Variational preparation of finite-temperature states on a quantum computer

Sagastizabal, R.; Klaver, B. A.; Rol, M. A.; Negîrneac, V.; Zou, X.; Muthusubramanian, N.; Beekman, M.; Zachariadis, C.; Ostroukh, V. P.; Haider, N.

ARTICLE OPEN

Variational preparation of finite-temperature states on a quantum computer

R. Sagastizabal[1,2,8], S. P. Premaratne [ID][3,8], B. A. Klaver[1,2], M. A. Rol [ID][1,2,6], V. Negîrneac [ID][1,4], M. S. Moreira [ID][1,2], X. Zou[3], S. Johri[3,7], N. Muthusubramanian [ID][1,2], M. Beekman[1,5], C. Zachariadis[1,2], V. P. Ostroukh [ID][1,2], N. Haider[1,5], A. Bruno[1,2], A. Y. Matsuura[3] and L. DiCarlo [ID][1,2✉]

The preparation of thermal equilibrium states is important for the simulation of condensed matter and cosmology systems using a quantum computer. We present a method to prepare such mixed states with unitary operators and demonstrate this technique experimentally using a gate-based quantum processor. Our method targets the generation of thermofield double states using a hybrid quantum-classical variational approach motivated by quantum-approximate optimization algorithms, without prior calculation of optimal variational parameters by numerical simulation. The fidelity of generated states to the thermal-equilibrium state smoothly varies from 99 to 75% between infinite and near-zero simulated temperature, in quantitative agreement with numerical simulations of the noisy quantum processor with error parameters drawn from experiment.

INTRODUCTION

The potential for quantum computers to simulate other quantum mechanical systems is well known[1], and the ability to represent the dynamical evolution of quantum many-body systems has been demonstrated[2]. However, the accuracy of these simulations depends on efficient initial state preparation within the quantum computer. Much progress has been made on the efficient preparation of non-trivial quantum states, including spin-squeezed states[3] and entangled cat states[4]. Studying phenomena like high-temperature superconductivity[5] requires preparation of thermal equilibrium states, or Gibbs states. Producing mixed states with unitary quantum operations and measurements is not straightforward, and has only recently begun to be explored[6,7]. In this work, we demonstrate the use of a variational quantum-classical algorithm to realize Gibbs states using (ideally unitary) gate control on a transmon quantum processor.

Our approach is mediated by the generation of thermofield double (TFD) states, which are pure states sharing entanglement between two identical quantum systems with the characteristic that when one of the systems is considered independently (by tracing over the other), the result is a mixed state representing equilibrium at a specific temperature. TFD states are of interest not only in condensed matter physics but also for the study of black holes[8,9] and traversable wormholes[10,11]. We use a variational protocol[12] motivated by quantum-approximate optimization algorithms (QAOA) that relies on alternation of unitary intra- and inter-system operations to control the effective temperature, eliminating the need for a large external heat bath. Other methods have been studied for generation of Gibbs states, such as quantum metropolis sampling[13] and imaginary time evolution using variational quantum simulation[14,15]. However, the advantage of QAOA compared to these proposals is that the form of the ansatz is relatively straightforward and low-depth, whereas the metropolis sampling involves phase estimation which leads to a high-depth circuit, and the imaginary time evolution proposal does not have a clear proposal for the form of the ansatz. Recently, verification of TFD state preparation was demonstrated on a trapped-ion quantum computer[6]. Our work experimentally demonstrates the generation of finite-temperature states in a superconducting quantum computer by variational preparation of TFD states in a hybrid quantum-classical manner.

RESULTS

Theory

Consider a quantum system described by Hamiltonian H with eigenstates $|j\rangle$ and corresponding eigenenergies E_j:

$$H|j\rangle = E_j|j\rangle. \tag{1}$$

The Gibbs state ρ_{Gibbs} of the system is

$$\rho_{\text{Gibbs}}(\beta) = \frac{1}{Z}\sum_j \exp(-\beta E_j)|j\rangle\langle j|, \tag{2}$$

where $\beta = 1/k_B T$ is the inverse temperature, k_B is the Boltzmann constant, and

$$Z = \sum_j \exp(-\beta E_j) \tag{3}$$

is the partition function. Except in the limit $\beta \to \infty$, the Gibbs state is a mixed state and thus impossible to generate strictly through unitary evolution. To circumvent this, we define the TFD state[12] on two identical systems A and B as

$$|\text{TFD}(\beta)\rangle = \frac{1}{\sqrt{Z}}\sum_j \exp\left(\frac{-\beta E_j}{2}\right)|j\rangle_B |j\rangle_A. \tag{4}$$

Tracing out either system yields the desired Gibbs state in the other.

[1]QuTech, Delft University of Technology, Delft, The Netherlands. [2]Kavli Institute of Nanoscience, Delft University of Technology, Delft, The Netherlands. [3]Intel Labs, Intel Corporation, Hillsboro, OR, USA. [4]Instituto Superior Técnico, Lisbon, Portugal. [5]Netherlands Organisation for Applied Scientific Research (TNO), Delft, The Netherlands. [6]Present address: Orange Quantum Systems, Delft, The Netherlands. [7]Present address: IonQ, College Park, MD, USA. [8]These authors contributed equally: R. Sagastizabal, S. P. Premaratne. ✉email: l.dicarlo@tudelft.nl
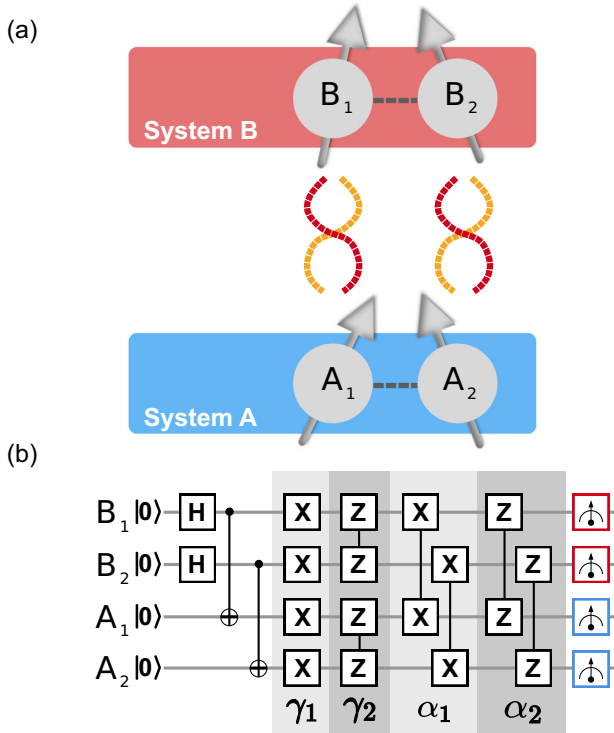
(a)

(b)

Fig. 1 Principle and generation of Thermofield Double state.
a Two identical systems A and B are variationally prepared in an ideally pure, entangled joint state such that tracing out one system yields the Gibbs state on the other. **b** Corresponding qubits in the two systems are first pairwise entangled to produce the $\beta = 0$ TFD state. Next, intra- and inter-system Hamiltonians are applied with optimized variational angles $(\boldsymbol{\alpha}, \boldsymbol{\gamma})$ to approximate the TFD state corresponding to the desired temperature.

To prepare the TFD states, we follow the variational protocol proposed by[12] and consider two systems each of size n. In the first step of the procedure, the TFD state at $\beta = 0$ is generated by creating Bell pairs $|\Phi_i^+\rangle = (|0\rangle_{Bi}|0\rangle_{Ai} + |1\rangle_{Bi}|1\rangle_{Ai})/\sqrt{2}$ between corresponding qubits i in the two systems. Tracing out either system yields a maximally mixed state on the other, and vice versa. The next steps to create the TFD state at finite temperature depend on the relevant Hamiltonian. Here, we choose the transverse field Ising model in a one-dimensional chain of n spins[16], with $n = 2$ [Fig. 1(a)]. We map spin up (down) to the computational state $|0\rangle (|1\rangle)$ of the corresponding transmon. The Hamiltonian describing system A is

$$H_A = ZZ_A + gX_A, \tag{5}$$

where $ZZ_A = Z_{A2}Z_{A1}$, $X_A = X_{A2} + X_{A1}$, and g is proportional to the transverse magnetic field. The Hamiltonian for system B is the same. We focus on $g = 1$, where a phase transition is expected in the transverse field Ising model at large n[17]. We use a QAOA-motivated variational ansatz[12,18], where intra-system evolution is interleaved with a Hamiltonian enforcing interaction between the systems:

$$H_{BA} = XX_{BA} + ZZ_{BA}, \tag{6}$$

where $XX_{BA} = X_{B2}X_{A2} + X_{B1}X_{A1}$, and analogously for ZZ_{BA}. For single-step state generation, the unitary operation describing the TFD protocol is

$$U(\boldsymbol{\alpha}, \boldsymbol{\gamma}) = U_{\text{inter}}(\boldsymbol{\alpha})U_{\text{intra}}(\boldsymbol{\gamma}), \tag{7}$$

where

$$
\begin{aligned}
U_{\text{intra}}(\boldsymbol{\gamma}) &= \exp[-i\gamma_2(ZZ_B + ZZ_A)/2] \\
&\quad \times \exp[-i\gamma_1(X_B + X_A)/2],
\end{aligned} \tag{8}
$$

$$U_{\text{inter}}(\boldsymbol{\alpha}) = \exp(-i\alpha_2 ZZ_{BA}/2)\exp(-i\alpha_1 XX_{BA}/2). \tag{9}$$

The variational parameters $\boldsymbol{\gamma} = (\gamma_1, \gamma_2)$, $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$ are optimized by the hybrid classical-quantum algorithm to generate states closest to the ideal TFD states. A single step of intra- and inter-system interaction ideally produces the state $|\psi(\boldsymbol{\alpha}, \boldsymbol{\gamma})\rangle = U(\boldsymbol{\alpha}, \boldsymbol{\gamma})(|\Phi_2^+\rangle \otimes |\Phi_1^+\rangle)$[19].

The variational algorithm extracts the cost function after each state preparation. We engineer a cost function $\mathcal{C}$ to be minimized when the generated state is closest to an ideal TFD state[19]. Following recent work on the concentration of control parameters for QAOA[20,21], we expect engineering the cost function based on the target state for small-sized systems to lead to a general expression for the cost function of an arbitrary-sized system. The engineered cost function is given by:

$$
\begin{aligned}
C(\beta) &= \langle X_A\rangle + \langle X_B\rangle + 1.57(\langle ZZ_A\rangle + \langle ZZ_B\rangle) \\
&\quad - \beta^{-1.57}(\langle XX_{BA}\rangle + \langle ZZ_{BA}\rangle).
\end{aligned} \tag{10}
$$

We compare the performance of this engineered cost function $\mathcal{C}_{1.57}$ to that of the non-optimized cost function $\mathcal{C}_{1.00}$, using the reduction of infidelity to the Gibbs state as the ultimate metric of success [see Supplementary Note 1]. The engineered cost function achieves an average improvement of 54% across the β range covered ($[10^{-2}, 10^2]$ in units of $1/g$), as well as a maximum improvement of up to 98% for intermediate temperatures ($\beta \sim 1$). Our choice of the class of cost functions to optimize lets us trade off a slight decrease in low-temperature performance with a significant increase in performance at intermediate temperatures. See ref. [19] for further details on the theory.

The quantum portion of the algorithm prepares the state according to a given set of angles $(\boldsymbol{\alpha}, \boldsymbol{\gamma})$, performs the measurements, and returns these values to the classical portion. The classical portion then evaluates the cost function according to the returned measurements, performs classical optimization, generates and returns the next set of variational angles to evaluate the quantum portion.

Experiment

We implement the algorithm using four of seven transmons in a monolithic quantum processor [Fig. 2(a)]. The four transmons (labeled A_1, A_2, B_1, and B_2) have square connectivity provided by coupling bus resonators, and are thus ideally suited for implementing the circuit in Fig. 1(b). Each transmon has a microwave-drive line for single-qubit gating, a flux-bias line for two-qubit controlled-Z (CZ) gates, and a dispersively coupled resonator with dedicated Purcell filter[22,23]. The four transmons can be simultaneously and independently read out by frequency multiplexing, using the common feedline connecting to all Purcell filters. All transmons are biased to their flux-symmetry point (i.e., sweetspot[24]) using static flux bias to counter residual offsets. Device details and a summary of measured transmon parameters are provided in Supplementary Note 3. Details on the experimental setup can be found in Supplementary Note 4.

In order to realize the theoretical circuit in Fig. 1(b), we first map it to the optimized depth-13 equivalent circuit shown in Fig. 2(b), which conforms to the native gate set in our control architecture. This gate set consists of arbitrary single-qubit rotations about any equatorial axis of the Bloch sphere, and CZ gates between nearest-neighbor transmons. Conveniently, all variational angles are mapped to either the axis or angle of single-qubit rotations. Further details on the compilation steps are reported in Methods and Supplementary Note 2. Bases pre-rotations are added at the
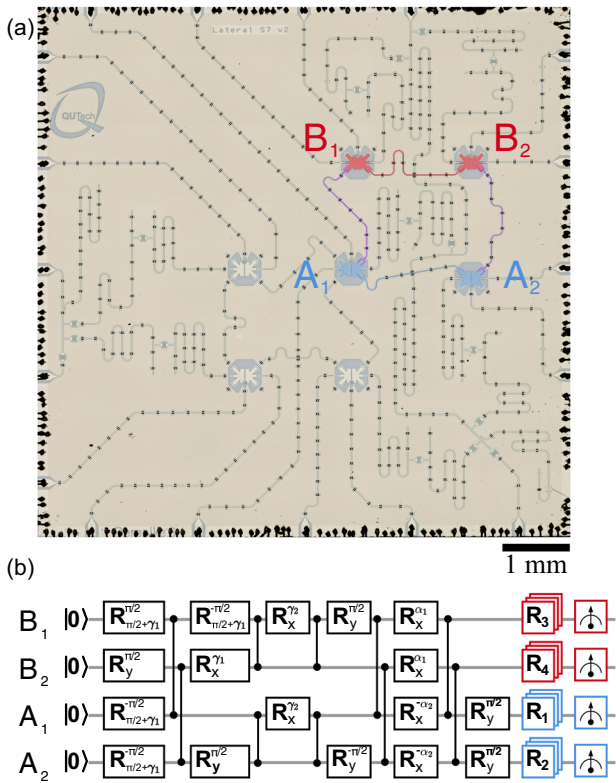
Fig. 2 Device and optimized quantum circuit. a Optical image of the transmon processor was used in this experiment, with false color highlighting the four transmons employed and the dedicated bus resonators providing their nearest-neighbor coupling. Scale bar corresponds to 1 mm. **b** Optimized circuit equivalent to that in Fig. 1 (**b**) and conforming to the native gate set in our architecture. All variational parameters are mapped onto rotation axes and angles of single-qubit gates. Tomographic pre-rotations $R_1–R_4$ are added to reconstruct the terms in the cost function $\mathcal{C}$ and to perform two-qubit state tomography of each system following optimization.

end of the circuit to first extract all the terms in the cost function $\mathcal{C}$ and finally to perform two-qubit state tomography of each system.

Prior to implementing any variational optimizer, it is helpful to build a basic understanding of the cost-function landscape. To this end, we investigate the cost function $\mathcal{C}$ at $\beta = 0$ using two-dimensional cuts Fig. 3: we sweep γ while keeping $\boldsymbol{a} = 0$ to study the effect of U_{intra} and vice versa to study the effect of U_{inter}. Note that owing to the $\beta^{-1.57}$ divergence, the cost function reduces to $-\langle H_{BA} \rangle$ in the $\beta = 0$ limit. Consider first the landscape for an ideal quantum processor, which is possible to compute for our system size. The γ landscape at $\boldsymbol{a} = 0$ is π-periodic in both directions due to the invariance of $|\mathrm{TFD}(\beta = 0)\rangle$ under bit-flip (X) and phase-flip (Z) operations on all qubits. The cost function is minimized to -4 at even multiples of $\pi/2$ on γ_1 and γ_2: $|\mathrm{TFD}(\beta = 0)\rangle$ is a simultaneous eigenstate of XX_{BA} and ZZ_{BA} with eigenvalue $+2$ due to the symmetry of the constituting Bell states $|\Phi_i^+\rangle$. In turn, the cost function is maximized to $+4$ at odd multiples of $\pi/2$, at which the $|\Phi_i^+\rangle$ are transformed to singlets $|\Psi_i^-\rangle = \left(|0\rangle_{Bi} |1\rangle_{Ai} - |1\rangle_{Bi} |0\rangle_{Ai} \right)/\sqrt{2}$. The $\boldsymbol{a}$ landscape at $\gamma = 0$ is constant, reflecting that $|\mathrm{TFD}(\beta = 0)\rangle$ is a simultaneous eigenstate of XX_{BA} and ZZ_{BA} and thus also of any exponentiation of these operators. The corresponding experimental landscapes show qualitatively similar behavior. The γ landscape clearly shows the π periodicity with respect to both angles, albeit with reduced contrast. The $\boldsymbol{a}$ landscape is not strictly constant, showing weak structure particularly with respect to a_2. These experimental deviations reflect underlying errors in our noisy intermediate-scale quantum

(NISQ) processor, which include transmon decoherence, residual ZZ coupling at the bias point, and leakage during CZ gates. We discuss these error sources in detail further below.

The task of the variational algorithm is to balance the mixture of the states at each β, in order to generate the corresponding Gibbs state. Although thermal states are well understood, it is challenging to accurately generate them in NISQ devices for studies of finite temperature systems. When working with small systems, it is possible and tempting to predetermine the variational parameters at each β by a prior classical simulation and optimization for an ideal or noisy quantum processor. We refer to this common practice[6,25] as *cheating*, since this approach does not scale to larger problem sizes and skips the main quality of variational algorithms: to arrive at the parameters variationally. Here, we avoid cheating altogether by starting at $\beta = 0$, with initial guess the obvious optimal variational parameters for an ideal processor ($\gamma = \boldsymbol{a} = 0$), and using the experimentally optimized $(\boldsymbol{a}, \gamma)$ at the last β as an initial guess when stepping β in the range $[0, 5]$ (in units of $1/g$). This approach only relies on the assumption that solutions (and their corresponding optimal variational angles) vary smoothly with β. At each β, we use the Gradient-Based Random-Tree optimizer of the *scikit-optimize*[26] Python package to minimize $\mathcal{C}$, using 4096 averages per tomographic pre-rotation necessary for the calculation of $\mathcal{C}$. After 200 iterations, the optimization is stopped. The best point is remeasured two times, each with 16384 averages per tomographic pre-rotation needed to perform two-qubit quantum state tomography of each system. A new optimization is then started for the next β, using the previous solution as the initial guess.

To begin comparing the optimized states ρ_{Exp} produced in an experiment to the target Gibbs states ρ_{Gibbs}, we first visualize their density matrices (in the computational basis) for a sampling of the β range covered (Fig. 4). Starting from the maximally-mixed state I $I/4$ at $\beta = 0$, the Gibbs state monotonically develops coherences (off-diagonal terms) between all states as β increases. Coherences between states of equal (opposite) parity have 0 (π) phase throughout. Populations (diagonal terms) monotonically decrease (increase) for even (odd) parity states. By $\beta = 5$, the Gibbs state is very close to the pure state $|\Upsilon\rangle\langle\Upsilon|$, where $|\Upsilon\rangle \approx \sqrt{0.36}(|01\rangle + |10\rangle) - \sqrt{0.14}(|00\rangle + |11\rangle)$. The noted trends are reproduced in ρ_{Exp}. However, the matching is evidently not perfect, and to address this we proceed to a quantitative analysis.

We employ two metrics to quantify experimental performance: the fidelity F of ρ_{Exp} to ρ_{Gibbs} and the purity P of ρ_{Exp}, given by

$$F = \mathrm{Tr}\left(\sqrt{\rho_{\mathrm{Gibbs}}^{1/2} \rho_{\mathrm{Exp}} \rho_{\mathrm{Gibbs}}^{1/2}} \right), \quad (11)$$

$$P = \mathrm{Tr}\left(\rho_{\mathrm{Exp}}^2 \right). \quad (12)$$

At $\beta = 0$, $F = 99\%$ and $P = 0.262$, revealing a very close match to the ideal maximally-mixed state. However, F smoothly worsens with increasing β, decreasing to 92% at $\beta = 1$ and 75% by $\beta = 5$. Simultaneously, P does not closely track the increase of purity of the Gibbs state. By $\beta = 5$, the Gibbs state is nearly pure, but P peaks at 0.601.

In an effort to quantitatively explain these discrepancies, we perform a full density-matrix simulation of a four-qutrit system using *quantumsim*[27]. Our simulation incrementally adds calibrated errors for our NISQ processor, starting from an ideal processor (model 0): transmon relaxation and dephasing times at the bias point (model 1), increased dephasing from flux noise during CZ gates (model 2), crosstalk from residual ZZ coupling at the bias point (model 3), and transmon leakage to the second-excited state during CZ gates (model 4). The experimental input parameters for each increment and details of modeling are described in Methods and Supplementary Notes 5–9. The added curves in Fig. 5 clearly show that model 4 quantitatively matches the observed
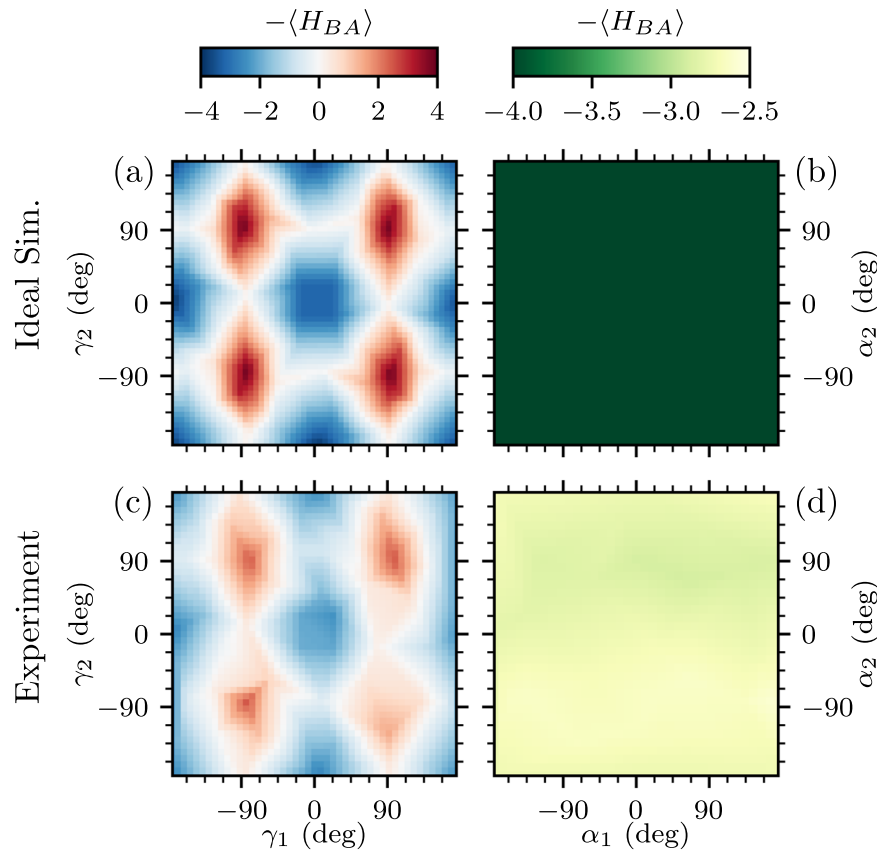
Fig. 3 Landscape of the cost function for infinite temperature. Panels (**a**) and (**c**): Landscape cuts obtained in noiseless simulation and in experiment, respectively, when varying control angles γ while keeping $\alpha = 0$ (the ideal solution value). Panels (**b**) and (**d**): Corresponding cuts obtained when varying control angles α while keeping $\gamma = 0$ (the ideal solution value). See text for details. These landscape cuts are sampled at 100 points and interpolated using the Python package *adaptive*[41].

dependence of F and P over the full β range, and identifies leakage from CZ gates as the dominant error.

DISCUSSION

The power of variational algorithms relies on their adaptability: the optimizer is meant to find its way through the variational parameter space, adapting to mitigate coherent errors as allowed by the chosen parametrization. For completeness, we compare in Fig. 5 the performance achieved with our variational strategy to that achieved by cheating, i.e., using the pre-calculated optimal (α, γ) for an ideal processor. Our variational approach, whose sole input is the obvious initial guess at $\beta = 0$, achieves comparable performance at all β. This aspect is crucial when considering the scaling with problem size, as classical pre-simulations will require prohibitive resources beyond ~50 qubits, but variational optimizers would not. Given the dominant role of leakage as the error source, which cannot be compensated by the chosen parametrization, it is unsurprising in hindsight that both approaches yield nearly identical performance.

In summary, we have presented the generation of finite-temperature Gibbs states in a quantum computer by variational targeting of TFD states in a hybrid quantum-classical manner. The algorithm successfully prepares mixed states for the transverse field Ising model with Gibbs-state fidelity ranging from 99% to 75% as β increases from 0 to $5/g$. The loss of fidelity with decreasing simulated temperature is quantitatively matched by a numerical simulation with incremental error models based on experimental input parameters, which identifies leakage in CZ gates as dominant. This work demonstrates the suitability of

variational algorithms on NISQ processors for the study of finite-temperature problems of interest, ranging from condensed-matter physics to cosmology. Our results also highlight the critical importance of continuing to reduce leakage in two-qubit operations when employing weakly-anharmonic multi-level systems such as the transmon.

During the preparation of this manuscript, we became aware of related experimental work[25] on a trapped-ion system, applying a non-variationally prepared TFD state to the calculation of a critical point.

METHODS

Quantum circuit

We map the theoretical circuit in Fig. 1(b) to an equivalent circuit conforming to the native gate set in our control architecture and exploiting virtual Z-gate compilation[28] to minimize circuit depth. Single-qubit rotations $R_{XY}(\phi, \theta)$, by arbitrary angle θ around any equatorial axis $\cos(\phi)\mathbf{x} + \sin(\phi)\mathbf{y}$ on the Bloch sphere, are realized using 20ns DRAG pulses[29,30]. Two-qubit CZ gates are realized by baseband flux pulsing[31,32] using the Net Zero scheme[33,34], completing in 80ns. In the optimized circuit [Fig. 2(b)], CZ gates only appear in pairs. These pairs are simultaneously executed and tuned as one block. Single-qubit rotations R_1-R_4 are used to change the measurement bases, as required to measure $\mathcal{C}$ during optimization and to perform two-qubit tomography[35] in each system to extract F and P. A summary of single- and two-qubit gate performance [see Supplementary Note 5] and a step-by-step derivation of the optimized circuit is provided [see Supplementary Note 2].
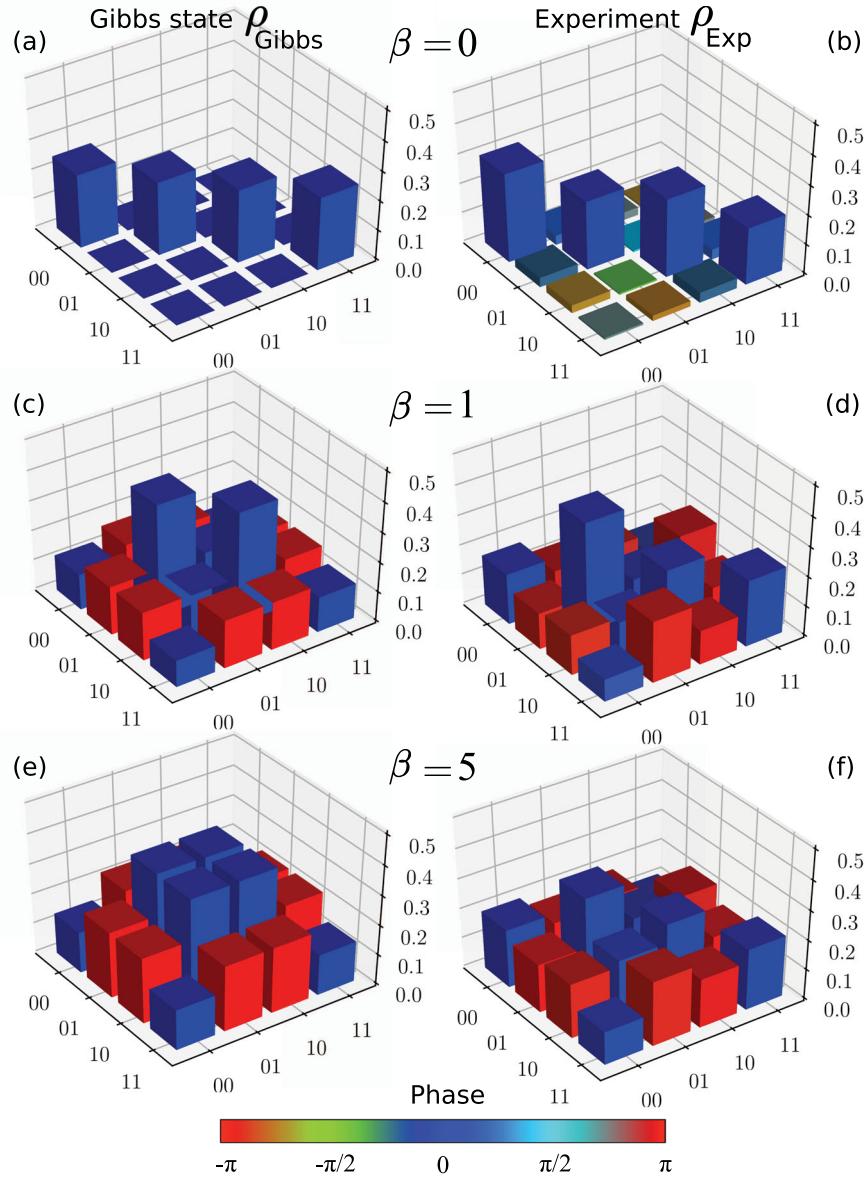
Fig. 4 Qualitative comparison of optimized states to the Gibbs state. Visualization of the density matrices (in the computational basis) for the targeted Gibbs states ρ_{Gibbs} (left) and the optimized experimental states ρ_{Exp} (right) at (**a, b**) $\beta = 0$, **a–d** $\beta = 1$ and (**e–f**) $\beta = 5$. As β increases, the Gibbs state monotonically develops coherence between all states, with phase 0 (π) for states with the same (opposite) parity. Populations in even (odd) parity states decrease (increase). The optimized experimental states show qualitatively similar trends.

Modeling and simulations

Noiseless simulations were performed prior to experiments, for verification of algorithm convergence. During simulations, experimental conditions were maintained exactly for the algorithm and the control software, while the outcome from readout hardware (Zurich Instruments UHFQC) was replaced with a simulated readout [see Supplementary Note 10 for details].

The models used to simulate the performance of the algorithm are incremental: model k contains all the noise mechanisms in model $k - 1$ plus one more, which we use for labeling in Fig. 5. Model 0 corresponds to an ideal quantum processor without any error. Model 1 adds the measured relaxation and dephasing times measured for the four transmons at their bias point. These times are tabulated in Supplementary Table 1. Model 2 adds the increased dephasing that flux-pulsed transmons experience during CZ gates. For this, we extrapolate the echo coherence time T_2^{echo} to the CZ flux-pulse amplitude using a $1/f$ noise model[36,37] with amplitude $\sqrt{A} = 1\mu\Phi_0$. This noise model is implemented following[38]. Model 3 adds the idling crosstalk due to residual ZZ coupling between transmons. This model expands on the implementation of idling evolution used for coherence times: the circuit gates are simulated to be instantaneous, and

the idling evolution of the system is trotterized. In this case, the residual ZZ coupling operator uses the measured residual ZZ coupling strengths at the bias point [see Supplementary Note 6 for details]. Finally, model 4 adds leakage to the CZ gates, based on randomized benchmarking with modifications to quantify leakage[33,39], and implemented in simulation using the procedure described in ref. [38].

Leakage in transmons

Leakage to transmon second-excited states is found essential to quantitatively match the performance of the algorithm by simulation. To reach this conclusion it was necessary to thoroughly understand how leakage affects the two-qubit tomographic reconstruction procedure employed. The readout calibration only considers computational states of the two transmons involved. Moreover, basis pre-rotations only act on the qubit subspace, leaving the population in leaked states unchanged. Using an overcomplete set of basis pre-rotations for state tomography, comprising both positive (X, Y, Z) and negative $(-X, -Y, -Z)$ bases for each transmon, leads to the misdiagnosis of a leaked state as a maximally
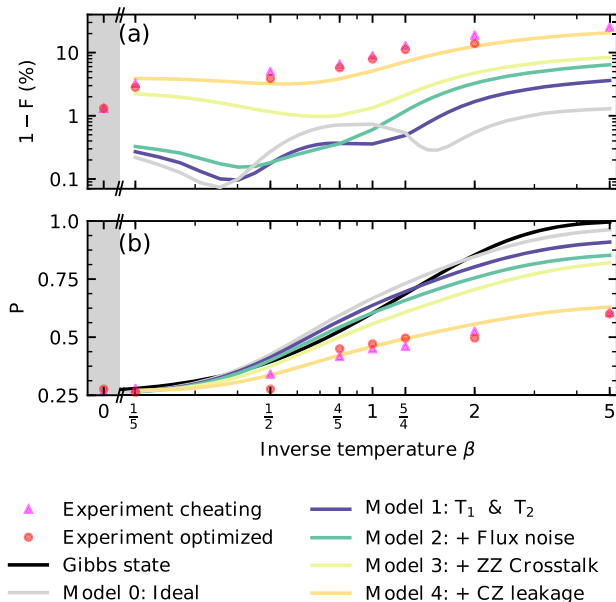
Fig. 5 **Performance of the variational algorithm. a** Fidelity to the Gibbs state as a function of inverse temperature β for experimental states obtained by optimization and cheating. **b** Purity of experimental states as a function of β, and comparison to the purity of the Gibbs state. Added curves in both panels are obtained by numerical simulation of a noisy quantum processor with incremental error models based on calibrated error sources for our device: qubit relaxation and dephasing times, increased dephasing from flux noise during CZ gates, residual ZZ crosstalk at the bias point, and leakage during CZ gates. Leakage is identified as the dominant error source.

mixed state qubit state for that transmon. See Supplementary Note 8 for further details.

DATA AVAILABILITY

The data from experiments and simulations in this work are available at the online repository `DiCarloLab-Delft/Gibbs_States_Data`[40].

CODE AVAILABILITY

The code that is deemed central to the conclusions is available from the corresponding author upon reasonable request.

REFERENCES

1. Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488 (1982).
2. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073 (1996).
3. Hosten, O., Engelsen, N. J., Krishnakumar, R. & Kasevich, M. A. Measurement noise 100 times lower than the quantum-projection limit using entangled atoms. *Nature* **529**, 505–508 (2016).
4. Vlastakis, B. et al. Deterministically encoding quantum information using 100-photon Schrödinger cat states. *Science* **342**, 607–610 (2013).
5. Lee, P. A., Nagaosa, N. & Wen, X.-G. Doping a mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.* **78**, 17–85 (2006).
6. Zhu, D. et al. Generation of thermofield double states and critical ground states with a quantum computer. *P. Natl Acad. Sci. USA* **117**, 25402–25406 (2020).
7. Chowdhury, A. N., Low, G. H. & Wiebe, N. A variational quantum algorithm for preparing quantum gibbs states (2020). Preprint at https://arxiv.org/abs/2002.00055.

8. Israel, W. Thermo-field dynamics of black holes. *Phys. Rev. A* **57**, 107 – 110 (1976).
9. Maldacena, J. Eternal black holes in anti-de sitter. *J. High Energy Phys.* **2003**, 021–021 (2003).
10. Maldacena, J., Stanford, D. & Yang, Z. Diving into traversable wormholes. *Fortschritte der Physik* **65**, 1700034 (2017).
11. Gao, P., Jafferis, D. L. & Wall, A. C. Traversable wormholes via a double trace deformation. *J. High Energy Phys.* **2017**, 151 (2017).
12. Wu, J. & Hsieh, T. H. Variational thermal quantum simulation via thermofield double states. *Phys. Rev. Lett.* **123**, 220502 (2019).
13. Temme, K., Osbourne, T., Vollbrecht, K., Poulin, D. & Verstraete, F. Quantum metropolis sampling. *Nature* **471**, 87 (2011).
14. Yuan, X., Endo, S., Zhao, Q., Li, Y. & S.C., B. Theory of variational quantum simulation. *Quantum* **3**, 191 (2019).
15. Motta, M. et al. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat. Phys.* **16**, 205–210 (2020).
16. Ho, W. W. & Hsieh, T. H. Efficient variational simulation of non-trivial quantum states. *SciPost Phys.* **6**, 29 (2019).
17. de Alcantara Bonfim, O. F., Boechat, B. & Florencio, J. Ground-state properties of the one-dimensional transverse Ising model in a longitudinal magnetic field. *Phys. Rev. E* **99**, 012122 (2019).
18. Hadfield, S. et al. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithm* **12** (2019).
19. Premaratne, S. P. & Matsuura, A. Y. Engineering a cost function for real-world implementation of a variational quantum algorithm. In *Proc. 2020 IEEE Int. Conf. Quantum Comp. Eng.*, 278–285 (2020).
20. Brandão, F. G. S. L., Broughton, M., Farhi, E., Gutmann, S. & Neven, H. For Fixed Control Parameters the Quantum Approximate Optimization Algorithm's Objective Function Value Concentrates for Typical Instances Preprint at https://arxiv.org/abs/1812.04170v1 (2018).
21. Streif, M. & Leib, M. Training the Quantum Approximate Optimization Algorithm without access to a Quantum Processing Unit. *Quantum Sci. Technol.* **5**, 034008. Preprint at https://iopscience.iop.org/article/10.1088/2058-9565/ab8c2b (2020).
22. Heinsoo, J. et al. Rapid high-fidelity multiplexed readout of superconducting qubits. *Phys. Rev. App.* **10**, 034040 (2018).
23. Bultink, C. C. et al. Protecting quantum entanglement from leakage and qubit errors via repetitive parity measurements. *Sci. Adv.* **6** (2020).
24. Schreier, J. A. et al. Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B* **77**, 180502(R) (2008).
25. Francis, A. et al. Many body thermodynamics on quantum computers via partition function zeros Preprint at https://arxiv.org/abs/2009.04648 (2020).
26. Head, T., Kumar, M., Nahrstaedt, H., Louppe, G. & Scherbatyi, I. Scikit-optimize. https://doi.org/10.5281/zenodo.4014775 (2020).
27. Tarasinski, B. M., Ostroukh, V. P., Bonet-Monroig, X., O'Brien, T. E. & Varbanov, B. Quantumsim. https://gitlab.com/quantumsim (2016).
28. McKay, D. C., Wood, C. J., Sheldon, S., Chow, J. M. & Gambetta, J. M. Efficient z gates for quantum computing. *Phys. Rev. A* **96**, 022330 (2017).
29. Motzoi, F., Gambetta, J. M., Rebentrost, P. & Wilhelm, F. K. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Phys. Rev. Lett.* **103**, 110501 (2009).
30. Chow, J. M. et al. Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Phys. Rev. A* **82**, 040305 (2010).
31. Strauch, F. W. et al. Quantum logic gates for coupled superconducting phase qubits. *Phys. Rev. Lett.* **91**, 167005 (2003).
32. DiCarlo, L. et al. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature* **460**, 240 (2009).
33. Rol, M. A. et al. Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits. *Phys. Rev. Lett.* **123**, 120502 (2019).
34. Rol, M. A. et al. Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor. *App. Phys. Lett.* **116**, 054001 (2020).
35. Sagastizabal, R. et al. Error mitigation by symmetry verification on a variational quantum eigensolver. *Phys. Rev. A* **100**, 010302(R) (2019).
36. Wellstood, F., Urbina, C. & Clarke, J. Excess noise in dc squids from 4.2 k to 0.022 k. *IEEE Trans. Magn.* **23**, 1662–1665 (1987).
37. Paladino, E., Galperin, Y. M., Falci, G. & Altshuler, B. L. 1/f noise: Implications for solid-state quantum information. *Rev. Mod. Phys.* **86**, 361–418 (2014).
38. Varbanov, B. M. et al. Leakage detection for a transmon-based surface code. *npj Quantum Inf.* **6**, 102 (2020).
39. Asaad, S. et al. Independent, extensible control of same-frequency superconducting qubits by selective broadcasting. *npj Quantum Inf.* **2**, 16029 (2016).
40. Sagastizabal, R. et al. Gibbs States Data. https://github.com/DiCarloLab-Delft/Gibbs_States_Data (2021).
41. Nijholt, B., Weston, J., Hoofwijk, J. & Akhmerov, A. Adaptive: parallel active learning of mathematical functions (2019). https://doi.org/10.5281/zenodo.1182437.

ACKNOWLEDGEMENTS

AUTHOR CONTRIBUTIONS

M.B, N.H., and L.D.C. designed the device. A.B., N.M., and C.Z. fabricated it. R.S. and B. A.K. performed the measurements and analysed the data with contributions from M. A.R and V.N. M.S.M. designed control electronics. S.P., X.Z., and S.J. developed the algorithm and theory. S.P., B.A.K., V.O., and R.S. performed the numerical simulations. A.M. and L.D.C. supervised theory and experiment, respectively. R.S. and S.P. wrote the manuscript with contributions from A.M. and L.D.C., and feedback from all co-authors. R.S. and S.P. contributed equally to this work as co-first authors.

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/s41534-021-00468-1.

Correspondence and requests for materials should be addressed to L.D.

Reprints and permission information is available at http://www.nature.com/reprints

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