by $\hat{h}_m$ (this follows from Uhlmann's theorem [21]; see SI). The unitary $e^{-i\Delta\tau A[m]}$ can then be determined by measurements and solving the least squares problem in this domain (Fig. 1). For example, for a nearest-neighbor

local Hamiltonian on a d-dimension cubic lattice, the domain size D is bounded by $O(C^d)$. In many physical systems, we expect the maximum correlation length throughout the Trotter steps to increase with β and saturate for $C_{\max} \ll N$ [22]. Fig. 1 shows the mutual information between qubits i and j as a function of imaginary time in the 1D and 2D ferromagnetic transverse field Ising models computed by tensor network simulation (see SI), demonstrating a monotonic increase and clear saturation.

The above replacement of imaginary time evolution steps by unitary updates can be extended to more general Hamiltonians, such as ones with long-range interactions and fermionic Hamiltonians. For example, for a Hamiltonian with long-range pairwise terms, the action of $e^{-\Delta\tau \hat{h}[m]}$ (if $\hat{h}[m]$ acts on qubits i and j) can be emulated by a unitary constructed in the neighborhoods of i and j, over a domain of $(2C\log(1/\delta))^k$ sites (see SI). The assumption of finite correlation length, however, is less natural for such Hamiltonians. For fermions, the locality of the corresponding qubit Hamiltonian depends on the spin mapping. In principle, a geometric k-local fermionic Hamiltonian can be mapped to a geometric local qubit Hamiltonian [23], allowing the above techniques to be directly applied. Alternatively, we conjecture that by using a fermionic unitary, where the Pauli basis in Eq. (4) is replaced by the fermionic operator basis $\{1, \hat{a}, \hat{a}^\dagger, \hat{a}^\dagger \hat{a}\}$, the unitary update can be constructed over a domain size $D \sim O(C^d)$ where C is the fermionic correlation length.

Cost of QITE. The number of measurements and classical storage at a given time step (starting propagation from a product state) is bounded by $\exp(O(C^d))$ (with C the correlation length at that time step), since each unitary at that step acts on at most $O(C^d)$ sites; classical solution of the least squares problem has a similar scaling $\exp(O(C^d))$, as does the synthesis and application as a quantum circuit (composed of two-qubit gates) of the unitary $e^{-i\Delta\tau A[m]}$. Thus, space and time requirements are bounded by exponentials in C^d, but are polynomial in N when one is interested in a local approximation of the state (or quasi-polynomial for a global approximation); the polynomial in N comes from the number of terms in H; see SI for details).

The exponential dependence on C^d can be greatly reduced in many cases. Suppose the Hamiltonian $A[m]$ of the unitary update has a locality structure, i.e. it is (approximately) a p-local Hamiltonian (i.e. in Eq. (4), all $a[m]_{i_1 \ldots i_k}$ coefficients are zero except for those where at most p of the σ_i operators are different from the identity). Then the cost of tomography becomes only $C^{O(dp)}$, while the cost of finding and implementing the unitary is $O(pC^d T_e)$, with T_e the cost to compute one entry of $A[m]$ [24]. If we assume further that $A[m]$ is geometric local, the cost of tomography is reduced further to $O(pC^d)$. However, even if C is too large to construct the unitaries exactly, we can still run the algorithm as a heuristic, truncating the unitary updates to domain sizes that fit the computational budget. This gives the inexact QITE algorithm, described further below.

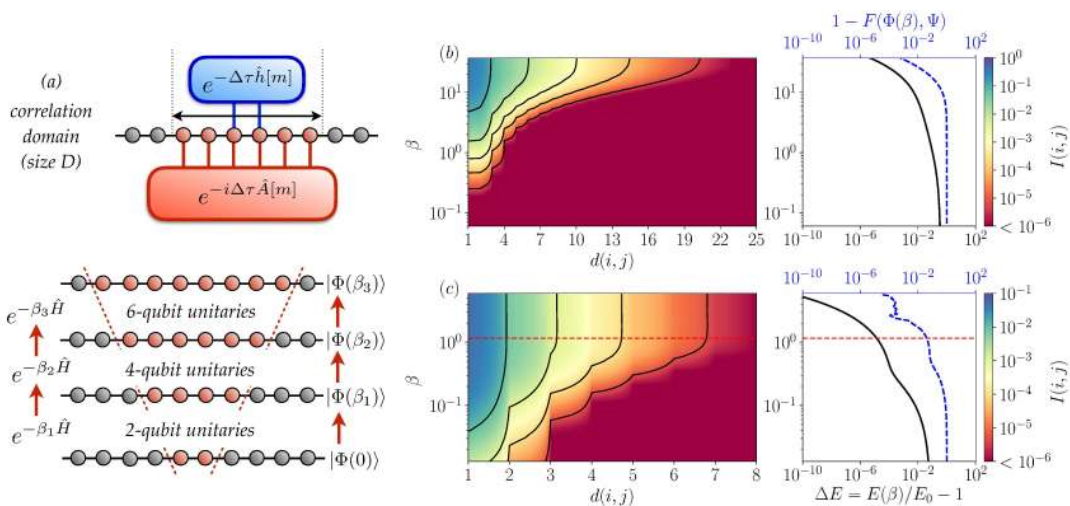*Comparison to classical implementations*. Compared to a di-

FIG. 1: (color online) (a) Schematic of the QITE algorithm. Top: imaginary-time evolution under a geometric k-local operator $\hat{h}[m]$ can be reproduced by a unitary operation acting on $D > k$ qubits. Bottom: exact imaginary-time evolution starting from a product state requires unitaries acting on a domain D that grows with correlations. (b,c) Left: mutual information $I(i,j)$ between qubits i, j as a function of distance $d(i,j)$ and imaginary time β, for a 1D (b) and a 2D (c) FM transverse-field Ising model, with $h = 1.25$ (1D) and $h = 3.5$ (2D). $I(i,j)$ saturates at longer times. Right: relative error in the energy ΔE and fidelity $F = |\langle \Phi(\beta)|\Psi\rangle|^2$ between the finite-time state $\Phi(\beta)$ and infinite-time state Ψ as a function of β. The noise in the 2D fidelity error at large β arises from the approximate nature of the algorithm used. See SI for details.

rect classical implementation of imaginary time evolution, the cost of a QITE time-step (for bounded correlation length C) is linear in N in space and polynomial in N in time, thus giving an exponential reduction in space and time. We can also compare to other classical algorithms. As QITE defines a quantum circuit for the imaginary time evolution, we could attempt to use it for a faster classical simulation. If we are only interested in local observables, we can apply the circuit in the Heisenberg picture in a classical emulation. However, this gives an extra exponential dependence on the number of previous time-steps: After the unitaries associated to $(e^{-\Delta\tau\hat{h}_1}e^{-\Delta\tau\hat{h}_2}\ldots)^l$ have been applied, the cost of applying the next unitary scales as $\exp(O(lD))$, with D the domain size of the unitaries, instead of $\exp(O(D))$ in QITE. Alternatively, if $|\Psi\rangle$ is represented by a tensor network in a classical simulation, then $e^{-\Delta\tau\hat{h}[m]}|\Psi\rangle$ can be represented as a classical tensor network with increased bond dimension [25, 26]. However, the bond dimension will scale as $\exp(O(lD))$. Apart from the extra exponential dependence on l, a further potential drawback in this approach is that we cannot guarantee contracting the resulting classical tensor network for an observable is efficient; it is a #P-hard problem in the worst case in 2D (and even in the average case for Gaussian distributed tensors) [27, 28]. Finally, we can compare QITE with bounded C with the classical heuristic of truncating the problem size at the correlation length C_0 of the ground-state and solving by exact diagonalization, which can be done in time $\exp(O(C_0^d))$ in d spatial dimensions. While this is a competitive strategy in many cases, it may not converge to the correct ground-state when there is frustration in the Hamiltonian, for example in glassy models.

Inexact QITE. Given limited resources, for example on near-term devices, we can choose to measure and construct the uni-

tary over a domain D smaller than induced by correlations, to fit the computational budget. For example, if $D = 1$, this gives a mean-field approximation of the imaginary time evolution. While the unitary is no longer an exact representation of the imaginary time evolution, there is no issue of a local minimum in its construction, although the energy is no longer guaranteed to decrease at every step. In this case, one can apply inexact imaginary time evolution until the energy stops decreasing; the energy will still be a variational upper bound. One can also use the quantum Lanczos algorithm, described later.

QITE experiments. To illustrate the QITE algorithm, we have carried out exact classical emulations (assuming perfect expectation values and perfect gates) for several Hamiltonians: short-range 1D Heisenberg; 1D AFM transverse-field Ising; long-range 1D Heisenberg with spin-spin coupling $J_{ij} = (|i-j|+1)^{-1}; i \neq j$; 1D Hubbard at half-filling (mapped by Jordan-Wigner transformation to a spin model); a 6-qubit MAXCUT [10–12] instance, and a minimal basis 2-qubit dihydrogen molecular Hamiltonian [29]. To assess the feasibility of implementation on near-term quantum devices, we have also carried out noisy classical emulation (sampling expectation values and with an error model) using the Rigetti quantum virtual machine (QVM) and a physical simulation using the Rigetti Aspen-1 QPUs, for a single qubit field model ($2^{-1/2}(X + Z)$)[30] and a 1D AFM transverse-field Ising model. We carry out QITE using different fixed domain sizes D for the unitary or fermionic unitary (see SI for descriptions of simulations and models).

Figs. 2 and 3 show the energy obtained by QITE as a function of β and D for the various models. As we increase D, the asymptotic ($\beta \to \infty$) energies rapidly converge to the exact
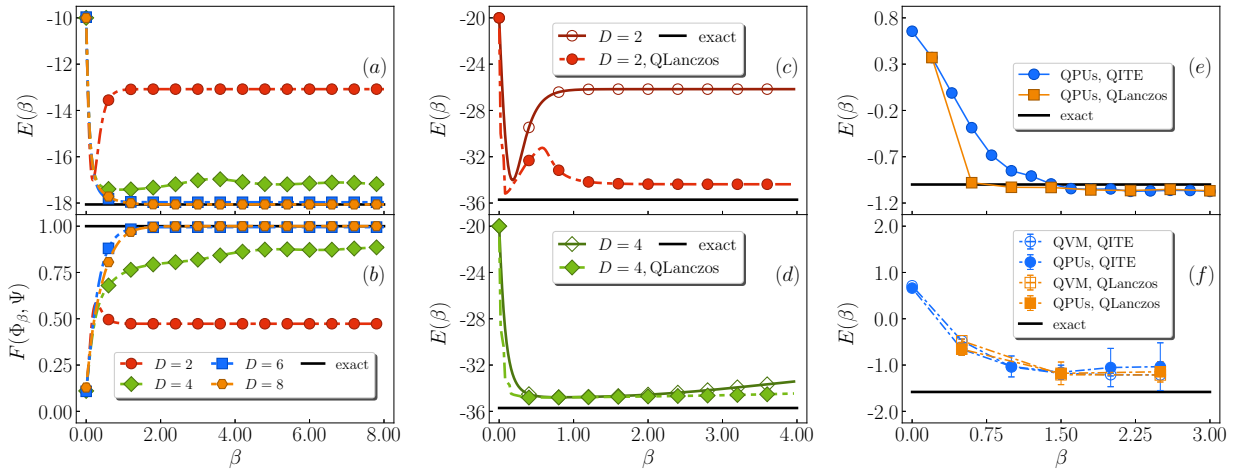
FIG. 2: Left: QITE energy $E(\beta)$ (a) and fidelity F (b) between finite-time state $\Phi(\beta)$ and exact ground state Ψ as function of imaginary time β, for a 1D 10-site Heisenberg model, showing the convergence with increasing unitary domains of $D = 2 - 8$ qubits. Middle: QITE (dashed red, dot-dashed green lines) and QLanczos (solid red, solid green lines) energies as function of imaginary time β, for a 1D Heisenberg model with $N = 20$ qubits, using domains of $D = 2$ (c) and 4 qubits (d), showing improved convergence of QLanczos over QITE. Black line is the exact ground-state energy/fidelity. Right: QITE and QLanczos energy $E(\beta)$ as a function of imaginary time β for (e) 1-qubit field model using the QVM and QPU (qubit 14 on Aspen-1), (f) 2-qubit AFM transverse field Ising model using the QVM and QPU (qubit 14, 15 on Aspen-1). Black line is the exact ground-state energy (see SI for details).

ground-state. For small D, the inexact QITE tracks the exact QITE for a time until the correlation length exceeds D. Afterwards, it may go down or up. The non-monotonic behavior is strongest for small domains; in the MAXCUT example, the smallest domain $D = 2$ gives an oscillating energy; the first point at which the energy stops decreasing is a reasonable estimate of the ground-state energy. In all models, increasing D past a maximum value (less than N) no longer affects the asymptotic energy, showing that the correlations have saturated (this is true even in the MAXCUT instance).

Figs. 2e and 2f show the results of running the QITE algorithm on Rigetti's QVM and Aspen-1 QPUs for 1- and 2-qubits, respectively. The error bars are due to gate, readout, incoherent and cross-talk errors. Sufficient samples were used to ensure that sampling error is negligible. Encouragingly for near-term simulations, despite these errors it is possible to converge to a ground-state energy close to the exact energy for the 1-qubit case. This result reflects a robustness that is sometimes informally observed in imaginary time evolution algorithms in which the ground state energy is approached even if the imaginary time step is not perfectly implemented. In the 2-qubit case, although the QITE energy converges, there is a systematic shift which is reproduced on the QVM using available noise parameters for readout, decoherence and depolarizing noise [31]. Remaining discrepancies between the emulator and hardware are likely attributable to cross-talk between parallel gates not included in the noise model (see SI). However, reducing decoherence and depolarizing errors in the QVM or using different sets of qubits with improved noise characteristics (see SI) all lead to improved convergence to the exact ground-state energy.

Quantum Lanczos algorithm. Given the QITE subroutine,
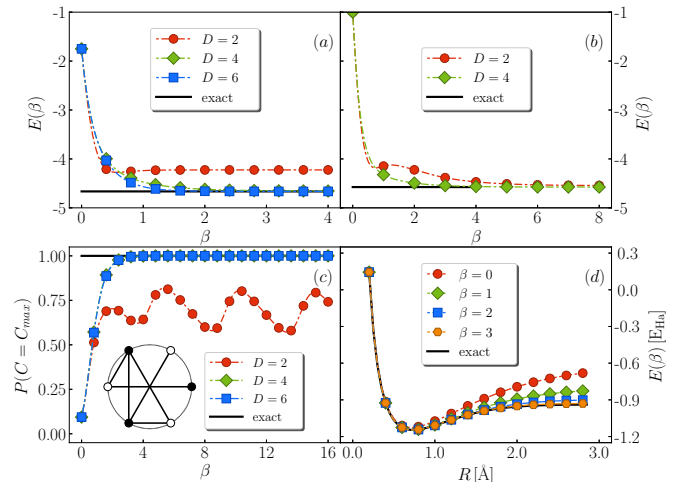

FIG. 3: (a) QITE energy $E(\beta)$ as a function of imaginary time β for a 6-site 1D long-range Heisenberg model, for unitary domains $D = 2 - 6$; (b) a 4-site 1D Hubbard model with $U/t = 1$, for unitary domains $D = 2, 4$. (c) Probability of MAXCUT detection, $P(C = C_{max})$ as a function of imaginary time β, for the 6-site graph in the panel. (d) QITE energy for the H_2 molecule in the STO-6G basis as a function of bond-length R and β. Black line is the exact ground-state energy/probability of detection.

we now consider how to formulate a quantum Lanczos algorithm, which is an especially economical realization of a quantum subspace method [32, 33]. An important practical motivation is that the Lanczos algorithm typically converges much more quickly than imaginary time evolution, and often in physical simulations only tens of iterations are needed to converge to good precision. In addition, Lanczos provides a
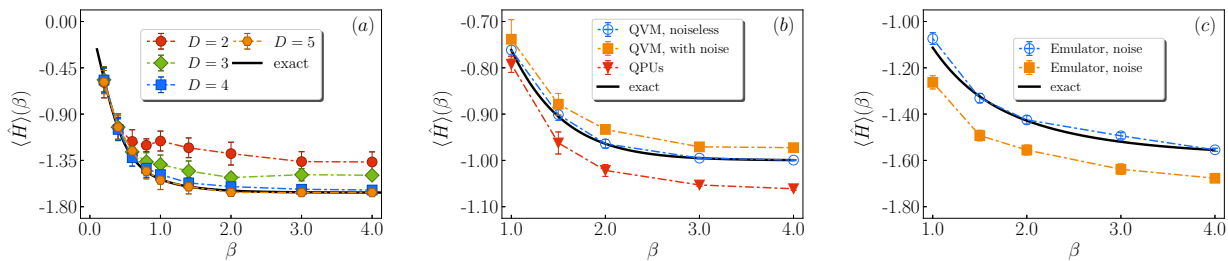
FIG. 4: Left: Thermal (Gibbs) average $\langle \hat{H} \rangle$ at temperature β from QMETTS for a 1D 6-site Heisenberg model (exact emulation). Black line is the exact thermal average without sampling error. Middle, Right: Thermal average $\langle \hat{H} \rangle$ at temperature β from QMETTS for (b) a 1 qubit field model using QVMs and QPUs, and (c) 2 qubit AFM transverse field Ising model using QVM.

natural way to compute excited states. Consider the sequence of imaginary time vectors $|\Phi_l\rangle = e^{-l\Delta\tau\hat{H}}|\Phi\rangle$, $l = 0, 1, \ldots n$, where $c_l = \|\Phi_l\|$. In QLanczos, we consider the vectors after even numbers of time steps $|\Phi_0\rangle, |\Phi_2\rangle \ldots$ to form a basis for the ground-state. (SI describes the equivalent treatment in terms of normalized imaginary time vectors). These vectors define an overlap matrix whose elements can be computed entirely from norms, $S_{ll'} = \langle\Phi_l|\Phi_{l'}\rangle = c^2_{(l+l')/2}$, where $c_{(l+l')/2}$ is the norm of another integer time step vector, and the overlap matrix elements for $n/2$ vectors can be accumulated for free after n steps of time evolution. The Hamiltonian matrix elements satisfy the identity $H_{ll'} = \langle\Phi_l|\hat{H}|\Phi_{l'}\rangle = \langle\Phi_{(l+l')/2}|\hat{H}|\Phi_{(l+l')/2}\rangle$. Although the Hamiltonian has $\sim n^2$ matrix elements, there are only $\sim n$ unique elements, and importantly, each is a simple expectation value of the energy during the imaginary time evolution. This economy of matrix elements is a property shared with the classical Lanczos algorithm. Whereas the classical Lanczos iteration builds a Krylov space in powers of $\hat{H}$, QLanczos builds a Krylov space in powers of $e^{-2\Delta\tau\hat{H}}$; in the limit of small $\Delta\tau$ these Krylov spaces are identical. Diagonalization of the QLanczos Hamiltonian matrix is guaranteed to give a ground-state energy that is lower than that of the last imaginary time vector Φ_n (while higher roots approximate excited states). Thus, as long as one is willing to take measurements of the energy during the imaginary time evolution process, one can use QLanczos to generate an improved ground state (or excited states).

With a limited computational budget, we can use inexact QITE to generate Φ_l, Φ'_l. However, in this case the above expressions for $S_{ll'}$ and $H_{ll'}$ in terms of expectation values are no longer exactly satisfied which can create numerical issues (e.g. the overlap may no longer be positive). To handle this as well as errors due to noise and sampling in real experiments, the QLanczos algorithm needs to be stabilized by ensuring that successive vectors are not nearly linearly dependent (see SI).

We demonstrate the QLanczos algorithm using classical emulation on the 1D Heisenberg Hamiltonian, as used for the QITE algorithm in Fig. 2 (see SI). Using exact QITE (large domains) to generate matrix elements, quantum Lanczos converges much more rapidly than imaginary time evolution. Using inexact QITE (small domains), convergence is

usually faster and also reaches a lower energy. We also assess the feasibility of QLanczos in presence of noise, using emulated noise on the Rigetti QVM as well as on the Rigetti Aspen-1 QPUs. In Fig. 2, we see that QLanczos also provides more rapid convergence than QITE with both noisy classical emulation as well as on the physical device for 1 and 2 qubits.

Quantum thermal averages. The QITE subroutine can be used in a range of other algorithms. For example, we discuss how to compute thermal averages $\text{Tr}[\hat{O}e^{-\beta\hat{H}}]/\text{Tr}[e^{-\beta\hat{H}}]$ using imaginary time evolution. Several procedures have been proposed for quantum thermal averaging, ranging from generating the finite-temperature state explicitly by equilibration with a bath [34], to a quantum analog of Metropolis sampling [35] that relies on phase estimation, as well as methods based on ancilla based Hamiltonian simulation with postselection [36] and approaches based on recovery maps [37]. However, given a method for imaginary time evolution, one can generate thermal averages of observables without any ancillae or deep circuits. This can be done by adapting to the quantum setting the classical minimally entangled typical thermal state (METTS) algorithm [19, 20], which generates a Markov chain from which the thermal average can be sampled. The QMETTS algorithm can be carried out as follows (i) start from a product state, carry out imaginary-time evolution (using QITE) up to time β (ii) measure the expectation value of the observable that one wants to produce a thermal average for (iii) measure a product operator such as $\hat{Z}_1\hat{Z}_2\ldots\hat{Z}_N$, to collapse back onto a random product state (iv) repeat (i). Note that in step (iii) one can measure in any product basis, and randomizing the product basis can be used to reduce the autocorrelation time and avoid ergodicity problems in sampling.

In Fig. 4 we show the results of quantum METTS (using exact classical emulation) for the thermal average $\langle\hat{H}\rangle$ as a function of temperature β, for the 6-site Heisenberg model for several temperatures and domain sizes; sufficiently large D converges to the exact thermal average at each β; error bars reflect only finite QMETTS samples. We also show an implementation of quantum METTS on the Aspen-1 QPU and QVM with a 1-qubit field model (Fig. 4b), and using the QVM for a 2-qubit AFM transverse field Ising model (Fig. 4c).

Conclusions. We have introduced quantum analogs of

imaginary time evolution (QITE) and the Lanczos algorithm (QLanczos), that can be carried out without ancillae or deep circuits, and which, for bounded correlation length, achieve exponential reductions in space and time per iteration relative to their classical counterparts. They provide new quantum routes to approximate ground-states of Hamiltonians in both physical simulations and in optimization that avoid some of the current disadvantages of phase estimation based approaches and variational algorithms. The QLanczos iteration appears especially powerful if sufficient sampling can be done, as in practice it obtains accurate estimates of ground-states from only a few iterations, and also provides an estimate of excited states. Additionally, further algorithms that use QITE and QLanczos as subroutines can be formulated, such as the quantum minimally entangled thermal states algorithm to compute thermal averages. Encouragingly, these algorithms appear useful in conjunction with near-term quantum architectures, and serve to demonstrate the power of quantum elevations of classical simulation techniques, in the continuing search for quantum supremacy.

<center>**SUPPLEMENTAL INFORMATION**</center>

<center>**Representing imaginary-time evolution by unitary maps**</center>

As discussed in the main text, we map the scaled non-unitary action of $e^{-\Delta\tau \hat{h}_m}$ on a state Ψ to that of a unitary $e^{-i\Delta\tau \hat{A}[m]}$, i.e.

$$|\Psi'\rangle \equiv c^{-1/2} e^{-\Delta\tau \hat{h}_m}|\Psi\rangle = e^{-i\Delta\tau \hat{A}[m]}|\Psi\rangle \quad . \tag{6}$$

where $c = \langle\Psi|e^{-2\Delta\tau \hat{h}_m}|\Psi\rangle$. $\hat{h}_m$ acts on k qubits; $\hat{A}$ is Hermitian and acts on a domain of D qubits around the support of $\hat{h}_m$, and is expanded as a sum of Pauli strings acting on the D qubits,

$$\hat{A}[m] = \sum_{i_1 i_2 \ldots i_D} a[m]_{i_1 i_2 \ldots i_D} \sigma_{i_1}\sigma_{i_2}\ldots\sigma_{i_D}$$
$$= \sum_I a[m]_I \sigma_I \tag{7}$$

where I denotes the index $i_1 i_2 \ldots i_D$. Define $|\Delta_0\rangle = \frac{|\Psi'\rangle - |\Psi\rangle}{\Delta\tau}$ and $|\Delta\rangle = -i\hat{A}[m]|\Psi\rangle$. Our goal is to minimize the difference $||\Delta_0 - \Delta||$. If the unitary $e^{-i\Delta\tau \hat{A}[m]}$ is defined over a sufficiently large domain D (related to the correlation length of $|\Psi\rangle$, see Section) then this error minimizes at ~ 0, for small $\Delta\tau$. Minimizing for real $a[m]$ corresponds to minimizing the quadratic function $f(a[m])$

$$f(a[m]) = f_0 + \sum_I b_I a[m]_I + \sum_{IJ} a[m]_I S_{IJ} a[m]_J \tag{8}$$

where

$$f_0 = \langle\Delta_0|\Delta_0\rangle \quad , \tag{9}$$
$$S_{IJ} = \langle\Psi|\sigma_I^\dagger \sigma_J|\Psi\rangle \quad , \tag{10}$$
$$b_I = i\langle\Psi|\sigma_I^\dagger|\Delta_0\rangle - i\langle\Delta_0|\sigma_I|\Psi\rangle \quad , \tag{11}$$

whose minimum obtains at the solution of the linear equation

$$\left(\mathbf{S} + \mathbf{S}^T\right)\mathbf{a}[m] = -\mathbf{b} \tag{12}$$

In general, $\mathbf{S} + \mathbf{S}^T$ may have a non-zero null-space. Thus, we solve Eq. (12) either by applying the generalized inverse of $\mathbf{S} + \mathbf{S}^T$ or by an iterative algorithm such as conjugate gradient.

For fermionic Hamiltonians, we replace the Pauli operators in Eq. (7) by fermionic field operators. For a number conserving Hamiltonian, such as the fermionic Hubbard Hamiltonian treated in Fig. 3 in the main text, we write

$$\hat{A}[m] = \sum_{i_1 i_2 \ldots i_D} a[m]_{i_1 i_2 \ldots i_D} \hat{f}_{i_1}^\dagger \ldots \hat{f}_{i_{D/2}}^\dagger \hat{f}_{i_{D/2+1}} \ldots \hat{f}_{i_D} \tag{13}$$

where $\hat{f}^\dagger$, $\hat{f}$ are fermionic creation, annihilation operators respectively.

<center>**Rigorous Run Time Bounds**</center>

Here we present a more detailed analysis of the running time of the algorithm. Consider a k-local Hamiltonian

$$H = \sum_{l=1}^m h_l \tag{14}$$

acting on a d-dimensional lattice with $\|h_i\| \leq 1$, where $\|*\|$ is the operator norm. In imaginary time evolution one typically applies Trotter formulae to approximate

$$\frac{e^{-\beta H}|\Psi_0\rangle}{\|e^{-\beta H}|\Psi_0\rangle\|} \tag{15}$$

for an initial state $|\Psi_0\rangle$ (which we assume to be a product state) by

$$\frac{\left(e^{-\beta h_1/n}\ldots e^{-\beta h_m/n}\right)^n |\Psi_0\rangle}{\left\|\left(e^{-\beta h_1/n}\ldots e^{-\beta h_m/n}\right)^n |\Psi_0\rangle\right\|}. \tag{16}$$

This approximation leads to an error which can be made as small as one wishes by increasing the number of time steps n.

Let $|\Psi_s\rangle$ be the state (after renormalization) obtained by applying s terms $e^{-th_i/n}$ from $\left(e^{-th_1/n}\ldots e^{-th_m/n}\right)^n$; with this notation $|\Psi_{mn}\rangle$ is the state given by Eq. (16). In the QITE algorithm, instead of applying each of the operators $e^{-th_i/n}$ to $|\Psi_0\rangle$ (and renormalizing the state), one applies local unitaries U_s which should approximate the action of the original operator. Let $|\Phi_s\rangle$ be the state after s unitaries have been applied.

Let C be an upper bound on the correlation length of $|\Psi_s\rangle$ for every s: we assume that for every s, and every pair of observables A and B separated by $\mathrm{dist}(A,B)$ sites,

$$C_s(A,B) = \langle\Psi_s|A\otimes B|\Psi_s\rangle - \langle\Psi_s|A|\Psi_s\rangle\langle\Psi_s|B|\Psi_s\rangle$$
$$\leq \|A\|\|B\|e^{-\mathrm{dist}(A,B)/C}. \tag{17}$$

Theorem 1. *For every $\varepsilon > 0$, there are unitaries U_s each acting on*

$$k(2C)^d \ln^d(2\sqrt{2}nm\varepsilon^{-1}) \tag{18}$$

qubits such that

$$\||\Psi_{mn}\rangle - |\Phi_{mn}\rangle\| \leq \varepsilon \tag{19}$$

Proof. We have

$$\||\Psi_s\rangle - |\Phi_s\rangle\| = \||\Psi_s\rangle - U_s|\Phi_{s-1}\rangle\|$$
$$\leq \||\Psi_s\rangle - U_s|\Psi_{s-1}\rangle\| + \||\Psi_{s-1}\rangle - |\Phi_{s-1}\rangle\|. \tag{20}$$

To bound the first term we use our assumption that the correlation length of $|\Psi_{s-1}\rangle$ is smaller than C. Consider a region R_v of all sites that are a distance at most v (in the Manhattan distance on the lattice) of the sites in which h_{i_s} acts. Let $\mathrm{tr}_{\backslash R_v}(|\Psi_s\rangle\langle\Psi_s|)$ be the reduced state on R_v, obtained by partial tracing over the complement of R_v in the lattice. Since

$$|\Psi_s\rangle = \frac{e^{-\beta h_{i_s}/n}|\Psi_{s-1}\rangle}{\|e^{-\beta h_{i_s}/n}|\Psi_{s-1}\rangle\|}, \tag{21}$$

it follows from Eq. (17) and Lemma 9 of [38] that

$$\left\|\mathrm{tr}_{\backslash R_v}(|\Psi_s\rangle\langle\Psi_s|) - \mathrm{tr}_{\backslash R_v}(|\Psi_{s-1}\rangle\langle\Psi_{s-1}|)\right\|_1$$
$$\leq \|e^{h_{i_s}/n}\|^{-1}e^{-\frac{v}{C}} \leq 2e^{-\frac{v}{C}}, \tag{22}$$

where we used that for $n \geq 2\beta$, $\|e^{-\beta h_{i_s}/n}\| \geq \|I - \beta h_{i_s}/n\| \geq 1 - \beta/n \geq 1/2$. Above $\|*\|_1$ is the trace norm.

The key result in our analysis is Uhlmann's theorem (see e.g. Lemmas 11 and 12 of [38]). It says that two pure states

with nearby marginals must be related by a unitary on the purifying system. In more detail, if $|\eta\rangle_{AB}$ and $|\nu\rangle_{AB}$ are two states s.t. $\|\eta_A - \nu_A\|_1 \leq \delta$, then there exists a unitary V acting on B s.t.

$$\||\eta\rangle_{AB} - (I\otimes V)|\nu\rangle_{AB}\| \leq 2\sqrt{\delta}. \tag{23}$$

Applying Uhlmann's theorem to $|\Psi_s\rangle$ and $|\Psi_{s-1}\rangle$, with $B = R_v$, and using Eq. (22), we find that there exists a unitary U_s acting on R_v s.t.

$$\||\Psi_s\rangle - U_s|\Psi_{s-1}\rangle\| \leq 2\sqrt{2}e^{-\frac{v}{2C}}, \tag{24}$$

which by Eq. (20) implies

$$\||\Psi_{nm}\rangle - |\Phi_{nm}\rangle\| \leq 2\sqrt{2}mne^{-\frac{v}{2C}}, \tag{25}$$

Choosing $\nu = 2C\ln(2\sqrt{2}nm\varepsilon^{-1})$ as the width of the support of the approximating unitaries, the error term above is ε. The support of the local unitaries is $k\nu^d$ qubits (as this is an upper bound on the number of qubits in R_ν). Therefore each unitary U_s acts on at most

$$k(2C)^d \ln^d(2\sqrt{2}nm\varepsilon^{-1}) \tag{26}$$

qubits. $\square$

Finding U_s: In the algorithm we claim that we can find the unitaries U_s by solving a least-square problem. This is indeed the case if we can write them as $U_s = e^{iA[s]/n}$ with $A[s]$ a Hamiltonian of constant norm. Then for sufficiently large n, $U_s = I + iA[s]/n + O((1/n)^2)$ and we can find $A[s]$ by performing tomography of the reduced state over the region where U_s acts and solving the linear problem given in the main text. Because we apply Uhlmann's Theorem to $|\Psi_{s-1}\rangle$ and

$$\frac{e^{-\beta h_{i_s}/n}|\Psi_{s-1}\rangle}{\|e^{-\beta h_{i_s}/n}|\Psi_{s-1}\rangle\|}, \tag{27}$$

using $e^{-\beta h_{i_s}/n} = I - \beta h_{i_s}/n + O((1/n)^2)$ and following the proof of the Uhlmann's Theorem, we find that the unitary can indeed be taken to be close to the identity, i.e. U_s can be written as $e^{iA[s]/n}$.

Total Run Time: Theorem 1 gives an upper bound on the maximum support of the unitaries needed for a Trotter update, while tomography of local reduced density matrices gives a way to find the unitaries. The cost for tomography is quadratic in the dimension of the region, so it scales as $\exp(O(k(2C)^d \ln^d(2\sqrt{2}nm\varepsilon^{-1})))$. This is also the cost to solve classically the linear system which gives the associated Hamiltonian $A[s]$ and of finding a circuit decomposition of $U_s = e^{iA[s]/n}$ in terms of two qubit gates. As this is repeated mn times, for each of the mn terms of the Trotter decomposition, the total running time (of both quantum and classical parts) is

$$mn\exp(O(k(2C)^d \ln^d(2\sqrt{2}nm\varepsilon^{-1}))). \tag{28}$$

This is exponential in C^d, with C the correlation length, and quasi-polynomial in n (the number of Trotter steps) and m (the number of local terms in the Hamiltonian. Note that typically $m = O(N)$, with N the number of sites). While this an exponential improvement over the $\exp(O(N))$ scaling classically, the quasi-polynomial dependence on m can still be prohibitive in practice. Below we show how to improve on that.

Local Approximation: If one is only interested in a local approximation of the state (meaning that all the local marginals of $|\Phi_{nm}\rangle$ are close to the ones of $e^{-\beta H}|\Psi_0\rangle$, but not necessarily the global states), then the support of the unitaries becomes independent of the number of terms of the Hamiltonian m (while for global approximation we have a polylogarithmic dependence on m):

Theorem 2. *For every $\varepsilon > 0$, there are unitaries U_s each acting on*

$$k(2C)^d \ln^d \left(2\sqrt{2}n(|S| + C\ln(8nC(2C)^{d+1}\varepsilon^{-1})^d \right) \quad (29)$$

qubits such that for every connected region S of size at most $|S|$,

$$\left\| tr_{\backslash S}(|\Psi_{mn}\rangle\langle\Psi_{mn}|) - tr_{\backslash S}(|\Phi_{mn}\rangle\langle\Phi_{mn}|) \right\|_1 \leq \varepsilon$$

Proof. Consider the unitaries U_s obtained in the proof of Theorem 1 satisfying Eq. (24).

Consider the replacement of the local term of the Trotter expansion by the unitary U_s for all local terms which are more than $2C\log(1/\delta)$ sites away from the region S. Because the correlation length is always smaller than C, we find by Lemma 9 of [38] that the total error on the reduced density matrix in region S can be bounded as

$$n \int_{2C\ln(1/\delta)}^{\infty} e^{-l/2C}l^d dl \leq 4nC(2C)^{d+1}\delta. \quad (30)$$

For the local terms which are at most a distance $2C\log(1/\delta)$ from the region S, in turn, the total error is bounded by the sum of each individual term, giving:

$$(|S| + C\log(1/\delta))^d n2\sqrt{2}e^{-\frac{\nu}{2C}} \quad (31)$$

Choosing $\delta = \varepsilon/(8nC(2C)^{d+1})$ and $\nu = 2C\ln(2\sqrt{2}n(|S| + C\ln(8nC(2C)^{d+1}\varepsilon^{-1})^d)$ gives the result. □

Non-Local Terms: Suppose the Hamiltonian has a term h_q acting on qubits which are not nearby, e.g. on two sites i and j. Then $e^{-\beta h_q/n}$ can still be replaced by a unitary, which only acts on sites i and j and qubits in the neighborhoods of the two sites. This is the case if we assume that the state has a finite correlation length and the proof is again an application of Uhlmann's theorem (we follow the same argument from the

proof of Theorem 1 but define R_v in that case as the union of the neighborhoods of i and j). Note however that the assumption of a finite correlation length might be less natural for models with long range interactions.

Scaling with temperature and increase of correlation length: Our discussion has been based on the assumption that the correlation length C is small on all intermediate states. Here we discuss the range of validity of the assumption.

Let us begin with an example where the correlation length can increase very quickly with number of local terms applied (this was communicated to us by Guang Hao Low). Consider a projection on two qubits $P_{i,i+1} = |0,0\rangle\langle0,0|_{i,i+1} + |1,1\rangle\langle1,1|_{i,i+1}$. Then

$$P_{1,2}P_{2,3}\ldots P_{n-1,n}|+\rangle^{\otimes n}, \quad (32)$$

with $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{(2)}$, is the GHZ state $(|0\ldots0\rangle + |1\ldots1\rangle)/\sqrt{2}$, which has correlation length $C = n$. While the projector $P_{i,i+1}$ cannot appear as a local term $e^{-\beta h_1/n}$ in the Trotter decomposition, this example show that we cannot expect a speed-of-sound bound on the spread of correlations for a circuit with non-unitary gates; indeed the example shows a depth two circuit can already create long range correlations.

However, we expect that generically the correlations do grow ballistically. Consider the state

$$|\psi_n\rangle := \frac{\left(e^{-\beta h_1/n}\ldots e^{-\beta h_m/n}\right)^n |\Psi_0\rangle}{\left\| \left(e^{-\beta h_1/n}\ldots e^{-\beta h_m/n}\right)^n |\Psi_0\rangle\right\|}. \quad (33)$$

after n rounds have been applied. Let us assume the Hamiltonian acts on a line, is translation invariant and has nearest-neighbor interactions. Then the state is a matrix product state of bond dimension at most 2^n. For matrix product states we can bound the correlations as follows (see e.g. Lemma 22 of [38])

$$C_s(A, B) = \langle\Psi_s|A \otimes B|\Psi_s\rangle - \langle\Psi_s|A|\Psi_s\rangle\langle\Psi_s|B|\Psi_s\rangle$$
$$\leq \|A\|\|B\|2^{2n}e^{-\Delta\text{dist}(A,B)}. \quad (34)$$

where we define the gap of the matrix-product-state as $\Delta := 1 - \lambda$, with λ the second largest eigenvalue of the transfer matrix of the matrix product state (normalized so that the largest eigenvalue is one). In the GHZ example above, the gap $\Delta = 0$ and that is the reason for the fast build up of correlations. Typically we expect the gap to be independent of n or decrease mildly as $1/\text{poly}(n)$.

From the above, we can replace a non-unitary local Trotter term applied to $|\psi_n\rangle$ by an unitary acting on $O(n/\Delta)$ qubits. Taking $n = O(\beta)$ to reach temperature β in the imaginary time evolution, the support of the unitaries would scale as $O(\beta/\Delta)$. Assuming Δ is a constant, we find a linear increase in temperature.

We also expect the linear growth of correlations/unitary support with inverse temperature also to hold generically in two dimensions, although there the analysis is more subtle as

rigorous results for the expected behaviour of the transfer operator (which becomes a one-dimensional tensor product operator) and its gap are not available.

Spreading of correlations

In the main text, we argued that the correlation volume V of the state $e^{-\beta H}|\Psi\rangle$ is bounded for many physical Hamiltonians and saturates at the ground-state with $V \ll N$ where N is the system size. To numerically measure correlations, we use the mutual information between two sites, defined as

$$I(i,j) = S(i) + S(j) - S(i,j) \tag{35}$$

where $S(i)$ is the von Neumann entropy of the density matrix of site i ($\rho(i)$) and similarly for $S(j)$, and $S(i,j)$ is the von Neumann entropy of the two-site density matrix for sites i and j ($\rho(i,j)$).

To compute the mutual information in Fig. 1 in the main text, we used matrix product state (MPS) and finite projected entangled pair state (PEPS) imaginary time evolution for the spin-1/2 1D and 2D FM transverse field Ising model (TFI)

$$H_{TFI} = -\sum_{\langle ij\rangle} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x \tag{36}$$

where the sum over $\langle i,j\rangle$ pairs are over nearest neighbors. We use the parameter $h = 1.25$ for the 1-D calculation and $h = 3.5$ for the 2-D calculations as the ground-state is gapped in both cases. It is known that the ground-state correlation length is finite.

MPS. We performed MPS imaginary time evolution (ITE) on a 1-D spin chin with $L = 50$ sites with open boundary conditions. We start from an initial state that is a random product state, and perform ITE using time evolution block decimation (TEBD) [39, 40] with a first order Trotter decomposition. In this algorithm, the Hamiltonian is separated into terms operating on even and odd bonds. The operators acting on a single bond are exponentiated exactly. One time step is given by time evolution of odd and even bonds sequentially, giving rise to a Trotter error on the order of the time step $\Delta\tau$. In our calculation, a time step of $\Delta\tau = 0.001$ was used.

We carry out ITE simulations with maximum bond dimension of $D = 80$, but truncate singular values less than 1.0e-8 of the maximum singular value. In the main text, the ITE results are compared against the ground state obtained via the density matrix renormalization group (DMRG)). This should be equivalent to comparing to a long-time ITE ground state. The long-time ITE ($\beta = 38.352$) ground state reached an energy per site of -1.455071, while the DMRG ground-state energy per site is -1.455076. The relative error of the nearest neighbor correlations is on the order of 10^{-4} to 10^{-3}, and about 10^{-2} for correlations between the middle site and the end sites (a distance of 25 sites). The error in fidelity between the two ground states was about 5×10^{-4}.

PEPS. We carried out finite PEPS [41–44] imaginary time evolution for the two-dimensional transverse field Ising model on a lattice size of 21×31. The size was chosen to be large enough to see the spread of mutual information in the bulk without significant effects from the boundary. The mutual information was calculated along the long (horizontal) axis in the center of the lattice. The standard Trotterized imaginary time evolution scheme for PEPS [45] was used with a time step $\Delta\tau = 0.001$, up to imaginary time $\beta = 6.0$, starting from a random product state. To reduce computational cost from the large lattice size, the PEPS was defined in a translationally invariant manner with only 2 independent tensors [46] updated via the so-called "simple update" procedure [47]. The simple update has been shown to be sufficiently accurate for capturing correlation functions (and thus $I(i,j)$) for ground states with relatively short correlation lengths (compared to criticality) [48, 49]. We chose a magnetic field value $h = 3.5$ which is detuned from the critical field ($h \approx 3.044$) but still maintains a correlation length long enough to see interesting behaviour.

Accuracy: Even though the simple update procedure was used for the tensor update, we still needed to contract the 21×31 PEPS at at every imaginary time step β for a range of correlation functions, amounting to a large number of contractions. To control the computational cost, we limited our bond dimension to $D = 5$ and used an optimized contraction scheme [50], with maximum allowed bond dimension of $\chi = 60$ during the contraction. Based on converged PEPS ground state correlation functions with a larger bond dimension of $D = 8$, our $D = 5$ PEPS yields $I(i, i+r)$ (where r denotes horizontal separation) at large β with a relative error of $\approx 1\%$ for $r = 1 - 4$, 5% or less for $r = 5 - 8$, and 10% or greater for $r > 8$. At smaller values of β (< 0.5) the errors up to $r = 8$ are much smaller because the bond dimension of 5 is able to completely support the smaller correlations (see Fig. 1, main text). While error analysis on the 2D Heisenberg model [48] suggests that errors with respect to $D = \infty$ may be larger, such analysis also confirms that a $D = 5$ PEPS captures the qualitative behaviour of correlation in the range $r = 5 - 10$ (and beyond). Aside from the bond dimension error, the precision of the calculations is governed by χ and the lattice size. Using the 21×31 lattice and $\chi = 60$, we were able to converge entries of single-site density matrices $\rho(i)$ to a precision of $\pm 10^{-6}$ (two site density matrices $\rho(i,j)$ had higher precision). For $\beta = 0.001 - 0.012$, the smallest eigenvalue of $\rho(i)$ fell below this precision threshold, leading to significant noise in $I(i,j)$. Thus, these values of β are omitted from Fig. 1 (main text) and the smallest reported values of I are 10^{-6}, although with more precision we expect $I \to 0$ as $r \to \infty$.

Finally, the energy and fidelity errors were computed with respect to the PEPS ground state *of the same bond dimension* at $\beta = 10.0$ (10000 time steps). The convergence of the quantities shown in Fig. 1 (main text) thus isolates the convergence of the imaginary time evolution, and does not include effects of other errors that may result from deficiencies in the wavefunction ansatz.

Simulation models

We here define, and give some background on, the models used in the QITE and QLanczos simulations.

1 qubit field model

$$\hat{H} = \alpha \hat{X} + \beta \hat{Z} \qquad (37)$$

This Hamiltonian has previously been used as a model for quantum simulations on physical devices in Ref. [30]. We used $\alpha = \frac{1}{\sqrt{2}}$ and $\beta = \frac{1}{\sqrt{2}}$. In simulations with this Hamiltonian, the qubit is assumed to be initialized in the Z basis.

1D Heisenberg and transverse field Ising model

The 1D short-range Heisenberg Hamiltonian is defined as

$$\hat{H} = \sum_{\langle ij \rangle} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j \quad , \qquad (38)$$

the 1D long-range Heisenberg Hamiltonian as

$$\hat{H} = \sum_{i \neq j} \frac{1}{|i-j|+1} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j \quad , \qquad (39)$$

and the AFM transverse-field Ising Hamiltonian as

$$\hat{H} = \sum_{\langle ij \rangle} \hat{S}_i^z \hat{S}_j^z + \sum_i h \hat{S}_i^x \quad . \qquad (40)$$

1D Hubbard model

The 1D Hubbard Hamiltonian is defined as

$$\hat{H} = -\sum_{\langle ij \rangle \sigma} a_{i\sigma}^\dagger a_{j\sigma} + U \sum_i \hat{n}_{i\uparrow}\hat{n}_{i\downarrow} \qquad (41)$$

where $\hat{n}_{i\sigma} = a_{i\sigma}^\dagger a_{i\sigma}$, $\sigma \in \{\uparrow, \downarrow\}$, and $\langle \cdot \rangle$ denotes summation over nearest-neighbors, here with open-boundary conditions. We label the n lattice sites with an index $i = 0 \ldots n-1$, and the $2n-1$ basis functions as $|\varphi_0\rangle = |0 \uparrow\rangle$, $|\varphi_1\rangle = |0 \downarrow\rangle$, $|\varphi_2\rangle = |1 \uparrow\rangle$, $|\varphi_3\rangle = |1 \downarrow\rangle \ldots$. Under Jordan-Wigner transformation, recalling that

$$\hat{n}_p = \frac{1-Z_p}{2} \quad ,$$
$$\hat{a}_p^\dagger \hat{a}_q + \hat{a}_q^\dagger \hat{a}_p = \frac{X_p X_q \prod_{k=q+1}^{p-1} Z_k \left(1 - Z_p Z_q\right)}{2} \quad , \qquad (42)$$

with $p = 0 \ldots 2n-2$ and $q < p$, the Hamiltonian takes the form

$$\hat{H} = -\sum_p \frac{X_p X_{p+2} Z_{p+1} \left(1 - Z_p Z_{p+2}\right)}{2}$$
$$+ U \sum_{p\,\text{even}} \frac{(1-Z_{2i})(1-Z_{2i+1})}{4} + \mu \sum_p \frac{(1-Z_p)}{2} \qquad (43)$$

H_2 molecule minimal basis model

We use the hydrogen molecule minimal basis model at the STO-6G level of theory. This is a common minimal model of hydrogen chains [51, 52] and has previously been studied in quantum simulations, for example in [29]. Given a molecular geometry (H-H distance R) we perform a restricted Hartree-Fock calculation and express the second-quantized Hamiltonian in the orthonormal basis of RHF molecular orbitals as [53]

$$\hat{H} = H_0 + \sum_{pq} h_{pq}\hat{a}_p^\dagger \hat{a}_q + \frac{1}{2}\sum_{prqs} v_{prqs}\hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_s \hat{a}_r \qquad (44)$$

where $a^\dagger$, a are fermionic creation and annihilation operators for the molecular orbitals.

The Hamiltonian (44) is then encoded by a Bravyi-Kitaev transformation into the 2-qubit operator

$$\hat{H} = g_0 I \otimes I + g_1 Z \otimes I + g_2 I \otimes Z + g_3 Z \otimes Z + g_4 X \otimes X + g_5 Y \otimes Y \quad , \qquad (45)$$

with coefficients g_i given in Table I of [29].

MAXCUT Hamiltonian

The MAXCUT Hamiltonian encodes the solution of the MAXCUT problem. Given a graph $\Gamma = (V, E)$, where V is a set of vertices and $E \subseteq V \times V$ is a set of links between vertices in V, a cut of Γ is a subset $S \subseteq V$ of V. The MAXCUT problem consists in finding a cut S that maximizes the number of edges between S and S^c (the complement of S). We denote the number of links in a given cut S as $C(S)$.

In Figure 3 of the main text, we consider a graph Γ with vertices $V = \{0, 1, 2, 3, 4, 5\}$ and links $E = \{(0,3), (1,4), (2,3), (2,4), (2,5), (4,5)\}$. It is easy to verify that $S = \{0, 2, 4\}$, $\{0, 1, 2\}$, $\{3, 4\}$ and their complements S^c are solutions of the MAXCUT problem, with weight $C_{max} = 5$.

The MAXCUT problem can be formulated as a Hamiltonian ground-state problem, by (i) associating a qubit to every vertex in V, (ii) associating to every partition $S =$ an element of the computational basis (here assumed to be in the z direction) of the form $|z_0 \ldots z_{n-1}\rangle$, where $z_i = 1$ if $i \in S$ and $z_i = 0$ if $i \in S^c$, and finding the minimal (most negative)

eigenvalue of the 2-local Hamiltonian

$$\hat{C} = - \sum_{(ij) \in E} \frac{1 - \hat{S}_i^z \hat{S}_j^z}{2} \quad . \tag{46}$$

The spectrum of $\hat{C}$ is a subset of numbers $C \in \{0, 1 \dots |E|\}$.

In the present work, we initialize the qubits in the state $|\Phi\rangle = |+\rangle^{\otimes n}$, where $|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}$, and evolve Φ in imaginary time. Measuring the evolved state at time β $|\Phi(\beta)\rangle$ will collapse it onto an element $|z_0 \dots z_{n-1}\rangle$ of the computational basis, which is also an eigenfunction of $\hat{C}$ with eigenvalue C. In Figure 3 in the main text, we illustrate the probability $P(|C| = C_{max})$ that such measurements yield a MAX-CUT solution. Note that, even in the presence of oscillations (with the smallest domain size $D = 2$) this probability remains above 60%.

Numerical simulation details

QITE stabilization

Sampling noise in the expectation values of the Pauli operators can affect the solution to Eq. (12) that sometimes lead to numerical instabilities. We regularize $\mathbf{S} + \mathbf{S}^T$ against such statistical errors by adding a small δ to its diagonal. To generate the data presented in Figures 2 and 4 of the main text, we used $\delta = 0.01$ for 1-qubit calculations and $\delta = 0.1$ for 2-qubit calculations.

QLanczos stabilization

In quantum Lanczos, we generate a set of wavefunctions for different imaginary-time projections of an initial state $|\Psi\rangle$, using QITE as a subroutine. The normalized states are

$$|\Phi_l\rangle = \frac{e^{-l \Delta\tau \hat{H}} |\Psi_T\rangle}{\|e^{-l \Delta\tau \hat{H}} \Psi_T\|} \equiv n_l \, e^{-l \Delta\tau \hat{H}} |\Psi_T\rangle \quad 0 \le l < L_{\max} \quad . \tag{47}$$

where n_l is the normalization constant. For the exact imaginary-time evolution and l, l' both even (or odd) the matrix elements

$$S_{l,l'} = \langle \Phi_l | \Phi_{l'} \rangle \quad , \quad H_{l,l'} = \langle \Phi_l | \hat{H} | \Phi_{l'} \rangle \tag{48}$$

can be computed in terms of expectation values (i.e. experimentally accessible quantities) only. Indeed, defining $2r = l + l'$, we have

$$S_{l,l'} = n_l n_{l'} \langle \Psi_T | e^{-l \Delta\tau \hat{H}} e^{-l' \Delta\tau \hat{H}} | \Psi_T \rangle = \frac{n_l n_{l'}}{n_r^2} \quad , \tag{49}$$

and similarly

$$\begin{aligned} H_{l,l'} &= n_l n_{l'} \langle \Psi_T | e^{-l \Delta\tau \hat{H}} \hat{H} e^{-l' \Delta\tau \hat{H}} | \Psi_T \rangle = \\ &= \frac{n_l n_{l'}}{n_r^2} \langle \Phi_r | \hat{H} | \Phi_r \rangle = S_{l,l'} \langle \Phi_r | \hat{H} | \Phi_r \rangle \quad . \end{aligned} \tag{50}$$

The quantities n_r can be evaluated recursively, since

$$\begin{aligned} \frac{1}{n_{r+1}^2} &= \langle \Psi_T | e^{-(r+1) \Delta\tau \hat{H}} e^{-(r+1) \Delta\tau \hat{H}} | \Psi_T \rangle = \\ &= \frac{\langle \Phi_r | e^{-2 \Delta\tau \hat{H}} | \Phi_r \rangle}{n_r^2} \quad , \end{aligned} \tag{51}$$

For inexact time evolution, the quantities n_r and $\langle \Phi_r | \hat{H} | \Phi_r \rangle$ can still be used to approximate $S_{l,l'}$, $H_{l,l'}$.

Given these matrices, we then solve the generalized eigenvalue equation $\mathbf{Hx} = E\mathbf{Sx}$ to find an approximation to the ground-state $|\Phi'\rangle = \sum_l x_l |\Phi_l\rangle$ for the ground state of $\hat{H}$. This eigenvalue equation can be numerically ill-conditioned, as S can contain small and negative eigenvalues for several reasons (i) as m increases the vectors $|\Phi_l\rangle$ become linearly dependent; (ii) simulations have finite precision and noise; (iii) S, H are computed approximately when inexact time evolution is performed.

To regularize the problem, out of the set of time-evolved states we extract a better-behaved sequence as follows (i) start from $|\Phi_{\text{last}}\rangle = |\Phi_0\rangle$ (ii) add the next $|\Phi_l\rangle$ in the set of time-evolved states s.t. $|\langle \Phi_l | \Phi_{\text{last}}\rangle| < s$, where s is a regularization parameter $0 < s < 1$ (iii) repeat, setting the $|\Phi_{\text{last}}\rangle = \Phi_l$ (obtained from (ii)), until the desired number of vectors is reached. We then solve the generalized eigenvalue equation $\tilde{\mathbf{H}}\mathbf{x} = E\tilde{\mathbf{S}}\mathbf{x}$ spanned by this regularized sequence, removing any eigenvalues of $\tilde{\mathbf{S}}$ less than a threshold ϵ. The exact emulated QLanczos calculations reported in the main text were stabilized with this algorithm (the source of error here is primarily (iii)) using stabilization parameter $s = 0.95$ and $\epsilon = 10^{-14}$. The stabilization parameters used in the QVM and QPU QLanczos calculations were $s = 0.75$ and $\epsilon = 10^{-2}$ (the main source of error in the simulations was (ii)). Note that the stabilization procedure is unlikely to fix all possible numerical instabilities, but was sufficient for all models and calculations performed in this work.

METTS algorithm

The METTS (minimally entangled typical thermal state) algorithm [54, 55] is a sampling method to calculate thermal properties based on imaginary time evolution. Consider the thermal average of an observable $\hat{O}$

$$\langle \hat{O} \rangle = \frac{1}{Z} \text{Tr}[e^{-\beta \hat{H}} \hat{O}] = \frac{1}{Z} \sum_i \langle i | e^{-\beta \hat{H}/2} \hat{O} e^{-\beta \hat{H}/2} | i \rangle \tag{52}$$

where $\{|i\rangle\}$ is an orthonormal basis set, and Z is the partition function. Defining $|\phi_i\rangle = P_i^{-1/2} e^{-\beta \hat{H}/2} |i\rangle$, we obtain

$$\langle \hat{O} \rangle = \frac{1}{Z} \sum_i P_i \langle \phi_i | \hat{O} | \phi_i \rangle \tag{53}$$

where $P_i = \langle i | e^{-\beta H} | i \rangle$. The summation in Eq.(53) can be estimated by sampling $|\phi_i\rangle$ with probability P_i / Z, and summing the sampled $\langle \phi_i | \hat{O} | \phi_i \rangle$.

In standard Metropolis sampling for thermal states, one starts from $|\phi_i\rangle$ and obtains the next state $|\phi_j\rangle$ from randomly proposing and accepting based an acceptance probability. However, rejecting and resetting in the quantum analog of Metropolis [56] is complicated to implement on a quantum computer, requiring deep circuits. The METTS algorithm provides an alternative way to sample $|\phi_i\rangle$ distributed with probability P_i/Z without this complicated procedure. The algorithm is as follows

1. Choose a classical product state (PS) $|i\rangle$.

2. Compute $|\phi_i\rangle = P_i^{-1/2}e^{-\beta H/2}|i\rangle$ and calculate observables of interest.

3. Collapse $|\phi_i\rangle$ to a new PS $|i'\rangle$ with probability $p(i \to i') = |\langle i'|\phi_i\rangle|^2$ and repeat Step 2.

In the above algorithm, $|\phi_i\rangle$ is named a minimally entangled typical thermal state (METTS). One can easily show that the set of METTS sampled following the above procedure has the correct Gibbs distribution [54]. Generally, $\{|i\rangle\}$ can be any orthonormal basis. For convenience when implementing METTS on a quantum computer, $\{|i\rangle\}$ are chosen to be product states.

On a quantum emulator or a quantum computer, the METTS algorithm is carried out as following

1. Prepare a product state $|i\rangle$.

2. Imaginary time evolve $|i\rangle$ with the QITE algorithm to $|\phi_i\rangle = P_i^{-1/2}e^{-\beta H/2}|i\rangle$, and measure the desired observables.

3. Collapse $|\phi_i\rangle$ to another product state by measurement.

In practice, to avoid long statistical correlations between samples, we used the strategy of collapsing METTS onto alternating basis sets [54]. For instance, for the odd METTS steps, $|\phi_i\rangle$ is collapsed onto the X-basis (assuming a Z computational basis, tensor products of $|+\rangle$ and $|-\rangle$), and for the even METTS steps, $|\phi_i\rangle$ is collapsed onto the Z-basis (tensor products of $|0\rangle$ and $|1\rangle$). The statistical error is then estimated by block analysis [57].

Implementation on emulator and quantum processor

We used pyQuil, an open source Python library, to express quantum circuits that interface with both Rigetti's quantum virtual machine (QVM) and the Aspen-1 quantum processing units (QPUs).

pyQuil provides a way to include noise models in the QVM simulations. Readout error can be included in a high-level API provided in the package and is characterized by p_{00} (the probability of reading $|0\rangle$ given that the qubit is in state $|0\rangle$) and p_{11} (the probability of reading $|1\rangle$ given that the qubit is in state $|1\rangle$). Readout errors can be mitigated by estimating

the relevant probabilities and correcting the estimated expectation values. We do so by using a high level API present in pyQuil. A general noise model can also be applied to a gate in the circuit by applying the appropriate Kraus maps. Included in the package is a high level API that applies the same decoherence error attributed to energy relaxation and dephasing to every gate in the circuit. This error channel is characterized by the relaxation time T_1 and coherence time T_2. We also include in our emulation our own high-level API that applies the same depolarizing noise channel to every single gate by using the appropriate Kraus maps. The depolarizing noise is characterized by p_1, the depolarizing probability for single-qubit gates and p_2, the depolarizing probability for two-qubit gates. We do not include all sources of error in our emulation. We applied the same depolarizing and dephasing channels to each gate operation for all qubits, when in reality, they can vary from qubit to qubit. In addition, noise due to crosstalk between qubits cannot be modeled using the QVM and is another source of discrepancy between the QVM and QPU results.

We investigate the influence of noise on the 2-qubit results obtained via the QVM using different noise parameters; Noise model 1: $p_{00} = 0.95$, $p_{11} = 0.95$, $T_1 = 10.5\mu$s, $T_2 = 14.0\mu$s, $p_1 = 0.001$, p_2=0.01, Noise model 2: $p_{00} = 0.99$, $p_{11} = 0.99$, $T_1 = 10.5\mu$s, T_2 =14.0μs, $p_1 = 0.001$, $p_2 = 0.01$ and, Noise model 3: $p_{00} = 0.99$, $p_{11} = 0.99$, $T_1 = 20.0\mu$s, T_2 =40.0μs, $p_1 = 0.0001$, p_2=0.001. Noise model 1 reflects realistic parameters that characterize the Aspen-1 QPUs we run our calculations on; p_{00}, p_{11}, T_1, and T_2 are reported values whereas p_1 and p_2 are values typically used to benchmark error mitigation algorithms [58]. We repeated 10 calculations for each noise model and note there is practically no variation from run to run. Fig. 5(a) shows that reducing the readout error does not greatly affect the converged ground state energy after readout error mitigation has been performed. However, reducing the other sources of error does improve the converged energy. Note that sufficient measurement samples are used such that the sampling variance is smaller than that due to noise.

We also ran 2-qubit simulations on different pairs of qubits on Aspen-1, with Q1 consisting of qubits 14, 15 and Q2 consisting of qubits 0, 1. These two pairs are reported to have different noise characteristics; Q1: $p_{00} = 0.95$, $p_{11} = 0.95$, $T_1 = 10.5\mu$s, T_2 =14.0μs, and, Q2: $p_{00} = 0.90$, $p_{11} = 0.90$, $T_1 = 6.5\mu$s, T_2 =8.0μs. Based on this, we expect simulations on Q2 to be worse. Note that in contrast to our QVM calculations, the results from the actual devices varied from run to run. Thus, we present the mean and standard deviation for 10 different runs on each pair. (Similarly, sufficient samples are taken when running the QVM such that the sampling variance is smaller than that due to noise). Fig. 5(b) indeed demonstrates that Q2 provides a less faithful implementation of the quantum algorithm.

Parameters used in QVM and QPUs simulations

In this section, we include the parameters used in our QPU and QVM simulations. Note that all noisy QVM simulations (unless stated otherwise in the text) were performed with noise parameters from noise model 1. We also indicate the number of samples used during measurements for each Pauli operator.

TABLE I: QPUs: 1-qubit QITE and QLanczos.

Trotter stepsize	nSamples	δ	s	ϵ
0.2	100000	0.01	0.75	10^{-2}

TABLE II: QPUs: 2-qubit QITE and QLanczos.

Trotter stepsize	nSamples	δ	s	ϵ
0.5	100000	0.1	0.75	10^{-2}

TABLE III: QPUs: 1-qubit METTS.

β	Trotter stepsize	nSamples	nMETTs	δ
1.5	0.15	1500	70	0.01
2.0	0.20	1500	70	0.01
3.0	0.30	1500	70	0.01
4.0	0.40	1500	70	0.01

TABLE IV: QVM: 2-qubit QITE and QLanczos.

Trotter stepsize	nSamples	δ	s	ϵ
0.5	100000	0.1	0.75	10^{-2}

TABLE V: QVM: 1-qubit METTS.

β	Trotter stepsize	nSamples	nMETTs	δ
1.0	0.10	1500	70	0.01
1.5	0.15	1500	70	0.01
2.0	0.20	1500	70	0.01
3.0	0.30	1500	70	0.01
4.0	0.40	1500	70	0.01

TABLE VI: QVM: 2-qubit METTS.

β	Trotter stepsize	nSamples	nMETTs	δ
1.0	0.10	30000	100	0.1
1.5	0.15	30000	100	0.1
2.0	0.20	30000	100	0.1
3.0	0.30	30000	100	0.1
4.0	0.40	30000	100	0.1

[*] Corresponding author: mariomotta31416@gmail.com
[†] Corresponding author: gkc1000@gmail.com

[1] R. P. Feynman, "Simulating physics with computers," *International Journal of Theoretical Physics*, vol. 21, no. 6, pp. 467–488, 1982.

[2] D. S. Abrams and S. Lloyd, "Simulation of many-body Fermi systems on a universal quantum computer," *Phys. Rev. Lett.*, vol. 79, pp. 2586–2589, 1997.

[3] S. Lloyd, "Universal quantum simulators," *Science*, vol. 273, no. 5278, pp. 1073–1078, 1996.

[4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, "Simulated quantum computation of molecular energies," *Science*, vol. 309, no. 5741, pp. 1704–1707, 2005.

[5] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, pp. 242 EP –, Sep 2017.

[6] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, "Extending the computational reach of a noisy superconducting quantum processor," *arXiv:1805.04492*, 2018.

[7] J. Kempe, A. Kitaev, and O. Regev, "The complexity of the local Hamiltonian problem," *SIAM Journal on Computing*, vol. 35, no. 5, pp. 1070–1097, 2006.

[8] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, "Quantum computation by adiabatic evolution." MIT-CTP-2936, 2000.

[9] A. Y. Kitaev, "Quantum measurements and the abelian stabilizer problem," 1995.

[10] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, "A quantum approximate optimization algorithm." MIT-CTP-4610, 2014.

[11] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, C. A. Ryan, D. Scarabelli, M. Scheer, E. A. Sete, P. Sivarajah, R. S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, B. R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, "Unsupervised machine learning on a hybrid quantum computer," 2017.

[12] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. lipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Möller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.

[13] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature Communications*, 2014.

[14] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New Journal of Physics*, vol. 18, no. 2, p. 023023, 2016.

[15] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, "Adapt-vqe: An exact variational algorithm for fermionic simulations on a quantum computer," *arXiv preprint arXiv:1812.11173*, 2018.

[16] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *arXiv preprint arXiv:1803.11173*, 2018.

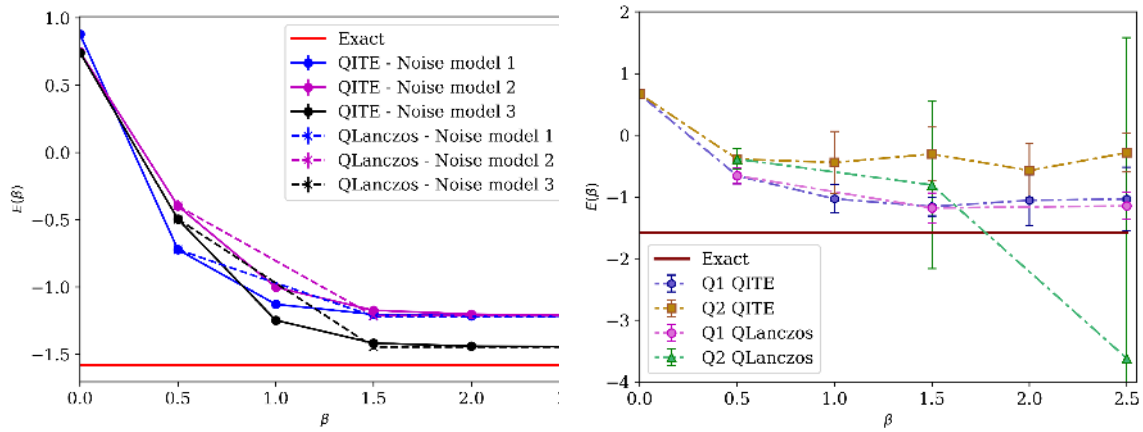[17] S. McArdle, T. Jones, S. Endo, Y. Li, S. Benjamin, and X. Yuan,

FIG. 5: QITE and QLanczos energy $E(\beta)$ as a function of imaginary time β for 2-qubit simulations (a) using the QVM with different noise models, (b) Q1 and Q2 on Aspen-1.

"Variational quantum simulation of imaginary time evolution," 2018.

[18] C. Lanczos, "An iteration method for the solution of the eigenvalue problem of linear differential and integral operators," *J. Res. Natl. Bur. Stand. B*, vol. 45, pp. 255–282, 1950.

[19] S. R. White, "Minimally entangled typical quantum states at finite temperature," *Phys. Rev. Lett.*, vol. 102, p. 190601, May 2009.

[20] E. M. Stoudenmire and S. R. White, "Minimally entangled typical thermal state algorithms," *New Journal of Physics*, vol. 12, no. 5, p. 055026, 2010.

[21] A. Uhlmann, "The "transition probability" in the state space of a*-algebra," *Reports on Mathematical Physics*, vol. 9, no. 2, pp. 273–279, 1976.

[22] M. B. Hastings and T. Koma, "Spectral gap and exponential decay of correlations," *Communications in Mathematical Physics*, vol. 265, no. 3, pp. 781–804, 2006.

[23] F. Verstraete and J. I. Cirac, "Mapping local Hamiltonians of fermions to local Hamiltonians of spins," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2005, no. 09, p. P09012, 2005.

[24] 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*, pp. 792–809, IEEE, 2015.

[25] G. Vidal, "Efficient simulation of one-dimensional quantum many-body systems," *Phys. Rev. Lett.*, vol. 93, p. 040502, Jul 2004.

[26] U. Schollwöck, "The density-matrix renormalization group in the age of matrix product states," *Annals of Physics*, vol. 326, no. 1, pp. 96 – 192, 2011.

[27] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, "Computational complexity of projected entangled pair states," *Phys. Rev. Lett.*, vol. 98, p. 140506, 2007.

[28] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, "Contracting projected entangled pair states is average-case hard," *arXiv preprint arXiv:1810.00738*, 2018.

[29] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-

Guzik, and J. M. Martinis, "Scalable quantum simulation of molecular energies," *Phys. Rev. X*, vol. 6, p. 031007, Jul 2016.

[30] H. Lamm and S. Lawrence, "Simulation of nonequilibrium dynamics on a quantum computer," *Phys. Rev. Lett.*, vol. 121, p. 170501, Oct 2018.

[31] R. Computing, "Quantum Cloud Services." https://qcs.rigetti.com/dashboard, accessed 2019-01-21.

[32] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, "Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states," *Phys. Rev. A*, vol. 95, p. 042308, Apr 2017.

[33] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, "Computation of molecular spectra on a quantum processor with an error-resilient algorithm," *Phys. Rev. X*, vol. 8, p. 011021, Feb 2018.

[34] B. M. Terhal and D. P. DiVincenzo, "Problem of equilibration and the computation of correlation functions on a quantum computer," *Phys. Rev. A*, vol. 61, p. 022301, 2000.

[35] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, "Quantum metropolis sampling," *Nature*, vol. 471, p. 87, 2011.

[36] A. N. Chowdhury and R. D. Somma, "Quantum algorithms for gibbs sampling and hitting-time estimation," *Quantum Information & Computation*, vol. 17, no. LA-UR-16-21218, 2017.

[37] F. G. Brandão and M. J. Kastoryano, "Finite correlation length implies efficient preparation of quantum thermal states," *Communications in Mathematical Physics*, vol. 365, no. 1, pp. 1–16, 2019.

[38] F. G. Brandão and M. Horodecki, "Exponential decay of correlations implies area law," *Communications in mathematical physics*, vol. 333, no. 2, pp. 761–798, 2015.

[39] G. Vidal, "Efficient simulation of one-dimensional quantum many-body systems," *Phys. Rev. B*, vol. 93, p. 040502, 2004.

[40] U. Schollwoeck, "The density-matrix renormalization group in the age of matrix product states," *Ann. Phys.*, vol. 326, pp. 96–192, 2011.

[41] T. Nishino and K. Okunishi, "Corner transfer matrix renormalization group method," *Journal of the Physical Society of Japan*, vol. 65, no. 4, pp. 891–894, 1996.

[42] F. Verstraete and J. I. Cirac, "Renormalization algorithms for quantum-many body systems in two and higher dimensions," *arXiv preprint cond-mat/0407066*, 2004.

[43] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, "Criticality, the area law, and the computational power of projected entangled pair states," *Physical review letters*, vol. 96, no. 22, p. 220601, 2006.

[44] R. Orús, "A practical introduction to tensor networks: Matrix product states and projected entangled pair states," *Annals of Physics*, vol. 349, pp. 117–158, 2014.

[45] F. Verstraete, V. Murg, and J. Cirac, "Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems," *Advances in Physics*, vol. 57, no. 2, pp. 143–224, 2008.

[46] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, "Classical simulation of infinite-size quantum lattice systems in two spatial dimensions," *Phys. Rev. Lett.*, vol. 101, p. 250602, Dec 2008.

[47] H. C. Jiang, Z. Y. Weng, and T. Xiang, "Accurate determination of tensor network state of quantum lattice models in two dimensions," *Phys. Rev. Lett.*, vol. 101, p. 090603, Aug 2008.

[48] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, "Algorithms for finite projected entangled pair states," *Phys. Rev. B*, vol. 90, p. 064425, Aug 2014.

[49] M. Lubasch, J. I. Cirac, and M.-C. Banuls, "Unifying projected entangled pair state contractions," *New Journal of Physics*, vol. 16, no. 3, p. 033014, 2014.

[50] Z. Y. Xie, H. J. Liao, R. Z. Huang, H. D. Xie, J. Chen, Z. Y. Liu, and T. Xiang, "Optimized contraction scheme for tensor-network states," *Phys. Rev. B*, vol. 96, p. 045128, Jul 2017.

[51] J. Hachmann, W. Cardoen, and G. K.-L. Chan, "Multireference correlation in long molecules with the quadratic scaling density matrix renormalization group," *The Journal of chemical physics*, vol. 125, no. 14, p. 144101, 2006.

[52] M. Motta, D. M. Ceperley, G. K.-L. Chan, J. A. Gomez, E. Gull, S. Guo, C. A. Jiménez-Hoyos, T. N. Lan, J. Li, F. Ma, A. J. Millis, N. V. Prokof'ev, U. Ray, G. E. Scuseria, S. Sorella, E. M. Stoudenmire, Q. Sun, I. S. Tupitsyn, S. R. White, D. Zgid, and S. Zhang, "Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods," *Phys. Rev. X*, vol. 7, p. 031059, 2017.

[53] A. Szabo and N. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Books on Chemistry, Dover Publications, 1996.

[54] E. M. Stoudenmire and S. R. White, "Minimally entangled typical thermal state algorithms," *New Journal of Physics*, vol. 12, p. 055026, 2010.

[55] S. R. White, "Minimally entangled typical quantum states at finite temperature," *Phys. Rev. Lett.*, vol. 102, p. 190601, 2009.

[56] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, "Quantum metropolis sampling," *Nature*, vol. 471, p. 87, 2011.

[57] H. Flyvbjerg, "Error estimates on averages of correlated data," *J. Chem. Phys*, vol. 91, p. 461, 1989.

[58] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Phys. Rev. Lett.*, vol. 119, p. 180509, Nov 2017.