nature portfolio

Peer Review File

Simple and efficient end-to-end quantum thermal and ground state preparation

Corresponding Author: Professor Zhiyan Ding

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors design and prove the convergence of an algorithm for preparing ground states and thermal states using an ancilla bath that interacts with the system of interest via a sequence of Hamiltonian evolutions designed to bring the system to thermal equilibrium. While a number of past works explore the ancilla-driven setting, the main advancement of the current work compared to previous similar results is that the authors managed to show it is possible to sample from a thermal distribution via an algorithm that simultaneously (i) requires only local operations, i.e. suitable for near-term devices, (ii) provides rigorous guarantees with resources that scale polynomially in the number of qubits and mixing time, and (iii) is universal and Hamiltonian-agnostic.

Overall, I find this result valid and sound, and it has significant consequences for both physics and quantum computing. This algorithm opens the possibility of studying equilibrium physics using rapidly developing quantum devices. I believe this work would be of broad interest to physicists beyond the field of quantum simulation, including those in thermodynamics, condensed matter physics, and other fields that involve quantum physics. Therefore, I think this work is a good fit for Nature Physics. At the same time I believe the text needs some work to make it publishable as I explain below.

Major comments:

1. The rigorous mixing time results (Theorems 2 and 3) cover free fermions and commuting Hamiltonians—both essentially non-interacting or classically simulable in some sense. How much of the technical machinery is specific to these cases? For example, I would like to note that very recent work (https://arxiv.org/abs/2601.16154) has extended the proof of polynomial-time efficiency of very similarly sounding algorithm to generic weakly interacting/high-temperature states (it probably appeared after the submission). Can the authors comment on this result and connect the two?

2. One of the critical points of the current work is hard to understand even after reading the Supplementary Material. In particular, the algorithm's efficiency ultimately rests on the mixing time being independent of σ. While the authors claim that they "need to carefully design the initial bath state and the interaction function," I find the explanation lacking essential details to make it clear for a broad readership. Is there intuition for which choices made this claim works and when it should fail? I believe the authors should carefully and in detail explain what innovations and elements of the algorithm make their result possible and how it stands out compared to previous work.

3. I believe the authors need explain some other important decisions regarding the algorithm's construction as well. For example, it is unclear how they choose the distribution of frequencies $g(\omega)$. The figure specifies it to be uniform on $[0, \omega_{\max}]$, but what is $\omega_{\max}$, and how should it be chosen based on the Hamiltonian -- for example, without the ability to compute the Hamiltonian's spectrum? A more unified discussion of how to choose g in practice would strengthen the paper. The same applies to the set of operators A: what are the minimal conditions on this set? I assume, for example, that A commuting with the Hamiltonian would not work. In other words, could the authors provide clear instructions for a quantum computation practitioner or physicist on how to apply this method to the Hamiltonian of their choice?

4. The paper mentions that the protocol provides a "benchmark for dissipative engineering," and the table in the Supplementary Information claims that it is suitable for early fault-tolerant devices. I think it would be valuable to give a better argument, e.g. discuss what system sizes, error rates, and which Hamiltonians would constitute a meaningful first application.

Minor comments:

1. I think the "Algorithm description" section needs some work to make it readable to the readers outside quantum computing field. For example, I am not sure that the Pauli matrix notation (X, Y, Z) is standard for a broad audience (e.g., below Eq. (1)). Also, the authors should be more careful with algorithm description, e.f. state that they start at time $-T$ etc.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)
Summary

This paper proposes a hardware-friendly algorithm to prepare quantum thermal and ground states via randomized system-bath interactions. The protocol requires only a single reusable ancilla qubit and relies entirely on the forward time evolution of the system Hamiltonian. In each step, the algorithm samples a coupling operator A_S and a bath frequency ω, evolves the joint system under a weak system-bath interaction modulated by a Gaussian time envelope, and subsequently traces out and resets the environment qubit. The authors rigorously prove that the target thermal or ground state is an approximate fixed point of this quantum channel and that the total Hamiltonian evolution time scales polynomially with the mixing time and other relevant parameters. Combined with their proof of polynomial mixing times for non-interacting spin and fermionic systems, the overall complexity is rigorously shown to be polynomial.

Strengths

1. Conceptual novelty and efficient implementation: To analyze the fixed point and mixing time, the authors show that this discrete channel can be effectively approximated by a continuous-time Lindbladian evolution. The resulting jump operators take the form of an operator Fourier transform of the Heisenberg-evolved coupling operators. While this structure shares similarities with recent Lindbladian-based state preparation algorithms [1] (which are inspired by the Davies generator), the conceptual contribution of this paper lies in how this Fourier transform is realized. In this work, the integral form of the jump operators naturally emerges from the Dyson series expansion of the weak system-bath interaction combined with the partial trace of a single, properly engineered environment qubit. Consequently, the algorithm circumvents the explicit implementation of the operator Fourier transform used to construct block-encodings for jump operators in previous works. This allows them to avoid any backward time evolution. This approach not only simplifies the quantum circuit implementation for early fault-tolerant devices, but it also mimics the physical thermalization process in nature, which does not perform explicit backward time evolution.

2. Rigorous performance guarantee: Because the effective Lindbladian is composed of a dissipative part that approximately satisfies the quantum detailed balance condition and a Lamb-shift part that approximately commutes with the system Hamiltonian, the authors can leverage existing mathematical tools to bound the mixing time. By mapping the channel's convergence to the spectral gaps of the underlying Davies-like generators, they rigorously establish polynomial mixing times for several physically relevant models (single-qubit toy models, non-interacting spin/fermionic systems, and local commuting spin systems).

Overall assessment

Overall, the authors propose a general framework to prepare ground and thermal states. The implementation is simple yet accompanied by rigorous end-to-end performance guarantees. Given the importance of state preparation and the scarcity of algorithms that achieve both hardware-friendly implementation and solid theoretical bounds, I consider this a strong contribution that will be of broad interest to the physics community, particularly for advancing the quantum simulation of many-body systems and condensed matter physics.

Comments to the Authors

1. Connection to the standard Lindblad derivation: The derivation of the effective Lindbladian in Appendix C shares structural similarities with the standard first-principle derivation of the Lindblad equation (which typically relies on the Born, Markov, and rotating-wave/secular approximations). However, the approximations used here seem to serve a different mathematical purpose (e.g., the Lindbladian naturally emerges from Dyson series truncation, while the secular approximation is used later purely as a proof artifact to establish detailed balance, rather than to guarantee CPTP). It would be beneficial for readers with a physics background if the authors could add a brief discussion explicitly comparing their algorithmic derivation with the first-principle derivation of the Lindblad equation.

2. Advantage of forward time evolution: The authors emphasize that their algorithm does not require backward time evolution. I agree this is conceptually novel. However, for most purely digital simulation algorithms based on Trotterization or QSVT, implementing the forward (e^{-iHt}) or backward (e^{iHt}) Hamiltonian evolution incurs the exact same gate complexity. The practical advantage of avoiding backward time evolution likely manifests in analog architectures, where Hamiltonian simulation might be implemented natively and reversing the sign of H is physically prohibitive. I suggest the authors elaborate slightly on this specific hardware context to strengthen the practical motivation for this feature.

Minor Comments

- In Eq. (9), $A_S(t)$ is introduced without an explicit definition. While it is presumably the Heisenberg-picture evolution of the coupling operator, it should be clearly stated.
- There is a missing space before the phrase "For instance" in the paragraph immediately following Table I.
- In Theorem 7, the parameter $\gamma(\omega)$ is defined but does not appear to be used within the theorem statement.
- In Theorem 7, the operator $V_{A_S, f, T}(\omega)$ is introduced without explanation. Although it is defined later in the proof, it would improve readability to define it within the theorem statement.

[1] Chen, C. F., Kastoryano, M., Brandão, F. G., & Gilyén, A. (2025). Efficient quantum thermal simulation. Nature, 646(8085), 561-566.

(Remarks on code availability)

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
I think the authors did a thorough job addressing my questions and concerns. They have incorporated recent progress in mixing time bounds, provided an updated explanation of their choice of parameters, and fixed minor issues in the manuscript. To reiterate, I believe this work has strong potential to be useful for Gibbs state simulation across multiple applications in physics.

Reviewer #2

(Remarks to the Author)
The authors have addressed my concerns in this revision.

The added discussion in Appendix C regarding the Lindblad derivation makes sense. Clarifying that the effective Lindbladian emerges from a finite-time evolution and a weak system-bath interaction—without relying on the traditional Markov and secular approximations—resolves my previous question. Furthermore, the added context clarifying that the practical advantage of forward-only time evolution specifically applies to analog simulators addresses my second point. Finally, I think the newly added intuition and practical guidelines for the parameter settings is a good addition. It is beneficial and will make the algorithm easier to understand and apply for a broader physics audience.

Overall, these revisions have improved the clarity of the manuscript. I have no further questions.

Open Access This Peer Review File 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 license, and indicate if changes were made.
In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.
The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit https://creativecommons.org/licenses/by/4.0/

****

The comments of the referees are given in *blue italics*. Our responses are given in black Roman with some sentences copied from the manuscript given in *black italics* or in shadow boxes.

****

Referee A:

We thank the referee for the invaluable comments and suggestions. Please find below the detailed responses.

Referee's comment

Overall, I find this result valid and sound, and it has significant consequences for both physics and quantum computing. This algorithm opens the possibility of studying equilibrium physics using rapidly developing quantum devices. I believe this work would be of broad interest to physicists beyond the field of quantum simulation, including those in thermodynamics, condensed matter physics, and other fields that involve quantum physics. Therefore, I think this work is a good fit for Nature Physics. At the same time I believe the text needs some work to make it publishable as I explain below.

Our Reply

We thank the referee for the positive comments, especially regarding our paper's scientific impact.

Referee's comment

The rigorous mixing time results (Theorems 2 and 3) cover free fermions and commuting Hamiltonians—both essentially non-interacting or classically simulable in some sense. How much of the technical machinery is specific to these cases? For example, I would like to note that very recent work (https://arxiv.org/abs/2601.16154) has extended the proof of polynomial-time efficiency of very similarly sounding algorithm to generic weakly interacting/high-temperature states (it probably appeared after the submission). Can the authors comment on this result and connect the two?

Our Reply

We thank the referee for the comment. Although our mixing-time result is established mainly for free-fermionic and commuting Hamiltonians, proving fast mixing even for these classes is already nontrivial and requires both a careful design of the system–bath interaction and new technical tools (see Appendix G for example). For these systems, our results rigorously justify that the algorithm achieves fast mixing and arbitrarily high accuracy for a broad class of physically relevant Hamiltonians.

We thank the reviewer for pointing out the very recent work (arXiv:2601.16154). We would like to highlight that a main goal of this paper is to provide a refined theoretical analysis of the algorithm proposed in our work. In particular, the authors substantially extend our mixing-time analysis for thermal-state preparation to a broader class of Hamiltonians, including local spin Hamiltonians above a certain temperature threshold, local weakly interacting fermionic Hamiltonians with sufficiently small interaction strength at constant inverse temperature β, and one-dimensional nearest-neighbor Hamiltonians at constant β. They also establish a novel monotonicity result for the spectral gap of the approximate Lindbladian dynamics. We believe this is a significant development in the mixing-time analysis

of the system–bath interaction framework, and these ideas have also been used in another recent paper [arXiv:2512.03457v2]. These follow up works further demonstrate the immediate impact of our work in the community. We have added citations and a brief discussion at the end of the related work section.

Referee's comment

One of the critical points of the current work is hard to understand even after reading the Supplementary Material. In particular, the algorithm's efficiency ultimately rests on the mixing time being independent of σ. While the authors claim that they "need to carefully design the initial bath state and the interaction function," I find the explanation lacking essential details to make it clear for a broad readership. Is there intuition for which choices made this claim works and when it should fail? I believe the authors should carefully and in detail explain what innovations and elements of the algorithm make their result possible and how it stands out compared to previous work.

Our Reply

We thank the reviewer for the comment. We have added a new subsection (Choice of interaction terms) after Theorem 4 in the Methods section to give a more detailed explanation of the intuition behind the algorithm and the role of its key design choices. A detailed comparison with previous work is also included in the itemized summary in the Related Work section. In addition, Appendix G, together with Remark 22, provides a simple single-ancilla example illustrating why these design choices are important for achieving fast mixing.

Here, we also briefly summarize the key intuition:

- H_E and ρ_E define a thermal single-ancilla bath. The transition energy induced by the system–bath interaction is determined by the parameter ω in H_E, while ρ_E is chosen as a thermal state at inverse temperature β to ensure the correct balance between upward and downward transitions, thereby enforcing an approximate detailed balance condition.

- The filter $\hat{f}$ determines the allowed transition window. To keep the Lamb-shift term from significantly perturbing the fixed point, one needs a large σ, which makes the fixed point closer to the target state but also concentrates the induced transitions more sharply around ω. To prevent the mixing time from deteriorating as σ increases, the induced transitions must still cover a sufficiently broad range, which is ensured by the choice of $g(\omega)$.

Referee's comment

I believe the authors need explain some other important decisions regarding the algorithm's construction as well. For example, it is unclear how they choose the distribution of frequencies $g(\omega)$. The figure specifies it to be uniform on $[0, \omega_{\max}]$, but what is $\omega_{\max}$, and how should it be chosen based on the Hamiltonian – for example, without the ability to compute the Hamiltonian's spectrum? A more unified discussion of how to choose g in practice would strengthen the paper. The same applies to the set of operators A: what are the minimal conditions on this set? I assume, for example, that A commuting with the Hamiltonian would not work. In other words, could the authors provide clear instructions for a quantum computation practitioner or physicist on how to apply this method to the Hamiltonian of their choice?

Our Reply

We thank the reviewer for the comment. We have added the discussion of these choices in the new subsection (Choice of interaction terms) after Theorem 4 in the Methods section.

The choice of $g(\omega)$ is flexible, as long as it provides a sufficiently broad energy-transition window. For example, one may take $g(\omega)$ to be uniform on $[0, \omega_{\max}]$ with $\omega_{\max} = \Theta(N)$ for an N-qubit local Hamiltonian, so that the induced transitions can cover nearly the full energy spectrum of the system. This choice does not require diagonalizing H. If $H = \sum_j h_j$ is a local Hamiltonian with $\|h_j\| = \mathcal{O}(1)$, one may simply take $\omega_{\max} = 2\sum_j \|h_j\| = \mathcal{O}(N)$, which upper bounds all Bohr frequencies.

When A_S induces only local transitions, one may instead choose $g(\omega)$ to be uniform on a smaller window with $\omega_{\max} = \mathcal{O}(1)$, which is expected to be sufficient in practice to cover the local energy changes generated by A_S.

The operator A_S in our algorithm plays the same role as a coupling operator in the Lindbladian literature. Its choice determines which transitions between eigenstates of the system Hamiltonian are enabled. At a minimum, one should not choose all coupling operators to commute with H, since then the interaction cannot induce transitions between distinct energy eigenspaces. More generally, for the corresponding effective Lindbladian with full-rank fixed point, a standard sufficient condition for uniqueness is that the commutant of the coherent term together with the jump operators is trivial. When users can provide additional prior knowledge of the system, A_S can be designed to favor transitions between the most relevant states. The Lindbladian literature further provides generic choices of A_S that satisfy this condition and induce efficient local transitions for broad classes of systems, such as $\mathcal{A} = \{\pm X_i, \pm Y_i, \pm Z_i\}$ for local spin systems and $\mathcal{A} = \{\pm c_i, \pm c_i^\dagger\}$ for fermionic systems. Many related works, including arXiv:1403.5224, arXiv:1505.07811, arXiv:2501.00443, arXiv:2601.16154 and our paper use these choices.

Referee's comment

The paper mentions that the protocol provides a "benchmark for dissipative engineering," and the table in the Supplementary Information claims that it is suitable for early fault-tolerant devices. I think it would be valuable to give a better argument, e.g. discuss what system sizes, error rates, and which Hamiltonians would constitute a meaningful first application.

Our Reply

We thank the referee for this helpful suggestion. We agree that the relevance to early fault-tolerant (EFT) devices deserves further clarification. Our intent was not to provide a hardware-specific resource estimate—e.g., precise logical qubit counts or code distances—which would be highly sensitive to architecture-specific compilation and is beyond the scope of this work.

Instead, our claim is based on the algorithmic efficiency and structural simplicity of the protocol. Specifically, the scheme requires only forward evolution under a system–bath Hamiltonian and a single reusable ancilla qubit. By avoiding costly requirements such as large ancilla registers, time-reversal, or the high T-gate overhead associated with controlled-long-time evolutions, the protocol is uniquely suited for regimes where logical resources are at a premium.

We have revised the manuscript to emphasize these points (in Applications of the algorithm.b) and to clarify that the most natural initial applications involve Hamiltonian families with native local structure—such as the free-fermion and commuting models in Section V—where the protocol provides a rigorous and resource-efficient path to ground and thermal state preparation.

Referee's comment

I think the "Algorithm description" section needs some work to make it readable to the readers outside quantum computing field. For example, I am not sure that the Pauli matrix notation (X, Y, Z) is standard for a broad audience (e.g., below Eq. (1)). Also, the authors should be more careful with algorithm description, e.f. state that they start at time $-T$ etc.

Our Reply

We thank the reviewer for the helpful comment. We have modified the manuscript accordingly (in Algorithm description section).

****

The referees' comments are given in *blue italics*. Our replies are given in black roman with some sentences copied from the manuscript given in *black italics* or in shadow boxes.

****

Referee B:

Referee's comment

Overall, the authors propose a general framework to prepare ground and thermal states. The implementation is simple yet accompanied by rigorous end-to-end performance guarantees. Given the importance of state preparation and the scarcity of algorithms that achieve both hardware-friendly implementation and solid theoretical bounds, I consider this a strong contribution that will be of broad interest to the physics community, particularly for advancing the quantum simulation of many-body systems and condensed matter physics.

Our Reply

We thank the reviewer for the generous praise of our work, particularly with respect to its technical innovations and scientific impact.

Referee's comment

Connection to the standard Lindblad derivation: The derivation of the effective Lindbladian in Appendix C shares structural similarities with the standard first-principle derivation of the Lindblad equation (which typically relies on the Born, Markov, and rotating-wave/secular approximations). However, the approximations used here seem to serve a different mathematical purpose (e.g., the Lindbladian naturally emerges from Dyson series truncation, while the secular approximation is used later purely as a proof artifact to establish detailed balance, rather than to guarantee CPTP). It would be beneficial for readers with a physics background if the authors could add a brief discussion explicitly comparing their algorithmic derivation with the first-principle derivation of the Lindblad equation.

Our Reply

We thank the reviewer for the suggestions. We have added the explanation in the Appendix C of our manuscript. Here, we briefly summarize the key points of the comparison:

The key difference lies in the timescale being analyzed. In our analysis, each application of the channel evolves the Hamiltonian only over the finite interval $[-T, T]$, while the dissipative correction appears perturbatively at order α^2. This makes it possible to expand the reduced channel as

$$\Phi(\rho) = \rho + \alpha^2 \mathcal{L}(\rho) + \mathcal{O}(\alpha^4)$$

where the $\mathcal{O}(\alpha^4)$ term is the higher-order remainder of the weak-coupling expansion. This argument does not rely on the decay rate of bath correlation functions. By contrast, the first-principles derivation of the Lindblad equation seeks an effective master equation valid on long system–bath timescales, and the Markov approximation is justified by rapid decay of the bath correlations.

Second, in our setting the generator $\mathcal{L}$ already has Lindblad form before any rotating-wave or secular approximation is invoked. Those approximations enter later only as analytical tools for comparing our generator with detailed-balance generators such as the Davies generator, rather than as steps needed to make the dynamics CPTP.

Referee's comment

Advantage of forward time evolution: The authors emphasize that their algorithm does not require backward time evolution. I agree this is conceptually novel. However, for most purely digital simulation algorithms based on Trotterization or QSVT, implementing the forward (e^{-iHt}) or backward (e^{iHt}) Hamiltonian evolution incurs the exact same gate complexity. The practical advantage of avoiding backward time evolution likely manifests in analog architectures, where Hamiltonian simulation might be implemented natively and reversing the sign of H is physically prohibitive. I suggest the authors elaborate slightly on this specific hardware context to strengthen the practical motivation for this feature.

Our Reply

We thank the reviewer for the helpful comments. The reviewer is correct that the complexity of implementing e^{-iHt} and e^{iHt} are the same for digital quantum simulators, and that any implementation advantage arises only in analog simulators. We have added further explanation in the manuscript in Methods/Algorithm Implementation section.

Referee's comment

1. In Eq. (9), $A_S(t)$ is introduced without an explicit definition. While it is presumably the Heisenberg-picture evolution of the coupling operator, it should be clearly stated.
2. There is a missing space before the phrase "For instance" in the paragraph immediately following Table I.
3. In Theorem 7, the parameter $\gamma(\omega)$ is defined but does not appear to be used within the theorem statement.
4. In Theorem 7, the operator $V_{A_S,f,T}(\omega)$ is introduced without explanation. Although it is defined later in the proof, it would improve readability to define it within the theorem statement.

Our Reply

We thank the reviewer for the comments. We have corrected all these mistakes in the revised manuscript.