Lawrence Berkeley National Laboratory

Title

Exploring finite temperature properties of materials with quantum computers

Permalink

Journal

ISSN

Authors

Powers, Connor
Bassman Oftelie, Lindsay
Camps, Daan
et al.

Publication Date

2023

DOI

Peer reviewed

scientific reports

OPEN

Exploring finite temperature properties of materials with quantum computers

Connor Powers[1,2]✉, Lindsay Bassman Oftelie[1], Daan Camps[1] & Wibe A. de Jong[1]

Thermal properties of nanomaterials are crucial to not only improving our fundamental understanding of condensed matter systems, but also to developing novel materials for applications spanning research and industry. Since quantum effects arise at the nano-scale, these systems are difficult to simulate on classical computers. Quantum computers can efficiently simulate quantum many-body systems, yet current quantum algorithms for calculating thermal properties of these systems incur significant computational costs in that they either prepare the full thermal state on the quantum computer, or they must sample a number of pure states from a distribution that grows with system size. Canonical thermal pure quantum (TPQ) states provide a promising path to estimating thermal properties of quantum materials as they neither require preparation of the full thermal state nor require a growing number of samples with system size. Here, we present an algorithm for preparing canonical TPQ states on quantum computers. We compare three different circuit implementations for the algorithm and demonstrate their capabilities in estimating thermal properties of quantum materials. Due to its increasing accuracy with system size and flexibility in implementation, we anticipate that this method will enable finite temperature explorations of relevant quantum materials on near-term quantum computers.

As the search for high-performance materials persists in research and industry alike, there is a need for better understanding the thermal properties of materials. In particular, exploring thermal properties of nanomaterials is of great interest, with applications ranging from energy production to nanoelectronics[1,2]. Quantum effects can dominate at the nano-scale, and the exponential growth of resources required to simulate quantum systems with classical computers makes the simulation of such materials quickly exceed the capabilities of even the largest classical supercomputers[3]. Quantum computers, by contrast, are able to simulate quantum many-body systems efficiently[4,5]. Therefore, quantum computers offer a promising route to studying thermal properties of quantum materials. While a plethora of simulations of systems at zero-temperature have been demonstrated on quantum devices in recent years[6–13], the landscape of quantum algorithms to calculate finite temperature properties remains more sparse[14–16]. The main challenge in exploring finite-temperature properties on quantum computers lies in the preparation of thermal states.

Current quantum algorithms for thermal state preparation fall into two main categories. The first comprises algorithms that initialize the qubits into the full thermal (i.e., mixed) state. In this case, the thermal average of an observable can be computed directly by measuring the observable in this state. Examples include algorithms that prepare the Gibbs state using phase estimation[17–19], which require quantum circuits that are too large for near-term quantum devices, otherwise known as noisy intermediate-scale quantum (NISQ) computers[20]. Other examples include variational quantum thermalizers[21,22] and methods that prepare thermofield double states[1,23], both of which rely on variational techniques. The variational nature of these algorithms necessitates the use of a cost function, which generally becomes hard to compute as system size increases. Such methods are therefore difficult to scale to large or complex systems. Still other methods for generating the full thermal state require a number of ancilla qubits that scales with system size or complexity[24,25], thus limiting the size of systems that can be simulated on current quantum hardware.

Algorithms in the second category prepare an ensemble of pure states, one pure state at a time, where each pure state has been sampled according to the correct thermal distribution. Existing examples rely on Monte Carlo sampling techniques, Markov chains, or both[26–31]. To calculate thermal averages, the desired observable is measured in each of the different pure states and the results are averaged over the ensemble. As pure states are much easier to prepare on a quantum computer than mixed states, this model for thermal state preparation is

[1]Lawrence Berkeley National Lab, Berkeley, CA, USA. [2]University of Maryland, College Park, MD, USA. ✉email: cdpowers@umd.edu

more promising for NISQ computers. However, the number of samples required generally grows with the size of the system being simulated, which can lead to significant resource requirements for large systems.

Canonical thermal pure quantum (TPQ) states[32] offer a promising way to estimate thermal averages on quantum computers. Thermal state approximation by TPQ states lies in a separate third category, as it neither incurs the quantum resources required to prepare a mixed state, nor relies on a number of samples that grows with system size. TPQ states are formed by applying a specific non-unitary transformation, which is a function of the system Hamiltonian and inverse temperature, to a random state. The resulting state is shown to be representative of the thermal equilibrium in that observables measured in this state will approximate thermal averages. Remarkably, the error in the expectation value of the observable is bounded by an exponentially decreasing function of system size N. Thus, at sufficiently large N, the expectation value of an observable in only a single TPQ state will yield a very close approximation to the thermal average. The effectiveness of canonical TPQ states has been demonstrated classically[32], but to our knowledge, has not been implemented on a quantum computer.

Here, we present an algorithm for generating canonical TPQ states on quantum computers, enabling the estimation of finite temperature properties of materials on NISQ devices. Our algorithm relies on a straightforward and scalable protocol for preparing the random state[33], which allows for circuit depths to be tuned to find a balance between desired accuracy and feasibility of execution on NISQ hardware. Furthermore, the algorithm is agnostic to implementation of the non-unitary transformation of the random state, which can be tailored to the resource constraints of different quantum devices. We compare three possible implementations for approximating the non-unitary transformation: (i) the quantum imaginary time evolution (QITE) algorithm[26], which can be better suited to devices constrained in qubit count, (ii) the dilated operator approach[34], which provides some advantages in resource requirements over the QITE algorithm when a single ancilla qubit is available, and (iii) the FABLE method[35] which is better suited for devices with limited coherence times. It is noted that Quantum Signal Processing may also be used to approximate this non-unitary operation, as discussed in Ref.[36].

We demonstrate our algorithm using each implementation for the Heisenberg model, a quintessential model used for studying a range of behaviors in materials[37–41]. We anticipate that this algorithm will facilitate estimations of finite temperature properties of materials on near-term quantum devices. Furthermore, since error in estimating thermal averages with TPQ states decreases with increasing system size, we believe this algorithm will only become increasingly useful as quantum hardware continues to grow in size in the coming years.

Theoretical framework

The (unnormalized) canonical TPQ state of a system of size N, governed by Hamiltonian H, at inverse temperature β is defined as[32]

$$|\beta, N\rangle = \hat{Q}|\Psi_R\rangle \equiv e^{-\beta H/2}|\Psi_R\rangle, \tag{1}$$

where $|\Psi_R\rangle = \sum_{i=1}^{2^N} c_i |i\rangle$ is a random state, defined by complex amplitudes c_i which are uniformly sampled from the unit hypersphere such that $\sum_{i=1}^{2^N} |c_i|^2 = 1$, and $|i\rangle$ is an arbitrary orthonormal basis. Defined this way, $|\Psi_R\rangle$ is characterized as a Haar-random pure state[42]. We define

$$\langle \hat{A} \rangle_{\beta,N}^{ens} \equiv \frac{\text{Tr}[e^{-\beta H} \hat{A}]}{\text{Tr}[e^{-\beta H}]}, \tag{2}$$

as the ensemble expectation value of an operator $\hat{A}$ for a system of size N at inverse temperature β, and

$$\langle \hat{A} \rangle_{\beta,N}^{TPQ} \equiv \frac{\langle \beta, N | \hat{A} | \beta, N \rangle}{\langle \beta, N | | \beta, N \rangle}, \tag{3}$$

as the corresponding expectation value of $\hat{A}$ in a single TPQ state. It is noted that $\hat{A}$ must be a low-degree polynomial of local operators for the following analysis to hold, but this group contains many prominently used observables including energy, magnetization, and a number of relevant correlation functions. The error between $\langle \hat{A} \rangle_{\beta,N}^{TPQ}$ and $\langle \hat{A} \rangle_{\beta,N}^{ens}$ is bounded by a value that becomes exponentially small with increasing system size N[32]. In practice, this means that for sufficiently large N, measuring the desired observable in a single TPQ state will provide a good approximation of the true thermal average. Indeed, Ref.[32] found less than a 1% error when estimating thermal properties using a single TPQ state for a quantum spin model with $N = 30$. At lower N, fidelity may be increased by averaging over the measured values from multiple TPQ states[32].

The first step in preparing a TPQ state on a quantum computer is the preparation of $|\Psi_R\rangle$. We can approximate this Haar-random state using random quantum circuits constructed in the manner proposed in Ref.[33] and illustrated in Fig. 1.

The random circuits we consider are composed of "blocks," where each block is composed of a layer of single-qubit rotation gates on every qubit followed by a layer of two-qubit entangling gates. The single-qubit rotation gates are selected from a finite set $\mathbb{A} = \{RX(\frac{\pi}{2}), RY(\frac{\pi}{2}), T\}$, with the constraint that no gate may be chosen for the same qubit two blocks in a row. In other words, if $R_{i,j}$ is the jth single-qubit gate acting on qubit i, then $R_{i,j} \in \mathbb{A} \setminus R_{i,j-1}$.

The two-qubit gate layers, as shown in Fig. 1b,c, follow a fixed pattern determined by the dimension of the Hamiltonian of interest; given a system dimension of k, there will be $2k$ different two-qubit gate layer patterns to loop through to ensure all coupling directions have been accounted for. In this way, the random circuits are easily generalizable to higher dimensions. For example, Fig. 1b shows the two 2-qubit gate patterns that must be looped through for a one-dimensional (1D) system. Similarly, Fig. 1c shows the four 2-qubit gate patterns to loop through for 2D systems.

Random circuits of this form can be defined by a single parameter d, which sets the number of blocks and ultimately controls the circuit depth. Such circuits can be seen as approximating successively higher t-designs as d is increased[43–46]. (A unitary t-design is an approximation of a Haar random unitary which accurately simulates the first t moments of Haar-random unitaries. Higher order t-designs generate states whose properties converge to those of Haar-random states). We can see this convergence by plotting the entropy of the resulting random state. This plot is useful for determining how large d needs to be for a given system size; the necessary d will be the point at which the state entropy is sufficiently converged to the value $\ln N - 1 + \gamma$ characteristic of Haar-random states, where $\gamma \approx 0.577$ is the Euler–Mascheroni constant. See Sect. I in the Supplementary Information (SI) for these convergence plots with and without simulated device noise. An N-qubit random circuit of this form with d layers has Nd single-qubit rotation gates and $O(Nd)$ two-qubit entangling gates. Rigorous bounds predict a polynomial scaling of the number of required layers with respect to system size N, but in practice sublinear scalings have been observed[46]. Here, we set $d = 20$, unless otherwise specified.

Once the random state $|\Psi_R\rangle$ has been generated, the non-unitary operator $\hat{Q}$ must be applied to generate a canonical TPQ state. Given that quantum computers can only perform unitary operations, this must be implemented by a unitary approximation of the non-unitary transformation. Several approximation methods exist and in practice, the method choice is informed by available quantum resources. Here, we demonstrate our algorithm with three different approaches, including QITE[26], a dilated operator approach[34], and the FABLE method[35].

With QITE, neither ancillary qubits nor post-selection of results are required. This however, comes at the cost of deeper circuits that require significant classical resources to generate. Circuit generation time with QITE grows quickly as the simulated system size is increased, though in practice this can be mitigated by employing the so-called "inexact QITE" method to truncate the domain-size[26]. When average correlation lengths of the simulated system are small, a domain-size significantly smaller than total system size can be chosen to dramatically decrease the computational complexity of generating the QITE circuit. However, this can result in reduced accuracy when correlation lengths of the system are larger than the selected domain size.

The unitary dilation method requires a single ancilla qubit as well as post-selection of results. Any experiment in which the ancilla qubit is measured to be in the '1' state must be discarded, which increases the number of shots required to generate good results. Overall, the unitary dilation method tends to generate shallower circuits than QITE, but at the price of requiring more shots.

Finally, FABLE requires a register of ancilla qubits that grows linearly with the simulated system size as well as post-selection of results. For a system size of N, FABLE requires $N + 1$ ancilla qubits, all of which must be measured to be in the '0' state (otherwise the result must be discarded). While its post-selection requirement based on such a large ancilla register will necessitate a very large number of shots, FABLE offers the most favorable circuit depths and circuit generation times of the three methods.

The algorithm shown in Fig. 2 summarizes our method for approximating thermal averages with TPQ states on quantum computers. First, the d-layer circuit is constructed to generate $|\Psi_R\rangle$. Layers of randomly chosen single-qubit rotation gates are alternated with layers of 2-qubit entangling gates following a pattern set by the Hamiltonian dimension k. Next, a unitary approximation of the non-unitary transformation $\hat{Q}$ is applied. The algorithm is flexible in how this non-unitary operation is implemented. The observable of interest is then measured in the resulting canonical TPQ state, and this process is repeated R times to construct and measure R distinct TPQ states. Finally, if applicable, the average of the measured thermal values is taken.

If QITE is chosen to implement line 10 in Fig. 2, the non-unitary operator $\hat{Q}$ is viewed as an evolution through imaginary time. To find a unitary approximation, this evolution is broken up into time-steps of imaginary time, and each time-step is approximated by a unitary sub-circuit[26]. Increasing the number of time-steps into which the imaginary-time evolution is broken will increase the transformation fidelity, i.e., the accuracy of the approximation. The size of the quantum circuit grows linearly with the number of time-steps, but a number of techniques have been developed to reduce the depth of these QITE circuits[47–50]. We also note that there exist other options for reducing the algorithmic error of QITE such as leveraging randomized compiling[51] or reinforcement learning[52]. The sub-circuit for each time-step is constructed based on a so-called domain-size, which can be set to be equal to or less than the simulated system size, and should be chosen based on the average correlation length within the simulated system. Accuracy of the QITE algorithm is improved by increasing the domain-size that the unitaries act upon, but the classical computational cost associated with finding these unitaries grows exponentially with the domain-size. When using QITE, the number of imaginary time-steps and the domain-size must be chosen appropriately by the user for the simulated system.

If line 10 in Fig. 2 is instead implemented with the dilated operator approach, an ancillary qubit initialized in state $|1\rangle$ must be added to create the augmented input state $\rho_{in} = (|1\rangle\langle1|) \otimes \rho_R$, where $\rho_R = |\Psi_R\rangle\langle\Psi_R|$. A dilated unitary operator $\hat{\Omega}$ is then constructed in the new 2^{N+1} dimensional Hilbert space that will approximate the action of $\hat{Q}$ on the original 2^N dimensional system Hilbert space[34]. $\hat{\Omega}$ is defined in terms of the non-unitary operator $\hat{Q}$ as

$$\hat{\Omega} \equiv \exp\left(i\epsilon \begin{pmatrix} 0 & -i\hat{Q} \\ i\hat{Q}^\dagger & 0 \end{pmatrix}\right), \tag{4}$$

where ϵ is a parameter that controls the performance of the operator. The dilated unitary operator $\hat{\Omega}$ is then applied to the augmented initial state ρ_{in}, and the ancillary qubit is measured. If it is measured in the state $|0\rangle$, the original N-qubit system of interest has been successfully transformed by an approximation to the non-unitary operator $\hat{Q}$. The probability that the ancillary qubit is measured in state $|0\rangle$, also called the probability of success, is denoted P_0. If the ancillary qubit is measured in state $|1\rangle$, the results of the entire circuit are discarded[34].

CNOT count				
N	QITE	Inexact QITE	Dilated operator	FABLE
2	14	20	41	16
3	97	963	218	64
4	–	1957	1025	256
5	–	2945	4474	1024

Table 1. CNOT gate count at various system sizes N for approximating the imaginary time evolution step of TPQ state preparation with the QITE, inexact QITE, dilated operator, and FABLE algorithms at $\beta = 1$.
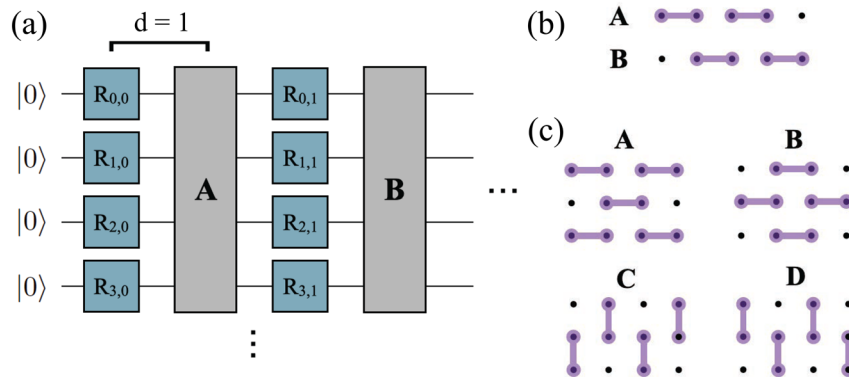
Figure 1. (**a**) Gate pattern for approximate Haar-random circuits. $R_{i,j}$ are randomly selected single-qubit rotation gates selected from a finite set with constraints (see main text). In 1D systems 2-qubit layers follow the pattern **ABAB...** with individual layer patterns specified in (**b**), while in 2D systems, 2-qubit layers follow the pattern **ABCDABCD...**, with individual layer patterns specified in (**c**).

To help quantify the performance of this dilated operator, it is useful to define a fidelity metric that measures how close the effective transformation on the N-qubit system is to the desired non-unitary transformation under $\hat{Q}$. This fidelity can be defined as:

$$F = \mathrm{Tr}\sqrt{\sqrt{\rho_0}|\beta, N\rangle\langle\beta, N|\sqrt{\rho_0}}. \tag{5}$$

Here, ρ_0 is the density matrix of the N-qubit system after tracing out the ancillary qubit, provided that it is measured in the state $|0\rangle$. The probability of success P_0 and the transformation fidelity F are ultimately controlled by the user's choice of ϵ. While a small ϵ is required to obtain an accurate approximation to the non-unitary operator (i.e. high transformation fidelity F), decreasing ϵ generally decreases the probability of success P_0, which means more shots need to be run until the ancilla is measured in the $|0\rangle$ state, indicating a successful transformation. See Sect. II of the SI for more information regarding the behavior of F and P_0 with varying ϵ.

While this method introduces a single ancillary qubit, its advantage is that increasing the accuracy of the approximate non-unitary transformation only requires increasing the average number of shots needed to see achieve a successful transformation. This stands in contrast to the QITE algorithm, where increasing the accuracy of the approximation requires a larger number of imaginary time-steps and subsequently deeper circuits. However, this is still considered a near-term approach since it requires the decomposition of the $N + 1$-qubit unitary into single and two-qubit gates; the decomposition of an m-qubit unitary generally requires a number of two-qubit entangling gates that is exponential in m. By counting the controlled single-qubit unitaries that need to be determined during the decomposition of multi-qubit unitaries via the column-by-column decomposition[53] technique used by Qiskit and recursively decomposing the multi-qubit diagonal gates that emerge during this process via Theorem 7 of Ref.[54], using this process to decompose an m-qubit unitary requires solving an upper bound of $2^{2m} + (m - 1)2^m - m$ (2×2) matrix equations and $2^{2m+1} - (m + 4)2^m + m + 2$ scalar equations before any circuit optimizations, yielding upper bounds of 4410 matrix equations and 7560 scalar equations for the simulations of 5-qubit systems seen in Fig. 4, and this decomposition method carries a classical time complexity of $O(m2^{3m})$[55]. If the dilated operator approach is preferred, then the only change in Fig. 2 is the addition of an ancillary qubit initialized in the $|1\rangle$ state which is not included in the random circuit construction. The non-unitary transformation approximation step then takes in the additional argument ϵ. It is noted that other recently-developed probabilistic methods of approximating the non-unitary imaginary time evolution operator with a single ancillary qubit can skirt this exponential scaling by leveraging forward and backward real-time evolution operators, for which efficient circuit implementations have previously been developed[56].

The third approach we follow to implement line 10 in Fig. 2 is a *block-encoding* of the non-unitary operator $\hat{Q}$, which is the embedding of the non-unitary in the leading principal block of a larger unitary U,

$$U = \begin{pmatrix} \hat{Q} & * \\ * & * \end{pmatrix}. \tag{6}$$

Here, $*$ indicate arbitrary matrix elements. Assuming a ancilla qubits are used to block encode the 2^N dimensional $\hat{Q}$, the operator U is constructed in the 2^{N+a} dimensional Hilbert space. Next, we apply U to the augmented input state $\rho_{\text{in}} = (|0\rangle^{\otimes a} \langle 0|^{\otimes a}) \otimes \rho_R$. If all a ancillary qubits are measured in the $|0\rangle$ state, the N-qubit operator $\hat{Q}$ has been successfully applied to ρ_R. The likelihood of a successful measurement can be increased through amplitude amplification.

We use FABLE [35] to generate a circuit for Eq. (6) based on the non-unitary operator $\hat{Q}$. FABLE circuits require $N+1$ ancilla qubits to encode an N-qubit operator and are fast to generate for small to medium-sized problems as the circuit generation algorithm scales as $\mathcal{O}(N4^N)$. nThe advantage of FABLE over the dilated operator approach

---

Input: H, $\hat{A}$, β, R, d, k
Output: $\langle \hat{A} \rangle_{\beta,N}^{TPQ} \approx \langle \hat{A} \rangle_{\beta,N}^{ens}$

```
1  thermal_vals = []
   /* Loop over R TPQ state realizations          */
2  for n=(0,R) do
      /* make random state preparation circuit     */
3     random_circ=empty_circuit()
4     prev_layer=[]
      /* Loop over d to add d blocks               */
5     for j=(0,d) do
         /* add single qubit gate layer            */
6        single_qub_layer = SingleQubLayer(prev_layer)
7        random_circ.append(single_qub_layer)
8        prev_layer = single_qub_layer
         /* Add two-qubit gate layer               */
9        random_circ.append(TwoQubLayer(k,j))
      /* Approximate transformation under Q̂        */
10    Q_circ = nonunitary_approx_circ(H,β,...)
11    TPQ_state_circ = random_circ + Q_circ
      /* measure observable of interest            */
12    thermal_val = measure(TPQ_state_circ, Â)
13    thermal_vals.append(thermal_val)
14 return average(thermal_vals)
```

Figure 2. Pseudocode algorithm for computation of thermal averages with canonical TPQ states on quantum computers.

---

is that no performance parameter ϵ is required and the circuits can be generated more efficiently while requiring fewer CNOT gates. The main disadvantage is that significantly more ancillary qubits are required.

In Tables 1, 2 and 3, we compare the three different methods utilized in this work for implementing the non-unitary evolution step of TPQ state preparation on quantum computers. Specifically, the three methods are compared across three metrics relevant to practical implementations on near-term devices: CNOT gate count of involved circuits (Table 1), ancillary qubit requirements (Table 2), and the time to classically generate the

Ancillary qubits				
N	QITE	Inexact QITE	Dilated operator	FABLE
2	0	0	1	3
3	0	0	1	4
4	0	0	1	5
5	0	0	1	6

Table 2. Ancillary qubit requirement at various system sizes N for approximating the imaginary time evolution step of TPQ state preparation with the QITE, inexact QITE, dilated operator, and FABLE algorithms at $\beta = 1$.

required circuits (Table 3). The latter is important because since the execution of the required quantum circuits can take on the order of milliseconds, the classical computational costs associated with generating the quantum circuits can dominate the overall cost of the quantum simulation. All data is generated for simulation of N-spin Heisenberg models (described in more detail in the following section). We first attempted to generate the QITE circuits with the XACC software package[57], however, this code requires the domain-size to be set equal to the system size, which made going to a system size of $N = 4$ prohibitively expensive for our chosen model. For $N > 3$, we therefore used the ArQTiC software package[58] to generate the QITE circuits. However, ArQTiC only allows a maximum domain-size of 3. Therefore, for systems with $N > 3$, a truncated domain-size of 3 had to be used, resulting in reduced accuracy of results. In Tables 1, 2 and 3, results from XACC are provided in the column labeled 'QITE', while results from ArQTiC are provided in the column labeled 'Inexact QITE', since we expect these results to have lowered accuracy due to domain-size truncation[26]. Since gate counts and/or circuit generation times of the QITE and dilated operator methods may vary between circuit realizations even for a given system size, data shown for these methods have been averaged over 10 circuit realizations. All data is collected for an inverse temperature of $\beta = 1$, and for QITE data, the imaginary time evolution operator is broken into 10 imaginary time steps.

As seen in Tables 1, 2 and 3, while the QITE algorithm has a clear advantage in that it requires no ancillary qubits, its scaling of circuit generation time is seen to be practically prohibitive even at relatively small system size. Therefore, this method seems to be best suited for simulating very small systems on equally small near-term devices. Switching to an inexact QITE with the domain-size truncated to 3 significantly alleviates the circuit generation time scaling with system size at the cost of increased entangling gate counts and reduced accuracy of results. The dilated operator method appears to bridge the resource requirement gaps of QITE and FABLE methods; it requires a minimal ancillary qubit requirement that is independent of simulated system size, and has more favorable circuit generation time scaling than QITE, but at a cost of increased CNOT gate counts. For small systems, the latter may be mitigated by some circuit synthesis techniques such as QFAST[59] and LEAP[60]. If sufficient ancillary qubit resources are available, the FABLE technique displays the best CNOT count and circuit generation time scalings of the three methods. Therefore, it is the best option when coherence times are a dominant limitation on experiments and ample qubit counts are available.

Demonstration

We now demonstrate our method by calculating the thermal energy of a Heisenberg spin model under an external magnetic field at various system sizes and inverse temperatures. The Hamiltonian of an N-spin Heisenberg model is given by

$$\hat{H} = \sum_{\alpha} \sum_{\langle i,j \rangle} J_{\alpha,ij} \sigma_i^{\alpha} \sigma_j^{\alpha} + h_x \sum_{i=1}^{N} \sigma_i^x, \tag{7}$$

where $J_{\alpha,ij}$ ($\alpha \in \{x, y, z\}$) gives the strength of the exchange coupling interaction between nearest neighbor spin pair $\langle i, j \rangle$ in the α-direction, h_x is the strength of an externally applied magnetic field in the x-direction, and σ^{α} are Pauli matrices. For results presented in this work, we set $J_x = 0.5$, $J_y = 1.25$, $J_z = 2.0$, and $h_x = 1.0$.

From TPQ state formalism it is known that, on average, the squared error from estimating a thermal value with a single TPQ state is bounded by a function that becomes exponentially small with increasing simulated system size[32]. At small system sizes, some of this error may be mitigated through averaging over multiple TPQ state realizations. For more details and numerical results, see Sect. III in the SI.

In Fig. 3, we present numerical results demonstrating the efficacy of the method in calculating thermal energies of 12-spin 1D and 2D Heisenberg models at various inverse temperatures. The random circuits are built through the aforementioned protocol and the non-unitary operation $\hat{Q}$ is numerically simulated. The thermal energies in both the 1D and 2D systems are closely approximated by averaging over just 10 canonical TPQ states.
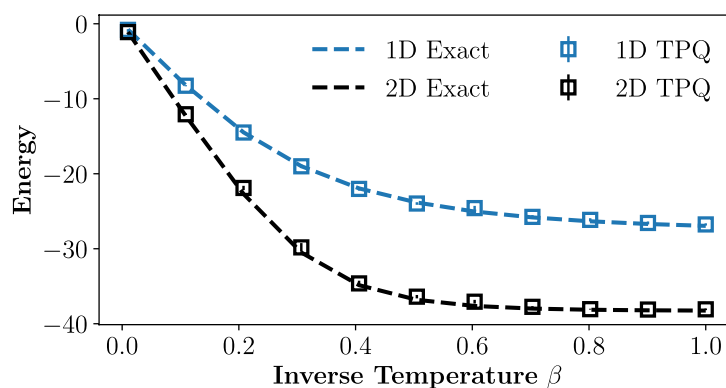
Figure 3. Thermal energy, as a function of inverse temperature β of a 1D Heisenberg model on a 12-qubit lattice (blue) and a 2D Heisenberg model on a 4×3 lattice (black). Shown are results averaged over 10 TPQ state realizations. The true ensemble thermal energies (dashed) are plotted for reference.
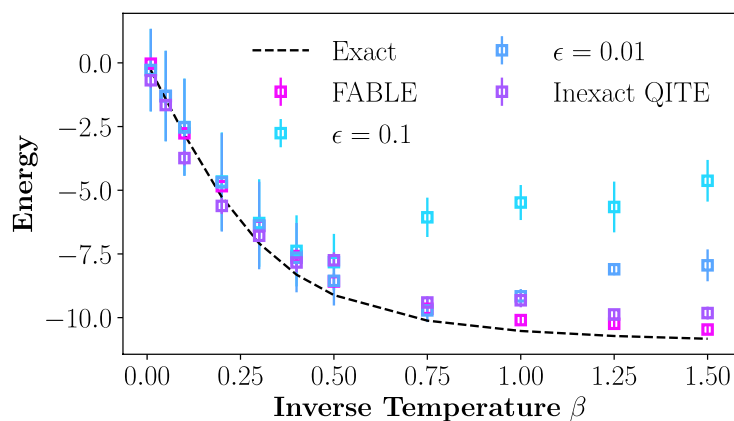
Figure 4. Thermal energy, as a function of inverse temperature β, calculated with TPQ states on a noiseless quantum simulator for a 5-qubit 1D Heisenberg model using FABLE, inexact QITE, and a dilated operator approach to approximate $\hat{Q}$.
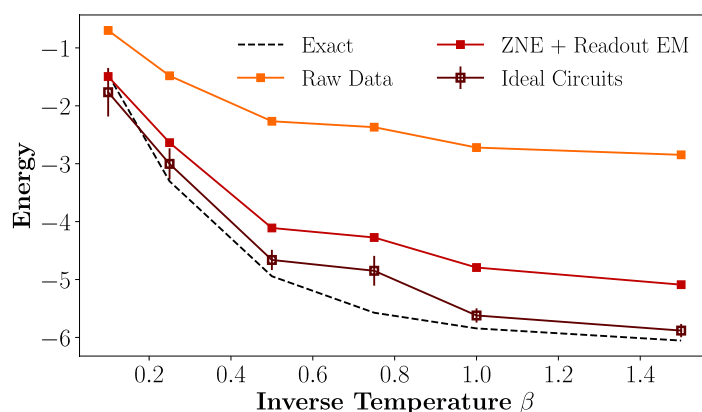
Figure 5. Thermal energy data from canonical TPQ states, as a function of inverse temperature β, of a 3-qubit 1D Heisenberg model executed on both a noiseless quantum simulator and on IBM's "ibmq_brooklyn" quantum computer. The QITE algorithm was performed with 10 imaginary time-steps for each value of β. Raw hardware data is presented in comparison to results after readout error mitigation (EM) and zero-noise extrapolation (ZNE) have been applied.

Circuit generation time (s)				
N	QITE	Inexact QITE	Dilated operator	FABLE
2	2.85	0.719	0.970	2.14×10^{-3}
3	1.44×10^2	3.71	1.53	6.39×10^{-3}
4	–	7.53	4.44	2.61×10^{-2}
5	–	11.51	17.4	8.71×10^{-2}

Table 3. Wall-clock circuit generation time at various system sizes N for approximating the imaginary time evolution step of TPQ state preparation with the QITE, inexact QITE, dilated operator, and FABLE algorithms at $\beta = 1$.

Next, we demonstrate quantum simulator results of our method using the three different quantum circuit implementations of $\hat{Q}$. In Fig. 4, the thermal energy of a 5-qubit 1D Heisenberg model as a function of inverse temperature is approximated using the dilated operator, inexact QITE, and FABLE approaches to implement $\hat{Q}$. Inexact QITE circuits were generated and simulated with the ArQTiC package at a truncated domain size of $D = 3$. Results using two different values of the dilated operator parameter ϵ are included to demonstrate its impact on performance. While smaller ϵ leads to significantly better results for larger β, it requires many more executions as the success probability decreases with smaller ϵ. After constructing the requisite dilated operator

to approximate the non-unitary transformation $\hat{Q}$, Qiskit was utilized to decompose it into a circuit with the basis gate set of the IBM's "ibmq_brooklyn" device. Presented results are averaged over $R = 10$ distinct TPQ states, and the error bars show uncertainty.

The dilated operator method appears to be consistently accurate at low β, while its performance begins degrading past some threshold inverse temperature. This threshold can be increased by decreasing the performance parameter ϵ. Results derived from using the FABLE technique have comparable accuracy to the dilated operator results at high temperatures and do not appear to exhibit a similar degradation at low temperatures. Results using the inexact QITE method exhibit some reduced accuracy compared to the other techniques at high temperatures, but still significantly outperform the dilated operator methods at low temperatures.

Finally, Fig. 5 compares results from circuits run both on a noiseless Qiskit quantum simulator and on IBM's "ibmq_brooklyn" quantum computer. Using the QITE version of the algorithm with no domain-size truncation, we calculated the thermal energy of a 3-site Heisenberg model as a function of inverse temperature β while employing readout error mitigation (EM) and zero-noise extrapolation[61,62] (ZNE) techniques to reduce error. The QITE circuits were generated using XACC[57] then compressed through numerical optimization using the QSearch[63] tool, reducing the total number of gates by around two orders of magnitude.

Figure 5 shows that while raw hardware data from generating and utilizing canonical TPQ states to measure thermal values has a significant amount of noise, this can be mitigated by standard techniques such as readout EM and ZNE. These error mitigation techniques can effectively accelerate the timeline in which this method for calculating thermal averages is achievable on near-term devices.

Conclusion

We have presented a new method for approximating finite temperature properties of materials on quantum computers using TPQ states. We demonstrated its efficacy through approximating thermal energies of Heisenberg models in one and two dimensions. To demonstrate flexibility in how the non-unitary step of the algorithm is implemented, we presented results from a quantum simulator derived from performing this transformation with the QITE algorithm, a dilated operator approach, and the recently-developed FABLE method. We also present hardware results executed on IBM's "ibmq_brooklyn" quantum computer, showing the efficacy of the method when combined with standard error mitigation techniques. Due to the increasing accuracy of thermal observables derived from TPQ states with system size, as well as the flexibility in the implementation of our algorithm, we expect our method for computing thermal properties of materials to be increasingly useful as higher quality qubits continue to be added to current quantum computers as we progress through the NISQ era and beyond.

Data availability

The datasets used and/or analysed in this work are available from the corresponding author on reasonable request.

References

1. Zhu, D. *et al.* Generation of thermofield double states and critical ground states with a quantum computer. *Proc. Natl. Acad. Sci.* **117**, 25402–25406 (2020).
2. Zhang, G. & Li, B. Impacts of doping on thermal and thermoelectric properties of nanomaterials. *Nanoscale* **2**, 1058–1068. https://doi.org/10.1039/C0NR00095G (2010).
3. De Raedt, H. *et al.* Massively parallel quantum computer simulator, eleven years later. *Comput. Phys. Commun.* **237**, 47–61 (2019).
4. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
5. Abrams, D. S. & Lloyd, S. Simulation of many-body Fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589 (1997).
6. Smith, A., Kim, M. S., Pollmann, F. & Knolle, J. Simulating quantum many-body dynamics on a current digital quantum computer. *NPJ Quantum Inf.* **5**, 106. https://doi.org/10.1038/s41534-019-0217-0 (2019).
7. Chiesa, A. *et al.* Quantum hardware simulating four-dimensional inelastic neutron scattering. *Nat. Phys.* **15**, 455–459 (2019).
8. Francis, A., Freericks, J. & Kemper, A. Quantum computation of magnon spectra. *Phys. Rev. B* **101**, 014411 (2020).
9. Arute, F. *et al.* Observation of separated dynamics of charge and spin in the fermi-hubbard model. Preprint at http://arxiv.org/abs/2010.07965 (2020).
10. Bassman, L. *et al.* Towards simulation of the dynamics of materials on quantum computers. *Phys. Rev. B* **101**, 184305. https://doi.org/10.1103/PhysRevB.101.184305 (2020).
11. Yeter-Aydeniz, K., Siopsis, G. & Pooser, R. C. Scattering in the ising model with the quantum Lanczos algorithm. *New J. Phys.* **23**, 043033 (2021).
12. Bassman, L. *et al.* Constant-depth circuits for dynamic simulations of materials on quantum computers. *Mater. Theory* **6**, 1–18 (2022).
13. Smith, A., Jobst, B., Green, A. G. & Pollmann, F. Crossing a topological phase transition with a quantum computer. *Phys. Rev. Res.* **4**, L022020 (2022).
14. Bassman, L., Klymko, K., Liu, D., Tubman, N. M. & de Jong, W. A. Computing free energies with fluctuation relations on quantum computers. Preprint at http://arxiv.org/abs/2103.09846 (2021).
15. Sun, S.-N. *et al.* Quantum computation of finite-temperature static and dynamical properties of spin systems using quantum imaginary time evolution. *PRX Quantum* **2**, 010317 (2021).
16. Bassman, L. *et al.* Simulating quantum materials with digital quantum computers. *Quantum Sci. Technol.* **6**, 043002 (2021).
17. Poulin, D. & Wocjan, P. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.* **103**, 220502. https://doi.org/10.1103/PhysRevLett.103.220502 (2009).
18. Riera, A., Gogolin, C. & Eisert, J. Thermalization in nature and on a quantum computer. *Phys. Rev. Lett.* **108**, 080402. https://doi.org/10.1103/PhysRevLett.108.080402 (2012).
19. Bilgin, E. & Boixo, S. Preparing thermal states of quantum systems by dimension reduction. *Phys. Rev. Lett.* **105**, 170405. https://doi.org/10.1103/PhysRevLett.105.170405 (2010).
20. Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018).

21. Verdon, G., Marks, J., Nanda, S., Leichenauer, S. & Hidary, J. Quantum hamiltonian-based models and the variational quantum thermalizer algorithm. Preprint at http://arxiv.org/abs/1910.02071 (2019).
22. Foldager, J., Pesah, A. & Hansen, L. K. Noise-assisted variational quantum thermalization. *Sci. Rep.* **12**, 1–11 (2022).
23. Wu, J. & Hsieh, T. H. Variational thermal quantum simulation via thermofield double states. *Phys. Rev. Lett.* **123**, 220502. https://doi.org/10.1103/PhysRevLett.123.220502 (2019).
24. Terhal, B. M. & DiVincenzo, D. P. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A* **61**, 022301. https://doi.org/10.1103/PhysRevA.61.022301 (2000).
25. Chowdhury, A. N. & Somma, R. D. Quantum algorithms for Gibbs sampling and hitting-time estimation. *Quantum Inf. Comput.* **17**, 41–64 (2017).
26. Motta, M. *et al.* Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat. Phys.* **16**, 205–210. https://doi.org/10.1038/s41567-019-0704-4 (2019).
27. Sun, S.-N. *et al.* Quantum computation of finite-temperature static and dynamical properties of spin systems using quantum imaginary time evolution. *PRX Quantum* **2**, 010317. https://doi.org/10.1103/PRXQuantum.2.010317 (2021).
28. Moussa, J. E. Low-depth quantum metropolis algorithm. Preprint at http://arxiv.org/abs/1903.01451 (2019).
29. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum metropolis sampling. *Nature* **471**, 87–90. https://doi.org/10.1038/nature09770 (2011).
30. Yung, M.-H. & Aspuru-Guzik, A. A quantum-quantum metropolis algorithm. *Proc. Natl. Acad. Sci.* **109**, 754–759. https://doi.org/10.1073/pnas.1111758109 (2012).
31. Lu, S., Bañuls, M. C. & Cirac, J. I. Algorithms for quantum simulation at finite energies. *PRX Quantum* **2**, 020321. https://doi.org/10.1103/PRXQuantum.2.020321 (2021).
32. Sugiura, S. & Shimizu, A. Canonical thermal pure quantum state. *Phys. Rev. Lett.* **111**, 010401. https://doi.org/10.1103/PhysRevLett.111.010401 (2013).
33. Richter, J. & Pal, A. Simulating hydrodynamics on noisy intermediate-scale quantum devices with random circuits. *Phys. Rev. Lett.* **126**, 230501. https://doi.org/10.1103/PhysRevLett.126.230501 (2021).
34. Gingrich, R. M. & Williams, C. P. Non-unitary probabilistic quantum computing. In *Proc. Winter International Symposium on Information and Communication Technologies, WISICT '04* 1–6 (Trinity College Dublin, 2004).
35. Camps, D. & Van Beeumen, R. Fable: Fast approximate quantum circuits for block-encodings. Preprint at http://arxiv.org/abs/2205.00081 (2022).
36. Coopmans, L., Kikuchi, Y. & Benedetti, M. Predicting Gibbs state expectation values with pure thermal shadows. Preprint at http://arxiv.org/abs/2206.05302 (2022).
37. Billoni, O. V., Cannas, S. A. & Tamarit, F. A. Spin-glass behavior in the random-anisotropy Heisenberg model. *Phys. Rev. B* **72**, 104407 (2005).
38. Gong, S.-S., Zhu, W. & Sheng, D. Emergent chiral spin liquid: Fractional quantum hall effect in a Kagome Heisenberg model. *Sci. Rep.* **4**, 1–6 (2014).
39. Jepsen, P. N. *et al.* Spin transport in a tunable Heisenberg model realized with ultracold atoms. *Nature* **588**, 403–407 (2020).
40. Tanaka, T. & Gohda, Y. Prediction of the curie temperature considering the dependence of the phonon free energy on magnetic states. *NPJ Comput. Mater.* **6**, 1–7 (2020).
41. Rodriguez-Nieva, J. F. Turbulent relaxation after a quench in the Heisenberg model. *Phys. Rev. B* **104**, L060302 (2021).
42. Liu, Z.-W. *On Quantum Randomness and Quantum Resources*. Ph.D. thesis, Massachusetts Institute of Technology (2018).
43. Brandao, F. G., Harrow, A. W. & Horodecki, M. Local random quantum circuits are approximate polynomial-designs. *Commun. Math. Phys.* **346**, 397–434 (2016).
44. Nakata, Y. *et al.* Quantum circuits for exact unitary *t*-designs and applications to higher-order randomized benchmarking. *PRX Quantum* **2**, 030339. https://doi.org/10.1103/PRXQuantum.2.030339 (2021).
45. Nakata, Y., Hirche, C., Koashi, M. & Winter, A. Efficient quantum pseudorandomness with nearly time-independent Hamiltonian dynamics. *Phys. Rev. X* **7**, 021006. https://doi.org/10.1103/PhysRevX.7.021006 (2017).
46. Boixo, S. *et al.* Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600. https://doi.org/10.1038/s41567-018-0124-x (2018).
47. Tan, K. C. Fast quantum imaginary time evolution. Preprint at http://arxiv.org/abs/2009.12239 (2020).
48. Nishi, H., Kosugi, T. & Matsushita, Y.-I. Implementation of quantum imaginary-time evolution method on nisq devices by introducing nonlocal approximation. *NPJ Quantum Inf.* **7**, 1–7 (2021).
49. Gomes, N. *et al.* Efficient step-merged quantum imaginary time evolution algorithm for quantum chemistry. *J. Chem. Theory Comput.* **16**, 6256–6266 (2020).
50. Gomes, N. *et al.* Adaptive variational quantum imaginary time evolution approach for ground state preparation. *Adv. Quantum Technol.* **4**, 2100114 (2021).
51. Ville, J.-L. *et al.* Leveraging randomized compiling for the qite algorithm. Preprint at http://arxiv.org/abs/2104.08785 (2021).
52. Cao, C., An, Z., Hou, S.-Y., Zhou, D. & Zeng, B. Quantum imaginary time evolution steered by reinforcement learning. *Commun. Phys.* **5**, 1–7 (2022).
53. Iten, R., Colbeck, R., Kukuljan, I., Home, J. & Christandl, M. Quantum circuits for isometries. *Phys. Rev. A* **93**, 032318. https://doi.org/10.1103/PhysRevA.93.032318 (2016).
54. Shende, V. V., Bullock, S. S. & Markov, I. L. Synthesis of quantum logic circuits. In *Proc. 2005 Asia and South Pacific Design Automation Conference* 272–275 (2005).
55. Iten, R. *et al.* Introduction to universalqcompiler. Preprint at http://arxiv.org/abs/1904.01072 (2019).
56. Kosugi, T., Nishiya, Y., Nishi, H. & Matsushita, Y.-I. Imaginary-time evolution using forward and backward real-time evolution with a single ancilla: First-quantized eigensolver algorithm for quantum chemistry. *Phys. Rev. Res.* **4**, 033121 (2022).
57. McCaskey, A. J., Lyakh, D. I., Dumitrescu, E. F., Powers, S. S. & Humble, T. S. XACC: A system-level software infrastructure for heterogeneous quantum–classical computing. *Quantum Sci. Technol.* **5**, 024002 (2020).
58. Bassman, L., Powers, C. & de Jong, W. A. Arqtic: A full-stack software package for simulating materials on quantum computers. *ACM Trans. Quantum Comput.* **3**, 1–17 (2021).
59. Younis, E., Sen, K., Yelick, K. & Iancu, C. Qfast: Conflating search and numerical optimization for scalable quantum circuit synthesis. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* 232–243 (2021).
60. Smith, E. *et al.* Leap: Scaling numerical optimization based synthesis using an incremental approach. Preprint at http://arxiv.org/abs/2106.11246 (2021).
61. Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509. https://doi.org/10.1103/PhysRevLett.119.180509 (2017).
62. Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050. https://doi.org/10.1103/PhysRevX.7.021050 (2017).
63. Davis, M., Smith, E. & Younis, E. *qsearch*. https://github.com/BQSKit/qsearch (2021).

Author contributions

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-023-28317-5.

Correspondence and requests for materials should be addressed to C.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.