

Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution

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 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 Science, California Institute of Technology, Pasadena, CA 91125, USA*

³*Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA 91125, USA*

The accurate computation of Hamiltonian ground, excited, and thermal states on quantum computers stands to impact many problems in the physical and computer sciences, from quantum simulation to machine learning. Given the challenges posed in constructing large-scale quantum computers, these tasks should be carried out in a resource-efficient way. In this regard, existing techniques based on phase estimation or variational algorithms display potential disadvantages; phase estimation requires deep circuits with ancillae, that are hard to execute reliably without error correction, while variational algorithms, while flexible with respect to circuit depth, entail additional high-dimensional classical optimization. Here, we introduce the quantum imaginary time evolution and quantum Lanczos algorithms, which are analogues of classical algorithms for finding ground and excited states. Compared to their classical counterparts, they require exponentially less space and time per iteration, and can be implemented without deep circuits and ancillae, or high-dimensional optimization. We furthermore discuss quantum imaginary time evolution as a subroutine to generate Gibbs averages through an analog of minimally entangled typical thermal states. Finally, we demonstrate the potential of these algorithms via an implementation using exact classical emulation as well as through prototype circuits on the Rigetti quantum virtual machine and Aspen-1 quantum processing unit.

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, though somewhat robust to coherent errors, are limited in accuracy by a fixed Ansatz, and involve high-dimensional noisy classical optimizations [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 $-\partial_\beta |\Phi(\beta)\rangle = \hat{H}|\Phi(\beta)\rangle$, $|\Psi\rangle = \lim_{\beta \rightarrow \infty} \frac{|\Phi(\beta)\rangle}{\|\Phi(\beta)\|}$ (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 family of approaches are variants of the iterative Lanczos method [18]. 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 Hamilto-

nian, 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), the quantum Lanczos (QLanczos) and the quantum analog of the minimally entangled typical thermal states (QMETTS) algorithm, to determine ground-states, ground and excited states and thermal states on a quantum computer. Under the assumption of finite correlation length, these methods rigorously use exponentially reduced space and time per propagation step or iteration, compared to their direct classical counterparts. Even when such assumptions do not hold, the inexact versions of the QITE and QLanczos algorithms remain valid heuristics that can be applied within a limited computational budget, and offer advantages over existing ground-state quantum algorithms, as they do not use deep circuits and converge to their solutions without non-linear optimization. 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 our algorithms potentially suitable for near-term quantum architectures. We demonstrate the algorithms on spin and fermionic Hamiltoni-

ans using exact classical emulation, and demonstrate proof-of-concept implementations on the Rigetti quantum virtual machine (QVM) and Aspen-1 quantum processing 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 neighboring qubits on an underlying graph) and a Trotter decomposition of the corresponding imaginary-time evolution,

$$e^{-\beta \hat{H}} = (e^{-\Delta\tau \hat{h}[1]} e^{-\Delta\tau \hat{h}[2]} \dots)^n + \mathcal{O}(\Delta\tau); \quad n = \frac{\beta}{\Delta\tau} \quad (1)$$

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

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

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 $\{\hat{\sigma}_i\}$ on k qubits,

$$\hat{A}[m] = \sum_{i_1 \dots i_k} a[m]_{i_1 \dots i_k} \hat{\sigma}_{i_1} \dots \hat{\sigma}_{i_k} \equiv \sum_I a[m]_I \hat{\sigma}_I. \quad (4)$$

Up to $\mathcal{O}(\Delta\tau)$, the coefficients $a[m]_I$ 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,

$$S_{I,I'} = \langle \Psi | \hat{\sigma}_I^\dagger \hat{\sigma}_{I'} | \Psi \rangle, \quad b_I = \frac{-i}{\sqrt{c}} \langle \Psi | \hat{\sigma}_I^\dagger \hat{h}[m] | \Psi \rangle. \quad (5)$$

In general, $\mathbf{S}$ has a null space; to ensure $\mathbf{a}[m]$ is real, we minimize $\|\bar{\Psi}' - (1 - i\Delta\tau \hat{A}[m])\Psi\|^2$ 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 $\exp(-L/C)$), $|\bar{\Psi}'\rangle$ can be generated by a unitary acting on a domain of width at most $\mathcal{O}(C)$ qubits surrounding the qubits acted on by $\hat{h}[m]$. This follows from Uhlmann's theorem [19], which states that two pure states with marginals close to each other must be related by a unitary transformation on the purifying subsystem (see SI). The unitary $e^{-i\Delta\tau \hat{A}[m]}$

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 $\mathcal{O}(C^d)$. In many physical systems, we expect the maximum correlation length throughout the Trotter steps to increase with β and saturate for $C_{\max} \ll N$ [20]. 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 fermions, in particular, the locality of the corresponding qubit Hamiltonian depends on the qubit mapping. In principle, a geometric k -local fermionic Hamiltonian can be mapped to a geometric local qubit Hamiltonian [21, 22], allowing above techniques to be applied directly. Alternatively, we conjecture that by constructing Eq. (4) with a local fermionic basis, the unitary update can be constructed over a domain size $D \sim \mathcal{O}(C^d)$, C being the fermionic correlation length (see SI).

Cost of QITE. The number of measurements and classical storage at a given time step (starting propagation from a product state) is bounded by $\exp(\mathcal{O}(C^d))$ (with C the correlation length at that time step), since each unitary at that step acts on at most $\mathcal{O}(C^d)$ sites; classical solution of the least squares problem has a similar scaling $\exp(\mathcal{O}(C^d))$, as does the synthesis and application as a quantum circuit (composed of two-qubit gates) of the unitary $e^{-i\Delta\tau \hat{A}[m]}$. 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, for example if $\hat{A}[m]$ has a locality structure, e.g. if it is (approximately) a p -local Hamiltonian (i.e. all $a[m]_{i_1 \dots i_k}$ in Eq. (4) are zero except for those where at most p of the $\hat{\sigma}_i$ operators differ from the identity) then the cost of tomography becomes only $C^{\mathcal{O}(dp)}$, while the cost of finding and implementing the unitary is $\mathcal{O}(pC^d T_e)$, T_e being the cost of computing one entry of $\hat{A}[m]$ [23]. If we assume further that $\hat{A}[m]$ is geometric local, the cost of tomography is reduced further to $\mathcal{O}(pC^d)$. However, it is important to note that 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 and studied in detail below.

Compared to a direct 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. Note that a finite correlation length C_0 in the ground-




FIG. 1: **Physical foundations of the quantum imaginary time evolution algorithm.** (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 $h = 1.25$, 50 qubits and $h = 3.5$, 21×31 qubits respectively. $I(i, j)$ 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).

state does not generally imply an efficient classical strategy. In the SI, we analyze multiple classical heuristics under the assumption of finite ground-state correlations, including: truncating the problem size at the ground-state correlation length C_0 , classical simulation in the Heisenberg representation, and tensor network calculations [24–27].

Inexact QITE. Given limited resources, for example on near-term devices, we can choose to measure and construct the unitary 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, and larger D gives successively better approximations to the ground-state. Importantly, 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.

Illustrative QITE calculations. To illustrate the QITE algorithm, we have carried out exact classical emulations (assuming perfect expectation values and gates) for several Hamiltonians (see SI): short-range 1D Heisenberg (with and without a field); 1D AFM transverse-field Ising; long-range 1D Heisenberg with spin-spin coupling $J_{ij} = (|i - j| + 1)^{-1}$; 1D Hubbard at half-filling; a 6-qubit MAXCUT [10–12] instance, and a minimal basis 2-qubit dihydrogen molecular Hamiltonian [28]. To assess the feasibility of implementation on near-term quantum devices, we have carried out noisy classical emulation (sampling expectation values and with an error model) using the Rigetti quantum virtual machine (QVM) and a phys-

ical simulation using the Rigetti Aspen-1 QPUs, for a single-qubit field model $(\hat{X} + \hat{Z})/\sqrt{2}$ [29] and a 1D AFM transverse-field Ising model. We also carried out measurement resource estimates for QITE on the short-range 1D Heisenberg (with field) model studied in Ref. [5] with VQE, and the 1D AFM transverse-field Ising model; we compared with resource estimates using the publicly available VQE implementation in IBM’s Qiskit. We carried out QITE using different fixed domain sizes D for the unitary or fermionic unitary (see SI for descriptions of simulations and models).

Fig. 2a-2f 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 ($\beta \rightarrow \infty$) energies rapidly converge to the exact 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. 2g, 2h show an estimate from classical emulation of the number of Pauli string expectation values to be measured in the QITE algorithm as well as the hardware-efficient VQE ansatz (using the optimization protocol in Ref. [5]) to obtain an energy accuracy of 1% in the 1D Heisenberg model with field $J = B = 1$ (Fig. 2g) and 1% or 2% in the 1D AFM transverse-field Ising model (Fig. 2h, the looser threshold was chosen to enable convergence of VQE). QITE is competitive with VQE for the 4-site model and requires significantly




FIG. 2: **Classical simulation and experimental implementation of QITE and QLanczos algorithms.** Column 1: (a) QITE energy $E(\beta)$ and (b) fidelity F between finite-time state $\Phi(\beta)$ and exact ground state Ψ as function of β , for a 1D 10-site Heisenberg model, showing convergence with increasing D . Column 2: QITE (dashed lines) and QLanczos (solid lines) energies as function of 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. Column 3: QITE and QLanczos energy as a function of β for (e) a 1-qubit model and (f) a 2-qubit AFM transverse field Ising model using QVMs and QPUs. Black lines denote the exact ground-state energy or maximum fidelity. Column 4: Estimate of the number of Pauli string expectation values (P_{total}) needed for QITE and VQE to converge within (g) 1% of the exact energy for a 4-site (left) and 6-site (right) 1D Heisenberg model with magnetic field, and (h) 1% (2%) of the exact energy for a 4-site (6-site) 1D AFM transverse-field Ising model. Error bars represent standard deviations computed from multiple runs.

fewer measurements in the 6-site model. While the number of measurements could potentially be reduced in VQE by different optimizers and Ansätze, the data suggests that QITE is a promising alternative to VQE on near-term devices.

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 [30]. 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, we now consider how to formulate a quantum Lanczos algorithm, which is an especially economical realization of a quantum subspace method [31, 32]. An important practical motivation is that the Lanczos algorithm typically converges




FIG. 3: **Application of QITE to long-range spin and fermionic models, and a combinatorial optimization problem.** (a) QITE energy as a function of 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 MAX-CUT 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.

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 natural way to compute excited states. Consider the sequence of imaginary time vectors $|\Phi_l\rangle = e^{-l\Delta\tau\hat{H}}|\Phi\rangle$, $l = 0, 1, \dots, n$,




FIG. 4: **Classical simulation and experimental implementation of the QMETTS algorithm.** 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. Error bars represent (block) standard deviations computed from multiple samples/runs.

where $c_l = \|\Phi_l\|$. In QLanczos, we consider the vectors after even numbers of time steps $|\Phi_0\rangle, |\Phi_2\rangle \dots$ 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, $S_{ll'} = \langle \Phi_l | \Phi_{l'} \rangle = c_{(l+l')/2}^2$, 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 in the basis of the Φ_l states, 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 lower than that of the last imaginary time vector Φ_n (while higher roots approximate excited states).

With a limited computational budget, we can use inexact QITE to generate $\Phi_l, \Phi_{l'}$. However, in this case the above expressions for $S_{ll'}$ and $H_{ll'}$ in terms of expectation values no longer exactly hold, 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, exact quantum Lanczos converges much more rapidly than imaginary time evolution. Convergence of inexact QITE (small domains), however, can both be faster and reach lower energies than inexact quantum Lanczos. 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 [33], to a quantum analog of Metropolis sampling [34] that relies on phase estimation, as well as methods based on ancilla based Hamiltonian simulation with post-selection [35] and approaches based on recovery maps [36]. 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 [37, 38], 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 $\hat{O}$ to produce its thermal average (iii) measure a product operator such as $\hat{Z}_1\hat{Z}_2 \dots \hat{Z}_N$, 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 In summary, the quantum analogs of imaginary-time evolution, Lanczos and METTS algorithms we have presented enable a new class of eigenstate and thermal state quantum simulations, that can be carried out without ancillae or deep circuits and that, for bounded correlation length, achieve exponential reductions in space and time per iteration relative to known classical counterparts. Encouragingly, these algo-

gorithms 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.

Acknowledgments. MM, GKC, FGSLB, ATKT, AJM were supported by the US NSF via RAISE-TAQS CCF 1839204. MJO'R was supported by an NSF graduate fellowship via grant No. DEG-1745301; the tensor network algorithms were developed with the support of the US DOD via MURI FA9550-18-1-0095. EY was supported by a Google fellowship. CS was supported by the US DOE via DE-SC0019374. GKC is a Simons Investigator in Physics and a member of the Simons Collaboration on the Many-Electron Problem. The Rigetti computations were made possible by a generous grant through Rigetti Quantum Cloud services supported by the CQIA-Rigetti Partnership Program. We thank GH Low, JR McClean, R Babbush for discussions, and the Rigetti team for help with the QVM and QPU simulations.

Author contributions and data availability.

MM, CS, GKC designed the algorithms. FGSLB established the mathematical proofs and error estimates. EY and MJO'R performed classical tensor network simulations. MM, CS, ATKT carried out classical exact emulations. ATKT and AJM designed and carried out the Rigetti QVM and QPU experiments. All authors contributed to the discussion of results and writing of the manuscript. The code used to generate the data presented in this study can be publicly accessed on GitHub at <https://github.com/mariomotta/QITE.git>

SUPPLEMENTAL INFORMATION

Representing imaginary-time evolution by unitary maps

In this section, we discuss how to emulate imaginary time evolution by measurement-assisted unitary circuits acting on suitable domains. As discussed in the main text, we map the scaled non-unitary action of $e^{-\Delta\tau\hat{h}[l]}$ on a state Ψ to that of a unitary $e^{-i\Delta\tau\hat{A}[l]}$, i.e.

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

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

$$\hat{A}[l] = \sum_{i_1 i_2 \dots i_D} a[l]_{i_1 i_2 \dots i_D} \hat{\sigma}_{i_1} \hat{\sigma}_{i_2} \dots \hat{\sigma}_{i_D} = \sum_I a[l]_I \hat{\sigma}_I, \quad (7)$$

where I denotes the index $i_1 i_2 \dots i_D$. Define

$$|\Delta_0\rangle = \frac{|\bar{\Psi}'\rangle - |\Psi\rangle}{\Delta\tau}, \quad |\Delta\rangle = -i\hat{A}[l]|\Psi\rangle. \quad (8)$$

Our goal is to minimize the difference $\| \Delta_0 - \Delta \|$. If the unitary $e^{-i\Delta\tau\hat{A}[l]}$ 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 $a[l]$ corresponds to minimizing the quadratic function $f(a[l])$

$$f(a[l]) = f_0 + \sum_I b_I a[l]_I + \sum_{IJ} a[l]_I S_{IJ} a[l]_J \quad (9)$$

where

$$f_0 = \langle\Delta_0|\Delta_0\rangle, \quad (10)$$

$$S_{IJ} = \langle\Psi|\hat{\sigma}_I^\dagger \hat{\sigma}_J|\Psi\rangle, \quad (11)$$

$$b_I = i \langle\Psi|\hat{\sigma}_I^\dagger|\Delta_0\rangle - i \langle\Delta_0|\hat{\sigma}_I|\Psi\rangle, \quad (12)$$

whose minimum obtains at the solution of the linear equation

$$(\mathbf{S} + \mathbf{S}^T) \mathbf{a}[l] = -\mathbf{b} \quad (13)$$

In general, $\mathbf{S} + \mathbf{S}^T$ may have a non-zero null-space. Thus, we solve Eq. (13) either by applying the generalized inverse of $\mathbf{S} + \mathbf{S}^T$ or by an iterative algorithm such as conjugate gradient.

Note that the above results for geometric local Hamiltonians can be extended to Hamiltonians with long-range terms. The primary difference is a modification of the domain over which the unitary acts. For example, for a spin Hamiltonian with long-range pairwise terms, the scaled action of $e^{-\Delta\tau\hat{h}[l]}$ (if $\hat{h}[l]$ acts on qubits i and j) can be emulated, to within accuracy ε , by a unitary constructed in the neighborhoods of i and j , with the domain size given by the result in Eq. (34), see also discussion below). For fermionic Hamiltonians, we replace the Pauli operators in Eq. (7) by fermionic field operators, i.e. $\sigma \in \{1, f, f^\dagger, f^\dagger f\}$, and conjecture that the analogous result holds for the domains for the fermionic operators as for spin operators. For a number conserving Hamiltonian, such as the fermionic Hubbard Hamiltonian treated in Fig. 3 in the main text, we retain only those terms with an equal number of creation and annihilation operators, to conserve particle number.

Real Hamiltonians and states

As the cost to construct the quantities $\mathbf{S}$, $\mathbf{b}$ and to solve the linear system (13) increases exponentially with D , it is natural to seek out simplifications that can be made rigorously. For example, as described above, to emulate the non-unitary action of a number-conserving fermionic Hamiltonian, we can consider $\hat{A}[l]$ to contain only fermionic operator strings that also conserve particle number. In the main text, we also considered the case where $\hat{A}[l]$ is itself approximately p -local, which removes the exponential dependence on D .

Another common scenario concerns Hamiltonians and states that only have real matrix elements and coefficients in the Z computational basis. Then, since $\langle\Psi|$ and $|\Delta_0\rangle$ are real in the computational basis and

$$b_I = -2 \text{Im} \left[\langle\Psi|\hat{\sigma}_I^\dagger|\Delta_0\rangle \right], \quad (14)$$

$b_I \equiv 0$ unless the matrix elements of $\hat{\sigma}_I^\dagger$ have non-zero imaginary part. When the Pauli basis is used, this means that $b_I \equiv 0$ unless $\hat{\sigma}_I^\dagger$ contains an odd number of $\hat{Y}$ operators.

The number of such operators for a domain size D , $y(D)$, is $2^D \frac{2^D - 1}{2} \simeq \frac{4^D}{2}$. This can be shown by induction over D . For $D = 1$, one has

$$y(D = 1) = 1 \quad (15)$$

Pauli strings with an odd number of $\hat{Y}$'s (i.e. just $\hat{Y}$). For $D = 2$, one has

$$y(D = 2) = 6 \quad (16)$$

such strings, namely

$$\hat{I}\hat{Y}, \hat{X}\hat{Y}, \hat{Z}\hat{Y}, \hat{Y}\hat{I}, \hat{Y}\hat{X}, \hat{Y}\hat{Z} \quad (17)$$

For $D = 3$, the number grows to

$$y(D = 3) = 28 \quad (18)$$

with strings

$$\begin{aligned} &\hat{I}\hat{I}\hat{Y}, \dots, \hat{I}\hat{Y}\hat{Z} \\ &\hat{X}\hat{I}\hat{Y}, \dots, \hat{X}\hat{Y}\hat{Z} \\ &\hat{Y}\hat{I}\hat{I}, \dots, \hat{Y}\hat{Z}\hat{Z} \\ &\hat{Z}\hat{I}\hat{Y}, \dots, \hat{Z}\hat{Y}\hat{Z} \end{aligned} \quad (19)$$

and so on. Pauli strings of length $D + 1$ containing an odd number of $\hat{Y}$ operators are obtained either by attaching a $\hat{Y}$ operator to a length- D string containing an even number of $\hat{Y}$ operators, or by attaching a $\hat{I}$, $\hat{X}$, $\hat{Z}$ operator to a length- D string containing an odd number of $\hat{Y}$ operators. Therefore, $O(D)$ obeys the recursion relation

$$y(D + 1) = 3y(D) + (4^D - y(D)) \quad (20)$$

This recursion relation is solved, with the initial condition $y(D = 1) = 1$, by

$$y(D) = 2^D \frac{2^D - 1}{2} \quad (21)$$

This result means that both $\mathbf{b}$ and $\mathbf{S} + \mathbf{S}^T$ can be assembled from $y(D)$ Pauli string expectation values, roughly half the number of measurements needed if one did not assume real Hamiltonians and states. Further, the dimension of $\mathbf{b}$ and $\mathbf{S} + \mathbf{S}^T$ is $y(D)$ and $y(D) \times y(D)$ respectively. Asymptotically, this reduces the cost of solving the linear system Eq. (13) by a factor of 1/8, assuming dense matrix techniques.

Rigorous run time bounds

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

$$H = \sum_{l=1}^m \hat{h}[l], \quad (22)$$

acting on a d -dimensional lattice with $\|h_i\| \leq 1$, where $\|*\|$ is the operator norm. Note that, if a quantum chemistry system is studied with an orthonormal basis of spatially localized states, such states can be approximately positioned on a lattice, and the results of this section apply on length scales larger than the size of the employed basis functions. In imaginary-time evolution one typically applies Trotter formulae to approximate

$$\frac{e^{-\beta \hat{H}} |\Psi_0\rangle}{\|e^{-\beta \hat{H}} |\Psi_0\rangle\|} \simeq \frac{\left(e^{-\Delta\tau \hat{h}[1]} \dots e^{-\Delta\tau \hat{h}[m]} \right)^n |\Psi_0\rangle}{\left\| \left(e^{-\Delta\tau \hat{h}[1]} \dots e^{-\Delta\tau \hat{h}[m]} \right)^n |\Psi_0\rangle \right\|} \quad (23)$$

for an initial state $|\Psi_0\rangle$ (which we assume to be a product state). 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 $e^{-\Delta\tau \hat{h}[i]}$ from $(e^{-\Delta\tau \hat{h}[1]} \dots e^{-\Delta\tau \hat{h}[m]})^n$; with this notation $|\Psi_{mn}\rangle$ is the state given by Eq. (23). In the QITE algorithm, instead of applying each of the operators $e^{-\Delta\tau \hat{h}[i]}$ to $|\Psi_0\rangle$ (and renormalizing the state), one applies local unitaries $\hat{U}_s$ 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 $\hat{A}$ and $\hat{B}$ acting on domains separated by $\text{dist}(A, B)$ sites,

$$\begin{aligned} C_s(\hat{A}, \hat{B}) &= \langle \Psi_s | \hat{A} \otimes \hat{B} | \Psi_s \rangle - \langle \Psi_s | \hat{A} | \Psi_s \rangle \langle \Psi_s | \hat{B} | \Psi_s \rangle \\ &\leq \|\hat{A}\| \|\hat{B}\| e^{-\text{dist}(A, B)/C}. \end{aligned} \quad (24)$$

Theorem 1. For every $\varepsilon > 0$, there are unitaries $\hat{U}_s$ each acting on

$$N_q = k (2C)^d \ln^d \left(2\sqrt{2} nm \varepsilon^{-1} \right) \quad (25)$$

qubits, such that

$$\| |\Psi_{mn}\rangle - |\Phi_{mn}\rangle \| \leq \varepsilon. \quad (26)$$

Proof. We have

$$\begin{aligned} \| |\Psi_s\rangle - |\Phi_s\rangle \| &= \| |\Psi_s\rangle - \hat{U}_s |\Phi_{s-1}\rangle \| \\ &\leq \| |\Psi_s\rangle - \hat{U}_s |\Psi_{s-1}\rangle \| + \| |\Psi_{s-1}\rangle - |\Phi_{s-1}\rangle \|. \end{aligned} \quad (27)$$

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 $\text{tr}_{R_v}(|\Psi_s\rangle\langle\Psi_s|)$ be the reduced state on R_v , obtained by partial tracing over the complement of R_v in the lattice. Since

$$|\Psi_s\rangle = \frac{e^{-\Delta\tau \hat{h}[i_s]} |\Psi_{s-1}\rangle}{\|e^{-\Delta\tau \hat{h}[i_s]} |\Psi_{s-1}\rangle\|}, \quad (28)$$

it follows from Eq. (24) and Lemma 9 of [39] that

$$\begin{aligned} & \left\| \text{tr}_{R_v}(|\Psi_s\rangle\langle\Psi_s|) - \text{tr}_{R_v}(|\Psi_{s-1}\rangle\langle\Psi_{s-1}|) \right\|_1 \\ & \leq \|e^{\Delta\tau\hat{h}[i_s]}\|^{-1} e^{-\frac{\nu}{2C}} \leq 2e^{-\frac{\nu}{2C}}, \end{aligned} \quad (29)$$

where we used that for $n \geq 2\beta$,

$$\|e^{-\Delta\tau\hat{h}[i_s]}\| \geq \|I - \Delta\tau\hat{h}[i_s]\| \geq 1 - \Delta\tau \geq 1/2. \quad (30)$$

Above, $\|\cdot\|_1$ is the trace norm. The key result in our analysis is Uhlmann's theorem (see e.g. Lemmas 11 and 12 of [39]). It states 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 with partial traces $|\eta\rangle_A$ and $|\nu\rangle_A$ over the complement of A , such that $\| |\eta\rangle_A - |\nu\rangle_A \|_1 \leq \delta$, then there exists a unitary $\hat{V}$ acting on B such that

$$\| |\eta\rangle_{AB} - (I \otimes \hat{V}) |\nu\rangle_{AB} \| \leq 2\sqrt{\delta}. \quad (31)$$

Applying Uhlmann's theorem to $|\Psi_s\rangle$ and $|\Psi_{s-1}\rangle$, with $B = R_v$, and using Eq. (29), we find that there exists a unitary $\hat{U}_s$ acting on R_v s.t.

$$\| |\Psi_s\rangle - \hat{U}_s |\Psi_{s-1}\rangle \| \leq 2\sqrt{2} e^{-\frac{\nu}{2C}}, \quad (32)$$

which by Eq. (27) implies

$$\| |\Psi_{nm}\rangle - |\Phi_{nm}\rangle \| \leq 2\sqrt{2} mn e^{-\frac{\nu}{2C}}, \quad (33)$$

Choosing $\nu = 2C \ln(2\sqrt{2}nm\epsilon^{-1})$ 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 $\hat{U}_s$ acts on at most

$$N_q = k(2C)^d \ln^d(2\sqrt{2}nm\epsilon^{-1}) \quad (34)$$

qubits. $\square$

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

$$|\Psi_{s-1}\rangle \text{ and } \frac{e^{-\Delta\tau\hat{h}[i_s]}|\Psi_{s-1}\rangle}{\|e^{-\Delta\tau\hat{h}[i_s]}|\Psi_{s-1}\rangle\|}, \quad (35)$$

using $e^{-\Delta\tau\hat{h}[i_s]} = I - \Delta\tau\hat{h}[i_s] + O(\frac{1}{n^2})$ 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. $\hat{U}_s$ can be written as $e^{i\frac{\hat{A}[s]}{n}}$.

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 $\exp(O(N_q))$. This is also the cost to solve classically the linear system which gives the associated Hamiltonian $\hat{A}[s]$ and of finding a circuit decomposition of $\hat{U}_s = e^{i\frac{\hat{A}[s]}{n}}$ 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

$$T = mn e^{O(N_q)} = mn e^{O(k(2C)^d \ln^d(2\sqrt{2}nm\epsilon^{-1}))}. \quad (36)$$

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 $e^{-\beta\hat{H}}|\Psi_0\rangle$, 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 $\epsilon > 0$, there are unitaries $\hat{U}_s$ each acting on*

$$N_q = k(2C)^d \ln^d \left(2\sqrt{2}n \left(|S| + C \ln^d(8nC(2C)^{d+1}\epsilon^{-1}) \right) \right) \quad (37)$$

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

$$\| \text{tr}_{\setminus S}(|\Psi_{mn}\rangle\langle\Psi_{mn}|) - \text{tr}_{\setminus S}(|\Phi_{mn}\rangle\langle\Phi_{mn}|) \|_1 \leq \epsilon. \quad (38)$$

Proof. Consider the unitaries $\hat{U}_s$ obtained in the proof of Theorem 1 satisfying Eq. (32).

Consider the replacement of the local term of the Trotter expansion by the unitary $\hat{U}_s$ 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 [39] that the total error η on the reduced density matrix in region S can be bounded as

$$\eta = n \int_{2C \ln(1/\delta)}^{\infty} e^{-l/2C} l^d dl \leq 4nC(2C)^{d+1}\delta. \quad (39)$$

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

$$\eta = (|S| + C \log(1/\delta))^d n 2\sqrt{2} e^{-\frac{\nu}{2C}}. \quad (40)$$

Choosing

$$\delta = \frac{\varepsilon}{(8nC(2C)^{d+1})}$$

$$\nu = 2C \ln \left(2\sqrt{2}n \left(|S| + C \ln (8nC(2C)^{d+1}\varepsilon^{-1})^d \right) \right) \quad (41)$$

gives the result. $\square$

Non-local Terms: Suppose the Hamiltonian has a term $\hat{h}[q]$ acting on qubits which are not nearby, e.g. on two sites i and j . Then $e^{-\Delta\tau\hat{h}[q]}$ 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 $\hat{P}_{i,i+1} = |0,0\rangle\langle 0,0|_{i,i+1} + |1,1\rangle\langle 1,1|_{i,i+1}$. Then

$$\hat{P}_{1,2}\hat{P}_{2,3}\dots\hat{P}_{n-1,n}|+\rangle^{\otimes n}, \quad (42)$$

with $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, is the GHZ state $(|0\dots 0\rangle + |1\dots 1\rangle)/\sqrt{2}$, which has correlation length $C = n$. While the projector $\hat{P}_{i,i+1}$ cannot appear as a local term $e^{-\Delta\tau\hat{h}[i]}$ 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

$$|\psi_n\rangle := \frac{\left(e^{-\Delta\tau\hat{h}[1]} \dots e^{-\Delta\tau\hat{h}[m]} \right)^n |\Psi_0\rangle}{\left\| \left(e^{-\Delta\tau\hat{h}[1]} \dots e^{-\Delta\tau\hat{h}[m]} \right)^n |\Psi_0\rangle \right\|}. \quad (43)$$

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 [39])

$$C_s(\hat{A}, \hat{B}) = \langle \Psi_s | \hat{A} \otimes \hat{B} | \Psi_s \rangle - \langle \Psi_s | \hat{A} | \Psi_s \rangle \langle \Psi_s | \hat{B} | \Psi_s \rangle$$

$$\leq \|\hat{A}\| \|\hat{B}\| 2^{2n} e^{-\Delta \text{dist}(A,B)}. \quad (44)$$

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 behavior 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

$$I(i, j) = S(i) + S(j) - S(i, j) \quad (45)$$

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)

$$\hat{H}_{TFI} = - \sum_{\langle ij \rangle} \hat{Z}_i \hat{Z}_j - h \sum_i \hat{X}_i \quad (46)$$

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 chain 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) [40, 41] 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 $1.0\text{e-}8$ 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 [42–45] 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 [46] 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 [47] updated via the so-called “simple update” procedure [48]. 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) [49, 50]. 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 behavior.

Accuracy: Even though the simple update procedure was used for the tensor update, we still needed to contract the 21×31 PEPS 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 [51], 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 [49] suggests that errors with respect to $D = \infty$ may be larger, such analysis also confirms that a $D = 5$ PEPS captures the qualitative behavior 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 $I \rightarrow 0$ as $r \rightarrow \infty$.

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.

Comparison to classical algorithms

In the main text, we noted that QITE provided an exponential speedup per iteration over the direct classical implementation of imaginary time evolution algorithm, given bounded correlation length C during the evolution. We now compare to some other possible classical algorithms.

We first note that a finite correlation length C_0 in the ground-state does not itself imply an efficient classical strategy. For example, a simple heuristic is to solve the problem locally, e.g. to truncate the problem size at correlation length C_0 of the ground-state and solve by exact diagonalization, which can be done in time $\exp(O(C_0 d))$ in d spatial dimensions. But this will not generally converge to the correct ground-state in a frustrated Hamiltonian, as this would efficiently solve NP-hard classical satisfiability problems even though these have $C_0 = 0$; physical examples include glassy models.

Similarly, as QITE defines a quantum circuit for the imaginary time evolution, we might 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 $(e^{-\Delta\tau\hat{h}[1]}e^{-\Delta\tau\hat{h}[2]}\dots)^l$ 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 $e^{-\Delta\tau\hat{h}[l]}|\Psi\rangle$ can be represented as a classical tensor network with increased bond dimension [24, 25]. However, the bond dimension will scale as $\exp(O(lD))$. Further, apart from the extra exponential dependence on l , another potential drawback in the tensor network 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) [24–27].

Simulation models

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

1 qubit field model

$$\hat{H} = \alpha \hat{X} + \beta \hat{Z} \quad (47)$$

This Hamiltonian has previously been used as a model for quantum simulations on physical devices in Ref. [29]. 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

$$\hat{H} = \sum_{\langle ij \rangle} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j, \quad (48)$$

the 1D short-range Heisenberg Hamiltonian in the presence of a field as

$$\hat{H} = J \sum_{\langle ij \rangle} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j + B \sum_i Z_i, \quad (49)$$

the 1D long-range Heisenberg Hamiltonian as

$$\hat{H} = \sum_{i \neq j} \frac{1}{|i - j| + 1} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j, \quad (50)$$

and the 1D AFM transverse-field Ising Hamiltonian as

$$\hat{H} = J \sum_{\langle ij \rangle} \hat{Z}_i \hat{Z}_j + h \sum_i \hat{X}_i. \quad (51)$$

1D Hubbard model

The 1D Hubbard Hamiltonian is defined as

$$\hat{H} = - \sum_{\langle ij \rangle \sigma} \hat{a}_{i\sigma}^\dagger \hat{a}_{j\sigma} + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} \quad (52)$$

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 $i = 0 \dots n - 1$, and the $2n - 1$ basis functions as $|\varphi_0\rangle = |0 \uparrow\rangle$, $|\varphi_1\rangle = |0 \downarrow\rangle$, $|\varphi_2\rangle = |1 \uparrow\rangle$, $|\varphi_3\rangle = |1 \downarrow\rangle \dots$. Under Jordan-Wigner transformation, recalling that

$$\begin{aligned} \hat{n}_p &= \frac{1 - \hat{Z}_p}{2}, \\ \hat{a}_p^\dagger \hat{a}_q + \hat{a}_q^\dagger \hat{a}_p &= \frac{\hat{X}_p \hat{X}_q \prod_{k=q+1}^{p-1} \hat{Z}_k (1 - \hat{Z}_p \hat{Z}_q)}{2}, \end{aligned} \quad (53)$$

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

$$\begin{aligned} \hat{H} &= - \sum_p \frac{\hat{X}_p \hat{X}_{p+2} \hat{Z}_{p+1} (1 - \hat{Z}_p \hat{Z}_{p+2})}{2} \\ &+ U \sum_{p \text{ even}} \frac{(1 - \hat{Z}_{2i})(1 - \hat{Z}_{2i+1})}{4} + \mu \sum_p \frac{(1 - \hat{Z}_p)}{2} \end{aligned} \quad (54)$$

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 [52, 53] and has previously been studied in quantum simulations, for example in [28]. 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 [54]

$$\hat{H} = H_0 + \sum_{pq} h_{pq} \hat{a}_p^\dagger \hat{a}_q + \frac{1}{2} \sum_{pqrs} v_{pqrs} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_s \hat{a}_r \quad (55)$$

where $a^\dagger, a$ are fermionic creation and annihilation operators for the molecular orbitals. The Hamiltonian (55) is then encoded by a Bravyi-Kitaev transformation into the 2-qubit operator

$$\hat{H} = g_0 + g_1 \hat{Z}_1 + g_2 \hat{Z}_2 + g_3 \hat{Z}_1 \hat{Z}_2 + g_4 \hat{X}_1 \hat{X}_2 + g_5 \hat{Y}_1 \hat{Y}_2, \quad (56)$$

with coefficients g_i given in Table I of [28].

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 and links

$$\begin{aligned} V &= \{0, 1, 2, 3, 4, 5\}, \\ E &= \{(0, 3), (1, 4), (2, 3), (2, 4), (2, 5), (4, 5)\}, \end{aligned} \quad (57)$$

respectively. 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 $|z_0 \dots z_{n-1}\rangle$, 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

$$\hat{C} = - \sum_{(ij) \in E} \frac{1 - \hat{Z}_i \hat{Z}_j}{2} . \quad (58)$$

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. (13) 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

$$|\Phi_l\rangle = \frac{e^{-l\Delta\tau\hat{H}}|\Psi_T\rangle}{\|e^{-l\Delta\tau\hat{H}}\Psi_T\|} \equiv n_l e^{-l\Delta\tau\hat{H}}|\Psi_T\rangle \quad 0 \leq l < L_{\max} . \quad (59)$$

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

$$S_{l,l'} = \langle \Phi_l | \Phi_{l'} \rangle , \quad H_{l,l'} = \langle \Phi_l | \hat{H} | \Phi_{l'} \rangle \quad (60)$$

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

$$S_{l,l'} = n_l n_{l'} \langle \Psi_T | e^{-l\Delta\tau\hat{H}} e^{-l'\Delta\tau\hat{H}} | \Psi_T \rangle = \frac{n_l n_{l'}}{n_r^2} , \quad (61)$$

and similarly

$$\begin{aligned} H_{l,l'} &= n_l n_{l'} \langle \Psi_T | e^{-l\Delta\tau\hat{H}} \hat{H} e^{-l'\Delta\tau\hat{H}} | \Psi_T \rangle = \\ &= \frac{n_l n_{l'}}{n_r^2} \langle \Phi_r | \hat{H} | \Phi_r \rangle = S_{l,l'} \langle \Phi_r | \hat{H} | \Phi_r \rangle . \end{aligned} \quad (62)$$

The quantities n_r can be evaluated recursively, since

$$\begin{aligned} \frac{1}{n_{r+1}^2} &= \langle \Psi_T | e^{-(r+1)\Delta\tau\hat{H}} e^{-(r+1)\Delta\tau\hat{H}} | \Psi_T \rangle = \\ &= \frac{\langle \Phi_r | e^{-2\Delta\tau\hat{H}} | \Phi_r \rangle}{n_r^2} , \end{aligned} \quad (63)$$

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 $\mathbf{H}\mathbf{x} = E\mathbf{S}\mathbf{x}$ 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 [55, 56] is a sampling method to calculate thermal properties based on imaginary time evolution. Consider the thermal average of an observable $\hat{O}$

$$\langle \hat{O} \rangle = \frac{1}{Z} \text{Tr}[e^{-\beta\hat{H}} \hat{O}] = \frac{1}{Z} \sum_i \langle i | e^{-\frac{\beta}{2}\hat{H}} \hat{O} e^{-\frac{\beta}{2}\hat{H}} | i \rangle , \quad (64)$$

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

$$\langle \hat{O} \rangle = \frac{1}{Z} \sum_i P_i \langle \phi_i | \hat{O} | \phi_i \rangle . \quad (65)$$

The summation in Eq. (65) 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 on an acceptance probability. However, rejecting and resetting in the quantum analog of Metropolis [57] 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 $|\phi_i\rangle = P_i^{-1/2} e^{-\frac{\beta}{2}\hat{H}}|i\rangle$ and calculate observables of interest.
3. Collapse the state $|\phi_i\rangle$ to a new PS $|i'\rangle$ with probability $p(i \rightarrow i') = |\langle i'|\phi_i\rangle|^2$ 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 [55]. 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 $|\phi_i\rangle = P_i^{-1/2} e^{-\frac{\beta}{2}\hat{H}}|i\rangle$, 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 [55]. 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 [58].

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 cross-talk 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:

$$\begin{aligned} p_{00} &= 0.95 & p_{11} &= 0.95 \\ T_1 &= 10.5 \mu s & T_2 &= 14.0 \mu s \\ p_1 &= 0.001 & p_2 &= 0.01 \end{aligned}$$

- Noise model 2:

$$\begin{aligned} p_{00} &= 0.99 & p_{11} &= 0.99 \\ T_1 &= 10.5 \mu s & T_2 &= 14.0 \mu s \\ p_1 &= 0.001 & p_2 &= 0.01 \end{aligned}$$

- Noise model 3:

$$\begin{aligned} p_{00} &= 0.99 & p_{11} &= 0.99 \\ T_1 &= 20.0 \mu s & T_2 &= 40.0 \mu s \\ p_1 &= 0.0001 & p_2 &= 0.001 \end{aligned}$$

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 [59]. 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:

$$\begin{aligned} p_{00} &= 0.95 & p_{11} &= 0.95 \\ T_1 &= 10.5 \mu s & T_2 &= 14.0 \mu s \end{aligned}$$

- Q2:

$$p_{00} = 0.90 \quad p_{11} = 0.90$$

$$T_1 = 6.5 \mu s \quad T_2 = 8.0 \mu s$$

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.




FIG. 5: Comparison of energies obtained using different noise models(NM) for (a) QITE and (b) QLanczos. Comparison of energies obtained using different pair of qubits for (c) QITE and (d) QLanczos. The performance of QITE and QLanczos improves as noise is reduced, indicating the potential of the algorithms.

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

Comparison of QITE and VQE

To address the feasibility of running QITE for larger systems on near-term devices, we compared the total number of Pauli string measurements needed for both VQE and QITE to obtain the ground state of two different spin models; (a) a 1D Heisenberg chain in a magnetic field with the parameters $J = B = 1$; the 4-site instance of this model was studied in Ref. [60], and (b) 1D AFM transverse-field Ising model ($J = h = 1/\sqrt{2}$). Specifically, we estimated how many expectation values of Pauli strings would need to be measured to obtain the ground state of a 4-site and 6-site instance. The state under the evolution of QITE or in VQE was said to be converged to the ground state if its energy was within 1% of the exact ground state energy for the Heisenberg model, and 1% or 2% for the Ising model (the relaxed criterion for the Ising model was chosen so that the VQE optimization could complete in a reasonable number of steps). The following section describes how we counted the total number of Pauli string measurements in VQE and QITE.

Counting Pauli strings in VQE

To perform the VQE calculations, we used the hardware-efficient variational Ansatz as described in [60]. This consists of first applying rotation unitaries represented by $U^{q,i}(\theta) = R_{z_{\theta_1^q,i}} R_{x_{\theta_2^q,i}} R_{z_{\theta_3^q,i}}$ to all qubits before applying layers of a certain depth d ; each layer begins by applying CZ gates between nearest-neighbors followed by applying $U^{q,i}(\theta) = R_{z_{\theta_1^q,i}} R_{x_{\theta_2^q,i}} R_{z_{\theta_3^q,i}}$ to all qubits again. Details of the circuit can be seen in Fig. 6.

As in [60], we also used the simultaneous perturbation stochastic approximation (SPSA) algorithm as the optimization protocol. The SPSA algorithm is commonly used because (i) it performs well in the presence of stochastic fluctuations and (ii) it requires only evaluating the objective function twice to update the variational parameters regardless of the number of parameters involved. The performance of the optimizer depends on the hyperparameters α and γ as described in [60] and we found that their reported values of $\alpha = 0.602$ and $\gamma = 0.101$ also gave the best results for us. Numerical evidence of this is provided later on.

Evaluating the objective function involves estimating the expectation value of the Pauli strings that appear in the Hamiltonian. To prevent sampling errors from influencing the comparison, we evaluated the expectation values exactly. We conducted the VQE calculations using Qiskit, a quantum emulator Python package provided by IBM. The package provides both the SPSA algorithm and a variational Ansatz which we modified to reproduce the exact Ansatz used in [60].

To count the number of Pauli strings needed for convergence, we ran VQE using different layer depths and determined the number of iterations N needed for the algorithm to converge to a state with an energy within a certain percent-

age (1% or 2%) of the exact ground state energy. Examples of converged VQE calculations for the 1D Heisenberg model are given in Fig. 7(a) and (b). In each iteration, the objective function was evaluated twice, and the evaluation of the objective function required measuring the expectation value of the M Pauli strings that appear in the Hamiltonian. Therefore, the total number of Pauli strings P_{total} is given as

$$P_{total} = 2 \times N \times M \quad (66)$$

We note that the results from one VQE trajectory can differ slightly from the next. For our VQE calculations, we always performed 10 trajectories and analyzed our data using the average trajectory. We summarize the VQE parameters that we found gave the lowest number of total Pauli measurements to converge to the ground state in Table VII. For the 6-site 1D AFM transverse field Ising model, VQE could not converge to within 1% and we instead used the 2% convergence criterion. We found that for $\alpha = 0.602, \gamma = 0.101$, our simulation results for the 6-site 1D Heisenberg model indicates that using a circuit depth of 20 requires the least number of total Pauli measurements. We also ran some tests to determine what values of α and γ gave the best result for the 6-site 1D Heisenberg model; the VQE calculation for the 6-site model conducted using $\alpha = 0.602, \gamma = 0.101$ and a circuit depth of 20 converged within 8400 optimization steps. We ran VQE calculations for different α and γ using the same circuit depth of 20 and a total of 9000 optimization steps. The data in table VIII clearly shows that $\alpha = 0.602$ and $\gamma = 0.101$ produced the best result for us.

Counting Pauli strings in QITE

To implement QITE, we used the second-order Trotter decomposition given by

$$e^{-\beta \hat{H}} = (e^{-\Delta\tau/2\hat{h}[1]} \dots e^{-\Delta\tau/2\hat{h}[K-1]} e^{-\Delta\tau\hat{h}[K]} \quad (67)$$

$$e^{-\Delta\tau/2\hat{h}[K-1]} \dots e^{-\Delta\tau\hat{h}[1]})^n + \mathcal{O}(\Delta\tau^2); n = \frac{\beta}{\Delta\tau}$$

to carry out the real time evolution. We initialized our state as: (a) $|0101\dots\rangle$ for the 1D Heisenberg model and (b) maximally-mixed state for the 1D AFM transverse-field Ising model. We converged to the ground state using a time step of $\Delta\tau = 0.1$ and a domain size D of 4, as seen in Figs. 7(c) and (d). To count the number of Pauli strings, we note that a domain size of 4 implies that to evaluate $e^{-\Delta\tau/2\hat{h}[i]}$ involves measuring $4^4 = 256$ Pauli strings (without using the real-valued nature of the Hamiltonian). Therefore, with a total number of Trotter steps T , the total number of Pauli strings P_{total} is given as

$$P_{total} = (2K - 1) \times T \times 256 \quad (68)$$

We summarize the parameters that we used for QITE to obtain the ground state in table IX. We had no trouble converging our ground state to arbitrary accuracy using QITE but we




FIG. 6: (i) VQE Ansatz that is composed of a sequence of interleaved single-qubit rotations $U^{q,i}(\theta)$ and entangling operations. (ii) The entangling operations consist of applying CZ gates between nearest neighbours.




FIG. 7: VQE calculations for (a) 4-site and (b) 6-site 1D Heisenberg model. k is the number of optimization steps, and d is the number of layers. QITE calculations for (c) 4-site and (d) 6-site for the same model. D is the domain size.

used the same convergence criterion as for VQE to facilitate comparison.

VQE and QITE

Data from Table X suggests that QITE is competitive with VQE with respect to the number of Pauli string measurements. In fact, for the 6 qubit system, the number of measurements needed in QITE was significantly less than in VQE, due largely to the SPSA iterations needed to reach convergence when optimizing the VQE energy. While it is likely that the VQE costs could be lowered by using a better optimizer, or a better VQE Ansatz, we also note that the counts for QITE can also be reduced by using the methods outlined in the main text and earlier sections that discussed how one can economize measurements in QITE. The widespread current implementation of VQE and the observed performance of

QITE suggest that it will be practical to implement the QITE protocol for intermediate system sizes on near-term devices.

TABLE VII: VQE simulation parameters for (a) 1D Heisenberg with applied field and (b) 1D AFM transverse field Ising. Conv. refers to the convergence criterion used. We note for the last case, the VQE optimization could not reach within 1% of the ground state energy, so we set the convergence criterion to 2 %

model	n-site	conv.	α	γ	d	N	P_{total}
a	4	1%	0.602	0.101	8	800	25,600
a	6	1%	0.602	0.101	20	8400	403,200
b	4	1%	0.602	0.101	12	800	12,800
b	6	2%	0.602	0.101	12	2890	69,360

TABLE VIII: Hyperparameters sweep for 6-site 1D Heisenberg model using a circuit depth of 20 for a total of 9000 optimization steps. The step at which the calculation converged is recorded under column T . '-' indicates that VQE failed to converge.

α	γ	T
0.400	0.066	-
0.400	0.101	-
0.400	0.133	-
0.602	0.066	-
0.602	0.101	8400
0.602	0.133	8800
0.800	0.066	-
0.800	0.101	-
0.800	0.133	-

TABLE IX: QITE simulation parameters for (a) 1D Heisenberg with applied field and (b) 1D AFM transverse field Ising. Conv. indicates the convergence criterion used. We used 2% for the final calculation to facilitate comparison with VQE which failed to converge to within 1%.

model	n-site	conv.	$\Delta\tau$	D	T	K	P_{total}
a	4	1%	0.1	4	7	4	12,544
a	6	1%	0.1	4	17	6	47,872
b	4	1%	0.2	4	7	4	12,544
b	6	2%	0.2	4	8	6	22,528

TABLE X: Total Pauli string expectation values in VQE and QITE for (a) 1D Heisenberg with applied field and (b) 1D AFM transverse field Ising. The total number of Pauli strings to be measured in QITE can be further reduced by using only Pauli strings with only an odd number of $\hat{Y}$ operators due to the real nature of the Hamiltonian. We show this reduced number in brackets.

model	n-site	VQE	QITE
a	4	25,600	12,544(5,880)
a	6	403,200	47,872(22,440)
b	4	12,800	12,544(5,880)
b	6	69,360	22,528(10,560)

* Corresponding author. ORCID 0000-0003-1647-9864. E-mail: mariomotta31416@gmail.com

† Corresponding author. ORCID 0000-0001-8009-6038. E-mail: gkc1000@gmail.com

- [1] R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **79**, 2586 (1997).
- [3] S. Lloyd, *Science* **273**, 1073 (1996).
- [4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
- [5] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature* **549**, 242 (2017).
- [6] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, *Nature* **567**, 491 (2019).
- [7] J. Kempe, A. Kitaev, and O. Regev, *SIAM J. Comput.* **35**, 1070 (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), [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [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. Scarbelle, 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), [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [12] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Ipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Möller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, *Quant. Sci. Tech.* **3**, 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, *Nat. Commun.* **5**, 4213 (2014), article.
- [14] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New J. Phys.* **18**, 023023 (2016).
- [15] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, *Nat. Commun.* **10**, 3007 (2019).
- [16] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nat. Commun.* **9**, 4812 (2018).
- [17] S. McArdle, T. Jones, S. Endo, Y. Li, S. Benjamin, and X. Yuan, “Variational quantum simulation of imaginary time evolution,” (2018), [arXiv:1804.03023](https://arxiv.org/abs/1804.03023).
- [18] C. Lanczos, *J. Res. Natl. Bur. Stand. B* **45**, 255 (1950).
- [19] A. Uhlmann, *Rep. Math. Phys.* **9**, 273 (1976).
- [20] M. B. Hastings and T. Koma, *Comm. Math. Phys.* **265**, 781 (2006).
- [21] S. B. Bravyi and A. Y. Kitaev, *Ann. Phys.* **298**, 210 (2002).
- [22] F. Verstraete and J. I. Cirac, *J. Stat. Mech.: Theor. Exp.* **2005**, P09012 (2005).
- [23] D. W. Berry, A. M. Childs, and R. Kothari, in *IEEE 56th Annual FOCS Symposium* (2015) pp. 792–809.
- [24] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [25] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [26] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [27] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, “Contracting projected entangled pair states is average-case hard,” (2018), [arXiv:1810.00738](https://arxiv.org/abs/1810.00738).
- [28] 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, *Phys. Rev. X* **6**, 031007 (2016).
- [29] H. Lamm and S. Lawrence, *Phys. Rev. Lett.* **121**, 170501 (2018).
- [30] “Rigetti computing: quantum cloud services,” <https://qcs.rigetti.com/dashboard>, accessed: 2019-01-21.
- [31] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, *Phys. Rev. A* **95**, 042308 (2017).
- [32] 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, *Phys. Rev. X* **8**, 011021 (2018).
- [33] B. M. Terhal and D. P. DiVincenzo, *Phys. Rev. A* **61**, 022301 (2000).
- [34] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, *Nature* **471**, 87 (2011).
- [35] A. N. Chowdhury and R. D. Somma, *Quantum Inf. Comput.* **17**, 41 (2017).
- [36] F. G. Brandão and M. J. Kastoryano, *Comm. Math. Phys.* **365**, 1 (2019).
- [37] S. R. White, *Phys. Rev. Lett.* **102**, 190601 (2009).
- [38] E. M. Stoudenmire and S. R. White, *New J. Phys.* **12**, 055026 (2010).
- [39] F. G. Brandão and M. Horodecki, *Communications in mathematical physics* **333**, 761 (2015).
- [40] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [41] U. Schollwöck, *Rev. Mod. Phys.* **77**, 259 (2005).
- [42] T. Nishino and K. Okunishi, *J. Phys. Soc. Jpn* **65**, 891 (1996).
- [43] F. Verstraete and J. I. Cirac, “Renormalization algorithms for quantum-many body systems in two and higher dimensions,”

- (2004), [arXiv:cond-mat/0407066](#) .
- [44] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, *Phys. Rev. Lett.* **96**, 220601 (2006).
 - [45] R. Orús, *Annals of Physics* **349**, 117 (2014).
 - [46] F. Verstraete, V. Murg, and J. Cirac, *Adv. Phys.* **57**, 143 (2008).
 - [47] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **101**, 250602 (2008).
 - [48] H. C. Jiang, Z. Y. Weng, and T. Xiang, *Phys. Rev. Lett.* **101**, 090603 (2008).
 - [49] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, *Phys. Rev. B* **90**, 064425 (2014).
 - [50] M. Lubasch, J. I. Cirac, and M.-C. Banuls, *New Journal of Physics* **16**, 033014 (2014).
 - [51] Z. Y. Xie, H. J. Liao, R. Z. Huang, H. D. Xie, J. Chen, Z. Y. Liu, and T. Xiang, *Phys. Rev. B* **96**, 045128 (2017).
 - [52] J. Hachmann, W. Cardoen, and G. K.-L. Chan, *J. Chem. Phys.* **125**, 144101 (2006).
 - [53] 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, *Phys. Rev. X* **7**, 031059 (2017).
 - [54] A. Szabo and N. Ostlund, *Modern Quantum Chemistry*, Dover Books on Chemistry (Dover Publications, 1996).
 - [55] E. M. Stoudenmire and S. R. White, *New J. Phys.* **12**, 055026 (2010).
 - [56] S. R. White, *Phys. Rev. Lett.* **102**, 190601 (2009).
 - [57] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, *Nature* **471**, 87 (2011).
 - [58] H. Flyvbjerg and H. G. Petersen, *J. Chem. Phys.* **91**, 461 (1989).
 - [59] K. Temme, S. Bravyi, and J. M. Gambetta, *Phys. Rev. Lett.* **119**, 180509 (2017).
 - [60] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature* **549**, 031007 (2017).

Publisher Correction: Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution

Mario Motta , Chong Sun, Adrian T. K. Tan, Matthew J. O'Rourke , Erika Ye, Austin J. Minnich ,
Fernando G. S. L. Brandão and Garnet Kin-Lic Chan 

Correction to: *Nature Physics* <https://doi.org/10.1038/s41567-019-0704-4>, published online 11 November 2019.

In the version of this Article originally published online, in the first equation in the sentence beginning “One popular approach is”, the term ‘ $-\partial_\beta|\Phi(\beta)\rangle$ ’ was repeated; the equation should have read “ $-\partial_\beta|\Phi(\beta)\rangle = \{|\Phi\rangle, \hat{H}|\Phi\rangle, \hat{H}^2|\Phi\rangle \dots\}|\Phi(\beta)\rangle$ ”. In the same sentence, in the expression “ $\langle\Phi(0)|\Psi \neq 0$ ”, the right ket bracket was missing and should have read ‘ $\langle\Phi(0)|\Psi\rangle \neq 0$ ’. All versions of this Article have been amended.

Published online: 21 November 2019

<https://doi.org/10.1038/s41567-019-0756-5>

© The Author(s), under exclusive licence to Springer Nature Limited 2019

Publisher Correction: Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution

Mario Motta , Chong Sun, Adrian T. K. Tan, Matthew J. O'Rourke , Erika Ye, Austin J. Minnich ,
Fernando G. S. L. Brandão and Garnet Kin-Lic Chan 

Correction to: *Nature Physics* <https://doi.org/10.1038/s41567-019-0704-4>, published online 11 November 2019; corrected online 21 November 2019.

In the version of this Article originally published online, and in the later correction, the first equation in the sentence beginning “One popular approach is” was incorrect; the correct equation is ‘ $-\partial_\beta|\Phi(\beta)\rangle = \hat{H}|\Phi(\beta)\rangle$ ’. This has now been amended in all versions of the Article.

Published online: 29 January 2020

<https://doi.org/10.1038/s41567-020-0798-8>

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Publisher Correction: Mathematical languages shape our understanding of time in physics

Nicolas Gisin

Correction to: *Nature Physics* <https://doi.org/10.1038/s41567-019-0748-5>, published online 6 January 2020.

In the version of this Comment originally published online, the final sentence in Box 1 contained the erroneous text “numbers and stochastic”; it has now been removed.

Published online: 16 January 2020

<https://doi.org/10.1038/s41567-020-0794-z>

© Springer Nature Limited 2020