



Mixing Time of Quantum Gibbs Sampling for Random Sparse Hamiltonians

Akshar Ramkumar ✉ 

Institute for Quantum Information and Matter, California Institute of Technology,
Pasadena, CA, USA

Mehdi Soleimanifar ✉ 

Institute for Quantum Information and Matter, California Institute of Technology,
Pasadena, CA, USA

Abstract

Providing evidence that quantum computers can efficiently prepare low-energy or thermal states of physically relevant interacting quantum systems is a major challenge in quantum information science. A newly developed quantum Gibbs sampling algorithm [11] provides an efficient simulation of the detailed-balanced dissipative dynamics of non-commutative quantum systems. The running time of this algorithm depends on the mixing time of the corresponding quantum Markov chain, which has not been rigorously bounded except in the high-temperature regime. In this work, we establish a $\text{polylog}(n)$ upper bound on its mixing time for various families of random $n \times n$ sparse Hamiltonians at any constant temperature. We further analyze how the choice of the jump operators for the algorithm and the spectral properties of these sparse Hamiltonians influence the mixing time. Our result places this method for Gibbs sampling on par with other efficient algorithms for preparing low-energy states of quantumly easy Hamiltonians.

2012 ACM Subject Classification Theory of computation → Quantum information theory; Theory of computation → Design and analysis of algorithms

Keywords and phrases Quantum algorithms, quantum Gibbs sampling, mixing time analysis

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.3

Related Version *Extended Version:* **arXiv preprint** [32]

Supplementary Material *Software (Source Code):* https://github.com/countableclouds/lindbladian_gap, archived at **swb**:1:dir:9db70c0155cd182d780b7562eca062cc64a051ad

Funding *Akshar Ramkumar:* AR acknowledges the Caltech SURF program and Zhuang Tang and Gebing Yi for their generous support.

Mehdi Soleimanifar: MS is supported by AWS Quantum Postdoctoral Scholarship and the National Science Foundation NSF CAREER award CCF-2048204.

Acknowledgements We thank Jiaoyang Huang, Luca Nashabeh, John Preskill, and Yongtao Zhan for valuable discussions. We are grateful to Chi-Fang Chen for insightful discussions and proposing this project in its early stages. We also thank Hongrui Chen and the authors of [12] for informing us about the overlap between our results and those in Section IV of their work. Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

1 Introduction

One of the main anticipated applications of quantum computers is the simulation and characterization of quantum systems in condensed matter physics [40], quantum chemistry [29], and high-energy physics [30, 4]. The problem of simulating the dynamics (time evolution) of an interacting quantum system under a local or sparse Hamiltonian H has largely been addressed, with efficient algorithms [22, 27, 5, 28, 20] that scale well with the number of



© Akshar Ramkumar and Mehdi Soleimanifar;
licensed under Creative Commons License CC-BY 4.0

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).

Editor: Bill Fefferman; Article No. 3; pp. 3:1–3:23



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

particles, simulation time, and required precision. However, the ability of quantum computers to evaluate the static features of quantum systems, such as their ground state or thermal properties, is less understood.

In this work, we focus on preparing the Gibbs (thermal) state $\rho_\beta = \frac{e^{-\beta H}}{\text{Tr}(e^{-\beta H})}$ of a quantum system, which represents the equilibrium state when the system is in contact with a thermal bath at a fixed temperature β^{-1} . This computational problem, known as Gibbs sampling or “cooling,” is valuable not only for simulating thermodynamic properties but also as a subroutine in quantum algorithms for optimization and learning [7, 2, 6]. However, to prepare the Gibbs state, quantum computers face challenges. In general, it is not believed that estimating the low-temperature properties of quantum systems can be solved efficiently by a quantum computer in the worst-case [24]. Fortunately, it has been hypothesized that this worst-case hardness of finding low-temperature states implied by arguments from complexity theory is due to pathological Hamiltonians, which are not apparent in many physical systems that normally occur in nature. This hypothesis is substantiated by the empirical success of natural cooling, such as using refrigerators, in reaching thermal equilibrium.

Quantum Gibbs sampling. Aiming to mimic nature’s cooling processes, a series of recent works have introduced quantum Markov Chain Monte Carlo (MCMC) algorithms, or quantum Gibbs samplers [11, 10, 36, 42, 31, 23, 43, 16, 19], as promising alternatives for tackling a range of classically intractable low-temperature simulation tasks on quantum computers. These algorithms are designed to replicate the success of classical Markov chains in preparing Gibbs states for classical Hamiltonians. The analysis of classical MCMC algorithms relies on the principle of detailed balance; however, achieving this in the quantum setting has been challenging and was only recently addressed by an algorithm in [11]. Part of the difficulty arises from a conflict between the finite energy resolution σ_E achievable by efficient quantum algorithms and the seemingly strict requirement to precisely distinguish energy levels to satisfy detailed balance. In this work, we focus primarily on this algorithm, referring to it as the CKG algorithm or the quantum Gibbs sampler when the context is clear. We give a detailed review of this algorithm in Section 4.1.3 and Appendix 4.2.1.

The Gibbs sampling algorithm provides a fully general method for preparing Gibbs states by evolving an initial state ρ_0 under a Lindbladian \mathcal{L}_β , which is efficiently implementable on a quantum computer and produces the state $\rho_t = e^{\mathcal{L}_\beta t}[\rho_0]$ after time t . The runtime of the quantum Gibbs sampler is governed by the *mixing time* of the corresponding quantum Markov chain, which is roughly the time required for ρ_t to approach the Gibbs state ρ_β . This in turn is bounded by the spectral gap $\lambda_{\text{gap}}(\mathcal{L}_\beta)$ of the Lindbladian by

$$t_{\text{mix}}(\mathcal{L}_\beta) \leq \frac{\mathcal{O}(\beta\|H\| + \log(n))}{\lambda_{\text{gap}}(\mathcal{L}_\beta)}.$$

The spectral gap is defined here to be λ_{min} , the smallest eigenvalue of $-\mathcal{L}_\beta$ for any eigenvector other than the fixed point ρ_β . Bounding the spectral gap, therefore, proves not only that \mathcal{L}_β has a unique fixed point, but also quantifies the rate of convergence. The mixing time varies based on the quantum system in question. Bounding this mixing time is challenging without access to fault-tolerant quantum computers, as we cannot run and benchmark the algorithm directly, making theoretical analysis essential. However, such analysis is hindered by a lack of technical tools for two key reasons. Firstly, the theory of convergence of quantum Markov chains is new, unlike the very mature twin field for classical Markov chains. Secondly, the Markov chain described by the algorithm is considerably complex, and depends on several parameters that we will discuss in more detail shortly: an energy resolution σ_E , a series of jump operators A^a for $a \in [M]$, and the inverse temperature β . The space of possibilities makes the algorithm’s performance more difficult to characterize.

This motivates the identification of quantum systems whose mixing times are tractable for analysis yet exhibit rich features that provide insights into the performance of the quantum Gibbs sampler for more general non-commuting Hamiltonians. In line with this, the mixing time of the CKG algorithm has recently been bounded for local Hamiltonians, showing a polynomial scaling with system size at high enough temperatures [33].

Mixing time of sparse Hamiltonians. In this work, we consider an alternative approach by characterizing the mixing time of a family of *sparse* Hamiltonians of the form

$$\mathbf{H} = \sum_{i,j \in [n]} H_{ij} |e_i\rangle \langle e_j|. \quad (1)$$

Such an operator can be understood as the Hamiltonian on a graph $G = (V, E)$ with $n = |V|$ vertices indexed by basis states $|e_i\rangle$, $i \in [n]$ and a set of edges E connecting vertices with $H_{ij} \neq 0$. When non-zero entries H_{ij} are all equal to 1, the Hamiltonian \mathbf{H} corresponds to the $n \times n$ adjacency matrix of the n -vertex graph. We define the degree d of the graph G as the sparsity of the underlying Hamiltonian and refer to Hamiltonians with constant or slowly increasing degrees $d = \text{polylog}(n)$ as *sparse*. Note that any $\log(n)$ -qubit Hamiltonian that consists of $m = \text{polylog}(n)$ terms each acting locally on $\kappa = O(1)$ qubits is a sparse Hamiltonian with degree $d \leq m2^\kappa \leq \text{polylog}(n)$. However, not all sparse Hamiltonians admit local qubit encodings.

Having defined sparse Hamiltonians, we now consider the dissipative dynamics of the system induced by a set of M jump operators $\mathbf{A}^a = \sum_{i,j \in [n]} A_{ij}^a |e_i\rangle \langle e_j|$, $a \in [M]$. We will soon explain how the jump operators \mathbf{A}^a relate to the Lindbladian \mathcal{L}_β . Briefly, the resulting dynamics can be understood as a combination of two processes: a continuous-time quantum walk of a single particle on the graph of states due to the coherent evolution of the Hamiltonian \mathbf{H} , which is combined with stochastic jumps on the graph determined by the jump operators \mathbf{A}^a .

Our interest in bounding the mixing time of the sparse Hamiltonians is multifaceted:

- (1) **Single-particle dynamics.** As stated earlier, bounding the mixing time of general interacting multipartite Hamiltonians is a challenging task. However, for simple choices of graphs G , the mixing time of the quantum Gibbs sampler may be easier to analyze, potentially leading to relevant techniques for tackling the case of interacting particles. In fact, we can think of the dynamics induced by the Hamiltonian \mathbf{H} (1) as the dynamics of a single-particle hopping on the graph G . This single-particle evolution on path graphs or grids is commonly analyzed in the tight-binding model in condensed matter physics. That being said, even in the simplified case of a single particle, the Hamiltonian \mathbf{H} is non-commuting, characterizing a continuous-time quantum walk that can yield exponential quantum advantage for certain oracular problems on graphs such as the glued trees [13].
- (2) **Chaotic Hamiltonians.** Our additional motivation for studying random sparse Hamiltonians stems from the fact that their spectra exhibit many of the same characteristics as *chaotic* Hamiltonians, such as the SYK model [34, 25, 26] and random p -spin models [37, 41]. Understanding whether chaotic Hamiltonians have a fast mixing time as they approach their thermal and low-energy states is a fundamental question in the study of quantum chaos [8, 1]. As a concrete step toward addressing this problem, we identify key spectral properties of random sparse Hamiltonians that can ensure a fast mixing time.
- (3) **Algorithmic applications.** Preparing quantum Gibbs states, and more broadly computing the matrix exponential of sparse matrices such as the adjacency or Laplacian of a graph, is a fundamental subroutine in solving various graph and optimization

problems. For instance, the Estrada index – defined as the trace of the matrix exponential of a graph’s adjacency matrix – measures subgraph centrality and provides structural insights [18]. Computing the matrix exponential is also related to matrix inversion and linear system solvers [35]. Moreover, quantum Gibbs sampling has been applied to solving semidefinite programs (SDPs) in optimization problems [21, 7, 6, 3], offering quantum speedups for these problems.

2 Our main results

Motivated by these considerations, we investigate the mixing time of quantum Gibbs samplers for sparse Hamiltonians and different choices of jump operators. Our study addresses two key questions regarding the performance of the quantum Gibbs sampler for sparse Hamiltonians. First, we ask

What choices of jump operators lead to a fast mixing time?

After exploring the effects of different jump operators \mathbf{A}^a , we then focus on the spectral properties of sparse Hamiltonians to understand:

What spectral property of the Hamiltonian determines its mixing time?

Answering these questions allows us to provide broad and intuitive insights on how the quantum Gibbs sampler operates for general families of sparse Hamiltonians.


2.1 Choice of jumps: graph-local vs unitary design

A natural set of jump operators for a given $n \times n$ Hamiltonian on a graph G are $\mathbf{A}^a = \frac{1}{\sqrt{n}}|e_a\rangle\langle e_a|$ or similar operators supported on a few neighboring vertices of G . Importantly, these are not “local” in the sense of multi-particle Hamiltonians, which refers to being composed of terms that act on a small number of qubits – often also geometrically close to one another. Utilizing graph-local jump operators also significantly simplifies the structure of the Lindbladian and the analysis of mixing times for certain graph families.

Moving beyond graph-local jumps, the Lindbladian \mathcal{L}_β of the quantum Gibbs sampler can still be efficiently implemented on a quantum computer with a much broader class of jumps. This is possible as long as each jump \mathbf{A}^a is efficiently implementable, the set of jumps M includes both \mathbf{A}^a and its adjoint $\mathbf{A}^{a\dagger}$, and $\sum_{a \in [M]} \|\mathbf{A}^{a\dagger} \mathbf{A}^a\|_\infty = 1$ (due to this normalization condition, we will sometimes speak of the jumping distribution \mathcal{A} , from which the jump operators \mathbf{A}^a are sampled with probability $\|\mathbf{A}^{a\dagger} \mathbf{A}^a\|$). This raises the question of whether there is an advantage in using *non-local* jumps that have a bounded spectral norm, or if more structured local jumps are sufficient to achieve a fast-mixing quantum MCMC. After all, *classical* continuous-time random walks are typically considered with local jumps on the graph vertices. However, in the context of graphs, we will see that the structured nature of graph-local jumps offers no advantage, but rather seems to cause a slowdown of the resulting algorithm.

Graph-local jumps. To this end, in the next theorem, we establish tight bounds on the spectral gap of the Lindbladian for cyclic graphs for graph-local jumps \mathbf{A}^a , with an approach similar to the one used in [38] to bound the spectral gap of a Davies generator.

► **Theorem 2.1** (Spectral gap of cyclic graphs with local jumps). *Fix temperature β^{-1} . There exists some constant energy resolution σ_E for which the spectral gap of the CKG Lindbladian \mathcal{L}_β for a cyclic graph with n vertices with jump operators $\mathbf{A}^a = \frac{1}{\sqrt{n}}|e_a\rangle\langle e_a|$ is asymptotically $\Theta(n^{-3})$.*



■ **Figure 1** Linear (above) and log-log (below) graphs of spectral gap with respect to system size. Gaps of ten random 4-regular graphs were averaged for each data point. For the cyclic graphs (one-dimensional lattices), the proven asymptotic decay aligns closely with the data.

In addition to theoretical analysis, we also generated data for cyclic graphs, path graphs, and random d -regular graphs with n vertices, as shown in Figure 1. These numerical results suggest spectral gaps of $o(n^{-1})$ for generic sparse graphs with graph-local jumps. We observed that increasing the constant d does improve the spectral gap decay, though it never improved past the asymptotic decay $O(n^{-1})$.

These results are all suboptimal, since for an $n \times n$ Hamiltonian \mathbf{H} , we expect an efficient result would be polynomial in the number of qubits, i.e. $\text{polylog}(n)$ rather than $\text{poly}(n)$. The poor performance can be attributed to two factors. (1) The operators $\mathbf{A}^{a\dagger}\mathbf{A}^a$ have L^1 norm $\frac{1}{n}$. (2) In the energy basis, many entries of \mathbf{A}^a are highly correlated.

The first drawback effectively scales the Lindbladian down by $\frac{1}{n}$, since the L^1 norm of $\mathbf{A}^\dagger\mathbf{A}$ can be as high as 1 when the operator norm $\|\mathbf{A}^\dagger\mathbf{A}\| = \frac{1}{n}$. However, the chosen jump operators are projectors, so their L^1 and operator norms are equal. In both Theorem 2.1 and in the data, a spectral gap even worse than $\frac{1}{n}$ is observed. This is due to the second drawback. The aforementioned correlations lead to off-diagonal terms in the Lindbladian, which in general have the potential to dampen the spectral gap, and in the case of the cyclic graph provably do so. It appears that more generally, the biases of an ensemble of local jumps can introduce off-diagonal terms to the Lindbladian that decrease the spectral gap. The same harmful correlations appear to exist in higher degree graphs in addition to cyclic ones, though to a lesser extent as evidenced by the improved spectral gap.

Unitary design jumps. To address some of these shortcomings, we next consider non-local jump operators, each independently drawn (along with its adjoint pair) according to a unitary 1-design $\mathcal{D}(U(n))$ on n vertices. More precisely, we define

► **Definition 2.2.** A set of jump operators $\{\mathbf{A}^a : a \in [M]\}$ is **drawn from a 1-design jumping distribution** if it is obtained by sampling $M/2$ jump operators i.i.d from a unitary 1-design $\mathcal{D}(U(n))$, normalizing each by $\frac{1}{\sqrt{M}}$, and including these operators along with their adjoint.

We include the adjoint of each randomly chosen jump since the CKG Lindbladian requires the set of jump operators to be closed under adjoint, $\{\mathbf{A}^a : a \in [M]\} = \{\mathbf{A}^{a\dagger} : a \in [M]\}$. When n is a power of 2, the unitary 1-design can be constructed as a tensor product of

random Pauli operators on $\log_2(n)$ qubits, in which case the jumps are self-adjoint and can be sampled and implemented efficiently. The efficiency of our results on a general system relies on the ability to efficiently implement some unitary 1-design.

As we will see, in our application, this 1-design sampling is effectively equivalent to sampling from a Haar-random distribution. This approach improves on the results given for graph-local jumps, and is able to achieve an efficient algorithm in the number of qubits for a graph (running time $\text{polylog}(n)$) for Gibbs sampling. This improved performance is in part because all the eigenvalues of a Haar random unitary have magnitude 1. Hence, it avoids the problem of $\mathbf{A}^\dagger \mathbf{A}$ having a relatively small L^1 norm given the constraint on its operator norm $\|\mathbf{A}^\dagger \mathbf{A}\|$. These jumps also avoid the second problem encountered for the graph-local jumps: Since the number of degrees of freedom of randomness is very large over the ensemble, any form of bias is mitigated. Indeed, the resulting Lindbladian over the full ensemble has no off-diagonal terms resulting from correlated elements of the jump operator.

Our results extend beyond cyclic graphs to any graph of bounded degree $d = O(1)$ where $\|\mathbf{H}\| \leq d$ at constant temperature, or more generally when $\beta\|\mathbf{H}\| = O(1)$. We refer to these sparse Hamiltonians as bounded degree and formally define them as:

► **Definition 2.3.** A *bounded degree system* is a sequence of temperatures $\beta(n)^{-1}$ and Hamiltonians $\mathbf{H}(n)$ for which $\beta(n)\|\mathbf{H}(n)\|$ is bounded from above by a constant independent of system size.

► **Theorem 2.4** (Constant spectral gap of Lindbladian in bounded degree systems). Let $\beta(n)^{-1}$ be a sequence of temperatures and $\mathbf{H}(n)$ a sequence of $n \times n$ Hamiltonians such that $\beta(n)\|\mathbf{H}(n)\| = O(1)$.

With any constant probability $1 - \xi$, the spectral gap of a Lindbladian \mathcal{L}_β with $\sigma_E = \beta^{-1}$ and M jump operators sampled from a 1-design jumping distribution for some $M = \Theta(\log(n))$, is bounded below by a constant, i.e. $\lambda_{\text{gap}} = \Omega(1)$.

Assume access to an efficient block-encoding of $\mathbf{H}(n)$. Then as a consequence and in the same setup, the Gibbs state of \mathbf{H} can be prepared with error ϵ in trace distance, in time $\text{poly}(\log(n), \log(\epsilon^{-1}))$.

While the examples of bounded degree Hamiltonians we consider are mostly graphs, the above theorem applies to preparing the Gibbs state for any Hamiltonian at a temperature β^{-1} such that $\beta^{-1} = O(\|\mathbf{H}\|)$.

2.2 Mixing time from spectral profile

Theorem 2.4 demonstrates that bounded-degree Hamiltonians with non-local jumps exhibit fast mixing times. However, it leaves open the case of Hamiltonians with unbounded degrees, such as those with $d \leq \text{polylog}(n)$. More formally, we define

► **Definition 2.5.** An *unbounded degree system* is a sequence of temperatures $\beta(n)^{-1}$ and Hamiltonians $\mathbf{H}(n)$ for which $\lim_{n \rightarrow \infty} \beta(n)\|\mathbf{H}(n)\| = \infty$. However, we still assume that $\beta\|\mathbf{H}\| = \text{polylog}(n)$, polynomial in the number of qubits.

In our next result, we show that again selecting the jumping distribution (including adjoints) to be M samples from a unitary 1-design for sufficiently large M and choosing the energy resolution $\sigma_E = \beta^{-1}$ yields an algorithm whose efficiency depends on its low-energy spectrum. In particular, the runtime scales inverse polynomially with the fraction of eigenvalues $\delta(n)$ of \mathbf{H} that are within $O(\beta^{-1})$ of the minimum eigenvalue.

► **Theorem 2.6** (Spectral gap of Lindbladian in unbounded degree systems). *Let $\beta(n)^{-1}$ be a sequence of temperatures and $\mathbf{H}(n)$ a sequence of $n \times n$ Hamiltonians, and let $\delta(n)$ be the fraction of eigenvalues of $\mathbf{H}(n)$ within $O(\beta^{-1})$ of λ_{\min} .*

With any constant probability $1 - \xi$, the spectral gap of a Lindbladian \mathcal{L}_β with $\sigma_E = \beta^{-1}$ and M jump operators sampled from a 1-design jumping distribution for some $M = \Theta(\delta(n)^{-2} \log(n) \log(\beta \|\mathbf{H}\|))$, is lower bounded by $\Omega(\delta(n))$.

Assume access to an efficient block-encoding of $\mathbf{H}(n)$. As a consequence, and in the same setup, the Gibbs state of \mathbf{H} can be prepared with error ϵ in trace distance, in time $\text{poly}(\delta(n)^{-1}, \log(n), \log(\epsilon^{-1}))$.

Note that if $\beta \|\mathbf{H}\|$ is bounded, then $\delta(n) = 1$. This result therefore generalizes Theorem 2.4.

2.3 Explicit examples of random sparse Hamiltonians

Having established a sufficient spectral condition for the fast mixing of random ensembles of sparse Hamiltonians, we now give explicit examples that satisfy this criterion. We also give one example, the hypercube, which does not, and for which local jumps in place of unitary design jumps achieve an exponential speedup. This example elucidates the potential of structured local jumps for speedups, in contrast to the case of the cyclic graph in which structured graph-local jumps yielded a slowdown.

Random regular graphs. The first example is when \mathbf{H} is the adjacency matrix of a randomly selected $\log(n)$ -regular graph, with $\text{polylog}(n)$ random 1-design jumps. In Section 4.4.1, we prove using Theorem 2.6 that this ensemble has a Lindbladian spectral gap of $\Omega(\log(n)^{-3/4})$ at constant temperature. This yields a polynomial algorithm to prepare the Gibbs state, given access to an efficient block-encoding of \mathbf{H} .

Random signed Pauli ensemble. The second example is the family of sparse Hamiltonians considered in [9], composed of random Pauli strings with random sign coefficients given by $\mathbf{H}_{PS} = \sum_{j=1}^m \frac{r_j}{\sqrt{m}} \sigma_j$, where σ is a random Pauli string on n_0 qubits (such that the size of Hamiltonian is $n \times n$ for $n = 2^{n_0}$), each r_j is sampled randomly from $\{-1, 1\}$, and $m = O(\frac{n_0^5}{\epsilon^4})$ for a parameter $1 \geq \epsilon \geq 2^{-o(n_0)}$. We show in Section 4.4.2 that the CKG Lindbladian has a spectral gap of $\Omega(\epsilon^{-3/2})$ when we choose $M = \tilde{O}(n_0^2 \epsilon^{-3})$ unitary 1-design jumps, inverse temperature $\beta = O(\epsilon^{-1})$, and $\sigma_E = \beta^{-1}$.

Hypercubes. The final example is the family of hypercubes. A hypercube with 2^d vertices and degree d can be interpreted as a Hamiltonian on d qubits $\sum_i \mathbf{X}_i$. At constant temperature, only an exponentially small fraction of eigenvalues lie near the minimum eigenvalue. As a result, the spectral profile implies a poor mixing time with unitary design jumps.

However, we show in Theorem 4.5 that by choosing local jumps $\frac{1}{\sqrt{d}} \mathbf{Z}_i$, the spectral gap is $\frac{1}{d}$, yielding an efficient algorithm for Gibbs sampling. Hypercubes therefore provide an example where not graph-local jumps, but rather local jumps on the qubits, ensure fast mixing. The crucial feature of local jumps that improves mixing time is that a local jump \mathbf{A}^a on a local Hamiltonian \mathbf{H} satisfies $\|[\mathbf{A}^a, \mathbf{H}]\| = O(1)$ – i.e., the jump operator only jumps between nearby eigenstates. This property is not held by graph-local jumps in general, so they displayed no improvement in the studied cases. In general, local jumps are the strongest candidates for fast-mixing on local Hamiltonians, though for which classes of local Hamiltonians fast-mixing can be achieved is still largely open. For most interesting classes of local Hamiltonians, the condition of Theorem 2.6 is unlikely to hold.

3 Proof sketch

3.1 Graph-local jumps

The proof of the mixing time for graph-local jumps in the cyclic graph involves two steps: First, in Appendix A, we derive a general expression for the CKG Lindbladian in the energy basis given by equation (5).

We then utilize this expression along with the fully known spectrum of the cyclic graph to show that the Lindbladian is block-diagonal in the energy basis of this graph, as demonstrated in Appendix B. One of these blocks corresponds to a classical Markov chain on the diagonal entries of the state in the energy basis, for which we establish a spectral gap lower bound using the canonical path method. For the remaining $n - 1$ blocks, we apply the Gershgorin circle theorem to bound their eigenvalues.

3.2 Unitary design jumps

Bounded-degree systems. To establish a lower bound on the spectral gap of the Lindbladian \mathcal{L}_β with unitary 1-design jumps, we consider a decomposition of the form $\mathcal{L}_\beta = \mathcal{L}_\mu + \delta\mathcal{L}$, where $\mathcal{L}_\mu = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}(U(n))}[\mathcal{L}_\beta]$ is the expected Lindbladian with the expectation taken over a single jump operator sampled from a unitary 1-design distribution $\mathcal{D}(U(n))$, and $\delta\mathcal{L}$ represents the remainder term. Due to the quadratic form of the Lindbladian given in expression (3) we see that $\mathbb{E}_{\mathbf{A} \sim \mathcal{H}(U(n))}[\mathcal{L}_\beta] = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}(U(n))}[\mathcal{L}_\beta]$. Here, $\mathcal{H}(U(n))$ is a Haar random distribution over jump operators.

Note that a CKG Lindbladian must have jump operators in adjoint pairs, so a Lindbladian with a single jump operator will not satisfy detailed balance. However, the expected Lindbladian over one Haar random jump operator is equal to the expected Lindbladian over the adjoint pair of a Haar random jump operator, by linearity of expectation.

The proof of Theorem 2.4 proceeds by first showing that this expected Lindbladian $\mathbb{E}_{\mathbf{A} \sim \mathcal{H}(U(n))}[\mathcal{L}_\beta]$ has a constant spectral gap, as long as $\beta(n)\|\mathbf{H}(n)\|$ is bounded by a constant as a function of system size. As before, we call such systems *bounded degree*, since for constant β bounded degree graphs satisfy the required property. Indeed, if such a system has degree bounded by d , it must have spectrum in $[-d, d]$ by Gershgorin's circle theorem, since every row consists of zeros and at most d ones. Adding phases to the edges of these bounded degree graphs remains feasible by a similar argument, so there is no constraint of stoquasticity.

The result is stated formally as follows:

► **Lemma 3.1** (Constant spectral gap of average Lindbladian for bounded degree systems). *Let $\beta(n)^{-1}$ be a sequence of temperatures and $\mathbf{H}(n)$ be a sequence of $n \times n$ Hamiltonians such that $\beta\|\mathbf{H}\| = O(1)$. The spectral gap of $\mathcal{L}_\mu = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}(U(n))}[\mathcal{L}_\beta]$, the expected CKG Lindbladian with energy resolution $\sigma_E = \beta^{-1}$ over the ensemble of one jump operator sampled from a unitary 1-design, is asymptotically $\Omega(1)$.*

To establish Lemma 3.1, we show that for any system, the average Lindbladian over a Haar random ensemble of jump operators decomposes as $\mathcal{L}_\mu = \mathcal{L}_{\text{classical}} + \mathcal{L}_{\text{dephasing}}$. For a density matrix in the energy basis, the evolution of $\mathcal{L}_{\text{classical}}$ is a classical continuous Markov chain of the diagonal. The evolution of this classical Lindbladian maps the diagonal, in the limit, to the Gibbs distribution. The spectral gap of $\mathcal{L}_{\text{classical}}$ can be analyzed with the large suite of techniques for classical Markov chains.

Meanwhile, $\mathcal{L}_{\text{dephasing}}$ damps the off-diagonal terms of the density matrix. In the limit as $t \rightarrow \infty$, the state therefore converges to a classical distribution on the diagonal in the energy basis, with no off-diagonal terms, as desired. The operator $\mathcal{L}_{\text{dephasing}}$ diagonalizes in the

energy basis of density matrices, with each off-diagonal element decaying at an independent rate. It is therefore simple to analyze as well. In summary, Lemma 3.1 establishes that the $\lambda_{\text{gap}}(\mathcal{L}_\mu) = \min(\lambda_{\text{gap}}(\mathcal{L}_{\text{classical}}), \lambda_{\min}(\mathcal{L}_{\text{dephasing}})) = \Omega(1)$.

However, this result does not imply that any given jump sampled from the unitary 1-design (with its adjoint pair) would yield a gapped Lindbladian. As a result, it does not yet yield an efficient Gibbs sampling algorithm. To obtain such a result in Theorem 2.4, we demonstrate that the remainder term $\delta\mathcal{L} = \mathcal{L}_\beta - \mathcal{L}_\mu$ has a small spectral norm when a Lindbladian is constructed from a sufficiently large number of jumps M , rather than just one. In particular, a Lindbladian sampled with $\Theta(\log(n))$ normalized jumps from any 1-design concentrates closely to its expectation, thereby establishing a spectral gap lower bound. Since this lower bound applies to any graph at constant temperature β^{-1} with bounded degree, it applies to the periodic lattices, path graphs, and k -regular graphs discussed in the previous section.

Unbounded degree systems. In the context of unbounded degree systems, 1-design unitaries can no longer, in general, achieve an algorithm that is efficient in $\log(n)$ at constant temperature. Indeed, $\mathcal{L}_\mu = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}(U(n))} [\mathcal{L}_\beta] = \mathcal{L}_{\text{classical}} + \mathcal{L}_{\text{dephasing}}$ does not necessarily have a constant spectral gap in general, as it did in the case of bounded degree systems. However, we may establish a condition on the spectrum of \mathbf{H} , with which we can recover a lower bound for the spectral gap:

► **Lemma 3.2** (Spectral gap of average Lindbladian for unbounded systems). *Let $\mathbf{H}(n)$ be a sequence of $n \times n$ Hamiltonians. For some C , let $\delta(n)$ be the proportion of eigenvalues λ_j of \mathbf{H} such that $\beta^{-1}(\lambda_j - \lambda_{\min}) \leq C$. The spectral gap of $\mathcal{L}_\mu = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}(U(n))} [\mathcal{L}_\beta]$, the expected CKG Lindbladian over the Haar random unitary ensemble of its jump operator at temperature β^{-1} with $\sigma_E = \Theta(\beta^{-1})$, is asymptotically $\Omega(\delta(n))$.*

The lemma expresses that if λ_{\min} is within $O(\beta^{-1})$ of $\delta(n)$ of the eigenvalues, the spectral gap is at least $\delta(n)$. Similarly to Theorem 2.4, using this result to obtain an efficient Gibbs sampling algorithm amounts to showing that a Lindbladian with enough independently sampled jump operators shares a similar asymptotic spectral gap to the average Lindbladian, using a concentration bound. When the average spectral gap from the above lemma is $\delta(n)$, the number of jump operators to concentrate around the expectation increases to $\delta(n)^{-2}$, along with an overhead of $\log(n) \log(\beta \|\mathbf{H}\|)^2$. This result is captured in Theorem 2.6, and results in an algorithm with runtime $\text{poly}(\delta(n)^{-1}, \log(n), \log(\epsilon^{-1}))$ for Gibbs sampling, where ϵ is the error in trace distance. This runtime bound relies on the standing assumption in this paper that $\log(\beta \|\mathbf{H}\|) = \text{poly}(\log(n))$.

4 Technical details

4.1 The quantum Gibbs sampler

4.1.1 Lindbladian evolution

The recently proposed CKG quantum MCMC algorithm addresses the problem of finding thermal states by imitating thermodynamic processes [11, 10]. In this process, a system of particles evolves in contact with a thermal bath at some fixed temperature β^{-1} . Due to interactions with the bath, the system is described by a probabilistic mixture of quantum states ρ . This state evolves in time, by approximation, with Markovian *dissipative* dynamics, $\frac{d\rho}{dt} = \mathcal{L}_\beta[\rho]$, given in terms of an operator \mathcal{L}_β known as the Lindbladian. This operator

involves a coherent term \mathbf{B} that describes the interaction among the particles in the system. There is also a term in \mathcal{L}_β that is specified by a series of *Lindblad operators* \mathbf{L}^j that drive the dissipative transitions. The dynamics of the coherent term \mathbf{B} are reversible, while the dissipative transitions drive all states toward some “stationary state”. These transitions can be understood as perturbations from the bath, and as all states converge to the Gibbs state, the information of the system is leaking via these perturbations to the bath. The expression, in terms of \mathbf{B} and \mathbf{L}^j , is:

$$\mathcal{L}_\beta[\cdot] = -i[\mathbf{B}, \cdot] + \sum_j \left(\mathbf{L}^j(\cdot)\mathbf{L}^{j\dagger} - \frac{1}{2}\{\mathbf{L}^{j\dagger}\mathbf{L}^j, \cdot\} \right).$$

The summands $\mathbf{L}^j(\cdot)\mathbf{L}^{j\dagger}$ are termed the transition part of the Lindbladian, and $-\frac{1}{2}\{\mathbf{L}^{j\dagger}\mathbf{L}^j, \cdot\}$ are the decay part of the Lindbladian. The choice of the Lindbladian operator \mathcal{L}_β can vary depending on the precise nature of interactions between the system and the bath. However, to prepare the Gibbs (thermal) state at temperature β^{-1} , the Lindbladian should satisfy

$$\frac{d\rho_\beta}{dt} = \mathcal{L}_\beta[\rho_\beta] = 0 \quad \text{where} \quad \rho_\beta := e^{-\beta\mathbf{H}}/\text{Tr}(e^{-\beta\mathbf{H}}), \quad (2)$$

and moreover ρ_β should be the unique stationary state of the Lindbladian. The long-term evolution of the system under this Lindbladian, as a result, would converge to the Gibbs state of the Hamiltonian \mathbf{H} at temperature $1/\beta$.

4.1.2 Detailed balance

To ensure that the Lindbladian \mathcal{L}_β converges to a state ρ_β , [11] designs a Lindbladian that satisfies Kubo-Martin-Schwinger (KMS) detailed balance with respect to ρ_β . KMS detailed balance is one of several ways of quantizing the notion of classical detailed balance for Markov chains. KMS detailed balance of \mathcal{L}_β is self-adjointness with respect to the inner product

$$\langle \sigma_1, \sigma_2 \rangle_{\rho_\beta^{-1}} = \text{Tr}(\sigma_1^\dagger \rho_\beta^{-1/2} \sigma_2 \rho_\beta^{-1/2}). \quad (\text{KMS Inner Product})$$

In particular, it is equivalent to the relation that

$$\mathcal{L}_\beta[\cdot] = \rho_\beta^{1/2} \mathcal{L}_\beta^\dagger \left[\rho_\beta^{-1/2}(\cdot) \rho_\beta^{-1/2} \right] \rho_\beta^{1/2} \quad (\text{Detailed Balance})$$

where $\mathcal{L}_\beta^\dagger$ is the adjoint Lindbladian with respect to the Hilbert-Schmidt inner product $\langle \sigma_1, \sigma_2 \rangle = \text{Tr}(\sigma_1^\dagger \sigma_2)$. The adjoint operator $\mathcal{L}_\beta^\dagger$, in the Heisenberg picture, describes the dynamics of observables under evolution by \mathcal{L}_β . The Lindbladian evolution is described by some quantum channel and therefore the observable I must always be fixed by $\exp(\mathcal{L}_\beta^\dagger)$. This implies that $\mathcal{L}_\beta^\dagger[I] = 0$. The detailed balance formula thereby implies that $\mathcal{L}_\beta[\rho_\beta] = 0$, as desired. Note that KMS detailed balance can be dually described as the self-adjointness of $\mathcal{L}_\beta^\dagger$ with respect to the inner product $\langle \sigma_1, \sigma_2 \rangle_{\rho_\beta} = \text{Tr}(\sigma_1^\dagger \rho_\beta^{1/2} \sigma_2 \rho_\beta^{1/2})$.

4.1.3 Construction and parameters

The quantum Gibbs sampler in [11] constructs a Lindbladian that satisfies two properties:

1. \mathcal{L}_β satisfies detailed balance with respect to ρ_β , and therefore $\mathcal{L}_\beta[\rho_\beta] = 0$.
2. The dynamics of \mathcal{L}_β can be efficiently implemented.

Their Lindbladian, which we term the CKG Lindbladian, can be simulated on a quantum computer with a cost per unit time $t = 1$ roughly equal to that of simulating the Hamiltonian dynamics of \mathbf{H} . The CKG Lindbladian is closely related to the Davies generator, which is a physically motivated Lindbladian that satisfies detailed balance, but that is not efficiently implementable in general. A full description of both Lindbladians are given in the Appendix.

The Gibbs sampling algorithm evolves an initial state ρ_0 according to the efficiently implemented Lindbladian \mathcal{L}_β , and produces the state $\rho_t = e^{\mathcal{L}_\beta t}[\rho_0]$ after time period t . The *mixing time* is roughly the time that it takes for the state ρ_t to approach the Gibbs state ρ_β . That is, $e^{\mathcal{L}_\beta t_{\text{mix}}}[\rho_0] \approx \rho_\beta$. The efficiency of the algorithm therefore scales linearly with the unit time simulation cost and the mixing time. The algorithm has several parameters in the Lindbladian's construction. In addition to the *inverse temperature* β , the algorithm specifies an *energy resolution* σ_E . A salient feature of [11]'s construction is that it can achieve detailed balance even though the algorithm only probes the energies of the Hamiltonian \mathbf{H} with approximate precision. σ_E quantifies this level of precision. The cost of the Lindbladian simulation depends linearly on σ_E^{-1} , but increasing the precision may also improve the mixing time. Taking $\sigma_E \rightarrow 0$ for absolute precision recovers the Davies generator – when distinguishing the energies of the system exactly is infeasible, this Lindbladian cannot be simulated efficiently.

A set of *jump operators* \mathbf{A}^a must also be specified for the Lindbladian. These operators are decomposed by frequency and reassembled in a particular way to construct the Lindblad operators that help \mathcal{L}_β satisfy detailed balance. They must appear in adjoint pairs: i.e., if $\mathbf{A} \in \{\mathbf{A}^a\}$, then $\mathbf{A}^\dagger \in \{\mathbf{A}^a\}$. The cost of simulation scales with the cost of implementing the oracle $|a\rangle \rightarrow |a\rangle \otimes \mathbf{A}^a$. In particular, the jump operators must be normalized when implemented for the algorithm, satisfying $\sum_a \|\mathbf{A}^{a\dagger} \mathbf{A}^a\| \leq 1$. CKG Lindbladians are linear in their jump operators – if \mathcal{L}_1 has one jump operator \mathbf{A}^1 and \mathcal{L}_2 has one jump operator \mathbf{A}^2 , then a Lindbladian \mathcal{L} with jump operators \mathbf{A}^1 and \mathbf{A}^2 satisfies $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2$. If \mathcal{L}_β was constructed from jumps \mathbf{A}^a , then jump operators $\sqrt{s}\mathbf{A}^a$ produce the Lindbladian $s\mathcal{L}_\beta$, scaling the mixing time by s . So we may therefore assume that $\sum_a \|\mathbf{A}^{a\dagger} \mathbf{A}^a\| = 1$ exactly, since renormalizing can only improve the spectral gap. In its normalized form, the set of jump operators can be understood as a jumping distribution over \mathbf{A}^a which we will notate $a \sim \mathcal{A}$, where each is sampled with probability $\|\mathbf{A}^{a\dagger} \mathbf{A}^a\|$.

4.2 Mathematical description

We begin with a description of the Davies generator, which is the limit of the CKG Lindbladian as $\sigma_E \rightarrow 0$. This generator was developed from a physical approximation of an open thermalizing quantum system, but at low temperatures it is unphysical and can be hard to implement. We then generalize the notions to the implementable CKG Lindbladian.

4.2.1 Davies generator

In the description of the Davies generator for a given system \mathbf{H} , there is a coherent term and there are jump operators \mathbf{A}^a . The \mathbf{A}^a terms must appear in adjoint pairs in the Davies generator. The dissipative part of the Lindbladian is expressed as follows:

$$\mathcal{L}_\beta[\cdot] = \sum_{a \in [M]} \int_{-\infty}^{\infty} \gamma(\omega) \left(\mathbf{A}_\omega^a(\cdot) \mathbf{A}_\omega^\dagger - \frac{1}{2} \{ \mathbf{A}_\omega^\dagger \mathbf{A}_\omega, \cdot \} \right) d\omega, \quad (3)$$

where \mathbf{A}_ω^a is the Operator Fourier Transform (OFT) of jump operator \mathbf{A}^a :

$$\mathbf{A}_\omega^a = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\mathbf{H}t} \mathbf{A}^a e^{-i\mathbf{H}t} e^{-i\omega t} dt.$$

The Davies' generator chooses Lindblad operators $\sqrt{\gamma(\omega)}\mathbf{A}_\omega^a$, each of which selects the energy transitions, or Bohr frequencies, in \mathbf{A}^a that are precisely ω . Because it requires certainty in energy, by Heisenberg's uncertainty principle of energy and time, in the general case simulating the evolution of the Davies generator efficiently is infeasible. In the above, γ is some function satisfying $\gamma(\omega) = \gamma(\omega) = e^{-\beta\omega}\gamma(-\omega)$. The Lindblad operators are scaled by $\gamma(\omega)$ precisely to satisfy KMS detailed balance. Since \mathbf{A}_ω^a represents jumps with Bohr frequency ω , the functional equation of γ ensures a desired ratio of jumps with Bohr frequency ω and $-\omega$. We choose the Metropolis filter, $\gamma(\omega) = \min(1, e^{-\beta\omega})$, though another common filter $\gamma(\omega) = \frac{1}{1+e^{-\beta\omega}}$ for "Glauber dynamics" could also be used for the same results.

The Davies generator satisfies detailed balance with respect to ρ_β , the thermal state. In some presentations of the Davies generator, it contains a coherent term $-i[\mathbf{H}, \cdot]$. If this term is included, the generator does not satisfy detailed balance, so we do not follow this convention. However, the term does not affect the fixed point of the generator, since ρ_β commutes with \mathbf{H} and therefore $-i[\mathbf{H}, \rho_\beta] = 0$.

4.2.2 CKG Lindbladian

The CKG Lindbladian is defined almost identically to the Davies generator, but is altered slightly so that it still obeys detailed balance, but is efficiently implementable.

$$\mathcal{L}_\beta[\cdot] = \underbrace{-i[\mathbf{B}, \cdot]}_{\text{coherent term}} + \sum_{a \in [M]} \int_{-\infty}^{\infty} \gamma(\omega) \left(\underbrace{\hat{\mathbf{A}}^a(\omega)(\cdot)\hat{\mathbf{A}}^a(\omega)^\dagger}_{\text{transition term}} - \underbrace{\frac{1}{2}\{\hat{\mathbf{A}}^a(\omega)^\dagger\hat{\mathbf{A}}^a(\omega), \cdot\}}_{\text{decay term}} \right) d\omega,$$

where $\hat{\mathbf{A}}^a(\omega)$ is now the Gaussian-supported OFT of jump operator \mathbf{A}^a :

$$\hat{\mathbf{A}}^a(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\mathbf{H}t} \mathbf{A}^a e^{-i\mathbf{H}t} e^{-i\omega t} f(t) dt.$$

To ensure that the jump operators do not have infinite precision in energy, a Gaussian supported OFT is performed instead to obtain $\hat{\mathbf{A}}^a(\omega)$, which selects a Gaussian band energies of around ω .

Here, $f(t) = e^{-\sigma_E^2 t^2} \sqrt{\sigma_E \sqrt{2/\pi}}$, with Fourier transform $\hat{f}(\omega) = \frac{1}{\sqrt{\sigma_E \sqrt{2\pi}}} \exp(-\frac{\omega^2}{4\sigma_E^2})$. As a result, the operator $\hat{\mathbf{A}}^a(\omega)$ can be shown to be equal to $\sum_\nu \hat{f}(\omega - \nu) \mathbf{A}_\nu^a$. The function $f(t)$ was chosen so that its squared Fourier transform $\hat{f}^2(\omega)$ is a Gaussian with standard deviation σ_E , which features prominently in the Lindbladian (since it consists of quadratic terms in $\hat{\mathbf{A}}^a(\omega)$). Taking $\sigma_E = \Theta(\beta^{-1})$ yields an efficient simulation algorithm with the assumption of a block-encoding of \mathbf{H} and a block-encoding for the jump operators $\sum_{a \in [M]} |a\rangle \otimes \mathbf{A}^a$, so σ_E is taken to be on the order of β^{-1} in this paper.

Since $\hat{\mathbf{A}}^a(\omega)$ is a noisy decomposition of \mathbf{A}^a into frequencies, it is not immediately clear whether there is a choice of function $\gamma(\omega)$ for which they can be recombined to achieve detailed balance. Indeed, as shown in [11], there is! The choice of $\gamma(\omega)$ is such that the transition part of \mathcal{L}_β , the summand $\sum_{a \in [M]} \int_{-\infty}^{\infty} \gamma(\omega) \hat{\mathbf{A}}^a(\omega)(\cdot)\hat{\mathbf{A}}^a(\omega)^\dagger d\omega$, still satisfies KMS detailed balance. [11] proved that there is a unique choice of \mathbf{B} , up to translation by a scalar, such that $-i[\mathbf{B}, \cdot] - \frac{1}{2} \sum_{a \in [M]} \int_{-\infty}^{\infty} \gamma(\omega) \{\hat{\mathbf{A}}^a(\omega)^\dagger \hat{\mathbf{A}}^a(\omega), \cdot\} d\omega$ also satisfies detailed balance. For the Davies generator, this coherent term \mathbf{B} is simply 0 (or corresponds to a Lamb shift that commutes with the Hamiltonian), and the decay term by itself already satisfies detailed balance. \mathbf{B} can be expressed in general as:

$$\mathbf{B} = \sum_{a \in [M]} \sum_{\nu_1, \nu_2} \frac{\tanh(-\beta(\nu_1 - \nu_2)/4)}{2i} (\mathbf{A}_{\nu_2}^a)^\dagger \mathbf{A}_{\nu_1}^a.$$

The choice of γ for this algorithm, for which the filter is efficiently implementable, is $\gamma(\omega) = \exp\left(-\beta \max\left(\omega + \frac{\beta\sigma_E^2}{2}, 0\right)\right)$. As $\sigma_E \rightarrow 0$ it converges to the Metropolis filter of the Davies generator. In particular, this γ is precisely Metropolis filter for the Davies generator shifted by $\beta\sigma_E^2$. The CKG Lindbladian requires a choice of γ for which $\alpha = \gamma * g$ satisfies the functional equation that was originally satisfied by γ in the Davies generator.

We also note that \mathcal{L}_β is bounded in operator norm.

► **Lemma 4.1.** *Consider the CKG Lindbladian \mathcal{L}_β with temperature β^{-1} , using the Metropolis filter, and with jump operators \mathbf{A}^a for which $\sum_a \|\mathbf{A}^{a\dagger} \mathbf{A}^a\| \leq 1$. This Lindbladian satisfies $\|\mathcal{L}_\beta\|_{\infty \rightarrow \infty} = O(\log(\beta \|\mathbf{H}\|))$, where $\|\cdot\|_{\infty \rightarrow \infty}$ is the operator norm of \mathcal{L}_β , with respect to the operator norm on the input and output vector spaces.*

Proof. The result follows from citing Proposition B.2 in [11] to bound the operator norm of the coherent term, and bounding the transition and decay terms manually. This proof is described in the arXiv version [32]. ◀

4.3 Spectral gap

Since the quantum MCMC algorithm was proposed recently, numerical and analytic characterizations of algorithm are limited. As for classical Markov chains, it has been shown that the mixing time of the algorithm can be characterized by the spectral gap $\lambda_{\text{gap}}(\mathcal{L}_\beta)$ of the Lindbladian. If the first eigenvalue $\lambda_1 = 0$ corresponds to eigenvector ρ_β , then the spectral gap is $\lambda_{\text{gap}}(\mathcal{L}_\beta) = \min_{j>1} |\lambda_j|$ [11]. Lindbladians are in general negative semidefinite like classical Markov chain generators, so $\lambda_{\text{gap}}(\mathcal{L}_\beta) = \min_j (-\lambda_j)$. More precisely, it holds that

$$\frac{\Omega(1)}{\lambda_{\text{gap}}(\mathcal{L}_\beta)} \leq t_{\text{mix}}(\mathcal{L}) \leq \frac{\log\left(\left\|\rho_\beta^{-1/2}\right\|\right)}{\lambda_{\text{gap}}(\mathcal{L}_\beta)} \leq \frac{\mathcal{O}(\beta \|\mathbf{H}\| + \log(\dim(\mathbf{H})))}{\lambda_{\text{gap}}(\mathcal{L}_\beta)}. \quad (4)$$

In particular, analytically bounding this spectral gap from below is sufficient to prove that ρ_β is the unique fixed point, and for obtaining an upper bound on the mixing time. For so-called rapid mixing, in which the mixing time is logarithmic in the number of qubits, the spectral gap bound often does not suffice. For our purposes of proving efficiency in the number of qubits, however, this issue is moot.

4.4 Unbounded degree systems

We now prove efficient Gibbs sampling results for certain unbounded sparse Hamiltonians.

4.4.1 Random $\log(n)$ -regular graphs

With high probability at constant temperature, a randomly selected $d = \log(n)$ -regular graph, with $\text{poly}(d)$ random 1-design jumps, has a Lindbladian spectral gap of $\Omega(d^{-3/4})$. This gives a polynomial algorithm to prepare the Gibbs state for most such graphs at constant temperature.

The gap of $\Omega(d^{-3/4})$ arises because a random d -regular graph, for $d \rightarrow \infty$, has one eigenvalue at d and the rest distributed from $-2\sqrt{d-1}$ to $2\sqrt{d-1}$ in a distribution that converges to a (normalized) semicircle. This semicircular distribution frequently appears in random matrix theory, for instance in the Gaussian unitary ensemble (GUE), which models the spectrum of many chaotic quantum systems. When the spectrum of a quantum system indeed follows this distribution, it implies that $\delta(n) = \Omega(d^{-3/4})$ of the eigenvalues lie within a constant of the minimum eigenvalue.

► **Theorem 4.2.** *With any constant probability $1 - \xi$, for a randomly selected $d = \log(n)$ -regular graph, there are $\delta = \Omega(d^{-3/4})$ eigenvalues within $O(1)$ of the minimum eigenvalue.*

As an immediate result of Theorem 4.2 and Theorem 2.6, we obtain an algorithm polynomial in d to prepare the Gibbs state of a d -regular graph. To prove the corollary, we use Theorem 2 in [17]. For this context, the following statement suffices.

► **Theorem 4.3.** *Let degree $d = \log(n)$. For sufficiently large n , there exists some $D > 0$ such that for any interval $I \subset \mathbb{R}$, $0 < \alpha < 1$, and $0 < \epsilon < \alpha$ such that $|I| > Dd^{-\alpha+\epsilon}$, with probability $1 - o(n^{-1})$ over all random d -regular graphs, $|\delta(n) - \mu| < d^{-\epsilon}|I|$, where $\mu = \int_I \rho_{sc}(x)dx$ and $\delta(n)$ is the fraction of eigenvalues of the d -regular graph in $I\sqrt{d-1}$.*

In the above, ρ_{sc} is the asymptotic distribution as $d \rightarrow \infty$ of a random d -regular graph is the semicircular distribution with radius $2\sqrt{d-1}$, $\rho_{sc}(x) = \frac{1}{2\pi(d-1)}\sqrt{4(d-1)-x^2}$. From this result, Theorem 4.2 follows. As described more explicitly in [32], the estimated density ρ_{sc} near the bottom of the semicircle can be estimated. Then, using the theorem, this can be related to the fraction of eigenvalues near the minimum eigenvalue as well, proving the result.

4.4.2 Pauli String Ensemble

We now mention another ensemble of Hamiltonians studied by [9] in the context of low-energy state preparation. In [9], efficient low energy state preparation with phase estimation is demonstrated under the same conditions as our efficient Gibbs sampling in Theorem 2.6. Indeed, if many eigenvectors are close to the ground-state energy, as we require, then performing phase estimation on the maximally mixed state has a high probability of measuring a low-energy state, so low-energy state preparation is possible as well. They study the following ensemble of Hamiltonians on n_0 qubits, $\mathbf{H}_{PS} = \sum_{j=1}^m \frac{r_j}{\sqrt{m}} \boldsymbol{\sigma}_j$, where $\boldsymbol{\sigma}$ is a random Pauli string on n_0 qubits, each r_j is sampled randomly from $\{-1, 1\}$, and $m = \left\lfloor c_2 \frac{n_0^5}{\epsilon^4} \right\rfloor$. The parameter ϵ satisfies $\epsilon \geq 2^{-n_0/c_1}$, and c_1, c_2 are absolute constants. The resulting spectrum is again close enough to a semicircular distribution to obtain an efficient Gibbs sampler for certain temperatures that depend on ϵ . As ϵ decreases, Gibbs sampling becomes efficient for even larger values of β (lower temperatures), since the ensemble's spectrum converges closer to a perfect semicircular distribution at the edge of the spectrum.

Using the results in their paper, we establish that Gibbs sampling is efficient in n_0 for certain values of ϵ and corresponding temperatures β^{-1} .

► **Theorem 4.4.** *Say that $\mathbf{H}(n_0)$ is sampled from the ensemble \mathbf{H}_{PS} on n_0 qubits, with $\epsilon = 2^{-o(n_0)}$ and $\epsilon \leq 1$. With any constant probability $1 - \xi$, for sufficiently large n_0 , $\delta = \Omega(\epsilon^{3/2})$ fraction of the eigenvalues lie within $O(\epsilon)$ of the minimum eigenvalue.*

Proof. We utilize two results from [9]. Firstly, they argue that $\Pr[\|\mathbf{H}(n_0)\| \geq 2(1 + \epsilon)] \leq \exp(-c_2 n_0)$ when $m \geq \frac{n_0^3}{\epsilon^4}$, which is satisfied in this case. With an arbitrary constant probability for sufficiently large n_0 , therefore, $\|\mathbf{H}(n_0)\| \leq 4$, since $\epsilon \leq 1$. The second result is that with probability $1 - \exp(-c_3 n_0^{1/3})$, at least $\Omega(\epsilon^{3/2})$ of the eigenvalues satisfy $\lambda_i \leq (1 - \epsilon)\lambda_{\min}$ where c_3 is an absolute constant. With any large constant probability, we therefore have that $|\lambda_i - \lambda_{\min}| \leq \epsilon\lambda_{\min} \leq 4\epsilon = O(\epsilon)$ for $\Omega(\epsilon^{3/2})$ of the eigenvalues. ◀

By Theorem 2.6, we obtain a Gibbs sampling algorithm that is $\text{poly}(\epsilon^{-1}, n_0)$ to prepare the Gibbs state at inverse temperature ϵ^{-1} . We may rephrase this result in terms of β . For any polynomially large β , it provides a Pauli string ensemble of Hamiltonians, \mathbf{H}_{PS} with $\epsilon = \beta^{-1}$, for which with high likelihood preparing the Gibbs state is efficient in n_0 , assuming access to a block-encoding of the Hamiltonian of interest.

4.4.3 Hypercube graphs

For hypercube with varying dimension at a constant temperature, using unitary 1-design jumps would yield an exponentially large runtime. The spectrum of a hypercube with dimension d and 2^d vertices consists of the integers $-d, -d+2, \dots, d-2, d$. The eigenvalue j has multiplicity $\binom{d}{\frac{d+j}{2}}$. In particular, for any constant C , only an exponentially small fraction of the eigenvalues $\delta(d)$ lie below $-d + C$. This leads to a naive algorithm with at worst exponential complexity in d .

However, a better result can be obtained considering the hypercube as a system of d qubits. The graph with dimension d has 2^d vertices, which can be considered length d bitstrings. Then, the adjacency matrix is the sum of Pauli X operators on each qubit, $\sum_{i=1}^d X_i$, since the hypercube has an edge between any two bitstrings of Hamming distance 1. Choosing d jump operators as $\frac{1}{\sqrt{d}}Z_a$, the mixing time can be improved to $\text{poly}(d, \log(\epsilon^{-1}))$:

► **Theorem 4.5** (Spectral Gap for Hypercube with Local Jumps). *For fixed β^{-1} , there exists some energy resolution σ_E such that the spectral gap of the CKG Lindbladian \mathcal{L}_β for a d -dimensional hypercube with jump operators $\mathbf{A}^a = \frac{1}{\sqrt{d}}Z_a$, is asymptotically $\Omega(d^{-1})$.*

Proof. The proof of this statement is given in the arXiv version [32] by showing that the Lindbladian, with this choice of jump operators, is the product of independent Lindbladians on each qubit. Each of their spectral gaps can then be calculated explicitly. ◀

In the case of the hypercube, the local jump operators $\frac{1}{\sqrt{d}}Z_i$ only jump between eigenstates whose eigenvalues differ by 1. This vastly improves the performance of the classical Markov chain and dephasing Markov chain within the Lindbladian. However, the Lindbladian does not consist only of these two terms, as it did in the limit of independently sampled 1-design jumps. Off-diagonal terms do exist, and the presence of Z_a for *every* index is necessary to ensure that these off-diagonal terms do not completely eliminate the spectral gap. In some way, there must be “enough uncorrelated” local energy jumps to dampen these off-diagonal terms. For more complicated local Hamiltonians, it is not clear how correlations may be suppressed while still maintaining locality.

5 Connection to previous work

Our results show that for a Hamiltonian \mathbf{H} with temperature β^{-1} such that some $\delta(n)$ fraction of the eigenstates are within $O(\beta^{-1})$ of the ground-state energy, the CKG quantum Gibbs sampler with 1-design jumps efficiently prepares the Gibbs state with trace distance at most ϵ . The running time scales polynomially with $\delta(n)^{-1}$, $\beta\|\mathbf{H}\|$, $\log(\epsilon^{-1})$, and the complexity of the block encoding of \mathbf{H} . This result is a baseline test that shows the CKG algorithm performs as well as other methods for preparing low-energy states of Hamiltonians. Indeed, our spectral condition is precisely the same as a condition that ensures easy quantum phase estimation of a near-ground state. Namely, performing quantum phase estimation on the maximally mixed state can prepare a random eigenstate, and with probability $\delta(n)$ it is within $O(\beta^{-1})$ of the minimum eigenvalue. Obtaining $O(\delta(n)^{-1})$ samples and taking the minimum energy can therefore prepare a near ground-state eigenvector. This approach is the basis of the previous analysis of random sparse Hamiltonians in [9].

Moreover, in [14], a quantum algorithm is presented that prepares the Gibbs state with a complexity that scales as $\text{poly}(\frac{n}{\delta(n)}, \log(\epsilon^{-1}))$. If $\delta(n)$ of the eigenstates are within $O(\beta^{-1})$ of the ground-state energy, then $\frac{n}{\delta(n)} = \Omega(\delta(n)^{-1})$, and therefore under such conditions, this algorithm efficiently prepares the Gibbs states as well. Effectively, the CKG Gibbs sampler

with “generic” 1-design jumps performs the same as previously developed algorithms – the algorithm is at least as powerful, but a potential advantage in cooling must arise from a smart (i.e., local and unbiased) choice of jump operators.

After completing our work, we also became aware of [12], where, among other contributions, the authors derive a new, efficient quantum Gibbs sampler algorithm that utilizes jump operators sampled from a Clifford-random circuit. This Gibbs sampler is shown to exhibit a spectral gap bound under the same condition on the spectral density considered in Theorem 2.6. In comparison, we show that under the conditions of Theorem 2.6, the spectral gap of the CKG Lindbladian with an ensemble of 1-design jumps is bounded with high probability.

Finally, our conditions on the spectrum and the structure of random unitary design jumps resemble previous works on chaotic Hamiltonians that apply the Eigenstate Thermalization Hypothesis (ETH) to prove the fast mixing of dissipative dynamics [8, 15]. In particular, in [8], the proposed algorithm implements a “rounded” Davies generator, yielding a physical Lindbladian that block-diagonalizes into components consisting of small-energy transitions. They propose their own version of ETH that relies on jump operators, for small Bohr frequencies ω , having independent Gaussian-distributed entries. The assumption that these entries are independent for the result is very strong, allowing them to conclude that their jump operators are both local *and* that distinct energy transitions are completely uncorrelated.

Our work shows fast mixing unconditionally for quantumly easy Hamiltonians, replacing the local jumps and ETH assumption for the rounded Davies generator with 1-design jump operators for the CKG Lindbladian. A similar ETH assumption to [8] would also yield fast-mixing for the CKG Lindbladian with local jumps, but more generally some approach must be taken to provably mitigate the correlations induced by implementing local jumps, in contrast with 1-design jump operators.

References

- 1 Eric R Anshuetz, Chi-Fang Chen, Bobak T Kiani, and Robbie King. Strongly interacting fermions are non-trivial yet non-glassy. *arXiv preprint arXiv:2408.15699*, 2024. doi:10.48550/arXiv.2408.15699.
- 2 Joran van Apeldoorn and András Gilyén. Improvements in quantum SDP-solving with applications. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 99:1–99:15, 2019. arXiv: 1804.05058 doi:10.4230/LIPIcs.ICALP.2019.99.
- 3 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. Earlier version in FOCS’17. arXiv: 1705.01843 doi:10.22331/q-2020-02-14-230.
- 4 Christian W. Bauer, Zohreh Davoudi, A. Baha Balantekin, Tanmoy Bhattacharya, Marcela Carena, Wibe A. de Jong, Patrick Draper, Aida El-Khadra, Nate Gemelke, Masanori Hanada, Dmitri Kharzeev, Henry Lamm, Ying-Ying Li, Junyu Liu, Mikhail Lukin, Yannick Meurice, Christopher Monroe, Benjamin Nachman, Guido Pagano, John Preskill, Enrico Rinaldi, Alessandro Roggero, David I. Santiago, Martin J. Savage, Irfan Siddiqi, George Siopsis, David Van Zanten, Nathan Wiebe, Yukari Yamauchi, Kübra Yeter-Aydeniz, and Silvia Zorzetti. Quantum simulation for high-energy physics. *PRX Quantum*, 4:027001, May 2023. doi:10.1103/PRXQuantum.4.027001.
- 5 Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114(9):090502, 2015. arXiv: 1412.4687 doi:10.1103/PhysRevLett.114.090502.

- 6 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 27:1–27:14, 2019. arXiv: 1710.02581 doi: 10.4230/LIPIcs.ICALP.2019.27.
- 7 Fernando G. S. L. Brandão and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 415–426, 2017. arXiv: 1609.05537 doi:10.1109/FOCS.2017.45.
- 8 Chi-Fang Chen and Fernando G. S. L. Brandão. Fast thermalization from the eigenstate thermalization hypothesis, 2021. doi:10.48550/arXiv.2112.07646.
- 9 Chi-Fang Chen, Alexander M. Dalzell, Mario Berta, Fernando G. S. L. Brandão, and Joel A. Tropp. Sparse random Hamiltonians are quantumly easy. *Phys. Rev. X*, 14:011014, February 2024. doi:10.1103/PhysRevX.14.011014.
- 10 Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation. arXiv: 2303.18224, 2023.
- 11 Chi-Fang Chen, Michael J. Kastoryano, and András Gilyén. An efficient and exact non-commutative quantum Gibbs sampler. *arXiv preprint arXiv:2311.09207*, 2023. doi: 10.48550/arXiv.2311.09207.
- 12 Hongrui Chen, Bowen Li, Jianfeng Lu, and Lexing Ying. A randomized method for simulating Lindblad equations and thermal state preparation. *arXiv preprint arXiv:2407.06594*, 2024. doi:10.48550/arXiv.2407.06594.
- 13 Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the 35th ACM Symposium on the Theory of Computing (STOC)*, pages 59–68, 2003. arXiv: quant-ph/0209131 doi:10.1145/780542.780552.
- 14 Anirban Narayan Chowdhury and Rolando D. Somma. Quantum algorithms for Gibbs sampling and hitting-time estimation. *Quantum Information and Computation*, 17(1&2):41–64, 2017. arXiv: 1603.02940 doi:10.26421/QIC17.1–2.
- 15 Zhiyan Ding, Chi-Fang Chen, and Lin Lin. Single-ancilla ground state preparation via Lindbladians. *Phys. Rev. Res.*, 6:033147, August 2024. doi:10.1103/PhysRevResearch.6.033147.
- 16 Zhiyan Ding, Bowen Li, and Lin Lin. Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition. *arXiv preprint arXiv:2404.05998*, 2024. doi:10.48550/arXiv.2404.05998.
- 17 Ioana Dumitriu and Soumik Pal. Sparse regular random graphs: Spectral density and eigenvectors. *The Annals of Probability*, 40(5), September 2012. doi:10.1214/11-aop673.
- 18 Ernesto Estrada and Juan A. Rodríguez-Velázquez. Subgraph centrality in complex networks. *Phys. Rev. E*, 71:056103, May 2005. doi:10.1103/PhysRevE.71.056103.
- 19 András Gilyén, Chi-Fang Chen, Joao F. Doriguello, and Michael J. Kastoryano. Quantum generalizations of Glauber and Metropolis dynamics. *arXiv preprint arXiv:2405.20322*, 2024. doi:10.48550/arXiv.2405.20322.
- 20 András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics [full version], 2018. arXiv: 1806.01838
- 21 Fernando G.S. L. Brandão, Richard Kueng, and Daniel Stilck França. Faster quantum and classical SDP approximations for quadratic binary optimization. *Quantum*, 6:625, January 2022. doi:10.22331/q-2022-01-20-625.
- 22 Jeongwan Haah, Matthew B. Hastings, Robin Kothari, and Guang Hao Low. Quantum algorithm for simulating real time evolution of lattice Hamiltonians. In *Proceedings of the 59th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 350–360, 2018. arXiv: 1801.03922 doi:10.1109/FOCS.2018.00041.

- 23 Jiaqing Jiang and Sandy Irani. Quantum Metropolis sampling via weak measurement. *arXiv preprint arXiv:2406.16023*, 2024. doi:10.48550/arXiv.2406.16023.
- 24 Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local Hamiltonian problem. In Kamal Lodaya and Meena Mahajan, editors, *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science*, pages 372–383, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. doi:10.48550/arXiv.quant-ph/0406180.
- 25 A Kitaev. Hidden correlations in the Hawking radiation and thermal noise, talk given at KITP. *Santa Barbara, California*, 12, 2015.
- 26 Alexei Kitaev. A simple model of quantum Holography (part 2). *Entanglement in strongly-correlated quantum matter*, page 38, 2015.
- 27 Guang Hao Low and Isaac L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118(1):010501, 2017. arXiv: 1606.02685 doi:10.1103/PhysRevLett.118.010501.
- 28 Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. arXiv: 1610.06546 doi:10.22331/q-2019-07-12-163.
- 29 Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, 2020. arXiv: 1808.10402 doi:10.1103/RevModPhys.92.015003.
- 30 John Preskill. Simulating quantum field theory with a quantum computer. arXiv: 1811.10085, 2018.
- 31 Patrick Rall, Chunhao Wang, and Pawel Wocjan. Thermal state preparation via rounding promises. *Quantum*, 7:1132, October 2023. doi:10.22331/q-2023-10-10-1132.
- 32 Akshar Ramkumar and Mehdi Soleimanifar. Mixing time of quantum Gibbs sampling for random sparse Hamiltonians. arXiv: 2411.04454, 2024.
- 33 Cambyse Rouzé, Daniel Stilck França, and Álvaro M Alhambra. Efficient thermalization and universal quantum computing with quantum Gibbs samplers. *arXiv preprint arXiv:2403.12691*, 2024. doi:10.48550/arXiv.2403.12691.
- 34 Subir Sachdev and Jinwu Ye. Gapless spin-fluid ground state in a random quantum Heisenberg magnet. *Physical Review Letters*, 70(21):3339–3342, May 1993. doi:10.1103/physrevlett.70.3339.
- 35 Sushant Sachdeva and Nisheeth K Vishnoi. Matrix inversion is as easy as exponentiation. *arXiv preprint arXiv:1305.0526*, 2013. doi:10.48550/arXiv.1305.0526.
- 36 Oles Shtanko and Ramis Movassagh. Preparing thermal states on noiseless and noisy programmable quantum processors. *arXiv preprint arXiv:2112.14688*, 2021. doi:10.48550/arXiv.2112.14688.
- 37 Brian Swingle and Mike Winer. Bosonic model of quantum Holography. *Physical Review B*, 109(9):094206, 2024. doi:10.48550/arXiv.2311.01516.
- 38 Kristan Temme. Lower bounds to the spectral gap of Davies generators. *Journal of Mathematical Physics*, 54(12), December 2013. doi:10.1063/1.4850896.
- 39 Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015. arXiv: 1501.01571 doi:10.1561/22000000048.
- 40 Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak, and Matthias Troyer. Solving strongly correlated electron models on a quantum computer. *Phys. Rev. A*, 92:062318, December 2015. doi:10.1103/PhysRevA.92.062318.
- 41 Michael Winer, Richard Barney, Christopher L Baldwin, Victor Galitski, and Brian Swingle. Spectral form factor of a quantum spin glass. *Journal of High Energy Physics*, 2022(9):1–47, 2022. doi:10.48550/arXiv.2203.12753.
- 42 Pawel Wocjan and Kristan Temme. Szegedy walk unitaries for quantum maps. *Communications in Mathematical Physics*, 402(3):3201–3231, 2023. doi:10.48550/arXiv.2107.07365.
- 43 Daniel Zhang, Jan Lukas Bosse, and Toby Cubitt. Dissipative quantum Gibbs sampling. *arXiv preprint arXiv:2304.04526*, 2023. doi:10.48550/arXiv.2304.04526.

A CKG Lindbladian in the energy basis

We consider a quantum system consisting of basis states $|e_i\rangle$ and a Hamiltonian \mathbf{H} . We choose some jump operators \mathbf{A}^a and denote $A_{lm}^a = \langle l | \mathbf{A}^a | m \rangle$. Notate the energy eigenstates as $|j\rangle$ with energy E_j . We independently calculate the three parts of the Lindbladian: the transition term \mathcal{L}_t , the decay term \mathcal{L}_d , and coherent term \mathcal{L}_c so that $\mathcal{L}_\beta = \mathcal{L}_c + \mathcal{L}_t + \mathcal{L}_d$. First, as mentioned above, the OFT of the jump operator \mathbf{A}^a is $\hat{\mathbf{A}}^a(\omega) = \sum_{lm} A_{lm}^a \hat{f}(\omega - \nu_{lm}) |l\rangle\langle m|$, where $\nu_{lm} = E_l - E_m$. To represent superoperators as linear maps, we vectorize operators with respect to the basis of operators $|m_1\rangle\langle m_2|$. In particular, $|\mathbf{m}\rangle$ with $\mathbf{m} = (m_1, m_2)$ will notate the basis operator $|m_1\rangle\langle m_2|$. Now, we may expand $\mathcal{L}_t, \mathcal{L}_d, \mathcal{L}_c$. The full calculations are performed in the arXiv version of this paper [32], yielding the following expressions:

$$\begin{aligned} \langle \mathbf{l} | \mathcal{L}_t | \mathbf{m} \rangle &= \sum_a A_{l_1 m_1}^a \overline{A_{l_2 m_2}^a} \theta(\nu_{l_1 m_1}, \nu_{l_2 m_2}), \\ \langle \mathbf{l} | \mathcal{L}_d | \mathbf{m} \rangle &= -\frac{1}{2} \left(\delta_{l_1 m_1} \sum_{a,j} \overline{A_{j m_2}^a} A_{j l_2}^a \theta(\nu_{j m_2}, \nu_{j l_2}) + \delta_{l_2 m_2} \sum_{a,j} \overline{A_{j l_1}^a} A_{j m_1}^a \theta(\nu_{j l_1}, \nu_{j m_1}) \right), \\ \langle \mathbf{l} | \mathcal{L}_c | \mathbf{m} \rangle &= \frac{1}{2} \left(\delta_{l_1 m_1} \tanh(\beta \nu_{m_2 l_2} / 4) \sum_{a,j} \overline{A_{j m_2}^a} A_{j l_2}^a \theta(\nu_{j m_2}, \nu_{j l_2}) - \right. \\ &\quad \left. \delta_{l_2 m_2} \tanh(\beta \nu_{l_1 m_1} / 4) \sum_{a,j} \overline{A_{j l_1}^a} A_{j m_1}^a \theta(\nu_{j l_1}, \nu_{j m_1}) \right). \end{aligned} \quad (5)$$

Note that they can all be taken in terms of $\alpha(\nu) = \theta(\nu, \nu)$ using the identity $\theta(\nu_1, \nu_2) = \alpha(\frac{\nu_1 + \nu_2}{2}) \exp\left(-\frac{(\nu_1 - \nu_2)^2}{8\sigma_E^2}\right)$.

B Graph-Local Jumps for Cyclic Graphs

In this section we prove Theorem 2.1. Consider a cyclic graph with n vertices with adjacency matrix \mathbf{H} , and eigenvectors $|j\rangle$.

The eigenbasis of a cyclic graph consists of vectors $|j\rangle = n^{-1/2} \sum_a \zeta_n^{-aj} |e_a\rangle$ with eigenvalues $2 \cos(\frac{2\pi j}{n})$. The jump operators on the graph are chosen to be graph-local $\mathbf{A}^a = n^{-1/2} |e_a\rangle\langle e_a|$, and therefore have coefficients $A_{lm}^a = n^{-3/2} \zeta_n^{a(l-m)}$. Now we observe that $\sum_a \zeta_n^{a(i-j)} = n \delta_{ij}$. We therefore have the relation that

$$\sum_a A_{l_1 m_1}^a \overline{A_{l_2 m_2}^a} = n^{-3} \sum_a \zeta_n^{a((l_1 - m_1) - (l_2 - m_2))} = n^{-2} \delta_{(l_1 - m_1)(l_2 - m_2)}.$$

As computed more explicitly in the arXiv version [32], we compute the components of the Lindbladian with these jump operators:

$$\begin{aligned} \langle \mathbf{l} | \mathcal{L}_t | \mathbf{m} \rangle &= n^{-2} \theta(\nu_{l_1 m_1}, \nu_{l_2 m_2}) \delta_{(l_1 - m_1)(l_2 - m_2)}, \\ \langle \mathbf{l} | \mathcal{L}_d | \mathbf{m} \rangle &= -\frac{n^{-2}}{2} \delta_{l_1 m_1} \delta_{l_2 m_2} (\alpha(\nu_{j m_2}) + \alpha(\nu_{j m_1})) \\ \langle \mathbf{l} | \mathcal{L}_c | \mathbf{m} \rangle &= 0. \end{aligned}$$

The above formulae imply that the Lindbladian is block diagonal in the eigenbasis. The coherent term vanishes and the decay term is fully diagonal. Setting $k = m_1 - m_2$ and $k' = l_1 - l_2$, the transition term $\langle \mathbf{l} | \mathcal{L}_t | \mathbf{m} \rangle$ is nonzero only if $k = k'$, due to the factor

$\delta_{(l_1-m_1)(l_2-m_2)} = \delta_{(l_1-l_2)(m_1-m_2)}$. There is therefore one block corresponding to each k , which we will denote \mathcal{L}^k . The block for $k = 0$ is the classical block of the Markov chain on the diagonal entries of the state. Finding its spectral gap, and then lower bounding the eigenvalues of the remaining $n - 1$ blocks, yields a bound for the spectral gap of the Lindbladian.

B.1 Spectral Gap Lower Bound

We will first show that the spectral gap of the classical block is asymptotically $\Omega(n^{-1})$. We have by explicit calculation that

$$\begin{aligned}\langle \mathbf{l} | \mathcal{L}_t^0 | \mathbf{m} \rangle &= \frac{1}{n^2} \alpha(\nu_{lm}) \\ \langle \mathbf{l} | \mathcal{L}_d^0 | \mathbf{m} \rangle &= -\frac{1}{n^2} \delta_{lm} \sum_j \alpha(\nu_{jm}).\end{aligned}$$

We will use the canonical path bound, a standard technique in the theory of classical Markov chains, to establish lower bounds on their spectral gaps. The canonical path lemma fixes a “canonical” path between each pair of vertices on a graph and obtains a corresponding spectral graph bound. For our purposes we let the canonical path between any two vertices to be the edge joining them, which obtains the following statement:

► **Lemma B.1.** *Say L^0 is a Markov chain generator with stationary state σ . Then, the spectral gap satisfies the following bound:*

$$\lambda \geq \min_{(l,m)} \frac{L_{lm}^0}{\sigma_l}.$$

Applying this bound in this case, and noting that the stationary state of this Markov chain is ρ_{ll} , we obtain the lower bound

$$\lambda \geq \min_{l \neq m} \frac{\alpha(\nu_{lm})n^{-2}}{\rho_{ll}}. \quad (6)$$

The first equality holds because every canonical path is length 1, so the only path containing the edge (l, m) is γ_{lm} . We may upper bound ρ_{ll} with $\rho_{ll} \leq \frac{e^{2\beta}}{\sum_i E_i} \leq n^{-1} \frac{e^{2\beta}}{e^{-2\beta}} = n^{-1} e^{4\beta}$. Moreover, $|\nu_{lm}| \leq 4$ since all energies lie in $[-2, 2]$, so α is bounded below by a positive constant C that is independent of n . We conclude that $\lambda \geq \min_{l \neq m} \frac{\alpha(\nu_{lm})n^{-2}}{\rho_{ll}} \geq \frac{Cn^{-2}}{e^{4\beta}n^{-1}} = \Omega(n^{-1})$, as desired.

Now, for the blocks with $k \neq 0$, we utilize the Gershgorin bound on the columns of the k th block of $-\mathcal{L}_\beta$, $-\mathcal{L}^k$, which states that the eigenvalues of \mathcal{L}^k must be larger than $\min_{\mathbf{m}} \left[\langle \mathbf{m} | \mathcal{L}^k | \mathbf{m} \rangle - \sum_{\mathbf{l} \neq \mathbf{m}} |\langle \mathbf{l} | \mathcal{L}^k | \mathbf{m} \rangle| \right]$. This approach is very similar to the one outlined in [38] for the Davies generator. By explicitly calculating and bounding this term in the arXiv version of this paper-[32], we obtain that the eigenvalues of \mathcal{L}^k are $\Omega(n^{-3})$. This completes the lower bound, showing that all the spectral gap in full must be $\Omega(n^{-3})$.

B.2 Spectral Gap Upper Bound

To prove the upper bound on the spectral gap, we consider the row vector v of length n^2 , that is 1 on the indices that correspond to the block $k = 1$, and 0 elsewhere. As an operator, it takes the value 1 on one offdiagonal with a fixed $l_1 - l_2 = k$. When calculating

$(v\mathcal{L}_\beta)_{\mathbf{m}} = (\mathcal{L}_\beta^\dagger v)_{\mathbf{m}}$, we obtain the same formula as found in the Gershgorin bound calculation above – indeed, the values on the diagonal are all positive, while the off-diagonal values are negative:

$$\begin{aligned} (\mathcal{L}_\beta^\dagger v)_{\mathbf{m}} = & \frac{1}{n^2} \sum_{\mathbf{l}} \left(\frac{\alpha(\nu_{l_1 m_1}) + \alpha(\nu_{l_2 m_2})}{2} - \alpha \left(\frac{\nu_{l_1 m_1} + \nu_{l_2 m_2}}{2} \right) \right) \\ & + \frac{1}{n^2} \sum_{\mathbf{l}} \alpha \left(\frac{\nu_{l_1 m_1} + \nu_{l_2 m_2}}{2} \right) \left(1 - \exp \left(\frac{-(\nu_{l_1 m_1} - \nu_{l_2 m_2})^2}{8\sigma_E^2} \right) \right). \end{aligned} \quad (7)$$

The previous lower bound shows that each of these values is nonnegative. Since $l_1 - l_2 = m_1 - m_2 = 1$, $\nu_{l_1 m_1} - \nu_{l_2 m_2} = \nu_{l_1 m_1} - \nu_{l_2 m_2}$ is $O(n^{-1})$. The first term in the expression, since it is composed of summands that are second differences in α , is n terms that are $O(n^{-2})$ scaled by n^{-2} – it is therefore $O(n^{-3})$. The summands of the second term can be similarly estimated to be $O(n^{-2})$, so it is also $O(n^{-3})$. The terms of $(\mathcal{L}_\beta^\dagger v)_{\mathbf{m}}$ are therefore nonnegative and are at most Cn^{-3} for some constant C .

To prove the upper bound, we make use of the inner product $\langle \cdot, \cdot \rangle_{\rho_\beta^{-1}}$ with respect to which $\mathcal{L}_\beta^\dagger$ is self-adjoint. Note that $\langle |i_1\rangle\langle i_2|, |j_1\rangle\langle j_2| \rangle_{\rho_\beta^{-1}} = \delta_{i_2 j_1} \delta_{i_1 j_2} (\rho_\beta)_{i_1 i_1}^{1/2} (\rho_\beta)_{i_2 i_2}^{1/2} \geq 0$. Hence, when $\langle \cdot, \cdot \rangle_{\rho_\beta^{-1}}$ is expressed in the energy basis as vMw for a matrix M , M has nonnegative elements. We therefore may upper bound $\langle (\mathcal{L}_\beta^\dagger v), v \rangle_{\rho_\beta^{-1}}$ by $Cn^{-3} \langle v, v \rangle$, since the coefficients of $(\mathcal{L}_\beta^\dagger v)$ and v are nonnegative and $(\mathcal{L}_\beta^\dagger v)$ is dominated by $Cn^{-3}v$ for some $C > 0$. We conclude that $\frac{\langle (\mathcal{L}_\beta^\dagger v), v \rangle_{\rho_\beta^{-1}}}{\langle v, v \rangle_{\rho_\beta^{-1}}}$ is $O(n^{-3})$. Since $\mathcal{L}_\beta^\dagger$ is self-adjoint with respect to this inner product, we obtain that v has Rayleigh quotient $O(n^{-3})$. v is also orthogonal to I , since $\langle v, \rho_\beta \rangle_{\rho_\beta^{-1}} = \text{Tr}(v \rho_\beta^{1/2} I \rho_\beta^{1/2}) = \text{Tr}(v \rho_\beta) = 0$, where the last equality holds since as an operator v is zero along the diagonal. v has no overlap with I , the fixed point of $\mathcal{L}_\beta^\dagger$, and therefore its Rayleigh quotient is an upper bound on the spectral gap. The spectral gap must therefore also be $O(n^{-3})$. This completes the proof of Theorem 2.1.

C Bounded Degree Systems with 1-Design Jumps

In this section we prove Lemma 3.1 and Theorem 2.4, demonstrating an improvement over the result in Theorem 2.1 for local jumps in cyclic graphs.

Proof of Lemma 3.1. To prove Lemma 3.1, we make use of the expressions (5) for the transition, decay, and coherent parts of a general Lindbladian, but with simply one Haar random jump (or equivalently any 1-design since the second moments of the operators are equal). The transition term is

$$\langle \mathbf{l} | \mathcal{L}_t | \mathbf{m} \rangle = A_{l_1 m_1} \overline{A_{l_2 m_2}} \theta(\nu_{l_1 m_1}, \nu_{l_2 m_2}).$$

The expectation of the product $A_{l_1 m_1} \overline{A_{l_2 m_2}}$ is zero if $l_1 \neq l_2$ or $m_1 \neq m_2$. The expectation of the norm squared of an element, on the other hand, is n^{-1} . By explicit calculation, we therefore obtain

$$\begin{aligned} \mathbb{E}[\langle \mathbf{l} | \mathcal{L}_t | \mathbf{m} \rangle] &= \frac{\delta_{l_1 l_2} \delta_{m_1 m_2}}{n} \alpha(\nu_{l_1 m_1}) \\ \mathbb{E}[\langle \mathbf{l} | \mathcal{L}_d | \mathbf{m} \rangle] &= -\frac{1}{2} \left(\frac{\delta_{l_1 m_1} \delta_{l_2 m_2}}{n} \sum_j (\alpha(\nu_{j m_2}) + \alpha(\nu_{j m_1})) \right), \\ \mathbb{E}[\langle \mathbf{l} | \mathcal{L}_c | \mathbf{m} \rangle] &= 0. \end{aligned}$$

The final Lindbladian \mathcal{L}_μ is therefore completely diagonal except for a “classical block” \mathcal{L}^0 of indices $|(m, m)\rangle$, whose off-diagonal terms are populated by the elements of \mathcal{L}_t . The spectral gap of this Lindbladian is therefore the minimum of the values along the diagonal, which are all positive, and the spectral gap of the classical block.

As in Section B, we use the Lemma B.1 to bound the spectral gap of the classical block. Using this lemma, with the canonical path being the edge between a pair of vertices, we obtain the bound

$$\lambda \geq \min_{l \neq m} \frac{\alpha(\nu_{lm})n^{-1}}{\rho_{ll}}. \quad (8)$$

We may upper bound ρ_{ll} with $\rho_{ll} \leq \frac{e^{-E_{\min}\beta}}{\sum_i E_i} \leq n^{-1} \frac{e^{-E_{\min}\beta}}{e^{-E_{\max}\beta}} = n^{-1}e^{-O(1)} = O(n^{-1})$ due to the fact that $\beta\|\mathbf{H}\| = O(1)$. Similarly, since α is the convolution of a Gaussian of radius $\sigma_E = \beta^{-1}$ with $\gamma(\omega) = \exp\left(-\beta \max\left(\omega + \frac{\beta\sigma_E^2}{2}, 0\right)\right)$, the assumption that $\beta\|\mathbf{H}\| = O(1)$ again yields that α evaluated at ν_{lm} is $\Omega(1)$. Indeed, within $O(\beta^{-1})$ of any value of ν_{lm} , $\gamma(\omega)$ is $\Omega(1)$, and as a consequence $\alpha(\nu_{lm}) = \Omega(1)$. This yields a lower bound on λ of $\min_{l \neq m} \frac{\Omega(n^{-1})}{O(n^{-1})} = \Omega(1)$.

Now, we lower bound the diagonal elements outside of the classical block. Since each such element is of the form

$$\mathbb{E}[\langle \mathbf{m} | \mathcal{L}_d | \mathbf{m} \rangle] = -\frac{1}{2} \left(\frac{1}{n} \sum_j \alpha(\nu_{jm_2}) + \frac{1}{n} \sum_j \alpha(\nu_{jm_1}) \right)$$

and we have already established that each α term is $\Omega(1)$, so the resulting diagonal values are all $\Omega(1)$. We conclude that the spectral gap of the Lindbladian is $\Omega(1)$. \blacktriangleleft

Proof of Theorem 2.4. We construct our Lindbladian by sampling $M = \Theta(\log(n))$ unnormalized jumps \mathbf{A}^a from the 1-design $\mathcal{D}(U(n))$ as in Definition 2.2, each with a corresponding Lindbladian \mathcal{L}_a (which has one jump \mathbf{A}^a along with its adjoint, normalized by 2). Then, we want to prove that with high probability, $\mathcal{L}_\beta = \frac{2}{M} \sum_{a=1}^{M/2} \mathcal{L}_a$, the Lindbladian with all M of these jumps now normalized by the number of jumps, has spectral gap bounded by a constant.

To prove the result, we shall make use of the matrix Bernstein’s inequality for our concentration bound:

► **Lemma C.1** (cf. [39]). *Say X_1, \dots, X_N are independent random $d \times d$ Hermitian matrices, such that $\mathbb{E}[X_i] = 0$ and $\|X_i\| \leq R$. Define $Y = \frac{1}{N} \sum_{i=1}^N X_i$, and say that $N\mathbb{E}[Y^2] \leq T$. Then $\Pr(\|Y\| \geq t) \leq 2d \exp(-\frac{3}{2} \frac{Nt^2}{3T + Rt})$.*

Call $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$, where $\mathcal{L}_\mu = \mathbb{E}_{\mathbf{A} \sim \mathcal{D}}[\mathcal{L}_\beta]$. Each of these operators has zero expectation. We have that $\delta\mathcal{L} = \frac{2}{M} \sum_{a=1}^{M/2} \delta\mathcal{L}_a$ is precisely the discrepancy between our Lindbladian $\mathcal{L}_\beta = \frac{2}{M} \sum_{a=1}^{M/2} \mathcal{L}_a$ and the expected Lindbladian \mathcal{L}_μ . We will apply Bernstein’s inequality for $X_a = \delta\mathcal{L}_a$, $Y = \delta\mathcal{L}$, and $N = \frac{M}{2}$. By Lemma 4.1, the CKG Lindbladian has operator norm $O(\log(\beta\|\mathbf{H}\|))$, which is $O(1)$ in this regime. We denote this upper bound $\frac{R}{2}$. We can now verify the condition $N\mathbb{E}[Y^2] \leq T$ in the statement of Lemma C.1, since we have that $\|\delta\mathcal{L}_a\| \leq \|\mathcal{L}_a\| + \|\mathcal{L}_\mu\| \leq R = \Theta(1)$, and

$$\frac{M}{2} \left\| \mathbb{E} \left[\left(\frac{\sum_{a=1}^{M/2} \delta\mathcal{L}_a}{M/2} \right)^2 \right] \right\| \leq \frac{2}{M} \sum_{a=1}^{M/2} \|\delta\mathcal{L}_a\|^2 \leq R^2.$$

Every operator that satisfies detailed balance is Hermitian in some fixed basis, and therefore each \mathcal{L}_a , as well as \mathcal{L}_μ , can be considered Hermitian. By linearity, the same holds true for $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$. The operators $\delta\mathcal{L}_a$ therefore satisfy the conditions of the matrix Bernstein inequality, and so their average $\delta\mathcal{L}$ satisfies

$$\Pr(\|\delta\mathcal{L}\| \geq t) \leq 2n^2 \exp\left(-\frac{3}{4} \frac{Mt^2}{3R^2 + Rt}\right),$$

where the n^2 is due to an overhead of the dimension of the Lindbladian.

For any constant t , there exists an $M = \Theta(\log(n))$ such that the term inside the exponential is at least $3 \log(n)$, since R is a constant. The probability that $\|\delta\mathcal{L}\| \leq t$ is then arbitrarily close to 1 for sufficiently large n and choice of $M = \Theta(\log(n))$. By Lemma 3.1, \mathcal{L}_μ has a constant spectral gap bounded below by some C . By Weyl's theorem, the eigenvalues of $\mathcal{L}_\mu + \delta\mathcal{L}$ may differ by at most t from those of \mathcal{L}_μ . Choosing $t \leq \frac{C}{2}$, it follows that \mathcal{L}_a , with any constant probability, has constant spectral gap. \blacktriangleleft

D Unbounded Degree Systems with 1-Design Jumps

Proof of Lemma 3.2. We follow the proof of Lemma 3.1. As in Lemma 3.1, we may obtain the following bound on the classical block:

$$\lambda \geq \min_{l \neq m} \frac{\alpha(\nu_{lm})n^{-1}}{\rho_{ll}} = \frac{\alpha(\nu_{lm}) \left(\frac{1}{n} \sum_j \exp(-\beta E_j) \right)}{\exp(-\beta E_l)}. \quad (9)$$

An explicit calculation, as in [32], therefore yields that

$$\lambda = \Omega(\delta(n))$$

by assumption that $\delta(n)$ of the eigenvalues are within $O(\beta^{-1})$ of λ_{\min} . Now bounding the diagonal elements outside of the classical block, we see that

$$\mathbb{E}[\langle \mathbf{m} | \mathcal{L}_d | \mathbf{m} \rangle] = -\frac{1}{2} \left(\frac{1}{n} \sum_j \alpha(\nu_{jm_2}) + \frac{1}{n} \sum_j \alpha(\nu_{jm_1}) \right).$$

Again, since $\delta(n)$ of eigenvalues are within $O(\beta^{-1})$ of λ_{\min} , the above sum is $\Omega(\delta(n))$, as desired. \blacktriangleleft

Proof of Theorem 2.6. We follow the proof of Theorem 2.4. Defining once again \mathcal{L}_a to be the Lindbladian with one jump operator \mathbf{A}^a and its adjoint (normalized by 2), and defining $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$, we can apply the matrix Bernstein's inequality to obtain

$$\Pr(\|\delta\mathcal{L}\| \geq t) \leq n^2 \exp\left(-\frac{3}{4} \frac{Mt^2}{3R^2 + Rt}\right),$$

since all the conditions of the inequality are again satisfied.

By Lemma 3.1, \mathcal{L}_μ has a constant spectral gap bounded below by some $C\delta(n)$. Selecting t to be $\frac{C}{2}\delta(n)$, there exists an $M = \Theta(\delta(n)^{-2} \log(\beta\|\mathbf{H}\|)^2 \log(n))$ for which the value inside the exponential is at least $3 \log(n)$. The probability that $\|\delta\mathcal{L}\| \leq t$ is therefore above any constant probability for sufficiently large n . By Weyl's theorem, the eigenvalues of $\mathcal{L}_\mu + \delta\mathcal{L}$ may differ by at most $t \leq \frac{C}{2}\delta(n)$. It follows that \mathcal{L}_a , with any constant probability, has spectral gap $\Omega(\delta(n))$. \blacktriangleleft