Quantum Gibbs Sampling Through the Detectability Lemma

Di Fang,[1,2] Jianfeng Lu,[1,3,4] Yu Tong,[1,2,5] and Chu Zhao[5]
[1]Department of Mathematics, Duke University, Durham, North Carolina 27708, USA
[2]Duke Quantum Center, Duke University, Durham, North Carolina 27701, USA
[3]Department of Physics, Duke University, Durham, North Carolina 27708, USA
[4]Department of Chemistry, Duke University, Durham, North Carolina 27708, USA
[5]Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina 27708, USA

(Received 1 May 2026; revised 25 August 2026; accepted 7 October 2026; published 5 November 2026)

Gibbs state preparation is an important subroutine in quantum computing. In this work we use the detectability lemma to improve Gibbs state preparation. Specifically, we design new Gibbs state preparation methods that do not rely on simulating Lindbladian evolution, thus avoiding the overhead from it. For local Lindbladians consisting of M terms, this approach reduces the cost by a factor of $\mathcal{O}(M)$. We also combine the detectability lemma operator and quantum singular value transformation to implement ground state projection operators of frustration-free Hamiltonians, resulting in a quadratic speedup in the spectral gap dependence. Applying this method to Lindbladians for the Gibbs state of local commuting Hamiltonians, we achieve quadratically better dependence on the Lindbladian spectral gap.

DOI: 10.1103/gfv4-qc3s

I. INTRODUCTION

Efficiently sampling from the Boltzmann distribution is a cornerstone problem in classical statistical physics and machine learning. Its quantum counterpart, the preparation of the quantum Gibbs state, is also central to the study of quantum phase transitions and the process of thermalization. Moreover, it serves as a subroutine for various quantum algorithms, ranging from quantum semi-definite programming [1,2] to simulating condensed matter systems at finite temperature [3–5].

In recent years, significant advances have been made in engineering Lindbladian dynamics to prepare Gibbs states [4,6–8]. The Lindbladian generators constructed in these works can be efficiently implemented on quantum computers using state-of-the-art Lindbladian solvers, such as those in Refs. [6,9–15], leading to efficient quantum algorithms for Gibbs state preparation if the Lindbladian dynamics converges quickly to the Gibbs state. In this approach, the overall complexity depends crucially on the mixing time of the Lindbladian dynamics. Consequently, much research has focused on establishing rigorous mixing time bounds for various physical systems [8,16–34].

However, constructing a Lindbladian dynamics and then simulating it on a quantum computer may not be the most efficient way to prepare Gibbs states. Simulation requires the quantum state produced by the algorithm to follow the dynamics closely along the entire trajectory of the dynamics. This requirement is overly restrictive for Gibbs state preparation: it is sufficient for the dynamics produced by the quantum algorithm to converge to a state that is close to the Gibbs state on a timescale comparable to the original dynamics. One algorithm following this idea is proposed in Ref. [34], which randomly applies a quantum channel corresponding to a Lindbladian term at each step, thus avoiding the overhead from accurately simulating the Lindbladian. Similar ideas have been used to improve Lindbladian simulation [27]. Discrete-time quantum processes satisfying an exact detailed balance condition without simulating Lindbladian evolution have also been developed [35]. A related quantum Metropolis scheme implements a quantum channel for which the Gibbs state is an approximate fixed point [36]. We see the same consideration in classical Markov chain Monte Carlo methods, where the objective is the design of an efficient transition kernel that ensures rapid convergence to the stationary distribution, rather than reproducing a specific physical stochastic process.

Another way to improve the runtime of the quantum algorithm is through improving the dependence on the spectral gap. While directly simulating the Lindbladian typically incurs a cost that scales linearly in inverse spectral gap, it is sometimes possible to design algorithms that run in time that scales as the *square root* of the inverse spectral gap, thus achieving a quadratic speedup. Examples include the Lindbladian constructed in Ref. [4], and certain Lindbladians corresponding to Hamiltonians whose

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

eigenvalues can either be exactly computed or satisfy a "rounding promise" [10]. These works follow the idea developed for quantizing a classical reversible Markov chain in Ref. [37], which also resulted in a quadratic improvement in the gap dependence. A recent result also extended such speedup to the continuous-time setting for classical Langevin dynamics [38].

Detectability lemma (DL) is a powerful technique to study Hamiltonian complexity and to prove entanglement area laws [39–41]. In the detectability lemma, the DL operator denotes the layered product of local ground-space projectors, which acts as an approximate projector onto the global ground space. In this work, we use the DL operator as an algorithmic tool to improve Gibbs state preparation along both of these directions. More specifically, we will construct efficiently implementable quantum channels that bypass Lindbladian simulation, thereby achieving a speedup by a factor of M, the number of local Lindbladian terms. These channels update the quantum state locally and are similar in spirit to classical Gibbs-sampling algorithms [42,43] and coordinate hit-and-run methods [44], where subsets of variables are updated conditional on the rest. We will also consider the parent Hamiltonian of the Lindbladian and use the DL to construct an efficiently implementable projection operator for its ground state. This allows us to achieve quadratically improved dependence on the spectral gap when an annealing path is available.

A. Main results

Our first main result is an algorithm to prepare the stationary state of a Lindbladian without simulating its time evolution. When the Lindbladian consists of M local terms, this approach achieves a linear dependence on M in its runtime. By comparison, state-of-the-art Lindbladian simulation algorithms, such as [13], need to first normalize the Lindbladian, thus resulting in an M factor appearing in the rescaled evolution time, making the simulation cost depending quadratically on M. Our method therefore achieves a factor of M speedup compared to the simulation-based method. Below we provide a precise statement of the result:

Theorem 1. (Corollary 3). Let $\mathcal{L}=\sum_{m=1}^{M}\mathcal{L}_m$, where each $\mathcal{L}_m$ is a Lindbladian that satisfies σ-KMS-detailed balance condition, and we assume $\mathcal{L}$ has a non-degenerate 0-eigenspace. We assume that each $\mathcal{L}_m$ commutes with all but at most g others. Then we can prepare σ up to ϵ error in trace distance using $\tilde{\mathcal{O}}(\frac{Mg^2}{\mathrm{gap}(\mathcal{L})}\log(\frac{1}{\sigma_{\min}\epsilon})\log^c(\frac{M}{\epsilon}))$ elementary single- and two-qubit gates, where $\sigma_{\min}$ is the minimum eigenvalue of σ, and c is a constant from the Solovay-Kitaev theorem.

In the above, the KMS-detailed balance condition is defined in Definition 2. $\mathcal{L}$ satisfying σ-KMS-detailed balance condition implies that σ is its stationary state. The algorithm consists of a sequence of local quantum channels that are applied repeatedly. The convergence toward the stationary state σ is analyzed using the detectability lemma (Lemma 1).

For the second direction of refinement, the DL operator can be used to obtain a quadratic improvement in the spectral gap dependence. For this, we require the parent Hamiltonian of the Lindbladian to have strict locality. To ensure this is true, we focus on preparing the Gibbs states of commuting local Hamiltonians.

Theorem 2. (Corollary 5). Consider commuting and bounded-degree local Hamiltonian H. For each $\nu\in[0,\beta]$, we construct σ_ν-detailed balanced $\mathcal{L}_\nu$ as stated in Theorem 6 and assume its irreducibility. Then, the corresponding parent Hamiltonian H_ν will be bounded-degree local and frustration free, and thus it satisfies the requirements of Theorem 6. Therefore there exists a quantum algorithm that prepares the purified Gibbs state at inverse temperature β with error δ and success probability at least $1-\delta$ with gate complexity

$$\mathcal{O}\left(\frac{M\beta\|H\|}{\sqrt{\gamma}}\log^2\left(\frac{\beta\|H\|}{\delta}\right)\log^c\left(\frac{M}{\sqrt{\gamma}}\frac{\beta\|H\|}{\delta}\right)\right) \quad (1)$$

using $\mathcal{O}(\log M)+1$ resettable ancilla qubits, where c is the exponent in the Solovay-Kitaev theorem, and $\gamma=\min_j\mathrm{gap}(\mathcal{L}_{\beta_j})$, $\beta_j=j\beta/K$, $K=\Theta(\beta\|H\|)$.

To prove the above theorem, we propose an algorithm for the frustration-free Hamiltonian ground state problem based on the DL operator, which may be of independent interest. By applying quantum singular value transformation (QSVT) [45] to the DL operator, we obtain a ground state projection operator whose cost scales as the inverse square root of the Hamiltonian spectral gap. Compared to previous methods that uses the block encoding of the original Hamiltonian to prepare the ground state [46,47], this approach achieves a quadratic speedup in the spectral gap dependence.

Theorem 3. (Corollary 4). Let $H=\sum_{m=1}^{M}H_m$, where each H_m satisfies $0\preceq H_m\preceq I$, is supported on at most k qubits, and overlaps with at most g other terms $H_{m'}$. We assume that H is frustration free in the sense that $P_H H_m=0$ for all $m=1,2,\ldots,M$, where P_H is its ground space projection operator. Then a $(1,\mathcal{O}(\log(M)),\epsilon)$-block encoding of P_H (see Definition 5) can be obtained using

$$\mathcal{O}(M\gamma^{-1/2}\log(1/\epsilon)\log^c(M\gamma^{-1/2}\log(1/\epsilon)/\epsilon))$$

elementary gates, where c is the exponent in the Solovay-Kitaev theorem.

B. Algorithmic overview

We briefly summarize the two algorithms that we develop in this work. The first algorithm takes as input a local KMS-detailed-balanced Lindbladian $\mathcal{L}=\sum_{m=1}^{M}\mathcal{L}_m$.

For every term, it implements the local channel $P_m^\dagger$, where $P_m = \lim_{t\to\infty} e^{t\mathcal{L}_m}$ and applies one sweep $\Phi^\dagger = P_M^\dagger \cdots P_1^\dagger$. Repeating this sweep makes the quantum system converge to the stationary state. The detectability lemma controls the contraction per sweep in terms of gap($\mathcal{L}$) and the interaction degree g, while the circuit directly implements local channels and never simulates the trajectory $e^{t\mathcal{L}^\dagger}$.

The second algorithm prepares a purified Gibbs state along a temperature schedule $0=\beta_0<\beta_1<\cdots<\beta_K=\beta$, with a properly chosen number of steps K. At each temperature, the KMS-detailed-balanced Lindbladian is mapped to a frustration-free parent Hamiltonian whose ground state is the purified Gibbs state. For commuting bounded-degree local Hamiltonians, this parent Hamiltonian is itself bounded-degree local. We compose block encodings of its local ground-space projectors to obtain a block encoding of the DL operator and then use QSVT to amplify the singular-value separation supplied by the detectability lemma. This produces an approximate global ground-space projector with inverse-square-root gap dependence. A second QSVT step applied to two consecutive approximate projectors implements the transition $|\psi_{\beta_{j-1}}\rangle \mapsto |\psi_{\beta_j}\rangle$. Annealing is then done by composing these transition operators.

C. Organization

The remainder of this paper is organized as follows. In Sec. II, we define the notation and provide the formal definitions for KMS-detailed balance. Section III reviews the detectability lemma and applies it to Lindbladians. In Sec. IV, we describe the implementation of Hamiltonian ground state projection operators via the detectability lemma with a quadratic speedup in the spectral gap dependence. Section V details the construction of the parent Hamiltonian of a KMS-detailed balanced Lindbladian. Section VI presents the state preparation algorithm via annealing with quadratically improved spectral gap dependence. Finally, in Sec. VII, we apply these techniques to commuting bounded-degree local Hamiltonians and discuss the resulting gate complexity.

II. LINDBLADIANS AND DETAILED BALANCE

We consider an n-qubit quantum system, with a Hilbert space $\mathcal{H}$. The operator algebra is denoted by $B(\mathcal{H})$ while the superoperator algebra (consisting of linear operators on the operator algebra) is denoted by $B(B(\mathcal{H}))$. A Lindbladian $\mathcal{L}$ is the generator of a quantum Markov semigroup $\{e^{t\mathcal{L}}\}_{t\geq 0}$. It takes a specific form

$$\mathcal{L}[X] = \sum_{j=1}^{r} L_j^\dagger X L_j - \frac{1}{2}\{L_j^\dagger L_j, X\},$$

where $\{\cdot\}$ denotes the anti-commutator, and each $L_j \in B(\mathcal{H})$ is called a *jump operator*. We consider a local Lindbladian

$$\mathcal{L} = \sum_{m=1}^{M} \mathcal{L}_m, \tag{2}$$

where each $\mathcal{L}_m$ is a Lindbladian that acts non-trivially only on k sites and overlaps with at most g other terms $\mathcal{L}_{m'}$. These Lindbladian terms do not have to commute with each other. $\|\mathcal{L}_m\|_\diamond \leq 1$ where $\|\cdot\|_\diamond$ denotes the diamond norm. We are working in the Heisenberg picture and therefore

$$\mathcal{L}_1(I) = \cdots = \mathcal{L}_M(I) = \mathcal{L}(I) = 0. \tag{3}$$

We will define the σ-KMS inner product for a quantum state $\sigma \in B(\mathcal{H})$:

Definition 1. (σ-KMS inner product). For any $X, Y \in B(\mathcal{H})$, we denote

$$\langle X, Y\rangle_\sigma = \mathrm{Tr}[X^\dagger \sqrt{\sigma} Y \sqrt{\sigma}].$$

The corresponding norm we denote by

$$\|X\|_\sigma = \langle X, X\rangle_\sigma^{1/2}.$$

With this we can generalize the detailed balance condition (DBC) to the quantum setting

Definition 2. (σ-KMS-detailed balance). A superoperator $\mathcal{A} \in B(B(\mathcal{H}))$ satisfies σ-KMS-detailed balance if it is self-adjoint under the σ-KMS inner product, i.e.,

$$\langle X, \mathcal{A}(Y)\rangle_\sigma = \langle \mathcal{A}(X), Y\rangle_\sigma,$$

for all $X, Y \in B(\mathcal{H})$.

From now on we assume that each $\mathcal{L}_m$ satisfies σ-KMS DBC, which then implies

$$\mathrm{Tr}[X\mathcal{L}_m^\dagger(\sigma)] = \mathrm{Tr}[\mathcal{L}_m(X)\sigma] = \langle \mathcal{L}_m(X), I\rangle_\sigma = \langle X, \mathcal{L}_m(I)\rangle_\sigma = 0,$$

for all $X \in B(\mathcal{H})$ [48]. This then implies $\mathcal{L}_m^\dagger(\sigma) = 0$ for all m. Therefore σ is a stationary state for the dynamical semigroup generated by $\mathcal{L}^\dagger$.

Because $\mathcal{L}$ is self-adjoint under the σ-KMS inner product, its spectrum must be real and located on the non-positive part of the real axis. 0 must be an eigenvalue, and here we assume it is non-degenerate. In particular, $\ker \mathcal{L} = \mathrm{span} I$. We denote the gap between 0 and the second largest eigenvalue of $\mathcal{L}$ by gap($\mathcal{L}$). More generally

Definition 3. (Spectral gap). For $\mathcal{A} \in B(B(\mathcal{H}))$ with real spectrum, we denote by gap($\mathcal{A}$) the gap between the two largest eigenvalues.

The state σ can then be prepared by evolving with $\mathcal{L}^\dagger$ for sufficiently long time. However, this strategy does not

necessarily translate to the most efficient quantum algorithm. Here we will consider another strategy. We approximately implement a quantum channel $\Phi^\dagger$, where Φ is defined through

$$\Phi = P_1 P_2 \cdots P_M, \quad P_m = \lim_{t\to\infty} e^{t\mathcal{L}_m}, \quad \forall\ 1 \le m \le M. \tag{4}$$

We will use the detectability lemma to show that $\Phi^\dagger$ works almost as well as $e^{t\mathcal{L}^\dagger}$, as a sufficient number of rounds of $\Phi^\dagger$ brings any initial state close to the stationary state σ. Note that in the above we are using an arbitrary ordering that does not depend on the underlying geometry of the quantum system.

III. THE DETECTABILITY LEMMA FOR LINDBLADIANS

The detectability lemma was originally formulated for the Hamiltonian situation. We will restate it here.

Lemma 1. (The detectability lemma [49]). Let $Q_1, Q_2, \ldots, Q_M$ be a set of projection operators and $H = \sum_{m=1}^M Q_m$. We assume that each Q_m commutes with all but at most g others. Given a state $|\psi\rangle$, define $|\phi\rangle = \prod_{m=1}^M (I - Q_m)|\psi\rangle$, where the product is in any order. Let $\epsilon_\phi = \langle\phi|H|\phi\rangle / \langle\phi|\phi\rangle$, then

$$\langle\phi|\phi\rangle \le \frac{1}{\epsilon_\phi / g^2 + 1}.$$

Note that the above lemma does not put any constraint on the locality on the operators Q_m, but only on their commutation relations. This lemma provides us a means to quantify how much the detectability lemma operator $\prod_{m=1}^M (I - Q_m)$ shrinks any quantum state that is orthogonal to the ground space of H, as stated in the following corollary

Corollary 1. (Corollary 3 of [49]). Under the same assumptions as Lemma 1, and assuming that the smallest positive eigenvalue of H is at least γ, let $|\psi^\perp\rangle$ be any quantum state that is orthogonal to $\ker(H)$, which can be trivial, then

$$\left\| \prod_{m=1}^M (I - Q_m)|\psi^\perp\rangle \right\|^2 \le \frac{1}{\gamma / g^2 + 1}.$$

We remark that the above corollary is slightly different from [49, Corollary 3], in that we did not assume the Hamiltonian H to be frustration free, i.e., its kernel can be trivial. This is because the proof of [49, Lemma 2] and [49, Corollary 3] in fact did not use the frustration-free property of H.

Coming back to the Lindbladian setting, we can similarly construct a Hamiltonian superoperator

$$H_{\mathcal{L}} = \sum_{m=1}^M (I - P_m), \tag{5}$$

for P_m defined in (4). Here I denotes the identity superoperator. This Hamiltonian superoperator $H_{\mathcal{L}}$ is self-adjoint under the σ-KMS inner product and can therefore be regarded as a Hamiltonian. $H_{\mathcal{L}}$ has a non-degenerate ground space span(I): on the one hand $H_{\mathcal{L}}(I) = 0$, and on the other hand if $H_{\mathcal{L}}(X) = 0$, we then have $P_m(X) = X$, which implies $\mathcal{L}_m(X) = 0$ for all m, thus forcing $X \propto I$.

Since $\mathcal{L}_m$ has all non-positive eigenvalues between 0 and 1, we have $I - P_m \succeq -\mathcal{L}_m$ (in the σ-KMS inner product sense), and consequently $H_{\mathcal{L}} \succeq -\mathcal{L}$. Additionally, because $H_{\mathcal{L}}$ and $-\mathcal{L}$ share the same ground space (span(I)), denoting the gap between 0 and the smallest positive eigenvalue of $H_{\mathcal{L}}$ by γ, we have

$$\gamma \ge \mathrm{gap}(\mathcal{L}). \tag{6}$$

With this we can directly apply Corollary 1 to $H_{\mathcal{L}}$ to obtain

Theorem 4. Let $\mathcal{L} = \sum_{m=1}^M \mathcal{L}_m$, where each $\mathcal{L}_m$ is a Lindbladian that satisfies σ-KMS DBC, and we assume $\mathcal{L}$ has a non-degenerate 0-eigenspace. We assume that each $\mathcal{L}_m$ commutes with all but at most g others. Then let the quantum channel Φ be defined as in (4). For any $X \in B(\mathcal{H})$ such that $\mathrm{Tr}[\sigma X] = 0$, we have

$$\|\Phi(X)\|_\sigma^2 \le \frac{\|X\|_\sigma^2}{\mathrm{gap}(\mathcal{L}) / g^2 + 1}.$$

Moreover $\mathrm{Tr}[\sigma\Phi(X)] = \mathrm{Tr}[\Phi^\dagger(\sigma)X] = \mathrm{Tr}[\sigma X] = 0$.

As a corollary of the above theorem, we can use the quantum channel Φ to prepare the Gibbs state σ from any initial state ρ_0.

Corollary 2. Under the same assumptions as Theorem 4, let $\rho_k = (\Phi^\dagger)^k(\rho_0)$ for any initial state ρ_0, then

$$\|\rho_k - \sigma\|_1 \le \frac{1}{(\mathrm{gap}(\mathcal{L}) / g^2 + 1)^{k/2} \sqrt{\sigma_{\min}}},$$

where $\sigma_{\min}$ is the smallest eigenvalue of σ.

Proof. Let $X \in B(\mathcal{H})$ satisfy $\|X\| \le 1$. Denote $\tilde{X} = X - \mathrm{Tr}[\sigma X] I$, and we have $\|\tilde{X}\|_\sigma \le \|X\|_\sigma \le \|X\| \le 1$. This ensures $\mathrm{Tr}[\sigma\tilde{X}] = 0$. We then compute

$$\begin{aligned} \mathrm{Tr}[X(\rho_k - \sigma)] &= \mathrm{Tr}[\tilde{X}(\rho_k - \sigma)] = \mathrm{Tr}[\tilde{X}\rho_k] \\ &= \mathrm{Tr}[\tilde{X}(\Phi^\dagger)^k(\rho_0)] = \mathrm{Tr}[\Phi^k(\tilde{X})\rho_0]. \end{aligned}$$

By the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \mathrm{Tr}[\Phi^k(\tilde{X})\rho_0]| &= |\langle \Phi^k(\tilde{X}), \sigma^{-1/2}\rho_0\sigma^{-1/2}\rangle_\sigma| \\ &\le \|\Phi^k(\tilde{X})\|_\sigma \|\sigma^{-1/2}\rho_0\sigma^{-1/2}\|_\sigma. \end{aligned}$$

For $\|\Phi^k(\tilde{X})\|_\sigma$ we directly apply Theorem 4 k times and use the fact that $\|\tilde{X}\|_\sigma \leq 1$ and $\mathrm{Tr}(\sigma\tilde{X}) = 0$. For $\|\sigma^{-1/2}\rho_0\sigma^{-1/2}\|_\sigma$ we have

$$\|\sigma^{-1/2}\rho_0\sigma^{-1/2}\|_\sigma = \mathrm{Tr}[\rho_0\sigma^{-1/2}\rho_0\sigma^{-1/2}]^{1/2} \leq \frac{1}{\sqrt{\sigma_{\min}}}.$$

Therefore we have

$$|\mathrm{Tr}[X(\rho_k - \sigma)]| \leq \frac{1}{(\mathrm{gap}(\mathcal{L})/g^2 + 1)^{k/2}\sqrt{\sigma_{\min}}}.$$

Because this is true for any $\|X\| \leq 1$, we have the desired inequality for the trace distance between ρ_k and σ through the duality between the Schatten 1- and ∞-norms. ∎

We will next discuss the gate complexity of preparing Gibbs states using the above approach. We assume that each $\mathcal{L}_m$ acts non-trivially only on $\mathsf{k} = \mathcal{O}(1)$ qubits. To prepare the Gibbs state σ that is the fixed point of $\mathcal{L}$ as discussed above, we repeatedly apply the quantum channel $\Phi^\dagger$. Each application of $\Phi^\dagger$ requires implementing all M channels P_m. Since each P_m acts non-trivially only on k qubits, we can compile a unitary circuit involving at most k qubits to implement it. Solovay-Kitaev Theorem then tells us that $\mathcal{O}(\log^c(1/\epsilon'))$ gates are needed to implement P_m to accuracy ϵ' in the diamond norm.

Corollary 2 tells us that in order to prepare the state to $\mathcal{O}(\epsilon)$ accuracy in trace distance, it suffices to have $k = \mathcal{O}(\frac{g^2}{\mathrm{gap}(\mathcal{L})}\log(1/(\sigma_{\min}\epsilon)))$. The accumulated unitary synthesis error in the whole procedure is at most $\mathcal{O}(kM\epsilon')$, for which we can choose $kM\epsilon' = \mathcal{O}(\epsilon)$ to ensure that it contributes only $\mathcal{O}(\epsilon)$ to the total error. Therefore the gate complexity for the whole procedure is upper bounded by

$$\begin{aligned}&\mathcal{O}\left(\frac{Mg^2}{\mathrm{gap}(\mathcal{L})}\log(1/(\sigma_{\min}\epsilon))\log^c(kM/\epsilon)\right)\\ &\quad = \tilde{\mathcal{O}}\left(\frac{Mg^2}{\mathrm{gap}(\mathcal{L})}\log\left(\frac{1}{\sigma_{\min}\epsilon}\right)\log^c\left(\frac{M}{\epsilon}\right)\right)\end{aligned} \tag{7}$$

We therefore arrive at the following corollary

Corollary 3. Let $\mathcal{L} = \sum_{m=1}^M \mathcal{L}_m$, where each $\mathcal{L}_m$ is a Lindbladian that satisfies σ-KMS DBC, and we assume $\mathcal{L}$ has a non-degenerate 0-eigenspace. We assume that each $\mathcal{L}_m$ commutes with all but at most g others. Then we can prepare σ up to ϵ error in trace distance using $\tilde{\mathcal{O}}(\frac{Mg^2}{\mathrm{gap}(\mathcal{L})}\log(\frac{1}{\sigma_{\min}\epsilon})\log^c(\frac{M}{\epsilon}))$ elementary single- and two-qubit gates, where $\sigma_{\min}$ is the minimum eigenvalue of σ, and c is a constant from the Solovay-Kitaev theorem.

Given the Lindbladian $\mathcal{L}$, the natural way to prepare σ is to simply simulate the time evolution $e^{t\mathcal{L}}$ to bring an arbitrary initial state close to σ. The required evolution time is $\mathcal{O}(\frac{1}{\mathrm{gap}(\mathcal{L})}\log(\frac{1}{\sigma_{\min}\epsilon}))$. Using state-of-the-art Lindbladian simulation algorithms such as [13], where the Lindbladian needs to be rescaled to have diamond norm $\mathcal{O}(1)$, the gate complexity scales as $\mathcal{O}(\frac{M^2}{\mathrm{gap}(\mathcal{L})}\log(\frac{1}{\sigma_{\min}\epsilon}))$ ignoring further poly-logarithmic factors. Therefore the method in Corollary 3 reduces the cost by a factor of M up to poly-logarithmic corrections.

This comparison with direct Lindbladian simulation uses the standard spectral-gap bound on the mixing time. If a stronger functional inequality or a model-specific argument proves a much shorter mixing time—for example, logarithmic mixing in the system size—then direct Lindbladian simulation can have a better guarantee than the gap-based bound above, and our present analysis does not imply an advantage over it. It remains an interesting open question whether such rapid-mixing estimates can be transferred directly to the discrete channel $\Phi^\dagger$ (4) that we implement in our algorithm.

IV. IMPLEMENTING GROUND STATE PROJECTION OPERATORS

In this section we will take a detour to discuss how one can implement the ground state projection operator of frustration-free Hamiltonians. In later sections we will discuss how this implementation leads to a quadratic speedup in terms of the spectral gap dependence in a simulated annealing procedure for quantum Gibbs states corresponding to commuting local Hamiltonians.

Definition 4. A Hamiltonian $H = \sum_{m=1}^M H_m$ acting on n qubits is called a bounded-degree k-local Hamiltonian if: (1) each term H_j acts non-trivially on at most k qubits, where $\mathsf{k} = \mathcal{O}(1)$, and (2) there exists a constant $g = \mathcal{O}(1)$ such that for every qubit $i \in \{1, \ldots, n\}$, at most g terms H_j act non-trivially on i.

Here and throughout this section, k denotes the locality of a Hamiltonian term; the ordinary italic k used in Sec. III denotes the number of DL-channel iterations.

We focus on such a bounded-degree local Hamiltonian on n qubits

$$H = \sum_{m=1}^M H_m, \tag{8}$$

where each H_m satisfies $0 \preceq H_m \preceq I$, is supported on at most k qubits, and overlaps with at most g other terms $H_{m'}$. The terms H_m do not necessarily commute with each other. Moreover, we assume that H is frustration free in the following sense: let P_H be the projection operator into the ground space of H, and then we have

$$P_H H_m = 0, \quad \forall\ m \in \{1, 2, \ldots, M\}. \tag{9}$$

Note that the above also implies $H_m P_H = 0$, and that the ground space of H is also its kernel. We also denote the dimension of the ground space by r.

For the above Hamiltonian, we will implement its ground projection operator using the detectability lemma operator

$$\mathrm{DL}(H) = \prod_{m=1}^{M} P_m. \tag{10}$$

Note that by applying Corollary 1 to $\sum_m (I - P_m)$, we have, for any $|\psi^\perp\rangle$ such that $P_H|\psi^\perp\rangle = 0$,

$$\|\mathrm{DL}(H)|\psi^\perp\rangle\|^2 \le \frac{1}{\gamma/g^2 + 1}, \tag{11}$$

where γ is the spectral gap separating 0 from the rest of the spectrum of H. Note that the above detectability lemma operator $\mathrm{DL}(H)$ can be implemented with gate complexity $\mathcal{O}(M)$ assuming that $\mathsf{k}, g = \mathcal{O}(1)$. Repeatedly applying $\mathrm{DL}(H)\mathcal{O}(\gamma^{-1}\log(1/\epsilon))$ times can take the resulting product ϵ-close to the ground projection operator P_H, thus resulting in a $\mathcal{O}(M\gamma^{-1}\log(1/\epsilon))$ gate complexity for implementing the projection operator. Below we will show how a $\gamma^{-1/2}$ scaling can be achieved using quantum singular value transformation (QSVT) [45].

To discuss QSVT we need to first introduce the notion of block encoding:

Definition 5. (Block encoding). An $(m+n)$-qubit unitary operator U is called an (α, m, ϵ)-block-encoding of an n-qubit operator A, if

$$\|A - \alpha(\langle 0^m| \otimes I)U(|0^m\rangle \otimes I)\| \le \epsilon. \tag{12}$$

An equivalent way to express (12) is

$$U = \begin{pmatrix} \tilde{A}/\alpha & * \\ * & * \end{pmatrix},$$

where $*$ can be any block matrices of the correct sizes and $\|\tilde{A} - A\| \le \epsilon$.

We first note that, because each P_m can be implemented via its (1, 1, 0)-block encoding, these block encodings can be composed to implement a $(1, M, 0)$-block encoding of $\mathrm{DL}(H)$. Using the compression gadget in Refs. [50,51] we can further reduce the number of ancilla qubit to $\mathcal{O}(\log(M))$, thus resulting in a $(1, \mathcal{O}(\log(M)), 0)$-block encoding of $\mathrm{DL}(H)$. Note that each implementation of P_m may involve exponentially small errors due to unitary synthesis, but these errors can be suppressed with polylogarithmic overhead, and their accumulation throughout the algorithm is well-controlled, as will be discussed in Corollary 4.

We then write down the singular value decomposition of $\mathrm{DL}(H)$:

$$\mathrm{DL}(H) = USV^\dagger, \tag{13}$$

where $S = \mathrm{diag}(s_1, s_2, \ldots, s_{2^n})$, $s_1 \ge s_2 \ge \cdots \ge s_{2^n}$. Since the ground space of H is r-dimensional, we have $s_1 = \cdots = s_{\mathsf{r}} = 1$. By (11), we have

$$0 \le s_j \le \frac{1}{\sqrt{\gamma/g^2 + 1}}, \qquad j = \mathsf{r}+1, \ldots, 2^n. \tag{14}$$

We denote by γ_* the lower bound of the gap separating the above s_j from 1:

$$\gamma_* = 1 - \frac{1}{\sqrt{\gamma/g^2 + 1}} = \Omega(\gamma). \tag{15}$$

The projection operator P_H can in fact be expressed in terms of U and V. We write

$$U = (\,U_1 \quad U_2\,), \qquad V = (\,V_1 \quad V_2\,), \tag{16}$$

where U_1 and V_1 are both of size $2^n \times \mathsf{r}$. In other words, the columns of U_1 and V_1 correspond to the singular values $s_1, \ldots, s_{\mathsf{r}}$. Then we claim

Lemma 2. For U_1 and V_1 defined above, we have $P_H = U_1 V_1^\dagger$.

Proof. For any $|\phi\rangle \in \mathrm{col}(V_1)$, we have $\|\mathrm{DL}(H)|\phi\rangle\| = \||\phi\rangle\|$, and therefore $\|P_m|\phi\rangle\| = \||\phi\rangle\|$ for all m. Since each P_m is a projection operator, we then have $P_m|\phi\rangle = |\phi\rangle$. Consequently

$$U_1 V_1^\dagger|\phi\rangle = \mathrm{DL}(H)|\phi\rangle = |\phi\rangle,$$

and $|\phi\rangle$ is in the ground space of H. As a result of the latter $P_H|\phi\rangle = |\phi\rangle$. Therefore P_H and $U_1V_1^\dagger$ agree in $\mathrm{col}(V_1)$.

If $|\phi\rangle \in \mathrm{col}(V_1)^\perp$, then $|\phi\rangle$ must be orthogonal to the ground space since the ground space is contained in $\mathrm{col}(V_1)$. Therefore $P_H|\phi\rangle = U_1V_1^\dagger|\phi\rangle = 0$. Combining this with the discussion in the previous paragraph we have $P_H = U_1V_1^\dagger$. ∎

For a polynomial $p(x)$ of degree d with fixed parity satisfying $|p(x)| \le 1$ for all $x \in [-1, 1]$, QSVT allows us to construct a $(1, \mathcal{O}(\log(M)), 0)$-block encoding of

$$\tilde{P}_H = Up(S)V^\dagger, \tag{17}$$

with d queries to the block encoding of $\mathrm{DL}(H)$. We will appropriately choose $p(x)$ so that $\tilde{P}_H$ approximates P_H. For this purpose, the polynomial in Ref. [52] suffices, but we will provide a simpler construction by simply rescaling the Chebyshev polynomial. We let

$$p(x) = \frac{T_\ell(x/(1-\gamma_*))}{T_\ell(1/(1-\gamma_*))}, \tag{18}$$

where T_ℓ is the ℓth Chebyshev polynomial of the first kind and is of degree ℓ with fixed parity. It satisfies $|p(x)| \le 1$

for $x \in [-1, 1]$, $p(1) = 1$, and when $x \in [-1 + \gamma_*, 1 - \gamma_*]$, we have

$$|p(x)| \leq \frac{1}{T_\ell(1/(1-\gamma_*))} \leq 2e^{-\ell\sqrt{1/(1-\gamma_*)-1}} \leq 2e^{-\ell\sqrt{\gamma_*}},$$

where the second inequality comes from [53, Lemma 13]. From this we have

$$\|p(S) - D\| \leq 2e^{-\ell\sqrt{\gamma_*}},$$

where $D = \mathrm{diag}(1, \ldots, 1, 0, \ldots, 0)$, with r diagonal entries being 1. The above implies

$$\begin{aligned}\|\tilde{P}_H - P_H\| &= \|\tilde{P}_H - U_1 V_1^\dagger\| = \|\tilde{P}_H - UDV^\dagger\| \\ &= \|p(S) - D\| \leq 2e^{-\ell\sqrt{\gamma_*}}. \end{aligned} \tag{19}$$

In order to ensure $\|\tilde{P}_H - P_H\| \leq \epsilon$, we only need

$$\ell = \mathcal{O}(\gamma_*^{-1/2} \log(1/\epsilon)) = \mathcal{O}(\gamma^{-1/2} \log(1/\epsilon)). \tag{20}$$

Therefore we have the following theorem:

Theorem 5. Let $H = \sum_{m=1}^M H_m$, where each H_m satisfies $0 \preceq H_m \preceq I$, is supported on at most k qubits and overlaps with at most g other terms $H_{m'}$. We assume that H is frustration free in the sense that $P_H H_m = 0$ for all $m = 1, 2, \ldots, M$, where P_H is its ground space projection operator. We assume that each ground projection operator of H_m is accessed through its $(1, \mathcal{O}(1), 0)$-block encoding. Then a $(1, \mathcal{O}(\log(M)), \epsilon)$-block encoding of P_H can be obtained using $\mathcal{O}(\gamma^{-1/2} \log(1/\epsilon))$ queries to each block encoding of the ground space projection operator of H_m.

We note that each block encoding of the ground space projection operator of H_m needs to be implemented with gates and has some errors. Suppose each introduces an ϵ' error in terms of the spectral norm, then the cumulative error in the block encoding of P_H is $\mathcal{O}(M\ell\epsilon')$. To ensure that the error in terms of the spectral norm is ϵ' for each small block encoding involving $\mathcal{O}(1)$ qubits, the Solovay-Kitaev theorem tells us that $\mathcal{O}(\log^c(1/\epsilon'))$ gates are needed, where $c = \log_{(1+\sqrt{5})/2}(2) + \delta$ for any $\delta > 0$. Therefore

Corollary 4. Let $H = \sum_{m=1}^M H_m$, where each H_m satisfies $0 \preceq H_m \preceq I$, is supported on at most k qubits and overlaps with at most g other terms $H_{m'}$. We assume that H is frustration free in the sense that $P_H H_m = 0$ for all $m = 1, 2, \ldots, M$, where P_H is its ground space projection operator. Then a $(1, \mathcal{O}(\log(M)), \epsilon)$-block encoding of P_H can be obtained using

$$\mathcal{O}(M\gamma^{-1/2} \log(1/\epsilon) \log^c(M\gamma^{-1/2} \log(1/\epsilon)/\epsilon))$$

elementary gates, where c is the exponent in the Solovay-Kitaev theorem.

V. THE PARENT HAMILTONIAN

The parent Hamiltonian is a coherent representation of the Lindbladian $\mathcal{L}$ which allows for a quadratic speedup in terms of the spectral gap dependence at the expense of doubling the Hilbert space. It is constructed such that its ground state provides access to the Gibbs state. The Hermiticity of $\mathcal{L}$ under the σ-KMS inner product results in the Hermiticity of this coherent representation, which is therefore viewed as a Hamiltonian. This representation enables us to prepare the Gibbs state by preparing the ground state of the parent Hamiltonian using the techniques in Sec. IV and obtain a quadratic speedup.

Define the superoperator $\Gamma_\sigma \colon \mathcal{B}(\mathcal{H}) \to \mathcal{B}(\mathcal{H})$ as

$$\Gamma_\sigma(\cdot) = \sigma^{\frac{1}{2}} \cdot \sigma^{\frac{1}{2}}. \tag{21}$$

Γ_σ recovers the σ-KMS inner product in the sense that

$$\langle X, Y\rangle_{\mathrm{KMS}} = \langle \Gamma_\sigma^{\frac{1}{2}} X, \Gamma_\sigma^{\frac{1}{2}} Y\rangle_{\mathrm{HS}}. \tag{22}$$

Define the superoperator $\mathfrak{H} \colon \mathcal{B}(\mathcal{H}) \to \mathcal{B}(\mathcal{H})$ as

$$\mathfrak{H} = \Gamma_\sigma^{\frac{1}{2}} \mathcal{L} \Gamma_\sigma^{-\frac{1}{2}}. \tag{23}$$

This superoperator is Hermitian under the usual Hilbert-Schmidt inner product, i.e.,

$$\langle \mathfrak{H} X, Y\rangle_{\mathrm{HS}} = \langle X, \mathfrak{H} Y\rangle_{\mathrm{HS}}, \quad \forall\ X, Y \in \mathcal{B}(\mathcal{H}). \tag{24}$$

By vectorizing the operator space $\mathcal{B}(\mathcal{H})$, the superoperator $\mathfrak{H}$ can be explicitly written as a Hamiltonian. Define the linear map $v \colon \mathcal{B}(\mathcal{H}) \to \mathcal{H}^{\otimes 2}$ that acts on the computational basis as follows:

$$v(|i\rangle\langle j|) = |i\rangle \otimes |j\rangle, \qquad i, j \in [d]. \tag{25}$$

By definition and linearity, we have

$$v(X) = v\left(\sum_{i,j} X_{ij} |i\rangle\langle j|\right) = \sum_{i,j} X_{ij} |i\rangle \otimes |j\rangle, \quad X \in \mathcal{B}(\mathcal{H}). \tag{26}$$

In particular, for a rank-1 operator $X = |\psi\rangle\langle\phi|$, we have

$$v(|\psi\rangle\langle\phi|) = |\psi\rangle \otimes |\phi^*\rangle, \tag{27}$$

where $|\phi^*\rangle$ is the entrywise complex conjugate of $|\phi\rangle$ in the computational basis. Indeed, writing $|\psi\rangle = \sum_i \psi_i |i\rangle$ and $|\phi\rangle = \sum_j \phi_j |j\rangle$ gives $X_{ij} = \psi_i \overline{\phi_j}$, so the preceding definition yields $v(X) = \sum_{i,j} \psi_i \overline{\phi_j} |i\rangle \otimes |j\rangle = |\psi\rangle \otimes |\phi^*\rangle$. Upon vectorization, the superoperator $\mathfrak{H}$ is represented as a Hamiltonian H acting on the doubled Hilbert space $\mathcal{H}^{\otimes 2}$,

$$\mathsf{H} = v\, \mathfrak{H}\, v^{-1}. \tag{28}$$

Observe that

$$\mathfrak{H}(\sqrt{\sigma}) = \Gamma_\sigma^{1/2}\mathcal{L}(I) = 0. \tag{29}$$

Therefore $\sqrt{\sigma}$ is the ground [54] state of the superoperator $\mathfrak{H}$ with eigenvalue 0. Or equivalently, the vectorized state $v(\sqrt{\sigma})$ is the ground state of the Hamiltonian H with eigenvalue 0. Moreover, assuming that the eigenvalue decomposition of σ is given by $\sigma = \sum_{i=1}^d \lambda_i |\psi_i\rangle\langle\psi_i|$,

$$v(\sqrt{\sigma}) = \sum_{i=1}^d \sqrt{\lambda_i}\, |\psi_i\rangle \otimes |\psi_i^*\rangle, \tag{30}$$

which is a purification of the Gibbs state σ, i.e., tracing over the second subsystem gives σ on the first subsystem. Since the map from the Lindbladian to the Hamiltonian is linear, given a Lindbladian that is a sum of M local Lindbladians, H is equal to the sum of M local terms, each of which can be classically computed in time that is upper bounded by some constant depending on locality. The problem of preparing the Gibbs state σ now reduces to preparing the ground state of the Hamiltonian H.

VI. STATE PREPARATION BY ANNEALING

To prepare the Gibbs state, we utilize a discretized temperature path similar to quantum annealing and prepare the Gibbs state by projecting the state consecutively to the ground space of the parent Hamiltonian of the next temperature. Denote the purified Gibbs state at inverse temperature ν by $|\psi_\nu\rangle$. Denote by $P_\nu := |\psi_\nu\rangle\langle\psi_\nu|$ the exact projection operator and $\tilde{P}_\nu$ the approximate projection operator. We first prepare the ground state $|\psi_{\beta_0}\rangle$ of the parent Hamiltonian at $\beta_0 = 0$, which is the purified maximally mixed state, or in other words the maximally entangled state on $\mathcal{H}^{\otimes 2}$. In the annealing process, we go through K temperature steps,

$$\beta_0 = 0 \to \beta_1 \to \beta_2 \to \cdots \to \beta_K = \beta, \quad \beta_{j-1} < \beta_j \ \forall\, j \in [K]. \tag{31}$$

Note that for each temperature, the associated parent Hamiltonian H_{β_j} is different, while the original Hamiltonian for the actual physical system as in the Gibbs state does not depend on temperature. For each temperature, using Corollary 4, we construct each approximate projection operator $\tilde{P}_{\beta_j}$. Then, we send the initial state to the final state using K transition operators, each constructed from two consecutive projectors,

$$O_{\beta_0}^{\beta_1} \to O_{\beta_1}^{\beta_2} \to \cdots \to O_{\beta_{K-1}}^{\beta_K}. \tag{32}$$

Denote the ideal transition operator by $O_{\beta_{j-1}}^{\beta_j} := |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|$. Error is introduced in the implementation of each $\tilde{P}_{\beta_j}$ as well as in the implementation of the transition operator. Denote the transition operator with error by $\tilde{O}_{\beta_{j-1}}^{\beta_j}$.

Denote the final state we obtain by $|\psi\rangle$,

$$|\psi\rangle = \frac{\prod_{j=1}^K \tilde{O}_{\beta_{j-1}}^{\beta_j} |\psi_{\beta_0}\rangle}{\| \prod_{j=1}^K \tilde{O}_{\beta_{j-1}}^{\beta_j} |\psi_{\beta_0}\rangle \|}. \tag{33}$$

$|\psi\rangle$ is obtained with success probability

$$p = \left\| \prod_{j=1}^K \tilde{O}_{\beta_{j-1}}^{\beta_j} |\psi_{\beta_0}\rangle \right\|^2. \tag{34}$$

For convenience, denote the unnormalized state by

$$|\tilde{\psi}_{\beta_i}\rangle = \prod_{j=1}^i \tilde{O}_{\beta_{j-1}}^{\beta_j} |\psi_{\beta_0}\rangle. \tag{35}$$

Let the error of each $\tilde{O}_{\beta_{j-1}}^{\beta_j}$ be upper bounded by $\delta/(2K)$, i.e., $\|\tilde{O}_{\beta_{j-1}}^{\beta_j} - O_{\beta_{j-1}}^{\beta_j}\| \le \delta/2K$. The error of the final state is then bounded by $\delta + \mathcal{O}(\delta^2)$ due to triangle inequality. More precisely,

$$\begin{aligned}
&\| |\tilde{\psi}_\beta\rangle - |\psi_\beta\rangle \| \\
&= \| \tilde{O}_{\beta_{K-1}}^{\beta_K} (|\tilde{\psi}_{\beta_{K-1}}\rangle - |\psi_{\beta_{K-1}}\rangle) + (\tilde{O}_{\beta_{K-1}}^{\beta_K} - O_{\beta_{K-1}}^{\beta_K}) |\psi_{\beta_{K-1}}\rangle \| \\
&\le \| |\tilde{\psi}_{\beta_{K-1}}\rangle - |\psi_{\beta_{K-1}}\rangle \| + \| (\tilde{O}_{\beta_{K-1}}^{\beta_K} - O_{\beta_{K-1}}^{\beta_K}) |\psi_{\beta_{K-1}}\rangle \| \\
&\le \cdots \le \sum_{j=1}^K \| \tilde{O}_{\beta_{j-1}}^{\beta_j} - O_{\beta_{j-1}}^{\beta_j} \| \le \frac{\delta}{2},
\end{aligned} \tag{36}$$

where from the first to the second line we additionally used the fact that $\|\tilde{O}_{\beta_{j-1}}^{\beta_j}\| \le 1$ since $\tilde{O}_{\beta_{j-1}}^{\beta_j}$ is encoded in a unitary. It immediately follows that $\| |\tilde{\psi}_\beta\rangle \| \ge 1 - \frac{\delta}{2}$. The error of the final state

$$\begin{aligned}
\| |\psi\rangle - |\psi_\beta\rangle \| &= \left\| \frac{|\tilde{\psi}_\beta\rangle}{\| |\tilde{\psi}_\beta\rangle \|} - |\psi_\beta\rangle \right\| \\
&= \left\| \frac{|\tilde{\psi}_\beta\rangle - \| |\tilde{\psi}_\beta\rangle \| |\psi_\beta\rangle}{\| |\tilde{\psi}_\beta\rangle \|} \right\| \\
&= \frac{\| |\psi_\beta\rangle (1 - \| |\tilde{\psi}_\beta\rangle \|) + (|\tilde{\psi}_\beta\rangle - |\psi_\beta\rangle) \|}{\| |\tilde{\psi}_\beta\rangle \|} \\
&\le \delta + \mathcal{O}(\delta^2).
\end{aligned} \tag{37}$$

On the other hand, the success probability

$$p = \| |\tilde{\psi}_\beta\rangle \|^2 \ge \left(1 - \frac{\delta}{2}\right)^2 \ge 1 - \delta. \tag{38}$$

We summarize the above in the following lemma.

Lemma 3. Let $\{\tilde{O}_{\beta_{j-1}}^{\beta_j}\}_{j=1}^K$ be a set of transition operators as in Eq. (32), the final state and the success probability

defined in Eq. (33) and Eq. (34). Then, if for each transition operator $\|\tilde{O}^{\beta_j}_{\beta_{j-1}} - |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|\| \le \delta/(2K)$, the final state $\||\psi\rangle - |\psi_\beta\rangle\| \le \delta + \mathcal{O}(\delta^2)$ and the success probability $p \ge 1-\delta$.

Hence it suffices to construct $\tilde{O}^{\beta_j}_{\beta_{j-1}}$ to precision $\delta/(2K)$. We use QSVT to construct $\tilde{O}^{\beta_j}_{\beta_{j-1}}$ from $\tilde{P}_{\beta_j}\tilde{P}_{\beta_{j-1}} \approx |\psi_{\beta_j}\rangle\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle\langle\psi_{\beta_{j-1}}|$, by using it to boost $|\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle|$ close to 1. Query complexity depends primarily on the overlap $|\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle|$, which is to be controlled by choosing an appropriate annealing path. The following lemma gives the query complexity of implementing $\tilde{O}^{\beta_j}_{\beta_{j-1}}$ with $\tilde{P}_{\beta_j}$, $\tilde{P}_{\beta_{j-1}}$.

Lemma 4. Given the block encodings of approximate projectors $\tilde{P}_{\beta_j}$ and $\tilde{P}_{\beta_{j-1}}$ where $\|\tilde{P}_{\beta_j} - |\psi_{\beta_j}\rangle\langle\psi_{\beta_j}|\|$ and $\|\tilde{P}_{\beta_{j-1}} - |\psi_{\beta_{j-1}}\rangle\langle\psi_{\beta_{j-1}}|\|$ are both upper bounded by $0 \le \mu < 1$. Assume there exists a constant $b > 0$ that lower bounds the overlap between these two states, $|\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle| \ge b$. Then, using one additional ancilla qubit and $l = \mathcal{O}(b^{-1}\log(\epsilon^{-1}))$ queries to the block encodings of $\tilde{P}_{\beta_j}$ and $\tilde{P}_{\beta_{j-1}}$, we can construct the block encoding of $\tilde{O}^{\beta_j}_{\beta_{j-1}}$ such that $\|\tilde{O}^{\beta_j}_{\beta_{j-1}} - |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|\| \le 4l\sqrt{3\mu} + \epsilon$.

Proof. Let U be the block encoding of $P_{\beta_j}P_{\beta_{j-1}} = |\psi_{\beta_j}\rangle\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle\langle\psi_{\beta_{j-1}}|$ such that $\langle\mathbf{0}|U|\mathbf{0}\rangle = P_{\beta_j}P_{\beta_{j-1}}$, where $|\mathbf{0}\rangle$ is the all-zero ancilla state. U can be obtained with one use of each of the block encodings of P_{β_j} and $P_{\beta_{j-1}}$. Let $\Pi = |\mathbf{0}\rangle\langle\mathbf{0}| \otimes I$, where I is the identity operator on the system. $\Pi U\Pi = W\Sigma V^\dagger$ has a single non-zero singular value $|\langle\psi_{\beta_j}|\psi_{\beta_{j-1}}\rangle| \ge b$. Following [45], we perform QSVT with an odd polynomial $p(\cdot)$ of degree $l = \mathcal{O}(b^{-1}\log(1/\epsilon))$ such that p is ϵ-close to 1 on the interval $[b,1]$. This lets us construct a circuit U_Φ with l uses of $U, U^\dagger$, such that $\Pi U_\Phi\Pi = p^{\mathrm{SV}}(\Pi U\Pi) = Wp(\Sigma)V^\dagger$. Because of our choice of p, we have $\|\Pi U_\Phi\Pi - |\mathbf{0}\rangle\langle\mathbf{0}| \otimes |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|\| \le \epsilon$.

Now we are given block encodings of approximate projectors $\tilde{P}_{\beta_j}, \tilde{P}_{\beta_{j-1}}$ rather than $P_{\beta_j}, P_{\beta_{j-1}}$. Let $\tilde{U}$ be the block encoding of $\tilde{P}_{\beta_j}\tilde{P}_{\beta_{j-1}}$. The robustness of QSVT [45, Lemma 22] shows that for matrices A and $\tilde{A}$ of operator norm at most 1, it holds that $\|p^{\mathrm{SV}}(A) - p^{\mathrm{SV}}(\tilde{A})\| \le 4l\sqrt{\|A-\tilde{A}\|}$. Since

$$\begin{aligned}
&\|\tilde{P}_{\beta_j}\tilde{P}_{\beta_{j-1}} - P_{\beta_j}P_{\beta_{j-1}}\| \\
&= \|(P_{\beta_j} + \delta P_{\beta_j})(P_{\beta_{j-1}} + \delta P_{\beta_{j-1}}) - P_{\beta_j}P_{\beta_{j-1}}\| \\
&= \|P_{\beta_j}\delta P_{\beta_{j-1}} + \delta P_{\beta_j}P_{\beta_{j-1}} + \delta P_{\beta_j}\delta P_{\beta_{j-1}}\| \\
&\le 2\mu + \mu^2 \le 3\mu,
\end{aligned} \tag{39}$$

where $\delta P_{\beta_j} = \tilde{P}_{\beta_j} - P_{\beta_j}$ and $\delta P_{\beta_{j-1}} = \tilde{P}_{\beta_{j-1}} - P_{\beta_{j-1}}$, we have $\|p^{\mathrm{SV}}(\Pi U\Pi) - p^{\mathrm{SV}}(\Pi\tilde{U}\Pi)\| \le 4l\sqrt{3\mu}$. Finally, by the triangle inequality,

$$\|p^{\mathrm{SV}}(\Pi\tilde{U}\Pi) - |\mathbf{0}\rangle\langle\mathbf{0}| \otimes |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|\| \le 4l\sqrt{3\mu} + \epsilon.$$

Let $\tilde{U}_\Phi$ be the circuit that implements $p^{\mathrm{SV}}(\Pi\tilde{U}\Pi)$, i.e., $\Pi\tilde{U}_\Phi\Pi = p^{\mathrm{SV}}(\Pi\tilde{U}\Pi)$. Thus $\tilde{U}_\Phi$ is the block encoding of $\tilde{O}^{\beta_j}_{\beta_{j-1}}$ and uses $\tilde{U}, \tilde{U}^\dagger$ a total of $\mathcal{O}(b^{-1}\log(1/\epsilon))$ times. ∎

By choosing $\epsilon \le \delta/(4K)$ and $\sqrt{\mu} \le \delta/(16\sqrt{3}lK)$, we have $\|\tilde{O}^{\beta_j}_{\beta_{j-1}} - |\psi_{\beta_j}\rangle\langle\psi_{\beta_{j-1}}|\| \le \delta/(2K)$. Also, when $\epsilon \le \delta/(4K)$, $l = \mathcal{O}(b^{-1}\log(K/\delta))$. Substituting into μ, we have $\mu = \mathcal{O}((\frac{b}{\frac{K}{\delta}\log\frac{K}{\delta}})^2)$. To control b, we adopt the following proposition from [4] which gives a lower bound on the overlap between two consecutive Gibbs states.

Lemma 5. (Proposition G.2 of [4]). Let $|\psi_\beta\rangle$ and $|\psi_{\beta+\delta\beta}\rangle$ be the ground states of the parent Hamiltonians H_β and $\mathsf{H}_{\beta+\delta\beta}$ respectively, or in other words, the purified Gibbs states of Hamiltonian H at temperature β and $\beta+\delta\beta$. Then

$$|\langle\psi_{\beta+\delta\beta}|\psi_\beta\rangle|^2 = 1 - \mathcal{O}(\delta\beta^2\|H\|^2). \tag{40}$$

By uniformly spacing the temperatures choosing $K = \Theta(\beta\|H\|)$, we can make b a constant. More precisely, choosing $K = \alpha\beta\|H\|$ for some constant $\alpha > 1$ gives $b^2 = 1 - \mathcal{O}(1/\alpha^2)$, which is a $\Theta(1)$ constant.

In total, we need

$$\mathcal{O}\left(\frac{M}{\sqrt{\gamma}}\log\frac{1}{\mu}\log^c\left(\frac{M}{\sqrt{\gamma}\mu}\log\frac{1}{\mu}\right)\right)\cdot\mathcal{O}\left(\frac{1}{b}\log\frac{K}{\delta}\right)K \tag{41}$$

uses of elementary gates, where c is the exponent in the Solovay-Kitaev theorem, and $\gamma = \min_j \mathrm{gap}(\mathsf{H}_{\beta_j})$. The first term comes from Corollary 4, which is the cost of constructing each approximate projector $\tilde{P}_{\beta_j}$ with error μ. The second term comes from Lemma 4, which is the cost of constructing each transition operator $O^{\beta_j}_{\beta_{j-1}}$ with error $\delta/(2K)$ given access to $\tilde{P}_{\beta_j}$ and $\tilde{P}_{\beta_{j-1}}$. The last term K is the number of steps. Substituting corresponding values, we have the following theorem.

Theorem 6. For a given Hamiltonian $H = \sum_{m=1}^M H_m$, we assume:

(1) There exists Lindbladians $\mathcal{L}_\nu$, for $\nu \in [0,\beta]$, where each $\mathcal{L}_\nu$ is irreducible and satisfies the σ_ν-detailed balance condition. Here $\sigma_\nu \propto e^{-\nu H}$ is the Gibbs state of H at inverse temperature ν.
(2) The parent Hamiltonian H_ν of each $\mathcal{L}_\nu$ is frustration-free and bounded-degree local (Definition 4).

Then there exists a quantum algorithm that prepares the purified Gibbs state at inverse temperature β with error δ and success probability at least $1-\delta$ using

$$\mathcal{O}\left(\frac{M\beta\|H\|}{\sqrt{\gamma}}\log^2\left(\frac{\beta\|H\|}{\delta}\right)\log^c\left(\frac{M}{\sqrt{\gamma}}\frac{\beta\|H\|}{\delta}\right)\right) \tag{42}$$

elementary gates, and

$$\mathcal{O}(\log M)+1 \tag{43}$$

resettable ancilla qubits, where c is the exponent in the Solovay-Kitaev theorem, and $\gamma=\min_j \mathrm{gap}(\mathcal{L}_{\beta_j})$, $\beta_j=j\beta/K$, $K=\Theta(\beta\|H\|)$.

The $\mathcal{O}(\log M)$ ancilla qubits above are used to coherently track the success of each projection operator involved in the DL operator (see the discussion under Definition 5 of block encoding).

VII. APPLICATION TO COMMUTING BOUNDED-DEGREE LOCAL HAMILTONIANS

Commuting Hamiltonians include classical spin models, such as Ising- and Potts-type systems, as well as commuting-projector and stabilizer Hamiltonians, including topologically ordered models such as the toric code. Preparing their Gibbs states provides coherent access to finite-temperature ensembles for estimating thermal observables and correlation functions, probing thermal phase structure and the thermal stability or loss of topological order, and serving as a subroutine in annealing-based estimation of partition-function ratios or free-energy differences. This class is also a standard test bed for studying thermalization and the stability of topological quantum memories against thermal noise [17,19,26,55–57].

In this section, we consider the KMS-DBC Lindbladian of the form in Ref. [8], which encompasses the Davies generator [58,59] and the algorithmic Lindbladian studied in Ref. [60]. The construction in Ref. [8] is not, in general, always irreducible. Nevertheless, we shall assume irreducibility, as this is the case for the Davies generator and the algorithmic Lindbladian in Ref. [60]. We will show that for commuting bounded-degree local Hamiltonians, the parent Hamiltonian is bounded-degree local and frustration free, therefore our method applies.

To begin with, we review the Lindbladian form in Ref. [8]. A jump operator L^a is assigned to each coupling operator A^a, defined as

$$L^a=\int_{-\infty}^{\infty} f^a(t)A^a(t)\mathrm{d}t, \quad A^a(t)=e^{iHt}A^a e^{-iHt}, \tag{44}$$

with a weighting function $f^a\colon \mathbb{R}\to\mathbb{C}$ given by

$$f^a(t)=\frac{1}{2\pi}\int_{-\infty}^{\infty} q^a(\omega)e^{-\beta\omega}e^{-it\omega}\mathrm{d}\omega. \tag{45}$$

The function $q^a(\omega)\colon \mathbb{R}\to\mathbb{C}$ satisfies $q^a(\omega)=\overline{q^a(-\omega)}$, and β is the inverse temperature. $A^a(t)$ is the time evolution of A^a under H. Note that the weighting function and hence the jump operators are in general temperature-dependent. For simplicity, we omit the subscript β representing the temperature for f^a and L^a for now. In addition to the jump operators, there is a coherent term

$$G=\sum_{a\in\mathcal{A}}\int_{-\infty}^{\infty} g(t)e^{iHt}\left((L^a)^\dagger L^a\right)e^{-iHt}\mathrm{d}t, \tag{46}$$

with $g(t)\colon \mathbb{R}\to\mathbb{C}$ defined as

$$g(t)=\frac{1}{2\pi}\int_{-\infty}^{\infty}-\frac{i}{2}\tanh\left(-\frac{\omega}{4}\right)\kappa(\omega)e^{-it\omega}\mathrm{d}\omega, \tag{47}$$

where $\kappa(\omega)$ satisfies

$$\kappa(\omega)=1 \quad \text{for } \omega\in[-2\|H\|,2\|H\|]. \tag{48}$$

Function f^a and g are all L^1-integrable. Finally, the total Lindbladian in the Heisenberg picture is defined as

$$\mathcal{L}(X)=i[G,X]+\sum_{a\in\mathcal{A}}\left((L^a)^\dagger X L^a-\frac{1}{2}\{(L^a)^\dagger L^a,X\}\right). \tag{49}$$

In our setting, to obtain the Lindbladian form of Eq. (2), let $G=\sum_{a\in\mathcal{A}}G^a$ in Eq. (46), with

$$G^a=\int_{-\infty}^{\infty} g(t)e^{iHt}\left((L^a)^\dagger L^a\right)e^{-iHt}\mathrm{d}t. \tag{50}$$

Letting

$$\mathcal{L}^a(X)=i[G^a,X]+(L^a)^\dagger X L^a-\frac{1}{2}\{(L^a)^\dagger L^a,X\}, \tag{51}$$

we obtain $\mathcal{L}=\sum_{a\in\mathcal{A}}\mathcal{L}^a$. For a specific temperature β, to explicitly denote the temperature dependence, we write $\mathcal{L}_\beta=\sum_{a\in\mathcal{A}}\mathcal{L}^a_\beta$. This Lindbladian is σ_β-detailed balance, as shown in Ref. [8]. We further assume that it is irreducible.

It follows from Eq. (28) that, for superoperator $K(\cdot)=A(\cdot)B^\dagger$, its vectorization $vKv^{-1}=A\otimes B^*$, where B^* is the entrywise complex conjugate. Subsequently, our parent Hamiltonian at temperature β can be written as

$$\mathsf{H}_\beta=\sum_{a\in\mathcal{A}} v\mathcal{L}^a_\beta v^{-1}=\sum_{a\in\mathcal{A}}\mathsf{H}^a_\beta, \tag{52}$$

where

$$\mathsf{H}_\beta^a = \left(\sigma_\beta^{\frac{1}{4}} \otimes (\sigma_\beta^\top)^{\frac{1}{4}}\right) \cdot (iG_\beta^a \otimes I - iI \otimes (G_\beta^a)^\top + (L_\beta^a)^\dagger \otimes (L_\beta^a)^\top - \frac{1}{2}(L_\beta^a)^\dagger L_\beta^a \otimes I - \frac{1}{2} I \otimes (L_\beta^a)^\top (L_\beta^a)^*) \cdot \left(\sigma_\beta^{-\frac{1}{4}} \otimes (\sigma_\beta^\top)^{-\frac{1}{4}}\right), \tag{53}$$

where the first and last term are the vectorization of Γ_σ defined in Eq. (21) and the middle term is the vectorization of $\mathcal{L}_\beta^a$. The parent Hamiltonian is frustration free, with $|\sqrt{\sigma_\beta}\rangle$ being the ground state for both H_β and H_β^a.

Consider a commuting Hamiltonian $H = \sum_m H_m$ with locality l and bounded degree q and a set of coupling operators $\{A^a\}$ with locality k_A and bounded degree g_A. The support of $A^a(t)$ is the union of the support of A^a and all Hamiltonian terms that intersect A^a. The number of such intersecting Hamiltonian terms is bounded. Then by locality, the support of $A^a(t)$ is bounded. By definition, L_β^a, G_β^a, both before and after applying by $\sigma_\beta^{\frac{1}{4}} \cdot \sigma_\beta^{-\frac{1}{4}} = e^{-\frac{1}{4}\beta H} \cdot e^{\frac{1}{4}\beta H}$, have the same support as $A^a(t)$. Therefore, the parent Hamiltonian is local. Furthermore, since $A^a(t)$ has bounded support, the number of Hamiltonian terms that intersect $A^a(t)$ is bounded, therefore the number of a' such that $\mathrm{supp}(A^{a'}(t)) \cap \mathrm{supp}(A^a(t)) \neq \varnothing$ is bounded. Thus, the parent Hamiltonian is bounded-degree local.

Consider the annealing path with inverse temperature $\nu \in [0, \beta]$. Each $\mathcal{L}_\nu$ and H_ν are constructed as above, thus they meet the requirements for our method to apply. To summarize, we have the following corollary.

Corollary 5. Consider commuting and bounded-degree local Hamiltonian H (Definition 4). For each $\nu \in [0, \beta]$, we construct σ_ν-detailed balanced $\mathcal{L}_\nu$ as above and assume its irreducibility. Then, the corresponding parent Hamiltonian H_ν will be bounded-degree local and frustration free, thus it satisfies the requirements of Theorem 6. Therefore there exists a quantum algorithm that prepares the purified Gibbs state at inverse temperature β with error δ and success probability at least $1 - \delta$ with gate complexity

$$\mathcal{O}\left(\frac{M\beta\|H\|}{\sqrt{\gamma}} \log^2\left(\frac{\beta\|H\|}{\delta}\right) \log^c\left(\frac{M}{\sqrt{\gamma}}\frac{\beta\|H\|}{\delta}\right)\right) \tag{54}$$

using $\mathcal{O}(\log M) + 1$ resettable ancilla qubits, where c is the exponent in the Solovay-Kitaev theorem, and $\gamma = \min_j \mathrm{gap}(\mathcal{L}_{\beta_j})$, $\beta_j = j\beta/K$, $K = \Theta(\beta\|H\|)$.

VIII. DISCUSSION AND OUTLOOK

In this work, we use the detectability lemma to improve quantum Gibbs sampling. We propose a quantum algorithm for preparing the stationary state without simulating Lindbladian evolution, thereby avoiding the associated overhead. We also present a quantum algorithm for preparing the Gibbs state of commuting bounded-degree local Hamiltonians, achieving a quadratic improvement in its dependence on the Lindbladian spectral gap. As a by-product, we also obtain an efficient quantum algorithm for preparing the ground state of frustration-free bounded-degree local Hamiltonians.

For future work, an important direction would be to extend the quadratic improvement in spectral-gap dependence to a broader class of Gibbs-state preparation problems. A natural way to achieve this would be to develop a stronger version of the detectability lemma that applies to quasi-local interactions. Beyond its conceptual interest, such a result could also lead to an improved Gibbs-sampling algorithm along the lines of Corollary 4.

One might also want to apply the Lieb-Robinson bound to truncate each quasi-local jump operator to a ball of radius R, apply the present algorithm to the resulting finite-range generator, and choose R so that the truncation tail is below the target accuracy [4,60,61]. A direct application, however, goes beyond the scope of our current analysis. A generic truncation can destroy exact KMS-detailed balance and the exact frustration freeness of the parent Hamiltonian, both of which enter our proof. Moreover, since local fixed-point projections are implemented via unitary synthesis in our current algorithm, such truncation leads to superpolynomial scaling with respect to the truncation error. How to handle the non-commuting case while preserving algorithmic efficiency in this detectability-lemma-based approach is therefore still an open problem.

More broadly, it would be very interesting to understand whether the detectability lemma can be used as a tool for analyzing the mixing times of Lindbladians [8,16–32]. Progress in this direction could be especially valuable if it were to yield a genuine quantum advantage. At present, however, many seemingly efficient quantum algorithms based on Lindbladian dynamics are challenged by strong classical algorithms [62–66]. Establishing a framework in which detectability-lemma techniques lead to provable quantum advantages in this setting would therefore be particularly compelling.

Another important direction is improving Lindbladian simulation. The recently derived commutator bound for Lindbladians [67] yields a better dependence on the number of terms for local Lindbladians, but its precision dependence is worse than that of [9]. Achieving the same spacetime-volume scaling for Lindbladian simulation as has been obtained for Hamiltonian simulation [68,69] remains a major open problem.

ACKNOWLEDGMENTS

Y. T. thanks András Gilyén for helpful discussion. C. Z. acknowledges support from NSF QLCI Grant No. OMA-2120757. J. F. acknowledges support by the National Science Foundation under Award No. DMS-2309378 D. F. acknowledges the support from the U.S. Department of Energy, Office of Science,

Accelerated Research in Quantum Computing Centers, Quantum Utility through Advanced Computational Quantum Algorithms, Grant No. DE-SC0025572, and National Science Foundation via the NSF CAREER Award No. DMS-2438074.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] F. G. Brandão, A. Kalev, T. Li, C. Y.-Y. Lin, K. M. Svore, and X. Wu, *Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning*, arXiv:1710.02581.

[2] J. Van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, *Quantum SDP-solvers: Better upper and lower bounds*, in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2017), pp. 403–414.

[3] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K..-L. Chan, *Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution*, Nat. Phys. **16**, 205 (2020).

[4] C.-F. Chen, M. J. Kastoryano, F. G. Brandão, and A. Gilyén, *Quantum thermal state preparation*, arXiv:2303.18224.

[5] J. Lloyd and D. A. Abanin, *Quantum thermal state preparation for near-term quantum processors*, arXiv:2506.21318.

[6] C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandao, and A. Gilyen, *Quantum thermal state preparation*, arXiv:2303.18224.

[7] P. Rall, C. Wang, and P. Wocjan, *Thermal state preparation via rounding promises*, Quantum **7**, 1132 (2023).

[8] Z. Ding, B. Li, and L. Lin, *Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition*, Commun. Math. Phys. **406**, 67 (2025).

[9] R. Cleve and C. Wang, *Efficient quantum algorithms for simulating Lindblad evolution*, in *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)* (Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Dagstuhl, Germany, 2017).

[10] P. Wocjan and K. Temme, *Szegedy walk unitaries for quantum maps*, Commun. Math. Phys. **402**, 3201 (2023).

[11] C.-F. Chen, M. J. Kastoryano, and A. Gilyen, *An efficient and exact noncommutative quantum Gibbs sampler*, arXiv:2311.09207.

[12] X. Li and C. Wang, *Simulating Markovian open quantum systems using higher-order series expansion*, in *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)* (Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Dagstuhl, Germany, 2023).

[13] Z. Ding, X. Li, and L. Lin, *Simulating open quantum systems using Hamiltonian simulations*, PRX Quantum **5**, 020332 (2024).

[14] M. Pocrnic, D. Segal, and N. Wiebe, *Quantum simulation of Lindbladian dynamics via repeated interactions*, arXiv:2312.05371.

[15] C.-F. Chen, M. Kastoryano, F. G. S. L. Brandão, and A. Gilyén, *Efficient quantum thermal simulation*, Nature (London) **646**, 561 (2025).

[16] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete, *Hypercontractivity of quasi-free quantum semigroups*, J. Math. Phys. **47**, 405303 (2014).

[17] Z. Ding, B. Li, L. Lin, and R. Zhang, *Polynomial-time preparation of low-temperature Gibbs states for 2d toric code*, arXiv:2410.01206.

[18] I. Bardet, Á. Capel, L. Gao, A. Lucia, D. Pérez-García, and C. Rouzé, *Entropy decay for Davies semigroups of a one dimensional quantum lattice*, Commun. Math. Phys. **405**, 42 (2024).

[19] R. Alicki, M. Fannes, and M. Horodecki, *On thermalization in Kitaev's 2D model*, J. Phys. A **42**, 065303 (2009).

[20] A. Ramkumar and M. Soleimanifar, *Mixing time of quantum Gibbs sampling for random sparse Hamiltonians*, arXiv:2411.04454.

[21] J. Kochanowski, A. M. Alhambra, A. Capel, and C. Rouzé, *Rapid thermalization of dissipative many-body dynamics of commuting Hamiltonians*, Commun. Math. Phys. **406**, 176 (2025).

[22] I. Bardet, Á. Capel, L. Gao, A. Lucia, D. Pérez-García, and C. Rouzé, *Rapid thermalization of spin chain commuting Hamiltonians*, Phys. Rev. Lett. **130**, 060401 (2023).

[23] C. Rouzé, D. S. França, and Á. M. Alhambra, *Optimal quantum algorithm for Gibbs state preparation*, Phys. Rev. Lett. **136**, 060601 (2026).

[24] Á. Capel, P. Gondolf, J. Kochanowski, and C. Rouzé, *Quasi-optimal sampling from Gibbs states via noncommutative optimal transport metrics*, Ann. Henri Poincaré (2025).

[25] M. J. Kastoryano and K. Temme, *Quantum logarithmic Sobolev inequalities and rapid mixing*, J. Math. Phys. **54**, 052202 (2013).

[26] K. Temme and M. J. Kastoryano, *How fast do stabilizer Hamiltonians thermalize?*, arXiv:1505.07811.

[27] H. Chen, B. Li, J. Lu, and L. Ying, *A randomized method for simulating Lindblad equations and thermal state preparation*, Quantum **9**, 1917 (2025).

[28] T. Rakovszky, B. Placke, N. P. Breuckmann, and V. Khemani, *Bottlenecks in quantum channels and finite temperature phases of matter*, arXiv:2412.09598.

[29] B. Li and J. Lu, *Quantum space-time Poincare inequality for Lindblad dynamics*, arXiv:2406.09115.

[30] T. Barthel and Y. Zhang, *Solving quasi-free and quadratic Lindblad master equations for open fermionic and bosonic systems*, J. Stat. Mech.:Theory Exp. **2022**, 113101 (2022).

[31] D. Fang, J. Lu, and Y. Tong, *Mixing time of open quantum systems via hypocoercivity*, Phys. Rev. Lett. **134**, 140405 (2025).10.1103/PhysRevLett.134.140405

[32] Z. Ding, C.-F. Chen, and L. Lin, *Single-ancilla ground state preparation via Lindbladians*, Phys. Rev. Res. **6**, 033147 (2024).

[33] Y. Tong and Y. Zhan, *Fast mixing of weakly interacting fermionic systems at any temperature*, PRX Quantum **6**, 030301 (2025).

[34] A. Bakshi, A. Liu, A. Moitra, and E. Tang, *A Dobrushin condition for quantum Markov chains: Rapid mixing and conditional mutual information at high temperature*, arXiv:2510.08542.
[35] A. Gilyén, C.-F. Chen, J. F. Doriguello, and M. J. Kastoryano, *Quantum generalizations of Glauber and metropolis dynamics*, arXiv:2405.20322.
[36] J. Jiang and S. Irani, *Quantum metropolis sampling via weak measurement*, arXiv:2406.16023.
[37] M. Szegedy, *Quantum speed-up of Markov chain based algorithms*, in *45th Annual IEEE Symposium on Foundations of Computer Science* (IEEE, New York, 2004), pp. 32–41.
[38] J. Leng, Z. Ding, Z. Chen, and L. Lin, *Operator-level quantum acceleration of non-Logconcave sampling*, Proc. Natl. Acad. Sci. U. S. A. **123**, e2512789123 (2026).
[39] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, *Quantum Hamiltonian complexity and the detectability lemma*, arXiv:1011.3445.
[40] I. Arad, A. Kitaev, Z. Landau, and U. Vazirani, *An area law and sub-exponential algorithm for 1d systems*, arXiv:1301.1162.
[41] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, *The detectability lemma and quantum gap amplification*, in *Proceedings of the forty-first annual ACM symposium on Theory of computing* (Association for Computing Machinery, New York, NY, USA, 2009), pp. 417–426.
[42] F. Ascolani, H. Lavenant, and G. Zanella, *Entropy contraction of the Gibbs sampler under log-concavity*, arXiv:2410.00858.
[43] J. S. Liu, W. H. Wong, and A. Kong, *Covariance structure and convergence rate of the Gibbs sampler with various scans*, J. R. Stat. Soc.: Series B (Methodological) **57**, 157 (1995).
[44] H. Narayanan and P. Srivastava, *On the mixing time of coordinate hit-and-run*, Comb. Probab. Comput. **31**, 320 (2022).
[45] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, *Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics*, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC '19* (ACM, New York, NY, USA, 2019), pp. 193–204.
[46] L. Lin and Y. Tong, *Near-optimal ground state preparation*, Quantum **4**, 372 (2020).
[47] Y. Ge, J. Tura, and J. I. Cirac, *Faster ground state preparation and high-precision ground energy estimation with fewer qubits*, J. Math. Phys. **60**, 022202 (2019).
[48] Here we use † to denote conjugation under the Hilbert-Schmidt inner product.
[49] A. Anshu, I. Arad, and T. Vidick, *Simple proof of the detectability lemma and spectral gap amplification*, Phys. Rev. B **93**, 205142 (2016).
[50] G. H. Low and N. Wiebe, *Hamiltonian simulation in the interaction picture*, arXiv:1805.00675.
[51] D. Fang, L. Lin, and Y. Tong, *Time-marching based quantum solvers for time-dependent linear differential equations*, Quantum **7**, 955 (2023).
[52] M. Thibodeau and B. K. Clark, *Nearly-frustration-free ground state preparation*, Quantum **7**, 1084 (2023).
[53] L. Lin and Y. Tong, *Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems*, Quantum **4**, 361 (2020).
[54] Strictly speaking, this state is the top eigenstate of the parent Hamiltonian, but they are equivalent up to flipping sign.
[55] M. J. Kastoryano and F. G. S. L. Brandão, *Quantum Gibbs samplers: The commuting case*, Commun. Math. Phys. **344**, 915 (2016).
[56] Y. Hwang and J. Jiang, *Gibbs state preparation for commuting Hamiltonian: Mapping to classical Gibbs sampling*, Quantum **10**, 2209 (2026).
[57] T.-C. Lu, T. H. Hsieh, and T. Grover, *Detecting topological order at finite temperature using entanglement negativity*, Phys. Rev. Lett. **125**, 116801 (2020).
[58] E. B. Davies, *Markovian master equations*, Commun. Math. Phys. **39**, 91 (1974).
[59] E. B. Davies, *Markovian master equations. ii*, Math. Ann. **219**, 147 (1976).
[60] C.-F. Chen, M. J. Kastoryano, and A. Gilyén, *An efficient and exact noncommutative quantum Gibbs sampler*, arXiv:2311.09207.
[61] E. H. Lieb and D. W. Robinson, *The finite group velocity of quantum spin systems*, Commun. Math. Phys. **28**, 251 (1972).
[62] A. Bakshi, A. Liu, A. Moitra, and E. Tang, *High-temperature Gibbs states are unentangled and efficiently preparable*, arXiv:2403.16850.
[63] B. T. McDonough, C. Yin, A. Lucas, and C. Zhang, *Lieb-Robinson bounds with exponential-in-volume tails*, PRX Quantum **6**, 040322 (2025).
[64] A. Ramkumar, Y. Cai, Y. Tong, and J. Jiang, *High-temperature fermionic Gibbs states are mixtures of Gaussian states*, arXiv:2505.09730.
[65] H. Chen, C. Rouzé, J. Chen, J. Jiang, S. O. Scalet, Y. Zhan, G. K. Chan, L. Ying, and Y. Tong, *Convergence of the cumulant expansion and polynomial-time algorithm for weakly interacting fermions*, arXiv:2512.12010.
[66] A. Zlokapa and B. T. Kiani, *SYK thermal expectations are classically easy at any temperature*, arXiv:2602.22619.
[67] X. Wang, S. Zhou, X. Wang, Y.-C. Zheng, S. Zhang, and T. Li, *Lindbladian simulation with commutator bounds*, arXiv:2603.28602.
[68] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, *Quantum algorithm for simulating real time evolution of lattice Hamiltonians*, SIAM J. Comput. **52**, FOCS18 (2021).
[69] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, *Theory of trotter error with commutator scaling*, Phys. Rev. X **11**, 011020 (2021).