



Preparing quantum statistical ensembles using mid-circuit measurements

John P. T. Stenger¹ · C. Stephen Hellberg¹ · Daniel Gunlycke¹

Received: 31 October 2023 / Accepted: 6 May 2024

This is a U.S. Government work and not under copyright protection in the US; foreign copyright protection may apply 2024

Abstract

We explore the relationship between entropy and quantum measurements and present a variational algorithm for preparing statistical ensembles on quantum computers using mid-circuit measurements. This algorithm optimizes both the entropy and variational parameters describing the state to obtain the minimum free energy of quantum systems in thermal equilibrium with some external heat bath. We demonstrate our algorithm on IBM-Q Lagos.

Keywords Quantum computing · Quantum algorithm · Entropy · Mid-circuit measurement

1 Introduction

Quantum statistical ensembles are needed in many application areas, including solving optimization problems [1, 2], training Boltzmann machines in quantum machine learning [3–6], and simulating equilibrium physics [7–11]. Simulating equilibrium physics is important to a wide variety of solid-state and chemistry problems including heat transfer, chemistry in solution, charge mobility, and many more. The ability to prepare arbitrary quantum statistical ensembles on quantum computers is, therefore, critical for conducting simulations in these areas.

Some of the earliest quantum algorithms for preparing mixed states include quantum walks via a metropolis algorithm [12–14], quantum rejection sampling [15, 16], dimension reduction [17], dynamical simulation [10, 18, 19], and variational techniques [5–10, 20–23]. All mixed-state preparation techniques require that a non-unitary operator is applied to the system. There are several techniques for effectively applying the required non-unitary operator, including state amplification via Grover's algorithm [14, 16, 17], quantum imaginary time evolution [10, 11], and simulating a thermal environment by coupling the system to a bath of auxiliary qubits [5, 6, 8, 15, 20, 22, 24],

✉ John P. T. Stenger
john.p.stenger2.civ@us.navy.mil

¹ U.S. Naval Research Laboratory, Washington, DC 20375, USA

[25]. It is possible to avoid the use of auxiliary qubits when the simulated environment can be modeled by the environment of the quantum computer itself. In this case, some degree of control over that environment is required [26, 27].

Herein, we present a variational quantum algorithm for the preparation and optimization of mixed states on quantum computers based on a different approach. Specifically, we prepare mixed states that describe canonical ensembles of open quantum systems in thermal equilibrium with some external heat bath. The needed statistical mixtures in our algorithm are obtained through mid-circuit measurements. The entropy in our method is a function of the variational parameters in the quantum circuit. Thus, the entropy does not need to be determined from the final measurements, which can be challenging [8, 22]. Furthermore, because we do not use auxiliary qubits, our method requires the least number of qubits possible. Having direct access to the entropy and using minimal quantum resources are two main advantages of our method.

This paper is organized as follows. Section 2 presents the method we use to generate the quantum statistical ensembles. Section 3.1 presents ansatzes for generating specific values of entropy. In Sect. 4, we present a demonstration of our algorithm for a two-site Hubbard model. Section 5 provides a deeper discussion about mid-circuit measurements. Lastly, we present some concluding remarks in Sect. 6

2 Method

2.1 Problem definition

Consider a quantum system described by a Hamiltonian \hat{H} connected to an external heat bath with temperature T . The relevant thermodynamic potential for this canonical ensemble is the Helmholtz free energy

$$F(\hat{\rho}) = E(\hat{\rho}) - TS(\hat{\rho}), \quad (1)$$

which is given by the internal energy

$$E(\hat{\rho}) = \text{Tr}(\hat{H}\hat{\rho}), \quad (2)$$

and the von Neumann entropy

$$S(\hat{\rho}) = -k \text{Tr}(\hat{\rho} \ln \hat{\rho}), \quad (3)$$

where $\hat{\rho}$ is the density operator that describes the state of the system and k is the Boltzmann constant. Our goal is to find the $\hat{\rho}$ that minimizes $F(\hat{\rho})$ using a quantum computer.

We assume that a map from the Fock space \mathcal{F} of the quantum system to the Hilbert space $\mathbb{H}^{\otimes Q}$ of a quantum register with Q qubits is provided (e.g., the Jordan–Wigner mapping). A basis for $\mathbb{H}^{\otimes Q}$ is the set of states $\{|n\rangle\}$ that is formed by all states of the form $|n\rangle = \bigotimes_{q=1}^Q |n_q\rangle$, where $n_q \in \{0, 1\}$ specifies the basis state $|n_q\rangle$ of qubit q . In this basis, a general density operator can be written as $\hat{\rho} = \sum_{nn'} p_{nn'} |n\rangle\langle n'|$, where

$p_{nn'}$ are complex coefficients. Density operators can be constructed on the quantum register using a combination of quantum logic gates and mid-circuit measurements. Details of this construction are described in Sect. 3.

2.2 Approach

We need a way to create mixed states with nonzero entropy. We accomplish this by introducing mid-circuit measurements. The challenge is that the measurement of a quantum system in any pure state collapses that state to another pure state, and as all pure states have zero entropy, no single measurement induces entropy by itself. However, quantum mechanics is a probabilistic theory, where repeated measurements of the same system in the same state do not necessarily produce the same outcome. Thus, repeated measurements of the same system can produce quantum ensembles with nonzero entropy.

When we describe performing operations on a quantum computer, we are always referring to performing repeated operations on identically prepared quantum states. Let us define

$$\hat{\rho}_0 = |0\rangle\langle 0|, \quad (4)$$

as the initial density operator of the identically prepared quantum systems. Measuring this density operator simply returns the same density operator. In order to generate entropy, we must rotate the density operator out of the $\{|n\rangle\}$ basis. We define

$$\hat{\rho}(\phi) = \hat{V}(\phi)\hat{\rho}_0\hat{V}^\dagger(\phi), \quad (5)$$

where $\hat{V}(\phi)$ is a unitary operator and ϕ is a collection of parameters. We assume there is a known relation $s(\phi)$ such that when $\hat{\rho}$ is measured the resulting density operator has the entropy $s(\phi)$. Let us define $M(\hat{\rho})$ to represent the measurement of $\hat{\rho}$. In principle any complete set of measurement operators can be used. In this work, we choose to perform measurements in the occupation basis $\{|n\rangle\}$. These mid-circuit measurements then yield statistical samples of the basis states $|n\rangle$, which we use to approximate the probability distribution. After the measurement, the density operator becomes,

$$\hat{\rho}_{\text{meas}}(\phi) = M[\hat{\rho}(\phi)]. \quad (6)$$

From the construction of $\hat{V}(\phi)$, we have $S(\rho_{\text{meas}}) = s(\phi)$. This density operator has nonzero entropy; however, it is diagonal in the $\{|n\rangle\}$ basis. We want the density operator to be diagonal in the eigenbasis of the Hamiltonian. Thus, we perform a second unitary operation $\hat{U}(\theta)$ with a second collection of parameters θ and define

$$\hat{\rho}(\phi, \theta) = \hat{U}(\theta)\hat{\rho}_{\text{meas}}(\phi)\hat{U}^\dagger(\theta). \quad (7)$$

Because the entropy cannot be changed by applying a unitary operator, we know that $S(\hat{\rho}) = s(\phi)$. To obtain $E(\hat{\rho})$, we use the standard measurement protocol [28] for calculating expectation values from quantum measurement samples. There are more sophisticated measuring techniques [28]. However, those techniques are beyond

the scope of this work. Once both $S(\hat{\rho})$ and $E(\hat{\rho})$ have been obtained, $F(\hat{\rho})$ can be calculated from Eq. (1).

The goal is to optimize both ϕ and θ in order to minimize the Helmholtz free energy. We know from statistical mechanics that the free energy is minimized by the canonical ensemble

$$\hat{\rho}^* = \frac{1}{\zeta} \sum_{\lambda} e^{-\beta E_{\lambda}} |\psi_{\lambda}\rangle\langle\psi_{\lambda}|, \quad (8)$$

where E_{λ} are the eigenvalues and $|\psi_{\lambda}\rangle$ are the eigenstates of \hat{H} , λ runs over the set of eigenstates, and ζ is the partition function

$$\zeta = \sum_{\lambda} e^{-\beta E_{\lambda}}, \quad (9)$$

where $\beta = (kT)^{-1}$. In Appendix A, we show that there exist a $V(\phi)$ and $\hat{U}(\theta)$ that generate the canonical ensemble.

2.3 Algorithm

Using the approach discussed above, we can define a variational quantum algorithm for optimizing ϕ and θ so that the Helmholtz free energy (1) is minimized. The algorithm is shown in Table 1.

In step 1 of the algorithm, we choose initial values for ϕ_0 and θ_0 . Certain choices could speed up the optimization and help to avoid local minima. One choice that is likely to speed up the optimization is to set ϕ_0 so that $V(\phi_0)$ generates the Fermi–Dirac distribution

$$\bar{n}_q = \frac{1}{1 + e^{\beta(\epsilon_q - \mu)}}, \quad (10)$$

over the basis states $|n_q\rangle$ of the qubits where ϵ_q are the Hartree–Fock energies and μ is the chemical potential.

In step 3, we initialize identical copies of the quantum register. Requiring identical copies of the quantum register is not unique to our algorithm. It is required for any quantum algorithm that aims to calculate the expectation value of an operator. For any such algorithm, identical copies are required in order to acquire enough shots for a statistically accurate measurement of the expectation value. In practice, the quantum operations acting on these copies can be run sequentially instead of simultaneously, allowing us to reuse the same quantum register for each run.

In step 12, we choose a new set of parameters ϕ'_i and θ'_i and use those parameters in step 13 to compare two values of the free energy. This is the classical optimization routine that is common to quantum variational algorithms. The parameters ϕ'_i and θ'_i must be related in some way to the original parameters ϕ_i and θ_i . Let $\phi_{i\alpha}, \theta_{i\alpha}, \phi'_{i\alpha}$, and $\theta'_{i\alpha}$ be parameters in the collections $\phi_i, \theta_i, \phi'_i$, and θ'_i , respectively. We can relate the parameters by a small offset

$$\begin{aligned} \phi'_{i\alpha} &= \phi_{i\alpha} + \Delta\phi_{i\alpha} \\ \theta'_{i\alpha} &= \theta_{i\alpha} + \Delta\theta_{i\alpha}, \end{aligned} \quad (11)$$

Table 1 Free-energy minimization algorithm

1	Choose initial values for ϕ_0 and θ_0
2	Iterate over i
3	Initialize identical copies of the quantum register to
	Form the initial ensemble $\hat{\rho}_0$
4	Apply a set of quantum gates that perform the first
	Unitary operation $V(\phi_i)$
5	Measure each state of the quantum ensemble
6	Apply a set of quantum gates that perform the second
	Unitary operation $\hat{U}(\theta_i)$
7	Measure each state a second time to calculate the
	Energy expectation value of the ensemble $E(\phi_i, \theta_i)$
8	Use $E(\phi_i, \theta_i)$ and $s(\phi_i)$ to calculate the free energy
	$F(\phi_i, \theta_i)$
9	If $F(\phi_i, \theta_i)$ has converged
10	Exit the loop
11	Else
12	Repeat (3)-(8) with a set of perturbed parameters ϕ'_i and
	θ'_i to calculate $F(\phi'_i, \theta'_i)$
13	Using a predefined-optimization scheme, compare
	$F(\phi_i, \theta_i)$ to $F(\phi'_i, \theta'_i)$ in order to find a new set of
	parameters ϕ_{i+1} and θ_{i+1}
14	Return the optimized free energy

where $|\Delta\phi_{i\alpha}|$ and $|\Delta\theta_{i\alpha}|$ are small. In our demonstration below, we use the SPSA optimizer in which $\Delta\phi_{i\alpha} = \pm\Delta\phi_{i\beta}$ and $\Delta\theta_{i\alpha} = \pm\Delta\theta_{i\beta}$ for all α and β . The parameters in the next iteration are then chosen by taking the difference of the free energies

$$\phi_{i+1\alpha} = \phi_{i\alpha} - \gamma \frac{F(\phi, \theta) - F(\phi', \theta')}{\Delta\phi_{i\alpha}}. \quad (12)$$

for some predefined-optimization rate γ .

3 Entropy circuits

3.1 Discrete entropy injection via mid-circuit measurements

We want to find a unitary operator $\hat{V}(\phi)$ such that when $\hat{\rho} = \hat{V}(\phi)^\dagger |0\rangle\langle 0| V(\phi)$ is measured, the resulting density operator $\hat{\rho}_{\text{meas}} = M(\hat{\rho})$ has an entropy $S(\hat{\rho}_{\text{meas}}) = s(\phi)$. Let us consider a simple example of how measurements can inject entropy into the quantum system. Consider a single qubit initialized to $\hat{\rho}_0 = |0\rangle\langle 0|$. Then, apply a

Hadamard operator $\hat{\mathcal{H}} = \frac{1}{\sqrt{2}}(\hat{\sigma}_x + \hat{\sigma}_z)$. This results in a density operator

$$\hat{\rho} = \frac{1}{2}(|0\rangle + |1\rangle)(\langle 0| + \langle 1|). \quad (13)$$

Next, perform the measurements

$$\hat{\rho}_{\text{meas}} = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|). \quad (14)$$

The entropy of the initial state is zero

$$S(\hat{\rho}_0) = 0, \quad (15)$$

while the entropy of the measured state is

$$S(\hat{\rho}_{\text{meas}}) = k \ln 2. \quad (16)$$

If instead of a single qubit, we have a set of Q qubits and we apply Hadamard operators to Q_H of them, then the entropy becomes

$$S(\hat{\rho}_{\text{meas}}) = Q_H k \ln 2. \quad (17)$$

Thus, we can define

$$\hat{V}^{\text{disc}}(\phi_0) = \prod_{i=0}^{\lfloor \frac{\phi_0}{\ln 2} \rfloor} \hat{\mathcal{H}}_i, \quad (18)$$

where ϕ_0 is a parameter, $\lfloor \cdot \rfloor$ denotes the floor function, and $\hat{\mathcal{H}}_i$ is the Hadamard operator acting on qubit i . This works as long as we only need to inject discrete amounts of entropy.

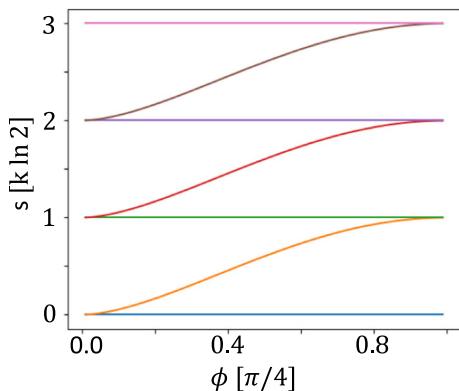
3.2 Continuous entropy injection via mid-circuit measurements

The construction in Sect. 3.1 is only able to generate entropy that is an integer multiple of $k \ln 2$. In general, we want to be able to generate any amount of entropy from $S_{\text{min}} = 0$ to $S_{\text{max}} = Nk \ln 2$, where N is the total number of qubits. This can be accomplished by applying a continuous rotation to one of the qubits. Let us define the operator

$$\hat{R}_i(\phi) = \hat{Z}_i \hat{R}_i^y(2\phi) = \hat{Z}_i \cos(\phi) + \hat{X}_i \sin(\phi), \quad (19)$$

where \hat{Z}_i and \hat{X}_i are the Pauli-z and Pauli-x matrices acting on qubit i , \hat{R}_i^y is the rotation operator around the y-axis of the Bloch sphere, and ϕ is an arbitrary angle. The \hat{Z}_i operator is used so that we can access the Hadamard operator, $\hat{R}_i(\pi/4) = \hat{\mathcal{H}}_i$.

Fig. 1 The value of $s(\phi)$ given input ϕ and Q_H . The straight lines from bottom to top have 0, 1, 2, and 3 \hat{H}_i operators and no \hat{R}_i operators. The curves from bottom to top have 0, 1, and 2 \hat{H}_i operators and 1 \hat{R}_i operator



Applying this operator to $\hat{\rho}_0 = |0\rangle\langle 0|$ gives us

$$\begin{aligned}\hat{\rho} &= \hat{R}_0(\phi)\hat{\rho}_0\hat{R}_0(\phi) \\ &= [\cos(\phi)|0\rangle + \sin(\phi)|1\rangle][\cos(\phi)\langle 0| + \sin(\phi)\langle 1|].\end{aligned}\quad (20)$$

Measuring generates an entropy

$$S[M(\hat{\rho})] = -k \cos^2(\phi) \ln [\cos^2(\phi)] - k \sin^2(\phi) \ln [\sin^2(\phi)], \quad (21)$$

that has a minimum of 0 and a maximum of $k \ln 2$. Therefore, we can generate an arbitrary amount of entropy by applying $Q_H = \lfloor \phi_0 / \ln 2 \rfloor$ Hadamard operators to the first Q_H qubits and a \hat{R}_i operator to the next qubit. We can apply these operators by setting $\hat{V}(\phi)$ to

$$\hat{V}(\phi) = \hat{V}^{\text{disc}}(\phi_0)R_{Q_H+1}(\phi), \quad (22)$$

where ϕ contains only a single parameter ϕ_0 and we define $\phi = \phi_0 - Q_H$ for convenience. The entropy is then given by

$$s(\phi)/k = -\cos^2(\phi) \ln [\cos^2(\phi)] - \sin^2(\phi) \ln [\sin^2(\phi)] + Q_H \ln(2). \quad (23)$$

Figure 1 shows plots of $s(\phi)$ for various combinations of Hadamard and rotation operations. We see that any input entropy can be generated from this form of $V(\phi)$. It should be noted that when $\hat{V}(\phi)$ is in this simple form, it does not need to be implemented on the quantum computer, see Appendix B for details.

3.3 Arbitrary distributions via mid-circuit measurements

Generating arbitrary entropy is not enough to generate an arbitrary density operator. Two density operators with the same entropy can have different eigenvalues. The eigenvalues of a density operator cannot be changed by applying unitary operators.

Therefore, a scheme to prepare a density operator with an arbitrary distribution of eigenvalues is required to solve for the minimum free energy in the most general case.

Preparing a density operator with arbitrary eigenvalues requires a $\hat{V}(\phi)$ composed of 2^N operators, in general. In order to avoid an exponential number of operators, $\hat{V}(\phi)$ has to be tailored to the specific problem. Bespoke operators are a common feature of variational quantum algorithms. See Appendix C for an example $\hat{V}(\phi)$ that balances complexity with control over the eigenvalues.

The entropy generated by measuring the density operator after applying an arbitrary $\hat{V}(\phi)$ may not be easily calculable. Fortunately, we do not have to calculate the entropy by hand. Instead, we can use the results of the mid-circuit measurements to obtain the entropy. The measurements provide us with a list of states $\{|n\rangle\}$ and the number of shots that were measured to be in each state $\{P_n\}$. While it is difficult to recover the entropy of the incoming ensemble, the entropy of the outgoing ensemble can be approximated using the results of the mid-circuit measurements. After the mid-circuit measurements, we approximate the density operator as $\hat{\rho} \approx \sum_n P_n / N_S |n\rangle \langle n|$, where N_S is the total number of shots. Therefore, the entropy is

$$S \approx - \sum_n \frac{P_n}{N_S} \ln \frac{P_n}{N_S}, \quad (24)$$

where the sum runs over only those states for which a shot is collected. The approximation becomes an equality for an infinite number of shots.

The number of shots required for an accurate approximation of S depends on the spread of the eigenvalues in the density operator. The distribution of eigenvalues in the canonical ensemble is controlled by the temperature. For both low $kT \ll E_{\max}$ and high $kT \gg E_{\max}$ temperatures, where E_{\max} is the width of the energy spectrum, the specific distribution of the eigenvalues of the density operator does not have a significant effect on the free energy. For low temperature, the majority of the weight is on the ground state. For high temperature, the weight tends toward an even distribution among all states. Therefore, we can use the form of $\hat{V}(\phi)$ in Sect. 3.2 for both the low and high entropy regimes. In the intermediate-temperature regime, we may need a more involved $\hat{V}(\phi)$ in which there is not a simple relation $s(\phi)$ for the entropy. Because the weights in the canonical ensemble (8) exponentially decay with energy, we expect to get a good approximation for the entropy using Eq. (24) with a subexponential number of shots, even in the intermediate regime.

3.4 Algorithm complexity

The overall complexity of the algorithm depends on many factors. For one, it depends on the circuit depth of $\hat{V}(\phi)$. As we have shown above, a full covering of the entropy can be achieved with a $V(\phi)$ that requires only Q gates. Further detail can be found in Appendix C.

The complexity also depends on the required depth of $\hat{U}(\phi)$ and on the number of optimization iterations required for the free energy to converge. These requirements are

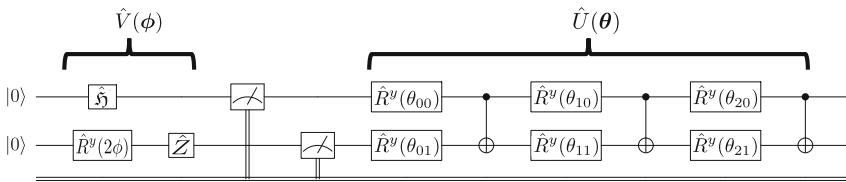


Fig. 2 Depiction of the circuit model we use in our demonstration. $\hat{V}(\phi)$ is composed of all of the operators before the initial measurement and $\hat{U}(\theta)$ is composed of all of the operators after the initial measurement. For $\hat{U}(\theta)$, we take the hardware efficient ansatz. In the figure, $\hat{V}(\phi)$ is composed of a Hadamard operator on the first qubit and a rotation on the second qubit; however, the optimizer is free to replace the Hadamard operator with a rotation on the first qubit if it removes the rotation from the second qubit. The rotation in $\hat{V}(\phi)$ is decomposed into native operators $\hat{R}_i(\phi) = \hat{R}_i^y(2\phi)\hat{Z}$, where $\hat{R}_i^y(\phi)$ is a rotation about the y-axis of the block sphere for qubit i and \hat{Z} is the Pauli-z matrix. The two qubit operators are controlled-NOT gates

open problems. However, it is generally believed that quantum variational algorithms require low depth circuits and can be optimized efficiently [29].

Finally, the complexity depends on the number shots required to achieve an accurate sampling of the probability distribution during both the mid-circuit and final measurements. The number of shots required to achieve an accuracy of ϵ goes as $N_s = \mathcal{O}(1/\epsilon^2)$ [29]. Notice that the \ln function in the definition of the entropy Eq. (3) means that a direct approximation of the entropy Eq. (24) can require high accuracy for some cases. However, we argue in Sect. 3.3 that we can use exact expressions for the entropy in most cases. We further argue that, for the cases where these expressions are not available, the entropy is well approximated without high accuracy.

4 Demonstrating the algorithm

We demonstrate our algorithm for the two-orbital Hubbard model described by the Hamiltonian

$$H_{\text{Hubbard}} = \mu(\hat{n}_0 + \hat{n}_1) + t(c_0^\dagger c_1 + c_1^\dagger c_0) + u\hat{n}_0\hat{n}_1, \quad (25)$$

where c_i^\dagger and c_i are creation and destruction operators, respectively, $\hat{n}_i = c_i^\dagger c_i$ is the number operator, μ is the chemical potential, t is the kinetic energy, and u is the fermion interaction strength. We scale all of the parameters by t . We set $\mu = -3.7t$ and $u = 0.7t$. We use the Jordan–Wigner transformation to convert Eq. (25) to a spin model that can be mapped to the quantum register.

Figure 2 shows the circuit model. Unless otherwise stated, we choose $\hat{V}(\phi)$ to be a list of Hadamard operators and one $\hat{R}_i(\phi)$ operator, as described in Eq. (19), and we take $\hat{U}(\theta)$ to be the hardware efficient ansatz. For $\hat{V}(\phi)$, we start with no Hadamard operators and a rotation operator on the first qubit. If the angle of the rotation operator reaches $\phi = \pi/4$ then a Hadamard operator is added to the first qubit, the rotation operator is moved to the second qubit, and the angle is reset to $\phi = 0$. The entropy associated with ϕ is determined by the relationship $s(\phi)$ from Eq (23). We are free

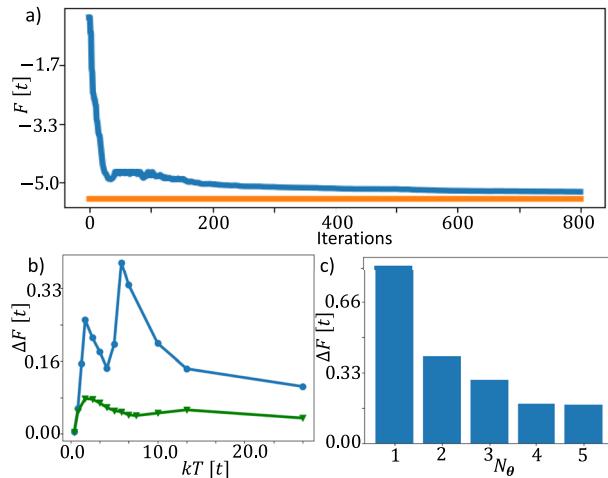


Fig. 3 Simulation of the free-energy algorithm. **a** shows the free energy as a function of optimization iterations. The blue curve is the simulated data and the orange line is the true minimum of the free energy. **b** shows the difference between the true minimum of the free energy and the simulated result as a function of temperature. The blue line with the blue dots shows results from the continuous-entropy method described in Sect. 3.2, where we use only one $\hat{R}_i(\phi)$ operator. The blue dots show the data points the blue line is plotted to guide the eye. The green line with the green triangles shows the results from the arbitrary-distribution method described in Appendix. B, where an $\hat{R}_i(\phi_i)$ operator is applied to each qubit. The green triangles show the data points and the green line is plotted to guide the eye. **c** shows the difference between the true minimum of the free energy and the simulated result for different numbers of layers of the ansatz. In **a** and **b**, $N_\theta = 4$ and in **a** and **c**, $kT = 3.3 t$

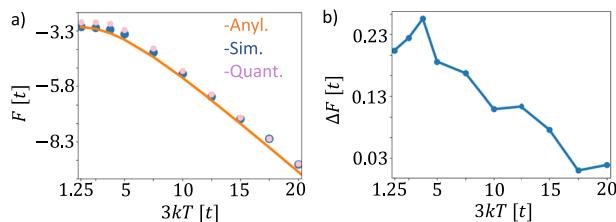


Fig. 4 Run of the free-energy algorithm on IBM-Q Lagos. **a** the free energy as a function of temperature. The orange curve is the true minimum of the free energy, the large blue data points are from classical simulation, and the small pink data points are taken from the quantum computer. For the quantum data, we use the optimized angles and entropy from the classical simulation. **b** the difference between the free energy measured on the quantum device and the free energy from the simulation. In both **a** and **b**, $N_\theta = 4$

to choose the number of angles in θ , which we define as $2N_\theta$ so that N_θ counts the number of layers in \hat{U} .

In Fig. 3, we show the results of a noise-free classical simulation of our algorithm. In Fig. 3a, we plot the free energy as a function of the optimization iterations. We observe an initial fast optimization for about 20 iterations then a second slow optimization with convergence around 600 iterations. In Fig. 3b, we plot the energy difference between the true-minimum free energy and the results of our simulated algorithm. The simulation is most accurate for small and large values of T . The decrease in accuracy

in the intermediate-temperature regime is expected, as discussed in Sect. 3.3. Even in this intermediate regime, the free energy from the simulation gets well within 10% of the true minimum of the free energy. Also, as predicted in Sect. 3.3, the free energy in the intermediate regime is more accurate when using the balanced ansatz described in Appendix C than when using the continuous-entropy ansatz described in Sect. 3.2. The accuracy of the algorithm is also dependent on N_θ . Figure 3c shows the free energy as a function of N_θ . We see that the final result improves as N_θ is increased from 1 to 4. At $N_\theta = 5$ and beyond, the algorithm does not improve.

We use the optimized parameter values from the classical simulation to test the circuit on a quantum computer. Figure 4 shows data from the quantum computer IBM-Q Lagos. See Appendix D for details about the quantum computer. Figure 4a shows the free energy as a function of temperature. Figure 4b shows the difference in free energy between the data from the simulation and the data from the quantum device. The error from the quantum device is less than 10% at $T = 0$ and decreases as the temperature is raised. This is due to the fact that at higher temperature, the density operator becomes more evenly spread across the eigenspectrum, so the free energy is tolerant of inexact quantum operators.

5 Discussion

While we focus on calculation of the Helmholtz free energy using a variational approach, we foresee the use of mid-circuit measurements to generate entropy as a component that can be used in a variety of algorithms. Mid-circuit measurements can be used, in general, to both increase and decrease the entropy as a subroutine in a quantum algorithm.

Above, we go into detail about how mid-circuit measurements can be used to increase entropy. More care is required in order to decrease entropy. As an expression of the second law of thermodynamics, we cannot reduce the entropy of the quantum system by measurement alone. If we want to reduce the entropy, we must perform post-selection using the measurement results. By sorting out shots with certain measurement outcomes, we artificially set some of the $P_m = 0$ in Eq. (24). This is analogous to the idea of a Maxwell Demon, which lowers the entropy of a system by selectively allowing only low energy particles through a screen. Postselection, like the Maxwell Demon, does not lower the entropy of the universe but simply redefines what we consider to be the system.

6 Conclusion

We present a method for generating quantum ensembles on quantum computers using mid-circuit measurements. The ability to generate mixed-state quantum ensembles is a critical component in simulating any physical system that is not isolated from its environment. Our work establishes a relationship between entropy and quantum measurements, which applies to all quantum algorithms that utilize mid-circuit measurements. We use this relationship to construct an algorithm for finding the free

energy of open systems on a quantum computer. We use the measurement apparatus as the environment for the open system. This allows us to avoid using auxiliary qubits. Furthermore, as entropy is unchanged after the mid-circuit measurement, we do not need to measure the final entropy, which can be challenging. Thus, we have two main advantages. First, because we do not use auxiliary qubits, the amount of quantum resources needed for our algorithm is greatly reduced. Second, the entropy does not need to be approximated because it is provided as an input parameter. Furthermore, our method is a natural extension of the variational quantum eigensolver and is practical for near-term quantum hardware.

Appendix A: Existence of an exact solution within the framework of our method

Here, we show that unitaries \hat{V}^* and \hat{U}^* , that generate the exact canonical ensemble Eq. (8), exist. These unitaries require information about the eigenvalues of the Hamiltonian and cannot be generated efficiently in general. However, it is important to demonstrate that the algorithm has the capacity to find the global minimum free energy. This is analogous to the variational quantum eigensolver algorithm in which there is some unitary operator that generates the exact ground state but in practice we use an ansatz that may not reach the global energy minimum.

Let us choose \hat{V}^* such that

$$\hat{\varrho} = \frac{1}{\zeta} \left(\sum_n e^{-\beta E_n/2} |n\rangle \right) \left(\sum_{n'} e^{-\beta E_{n'}/2} \langle n'| \right), \quad (\text{A1})$$

where $|n\rangle$ describes the state of the quantum register as in the previous section. As $\hat{\varrho}$ represents a normalized pure state, there is guaranteed to be a unitary operator \hat{V}^* that generates this state. Next, we measure $\hat{\varrho}$ and obtain

$$\hat{\varrho}_{\text{meas}} = \frac{1}{\zeta} \sum_n e^{-\beta E_n} |n\rangle \langle n|. \quad (\text{A2})$$

The next step is to apply the second unitary operator \hat{U}^* . We can build \hat{U}^* out of the eigenvectors of \hat{H} such that $|\psi_{\lambda=n}\rangle = \hat{U}^*|n\rangle$. Applying this \hat{U}^* to $\hat{\varrho}_{\text{meas}}$, we have

$$\hat{\varrho} = \frac{1}{\zeta} \sum_{\lambda} e^{-\beta E_{\lambda}} |\psi_{\lambda}\rangle \langle \psi_{\lambda}| \quad (\text{A3})$$

and, thus, $\hat{\varrho} = \hat{\varrho}^*$. Therefore, there exists a \hat{U}^* and a \hat{V}^* that generate the absolute minimum of the free energy.

Appendix B: Classical simulation of the measurement procedure

The algorithm is composed of two unitary operators: $\hat{V}(\phi)$, which is applied before the entropy-injecting measurement, and $\hat{U}(\theta)$, which is applied after the measurement. When $\hat{V}(\phi)$ is composed of only single-qubit operators, it can be efficiently simulated on a classical computer. The result of measuring $\hat{\rho} = \hat{V}^\dagger(\phi)|0\rangle\langle 0|\hat{V}(\phi)$ can simply be stated. Each qubit with a Hadamard operator has a fifty percent chance of being in the $|1\rangle$ state while each qubit with an $\hat{R}_i(\phi)$ operator has a $\cos^2(\phi)$ chance of being $|1\rangle$ state. As there are no entangling operators, orientation of each qubit is independent of all the other qubits.

Therefore, instead of generating $\hat{\rho}_{\text{meas}} = M(\hat{\rho})$ on the quantum computer, we can start each shot of the algorithm in a randomly generated state

$$\hat{\rho}_{\text{meas},s} = \prod_i \hat{X}_i^{n_{si}} |00\dots\rangle\langle 00\dots| \prod_i \hat{X}_i^{n_{si}}, \quad (\text{B1})$$

where $n_{si} \in \{0, 1\}$ is a random binary. We select each n_{si} from a probability distribution $P(n_{si} = 1)$ that depends on the operator that $\hat{V}(\phi)$ applies to qubit i :

$$\begin{aligned} \hat{\mathcal{H}}_i &\mapsto P(n_{si} = 1) = 1/2, \\ \hat{R}_i(\phi) &\mapsto P(n_{si} = 1) = \cos^2(\phi), \\ \hat{I}_i &\mapsto P(n_{si} = 1) = 0, \end{aligned} \quad (\text{B2})$$

where \hat{I}_i is the identity acting on qubit i . With enough shots we effectively generate the desired ensemble

$$\hat{\rho}_{\text{meas}} = \lim_{N_S \rightarrow \infty} \frac{1}{N_S} \sum_{s=1}^{N_S} \hat{\rho}_{\text{meas},s}. \quad (\text{B3})$$

While generating $\hat{\rho}_{\text{meas}}$ classically in this manner is equivalent to Eq. (22), it does not reproduce an arbitrary $\hat{V}(\phi)$, where $\hat{\rho}_{\text{meas}}$ can be exponentially difficult to approximate.

Appendix C: A balanced ansatz to control the density operator

While the full range of entropy can be reached using the methods in the main text, we cannot generate an arbitrary distribution of eigenvalues for the density operator using single-qubit rotations alone. In this section, we discuss how one can systematically increase control over the density operator. In general, the density operator has 2^N eigenvalues, where N is the number of qubits. However, we assume that we only ever need a subexponential number of control parameters. This assumption is based on the fact that the canonical ensemble (8) exponentially decays with energy. While this decay can be washed out at very high temperature, the high-temperature regime is well captured by the methods in the main text. It is only in the intermediate-temperature regime where deep control over the eigenvalues of the density operator is required.

In the main text, we take some data for a pre-measurement ansatz that applies single-qubit rotation $\hat{R}_i(\phi_i)$ operators to each qubit, where each ϕ_i is an independent parameter. Let us go into more detail about this ansatz as it is an excellent balance between control of the density operator and depth of the circuit. After mid-circuit measurements, the density operator becomes

$$M(\hat{\rho}) = \sum_{n_1, n_2, \dots} \prod_i \cos^{2(1-n_i)}(\phi_i) \sin^{2n_i}(\phi_i) \times |n_1 n_2 \dots\rangle \langle n_1 n_2 \dots|, \quad (\text{C1})$$

for $n_i \in \{0, 1\}$.

The measured density operator has an entropy of

$$S[M(\hat{\rho})] = - \sum_{n_1, n_2, \dots} \prod_i \cos^{2(1-n_i)}(\phi_i) \sin^{2n_i}(\phi_i) \times \ln \left[\prod_i \cos^{2(1-n_i)}(\phi_i) \sin^{2n_i}(\phi_i) \right]. \quad (\text{C2})$$

This entropy is exponentially difficult to calculate as there are 2^N terms in the summation. However, we can approximate the entropy from the mid-circuit measurement results as described in Sect. 3.3.

Notice that in Eq. (C1) the density operator $M(\hat{\rho})$ has 2^N nonzero eigenvalues. However, there are only N independent control parameters. Therefore, many of the eigenvalues are coupled. To decouple the eigenvalues we must apply multi-qubit rotations

$$\hat{R}_{\{i\}_n}^n(\phi) = \cos \phi_{\{i\}_n} \prod_{i \in \{i\}_n} \hat{Z}_i + \sin \phi_{\{i\}_n} \prod_{i \in \{i\}_n} \hat{X}_i, \quad (\text{C3})$$

where $\{i\}_n$ is a set of n indices. We can systematically increase the control over the eigenvalues of the density operator by successively adding additional rotation operators to the pre-measurement ansatz. In general, if we apply N_r unique rotations, as described by Eq. (C3), followed by mid-circuit measurements, we have a density operator of the form

$$M(\hat{\rho}) = \sum_{n_1, n_2, \dots} f(\phi_1, \phi_2, \dots, \phi_{N_r}) |n_1 n_2 \dots\rangle \langle n_1 n_2 \dots|, \quad (\text{C4})$$

where f is a sum of products of trig functions. We see that each additional rotation provides an additional control parameter to decouple the eigenstates of the density operator. In the limit $N_r = 2^N$, the eigenvalues are fully controlled; however, this requires exponentially many rotations.

Even for relatively few multi-qubit rotations, the circuit depth can become large as the rotations can require several CNOT operators. The number of CNOT operators depends on the exact arrangement of the indices and the layout of the qubits on the chip. Fortunately, there are two aspects of the algorithm that work in our favor. For one, the mid-circuit measurements restart the decoherence. This means that increasing the

Table 2 Specifications of IBM-Q Lagos during the execution of our algorithm

	q_0	q_1
Frequency (GHz)	5.24	5.10
T1 (μ s)	112.9	118.6
T2 (μ s)	41.0	89.9
Readout error	0.0075	0.0109
\sqrt{X} error	0.00025	0.00031
CX_{q_i, q_0} error		0.0076
CX_{q_i, q_1} error	0.0076	
CX_{q_i, q_2} error		0.0074
CX_{q_i, q_3} error		0.0083

depth of the pre-measurement ansatz does not increase the decoherence in the post-measurement ansatz. Secondly, because single-qubit rotations can span all possible values of the entropy, it is rare that we need more than a few two qubit rotations. As shown in Ref. [9], a quantum statistical ensemble with the correct entropy can often be transformed into a good approximation of the desired ensemble using only unitary operators.

Appendix D: Device specifications

We demonstrate our method on IBM-Q Lagos, which is a seven-qubit quantum computer. Lagos has the following qubit pairs connected: (q_0, q_1) , (q_1, q_2) , (q_1, q_3) , (q_3, q_5) , (q_4, q_5) , (q_5, q_6) , where q_i refers to qubit i . We use only qubits q_0 and q_1 . The specifications of the device at the time we ran our circuits can be found in Table 2. All of the operators were converted to z-rotations R^z , root-x operators \sqrt{X} , and controlled-not operators CX using the Qiskit transpile function. The R_z operators have no error as they are implemented by changing the measurement angles. These specifications are based on a calibration performed on August 25, 2022.

Acknowledgements This work has been supported by the Office of Naval Research (ONR) through the U.S. Naval Research Laboratory (NRL). We acknowledge quantum computing resources from IBM through a collaboration with the Air Force Research Laboratory (AFRL).

Author Contributions All authors contributed to the study conception and design. Material preparation, data collection, and analysis were performed by J.S. The first draft of the manuscript was written by J.S. and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Declarations

Conflict of interest The authors declare no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence,

and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by Simulated Annealing. *Science* **220**, 671 (1983)
2. Somma, R.D., Boixo, S., Barnum, H., Knill, E.: Quantum Simulations of Classical Annealing Processes. *Phys. Rev. Lett.* **101**, 130504 (2008). <https://doi.org/10.1103/PhysRevLett.101.130504>
3. Kieferová, M., Wiebe, N.: Tomography and generative training with quantum Boltzmann machines. *Phys. Rev. A* **96**, 062327 (2017). <https://doi.org/10.1103/PhysRevA.96.062327>
4. Amin, M.H., Andriyash, E., Rolfe, J., Kulchytskyy, B., Melko, R.: Quantum Boltzmann Machine. *Phys. Rev. X* **8**, 021050 (2018). <https://doi.org/10.1103/PhysRevX.8.021050>
5. Anschuetz, E.R., Cao, Y.: Realizing quantum boltzmann machines through eigenstate thermalization (2019) <https://doi.org/10.48550/ARXIV.1903.01359>
6. Zoufal, C., Lucchi, A., Woerner, S.: Variational quantum Boltzmann machines. *Quantum Mach. Intell.* (2021). <https://doi.org/10.1007/s42484-020-00033-7>
7. Islam, R., Ma, R., Preiss, P.M., Eric Tai, M., Lukin, A., Rispoli, M., Greiner, M.: Measuring entanglement entropy in a quantum many-body system. *Nature* **528**, 77 (2015). <https://doi.org/10.1038/nature15750>
8. Wu, J., Hsieh, T.H.: Variational Thermal Quantum Simulation via Thermofield Double States. *Phys. Rev. Lett.* **123**, 220502 (2019). <https://doi.org/10.1103/PhysRevLett.123.220502>
9. Martyn, J., Swingle, B.: Product spectrum ansatz and the simplicity of thermal states. *Phys. Rev. A* (2019). <https://doi.org/10.1103/physreva.100.032107>
10. Motta, M., Sun, C., Tan, A.T.K., O'Rourke, M.J., Ye, E., Minnich, A.J., Brandão, F.G.S.L., Chan, G.K.-L.: Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat. Phys.* **16**, 205 (2019). <https://doi.org/10.1038/s41567-019-0704-4>
11. Sun, S.-N., Motta, M., Tazhigulov, R.N., Tan, A.T., Chan, G.K.-L., Minnich, A.J.: Quantum Computation of Finite-Temperature Static and Dynamical Properties of Spin Systems Using Quantum Imaginary Time Evolution. *PRX Quantum* **2**, 010317 (2021). <https://doi.org/10.1103/PRXQuantum.2.010317>
12. Terhal, B.M., DiVincenzo, D.P.: Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A* **61**, 022301 (2000). <https://doi.org/10.1103/PhysRevA.61.022301>
13. Temme, K., Osborne, T.J., Vollbrecht, K.G., Poulin, D., Verstraete, F.: Quantum Metropolis sampling. *Nature* **471**, 87 (2011). <https://doi.org/10.1038/nature09770>
14. Yung, M.-H., Aspuru-Guzik, A.: A quantum–quantum Metropolis algorithm. *Proc. Natl. Acad. Sci.* **109**, 754 (2012). <https://doi.org/10.1073/pnas.1111758109>
15. Poulin, D., Wocjan, P.: Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer. *Phys. Rev. Lett.* (2009). <https://doi.org/10.1103/physrevlett.103.220502>
16. Wiebe, N., Kapoor, A., Svore, K.M.: Quantum Deep Learning. *Quantum Info. Comput.* **16**, 541–587 (2016)
17. Bilgin, E., Boixo, S.: Preparing Thermal States of Quantum Systems by Dimension Reduction. *Phys. Rev. Lett.* **105**, 170405 (2010). <https://doi.org/10.1103/PhysRevLett.105.170405>
18. Riera, A., Gogolin, C., Eisert, J.: Thermalization in Nature and on a Quantum Computer. *Phys. Rev. Lett.* **108**, 080402 (2012). <https://doi.org/10.1103/PhysRevLett.108.080402>
19. Kaplan, D.B., Klco, N., Roggero, A.: Ground States via Spectral Combing on a Quantum Computer (2017) <https://doi.org/10.48550/ARXIV.1709.08250>
20. Yuan, X., Endo, S., Zhao, Q., Li, Y., Benjamin, S.C.: Theory of variational quantum simulation. *Quantum* **3**, 191 (2019). <https://doi.org/10.22331/q-2019-10-07-191>
21. Chowdhury, A.N., Low, G.H., Wiebe, N.: A variational quantum algorithm for preparing quantum gibbs states (2020) <https://doi.org/10.48550/ARXIV.2002.00055>

22. Wang, Y., Li, G., Wang, X.: Variational Quantum Gibbs State Preparation with a Truncated Taylor Series. *Phys. Rev. Appl.* (2021). <https://doi.org/10.1103/physrevapplied.16.054035>
23. Liang, J.-M., Lv, Q.-Q., Wang, Z.-X., Fei, S.-M.: Assisted quantum simulation of open quantum systems. *iScience* **26**, 106306 (2023). <https://doi.org/10.1016/j.isci.2023.106306>
24. Wang, H., Ashhab, S., Nori, F.: Quantum algorithm for simulating the dynamics of an open quantum system. *Phys. Rev. A* **83**, 062317 (2011). <https://doi.org/10.1103/PhysRevA.83.062317>
25. Su, H.-Y., Li, Y.: Quantum algorithm for the simulation of open-system dynamics and thermalization. *Phys. Rev. A* **101**, 012328 (2020). <https://doi.org/10.1103/PhysRevA.101.012328>
26. Verstraete, F., Wolf, M.M., Cirac, J.I.: Quantum computation, quantum state engineering, and quantum phase transitions driven by dissipation (2008) <https://doi.org/10.48550/ARXIV.0803.1447>
27. Brandao, F.G., Kastoryano, M.J.: Finite correlation length implies efficient preparation of quantum thermal states (2016) <https://doi.org/10.48550/ARXIV.1609.07877>
28. García-Pérez, G., Rossi, M.A., Sokolov, B., Tacchino, F., Barkoutsos, P.K., Mazzola, G., Tavernelli, I., Maniscalco, S.: Learning to Measure: Adaptive Informationally Complete Generalized Measurements for Quantum Algorithms. *PRX Quantum* **2**, 040342 (2021). <https://doi.org/10.1103/PRXQuantum.2.040342>
29. Wecker, D., Hastings, M.B., Troyer, M.: Progress towards practical quantum variational algorithms. *Phys. Rev. A* **92**, 042303 (2015). <https://doi.org/10.1103/PhysRevA.92.042303>

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