

Classically compiling large-scale quantum circuits for approximate ground and thermal states

Simon Deconihout¹, Cenk Tüysüz², Enrico Rinaldi³, and Michele Grossi²

¹ École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

² European Organisation for Nuclear Research (CERN), Geneva, Switzerland

³ Quantinuum, London, United Kingdom

Introduction. Physical systems rarely sit exactly in their ground state. At finite temperature, a system weakly coupled to a large environment relaxes toward thermal equilibrium, and its equilibrium properties are described by the Gibbs state. Preparing Gibbs states on quantum computers has developed into an active research area. Efficient preparation with provable guarantees has been established in the high-temperature regime [1], and exact quantum Gibbs samplers have been constructed for general local Hamiltonians [2]. These algorithms target fault-tolerant devices, and the circuit depths and error rates they require place reliable execution beyond near-term hardware. A complementary line of work tailors algorithms to near-term devices through variational optimization. Variational quantum imaginary time evolution (VarQITE) [3, 4] has been adapted to Gibbs-state preparation [5–7]. Although these methods run on near-term hardware, they inherit limitations of variational training. Several cost functions exhibit barren plateaus, and even when barren plateaus are absent, the number of variational parameters often grows with system size, increasing the number of circuit evaluations and limiting scalability.

In this work, we introduce PP-VarQITE, a method that constructs variational circuits approximating ground and Gibbs states by combining VarQITE with Pauli propagation (Fig. 1). VarQITE evolves the circuit parameters along imaginary time using a metric matrix and a gradient vector built from expectation values of the parametrized circuit. We compute these expectation values classically with Pauli propagation, which removes the costly iterative quantum optimization and the shot overhead associated with barren plateaus. The ansatz is designed to remain shallow and executable on near-term hardware. As we detail below, PP-VarQITE prepares faithful ground and thermal states for up to 100 qubits (50 system qubits), reaching low-temperature regimes where direct operator-basis Pauli propagation loses accuracy. The classical Pauli-propagation step only evaluates the local expectation values required to update the variational parameters. It compiles the VarQITE trajectory, rather than simulating the final quantum state. After optimization, we obtain a shallow circuit, implementable on current hardware, that prepares the initial state for downstream algorithms such as sample-based quantum diagonalization [8].

Background. Consider a Hamiltonian $H = \sum_a h_a P_a$ acting on n qubits. At inverse temperature $\beta = (k_B T)^{-1}$, the thermal equilibrium state is $\rho_\beta = \exp(-\beta H)/Z$, where the partition function $Z = \text{Tr}[\exp(-\beta H)]$ ensures the normalization $\text{Tr}[\rho_\beta] = 1$. Since ρ_β is a mixed state, it cannot be represented as a state vector on the system Hilbert space alone. Nevertheless, ρ_β can be prepared by imaginary-time evolution of a purification on $2n$ qubits [5]. The system and ancilla qubits start in a maximally entangled state, whose reduced system state is the infinite-temperature Gibbs state $1/2^n$. Evolution under $H \otimes \mathbb{1}$ then cools the system, and tracing out the ancilla yields ρ_β .

Quantum imaginary-time evolution evolves an initial state according to the non-unitary operator $e^{-\tau H}$, suppressing high-energy components as the imaginary time τ increases. VarQITE [3, 4] approximates the evolved state by a unitary parametrized variational circuit that prepares the state

$$|\psi(\tau)\rangle \approx |\phi(\bar{\theta}(\tau))\rangle = U_1(\theta_1(\tau)) \dots U_L(\theta_L(\tau)) |\bar{0}\rangle, \quad (1)$$

on $2n$ qubits. Using McLachlan’s variational principle, the imaginary-time dynamics is projected onto the tangent space of the variational manifold, leading to the parameter update equation

$$M\dot{\theta} = V, \quad \text{where} \quad \begin{cases} M_{ij} = \text{Re}[\langle \partial_{\theta_i} \psi | \partial_{\theta_j} \psi \rangle], \\ V_i = -\text{Re}[\sum_a h_a \langle \partial_{\theta_i} \psi | P_a | \psi \rangle]. \end{cases} \quad (2)$$

The matrix M defines the local metric of the variational manifold, while V is the projection of the imaginary-time force generated by the Hamiltonian onto the variational directions.




Fig. 1. Overview of PP-VarQITE. The VarQITE parameter updates are evaluated classically with Pauli propagation, compiling a shallow variational circuit that prepares the thermal state of the target Hamiltonian.

Pauli propagation is a classical framework that tracks the evolution of observables under quantum channels rather than propagating full quantum states [9]. An observable is expanded in the Pauli basis, $O = \sum_i c_i P_i$, with Pauli strings P_i and complex coefficients c_i . Each elementary operation $\mathcal{E}_\ell$ maps a Pauli string P to $\mathcal{E}_\ell[P] = \sum_{j=1}^{N_{\mathcal{E}_\ell, P}} c_j^{(\ell)} P_j$, a sum of $N_{\mathcal{E}_\ell, P}$ non-zero terms. Iterating these transformations over all strings reconstructs the full evolution $\mathcal{E}[O]$ and yields expectation values of the form $\text{Tr}(\mathcal{E}[O] \rho)$. In practice, the expansion is truncated by discarding Pauli strings with coefficients below a cutoff, which sets the accuracy and classical cost.

PP-VarQITE. In standard VarQITE, the entries of M and V are estimated as expectation values measured on a quantum computer. This becomes a major bottleneck as the ansatz grows. For L variational parameters and a Hamiltonian with $|H|$ Pauli terms, a direct implementation requires evaluating $O(L^2)$ coefficients for M and $O(|H|)$ coefficients for V at each imaginary-time step. This scaling creates a trade-off between ansatz expressivity and the cost of the update rule. To reduce this measurement overhead, we replace the quantum estimation of these coefficients by a classical Pauli-propagation subroutine.

The effectiveness of VarQITE depends on the choice of ansatz. While hardware-efficient ansätze are easy to implement, they do not explicitly encode the structure of the target Hamiltonian and may contain trainable parameters that receive no update. To build an efficient ansatz, we use a selection rule derived from the evolution vector V . Since the Hamiltonian coefficients h_a are real and the derivative prefactors of Pauli rotations are imaginary, a parameter update is non-zero only if Pauli propagation generates an I/Z Pauli string with a non-zero imaginary coefficient when evaluated on the reference state $|0\rangle$. This condition identifies which Pauli rotations are active at initialization. The ansatz is constructed from the Hamiltonian decomposition by identifying, for each Pauli term P_a , the restricted set of Pauli-rotation generators that satisfy the first-step non-vanishing condition for the evolution vector V . This defines term-specific trainable subspaces. The resulting circuit has a three-block structure. Hamiltonian-guided trainable rotations come first, followed by a fixed Clifford Bell block and a post-Bell layer of trainable Pauli rotations. This final layer includes both Hamiltonian-guided rotations and additional rotations that build system-ancilla and interacting-site correlations during the evolution.

Numerical results. We benchmark PP-VarQITE on the 1D transverse-field Ising model, the J_1 – J_2 Heisenberg model, and the 2D Ising model, choosing the Hamiltonian parameters such that the exact ground-state energy is analytically known. As shown in Fig. 2, the prepared states reproduce two-point correlations and energy densities across inverse temperatures, with the J_1 – J_2 chain converging toward its exact ground-state energy for system sizes up to 50 qubits. The accuracy is retained at high β , including the 50-qubit J_1 – J_2 regime,




Fig. 2. (a) Finite-temperature benchmarking on the 1D transverse-field Ising model. Results for $n = 50$ sites. The top panel shows the two-point correlation C_{0j} versus distance j for inverse temperatures $\beta = 1, 2, 5$, with PP-VarQITE overlaying the matrix product state (MPS) reference. The bottom panel shows the energy density E/n versus β approaching the exact ground-state energy density E_0/n . **(b) Energy convergence for the J_1 - J_2 Heisenberg chain.** PP-VarQITE energy E versus inverse temperature β for the open J_1 - J_2 Heisenberg chain at the Majumdar-Ghosh point, $J_1 = 1$, $J_2 = J_1/2$, for system sizes $n = 10, 20, 30, 40, 50$. Each curve converges toward the exact ground-state energy $E_0 = -3n/2$.




Fig. 3. Spatial growth of correlations under cooling. Two-point correlation function of the 2D transverse-field Ising model on a 7×7 square lattice ($n = 49$ qubits), prepared by PP-VarQITE. Each node is a lattice site, colored on a logarithmic color scale by the connected correlation $C(r_{\text{ref}}, r) = \langle Z_{r_{\text{ref}}} Z_r \rangle - \langle Z_{r_{\text{ref}}} \rangle \langle Z_r \rangle$ between the reference site r_{ref} (ringed, bottom-right corner) and every site r . Panels show increasing inverse temperature $\beta = 0.1, 0.2, \dots, 0.5$ from left to right. As the system cools, correlations grow and spread outward from the reference site across the lattice.

whereas direct operator-basis propagation has been reported to lose accuracy beyond 30 qubits under the same coefficient truncation [10]. Figure 3 shows the spatial growth of correlations in the two-dimensional Ising model as the system cools, consistent with the expected spread of thermal correlations.

Conclusion. Across these benchmarks, PP-VarQITE prepares accurate thermal and ground states at a circuit cost within reach of near-term hardware. Matrix product states (MPS) also represent these states, but the bond dimension grows with the correlation length, and the compiled circuits become too deep for near-term hardware. For the 50-qubit transverse-field Ising model, PP-VarQITE reaches comparable accuracy at 940 CX gates and depth 88, against 1188 CX gates and depth 736 for the MPS-compiled circuit. At large system sizes and high β , the cost of any classical method grows and full convergence moves out of reach. PP-VarQITE can still compile a shallow circuit for an intermediate β , preparing a state with nontrivial overlap with the target state, exactly the input many quantum algorithms require [11].

References

1. C. Rouzé, D. S. França, and A. M. Alhambra, “Efficient thermalization and universal quantum computing with quantum Gibbs samplers,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.12691>
2. C.-F. Chen, M. J. Kastoryano, and A. Gilyén, “An efficient and exact noncommutative quantum Gibbs sampler,” 2023. [Online]. Available: <https://arxiv.org/abs/2311.09207>
3. S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, “Variational ansatz-based quantum simulation of imaginary time evolution,” *npj Quantum Information*, vol. 5, no. 1, p. 75, Sep. 2019.
4. X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, “Theory of variational quantum simulation,” *Quantum*, vol. 3, p. 191, 2019.
5. C. Zoufal, A. Lucchi, and S. Woerner, “Variational quantum Boltzmann machines,” *Quantum Machine Intelligence*, vol. 3, no. 1, p. 7, Feb. 2021.
6. A. Warren, L. Zhu, N. J. Mayhall, E. Barnes, and S. E. Economou, “Adaptive variational algorithms for quantum Gibbs state preparation,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.12757>
7. A. N. Chowdhury, G. H. Low, and N. Wiebe, “A variational quantum algorithm for preparing quantum Gibbs states,” 2020. [Online]. Available: <https://arxiv.org/abs/2002.00055>
8. J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki, and A. Mezzacapo, “Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer,” *Science Advances*, vol. 11, no. 25, 2025. [Online]. Available: <http://dx.doi.org/10.1126/sciadv.adu9991>
9. M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, “Pauli propagation: A computational framework for simulating quantum systems,” 2026. [Online]. Available: <https://arxiv.org/abs/2505.21606>
10. M. S. Rudolph, A. Angrisani, A. Wright, I. Sanderski, R. Puig, and Z. Holmes, “Thermal state simulation with Pauli and Majorana propagation,” 2026. [Online]. Available: <https://arxiv.org/abs/2602.04878>
11. M. Drudis, A. Baiardi, M. Chiurco, F. Tacchino, S. Woerner, and C. Zoufal, “Bowtie VarQTE: A resource-efficient quantum state preparation primitive,” 2026. [Online]. Available: <https://arxiv.org/abs/2605.20331>