Thermal State Preparation via Rounding Promises

Patrick Rall[1], Chunhao Wang[2], and Pawel Wocjan[3]

[1]IBM Quantum, MIT-IBM Watson AI Lab, Cambridge, Massachusetts 02142, USA
[2]Department of Computer Science and Engineering, Pennsylvania State University
[3]IBM Quantum, Thomas J Watson Research Center, Yorktown Heights, New York 10598, USA

A promising avenue for the preparation of Gibbs states on a quantum computer is to simulate the physical thermalization process. The Davies generator describes the dynamics of an open quantum system that is in contact with a heat bath. Crucially, it does not require simulation of the heat bath itself, only the system we hope to thermalize. Using the state-of-the-art techniques for quantum simulation of the Lindblad equation, we devise a technique for the preparation of Gibbs states via thermalization as specified by the Davies generator.

In doing so, we encounter a severe technical challenge: implementation of the Davies generator demands the ability to estimate the energy of the system unambiguously. That is, each energy of the system must be deterministically mapped to a unique estimate. Previous work showed that this is only possible if the system satisfies an unphysical 'rounding promise' assumption. We solve this problem by engineering a random ensemble of rounding promises that simultaneously solves three problems: First, each rounding promise admits preparation of a 'promised' thermal state via a Davies generator. Second, these Davies generators have a similar mixing time as the ideal Davies generator. Third, the average of these promised thermal states approximates the ideal thermal state.

Patrick Rall: patrickjrall@ibm.com
Chunhao Wang: cwang@psu.edu
Pawel Wocjan: pawel.wocjan@ibm.com

Contents

1 Introduction

Motivation. Preparing Gibbs states is a major task for quantum computers. There are several reasons for this. First, the Gibbs state is one of the most important states of matter. For quantum models comprised of many locally interacting particles, it describes a wide range of physical situations, relevant to condensed matter physics, high energy physics, quantum chemistry [Alh22]. Therefore, to use the quantum computer as a universal simulator of quantum systems, it is desirable to be able to prepare Gibbs states. Second, for general Hamiltonians, the Gibbs state is a crucial ingredient in some quantum algorithms such as those for solving semidefinite programs [VAGGdW20, BKL+19, vAG19, GSLW19] and for training quantum Boltzman machines [KW17]. Third, the problem of estimating the quantum partition function, which is connected to the problem of approximately preparing the Gibbs state, plays an important role in quantum complexity theory [BCGW21].

Background and overview of prior work. There are three main approaches to preparing Gibbs states.

The first one is a Grover-based approach in which an initial state is mapped onto a certain purification of the Gibbs state at inverse temperature β (see [PW09, CS17]). The resulting running time is dictated by the overlap between the initial and target states, which is exponentially small in the system size. The Gibbs state can be approximately prepared in time on the order of $\sqrt{D/\mathcal{Z}_\beta}$, where D denotes the dimension of the quantum system and $\mathcal{Z}_\beta$ the partition function at inverse temperature β. In [HMS+22] a quantum algorithm is presented for preparing a purification of the Gibbs state for the Hamiltonian $H_1 = H_0 + V$ at inverse temperature β starting from a purification of the thermal state of H_0.

The second approach is based on Davies generators, which is a differential equation that describes how nature thermalizes a quantum system to its thermal equilibrium [Dav76, Dav79]. Davies generators are special cases of Lindbladians [Lin76, BP02], which describe the most general continuous-time Markovian dynamics of an open quantum system, i.e., a quantum system that is weakly coupled to the environment and the dynamics of the environment are fast enough so that the information only flows from the system to the environment while no information is flowing back to the system. The method in [CB21] simulates a Davies generator by attaching a heat bath and simulating time evolution on the joint system while repeatedly refreshing the bath. Their algorithm in some sense follows the original derivation of the Davies generator as the limit of such joint system-bath Hamiltonian time evolution. When the system Hamiltonian satisfies the Eigenstate Thermalization Hypothesis (ETH), they show that the implemented quantum map not only converges to the desired Gibbs state, but also does so in polynomial time.

The third approach is based on the quantum Metropolis algorithm [TOV+11]. This approach also avoids the exponential scaling when the system Hamiltonian satisfies ETH [SC22, CB21].

The advantages of the second approach are two fold. First, for every quantum system that thermalizes fast in nature, it is expected that our algorithms can also prepare its corresponding Gibbs state efficiently without suffering from the exponential dependence on the number of qubits. Second, this approach fits well for many physics-motivated applications. For example, we can use our algorithm to prepare a "partially thermalized" (in the natural thermalization

process) thermal state, which might be of interest in some scenarios.

Main result. The present work examines how to approximately prepare Gibbs states for arbitrary system Hamiltonians by simulating time evolutions according to carefully engineered Lindbladians. These Lindbladians are derived from an ideal Davies generator having the Gibbs state as unique fixed point.

There are some similarities but also important differences to the work [CB21]. Obviously, both are based on Davies generators. In contrast, we do not approximate Lindblad evolution with the help of the Hamiltonian evolution of the system and a bath, a method with which it is *provably impossible* to achieve linear scaling in evolution time (see [CW17]). Instead, we rely on a method for directly simulating Lindbladian time evolution specified by any jump operators. Using this method can lead to a reduction in resources: specifically, we achieve linear scaling in the mixing time t_{mix}. In addition, a direct Lindblad simulation approach avoids the complication of dealing with the dynamics of the bath and the interaction Hamiltonians. We also seek to prove that our quantum map approximates the Gibbs state for any system Hamiltonian, that is, we do not need to make an assumption such as ETH.

The quantum algorithms for simulating Lindbladian time evolution in [CW17, LW22] assume that the jump operators have been suitably encoded. Unfortunately, it is not possible to construct jump operators of a Davies generator due to inherent imperfections of energy estimation of general Hamiltonians. However, we show how to construct a family of Lindbladians from the given Davies generators such that simulating them with the help of the simulation algorithms in [CW17, LW22] and taking the average of the resulting quantum states provides a good approximation of the Gibbs state. This is formulated in more detail in the theorem statement below.

Theorem 1 (Main result – informal statement)**.** *Assume we are given block encodings of the Hamiltonian H, coupling operators S_α, and a filter function G. With appropriately chosen S_α, these give rise to a Davies generator $\mathcal{L}$ that has the Gibbs state ρ_β for inverse temperature β as a unique fixed point. Assume that after time t_{mix} the time evolved state $\exp(t_{\mathrm{mix}}\mathcal{L})(\sigma_0)$ is ε-close to the Gibbs state ρ_β for any initial state σ_0.*

We engineer a certain family of 2^r many Lindbladians $\tilde{\mathcal{L}}^{(j)}$ from the above Davies generator $\mathcal{L}$. These Lindbladians have a new 'attenuated' mixing time t_{mix}, and their jump operators can be encoded efficiently with imperfect energy estimation. This makes it possible to simulate their time evolutions $\exp\left(t\tilde{\mathcal{L}}^{(j)}\right)$. We prove that the average

$$\frac{1}{2^r}\sum_j \exp\left(\tilde{t}_{\mathrm{mix}}\tilde{\mathcal{L}}^{(j)}\right)(\sigma_0) \tag{1}$$

is ε-close to the Gibbs state ρ_β for any initial state σ_0. Furthermore, we show how to implement the time evolution according to $\tilde{\mathcal{L}}^{(j)}$ to arbitrarily small failure probability $\delta_{\mathcal{L}}$, and that the total number of invocations of the block encoding of the Hamiltonian is bounded by:

$$O\left(\tilde{t}_{\mathrm{mix}}\cdot\gamma^{-1}\cdot\beta^3\varepsilon^{-7}\cdot\mathrm{polylog}(\tilde{t}_{\mathrm{mix}}/\delta_{\mathcal{L}})\cdot\log^2(\beta/\varepsilon)\right), \tag{2}$$

where γ is an attenuation coefficient that affects the attenuated mixing time $\tilde{t}_{\mathrm{mix}}$.

Remark. *After the first version of this manuscript was made public, recent work [CKBG23] resolved an open question raised in this manuscript (see Section 7) on utilizing the block encoding of the form* $\sum_j |j\rangle \otimes L_j$*. Using the Theorem III.1 of [CKBG23] together with slight adaptation of the simulation algorithm in [LW22], the complexity of our algorithm can be improved to* $\tilde{O}(\tilde{t}_{\mathrm{mix}} \cdot \gamma^{-1} \cdot \beta\varepsilon^{-2})$*. We also note that the additional factor of* $\log^2(\beta/\varepsilon)$ *can be removed via existing techniques from [Ral21].*

In Section 5 we give some numerical experiments indicating that the attenuated mixing time $\tilde{t}_{\mathrm{mix}}$ is on the order of the original mixing time t_{mix} for suitably chosen γ^{-1}.

We formulate in Section 7 some open research questions whose solution could lead to improvements of our current methods.

Technical overview. As mentioned above, the implementation of jump operators of a Davies generator requires perfect energy estimation. More precisely, what we mean by perfect is that energy estimation would have to be unambiguous: for each energy of the Hamiltonian, it must yield a unique energy estimate. Unfortunately, energy estimation unavoidably produces superpositions of different energy estimates for general Hamiltonians: if $|\psi\rangle$ is an eigenstate of the Hamiltonian with eigenvalue λ, energy estimation implements a map:

$$|\psi\rangle \quad \mapsto \quad |\psi\rangle \otimes \Big(\alpha\,|\tilde{\lambda}_1\rangle + \beta\,|\tilde{\lambda}_2\rangle\Big) \tag{3}$$

where $\tilde{\lambda}_1$ and $\tilde{\lambda}_2$ are two different estimates of λ. This superposition of estimates cause constant-size errors in quantum algorithms implementing Davies generators, and must be eliminated. With perfect estimation, the superposition on the second register is not present and there is always a unique estimate.

It is possible to construct an 'approximate Davies generator' from an energy estimation algorithm that produces superpositions of estimates. However, the resulting dynamics no longer correspond to the Davies generator of any particular Hamiltonian, making it challenging to analyze. To our knowledge, no method exists for rigorously proving the accuracy of an algorithm based on such an approximate Davies generator. We highlight some of the challenges of this task in Appendix C.

It was shown in [Ral21] that perfect energy estimation is possible when the Hamiltonian satisfies a 'rounding promise' assumption. The rounding promise prohibits the Hamiltonian from having any eigenvalues that induce a superposition of estimates. However, such an assumption on the Hamiltonian is extremely unphysical and will not be satisfiable in practice. The main technical idea of the present work is to shift the notion of a rounding promise away from the Hamiltonian itself, but rather to a family of states. The basic idea is very simple: if a quantum state has no support on any eigenstates whose eigenvalues have a superposition of estimates, then it is as if the Hamiltonian did not have such eigenvalues. This family of states is defined by a 'promised subspace'.

Since perfect energy estimation is possible on a promised subspace, implementation of Davies generators is possible as well. While the analysis involves a wide variety of error parameters, we find that all of them admit a mathematically rigorous treatment.

Our construction relies on just one assumption: that we can construct coupling operators for each promised subspace that ensure that the Davies generator converges in a reasonable amount of time. We give numerical evidence that projecting a coupling operator into a

promised subspace does not significantly reduce the Davies generator's convergence time. Notice how this assumption has nothing to do with the protocol's accuracy, only its convergence time.

We now summarize the different types of errors that occur in our implementation, and the rough ideas behind keeping them under control:

- **Lindblad simulation error**: There exists a quantum algorithm, e.g. [CW17], for simulating Lindblad evolution given a description of the jump operators of a Lindbladian $\mathcal{L}$. For evolution time t and accuracy δ_{sim}, this algorithm implements a quantum channel with accuracy δ_{sim} in terms of the diamond norm to the ideal channel $e^{\mathcal{L}t}$. In our case, we simulate the dynamics that approximates a special Lindbladian called the Davies generator.

- **Knowledge of mixing time**: Given a Hamiltonian H, coupling operators S_α, and a filter function G, the Davies generator is a certain Lindbladian whose fixed point is the desired thermal state $\rho_\beta = e^{-\beta H} / \operatorname{Tr} e^{-\beta H}$. We assume we are given the Hamiltonian and suitable coupling operators of the Davies generator together with a bound t_{mix} after which, the state is sufficiently close the desired thermal state.

- **Precision and failure probability of energy estimation**: The jump operators of a Davies generator depend on the coupling operators and on the Bohr frequencies (differences of the energies) of the Hamiltonian and the corresponding pairs of eigensubspaces. To realize these, we rely on energy estimation of the Hamiltonian.

  By restricting to the promised subspace, we can guarantee using techniques from [Ral21] that for every energy λ there exists a single unique estimate $\tilde{\lambda}$. There remain two other kinds of error, both of which can be dealt with rigorously.

  First, the energy estimation map is only implemented with a failure probability δ_{est} that can be exponentially suppressed. By phrasing this error as a deviation in spectral norm of the Lindbladian jump operators, we can show that this error is not blown up when the Davies generator is simulated for a long period of time.

  Second, the cost of energy estimation scales linearly with the precision of the estimate. To analyze this, we leverage a trick from [PW09]: instead of preparing the thermal state for the original Hamiltonian, one can interpret the resulting process as preparing the thermal state for a rounded Hamiltonian. Since the norm of the difference of these Hamiltonians is small, the corresponding thermal states are close to each other.

- **Preparation of an initial state on the promised subspace**: We implement a Davies generator whose dynamics are trivial outside the promised subspace. In order to prepare a promised thermal state, we must feed the Davies generator with a state that is exclusively supported on the promised subspace. We achieve this by taking an arbitrary input state and measuring a two-outcome 'left-right' POVM. Depending on the POVM outcome, we know that either a 'left' or a 'right' rounding promise is satisfied by the post-measurement state.

- **Approximation of the ideal Gibbs state**: Our goal is to prepare a Gibbs state supported on the entire Hilbert space. But our protocol only prepares promised Gibbs states, which

are only supported on the promised subspace. Thus, individual promised thermal states are generically far in trace distance from the ideal state. To deal with this, we show that there exists an ensemble of rounding promises such that the ensemble average of all the promised Gibbs states is an accurate approximation of the ideal thermal state. The basic idea of this ensemble is that the probability of any particular energy being excluded can be made arbitrarily small.

- **Leakage and attenuation errors in removal of rounding promise**: The constructed coupling operators of the Davies generators for the well-rounded Hamiltonians on the promised subspaces have two important types of errors, namely, 'leakage' and 'attenuation.'

  The leakage error δ_{leak} measures how much coupling operators and initial states 'leak' outside a promised space. Fortunately, δ_{leak} can be made exponentially small by using quantum singular value transformation.

  Blocks of coupling operators corresponding to some pairs of energies can be 'attenuated', i.e., multiplied by small positive numbers. However, as long as attenuation remains non-zero, the fixed-point remains the thermal state for the well-rounded Hamiltonian on the promised subspace. Unfortunately, the mixing speed can be negatively affected.

- **Mixing assumption for projected/attenuated coupling operators:** In principle, it could happen that the ideally projected coupling operators do not guarantee convergence to the thermal state of the well-rounded Hamiltonian on the promised subspace anymore. Moreover, even if they do, attenuation could increase the convergence time.

  However, we provide numerical evidence that these unfavorable situations do not typically occur. For the theoretical analysis of our quantum algorithm, we must assume that the attenuated coupling operators for the well-rounded Hamiltonian on the promised subspaces have similar mixing behavior as the original coupling operators.

2 Preliminaries

In this section, we first review some preliminaries about Gibbs states and Davies generators in order to establish notation and to rigorously define our goal: to prepare a Gibbs state by simulating the time evolution of a Davies generator. To do so, we leverage an algorithm for simulating general Lindblad evolution [LW22].

In order to implement the Davies generator, we require conditions under which the energy of a Hamiltonian can be estimated without producing superpositions of different energy estimates. So, in Section 2.2 we establish the notion of a rounding promise, and review a result from [Ral21] how a rounding promise can ensure that each energy is rounded to a unique energy when performing energy estimation.

Here are some notations and conventions used in this paper. We use $\mathbb{Z}_+$ to denote the set of all positive integers. Let $\mathcal{H}$ be a Hilbert space. We use $\mathrm{L}(\mathcal{H})$ to denote the collection of all linear operators (matrices) of the form: $A : \mathcal{H} \to \mathcal{H}$. For a matrix A, the *spectral norm* $\|A\|$ is the largest singular value of A, and the *trace norm* $\|A\|_1$ is the sum of the singular values. For a superoperator Λ, the *induced trace norm* of Λ, denoted by $\|\Lambda\|_1$, is defined as $\|\Lambda\|_1 = \max_{A \neq 0} \|\Lambda(A)\|_1 / \|A\|_1$. If Λ is acting on $\mathrm{L}(\mathcal{H})$ for some Hilbert space $\mathcal{H}$, then, the

diamond-norm of Λ, denoted by $\|\Lambda\|_\diamond$, is defined as $\|\Lambda\|_\diamond = \|\Lambda \otimes \mathcal{I}\|_1$, where $\mathcal{I}$ is the identity map acting on $\mathrm{L}(\mathcal{H})$. If A and B are matrices, then we say $A \geq B$ if $A - B$ is *positive semi-definite*. Finally, a *block encoding* of A is a unitary matrix U_A that, for some ancilla-count k, satisfies:

$$A = \left(\langle 0^k| \otimes I \right) U_A \left(|0^k\rangle \otimes I \right). \tag{4}$$

2.1 Gibbs states and Davies generators

We now define some fundamental concepts required to state our quantum algorithm and to analyze its performance.

Definition 2. *Let H be a Hamiltonian on the Hilbert space $\mathcal{H}$ with eigendecomposition*

$$H = \sum_i \lambda_i \Pi_i. \tag{5}$$

*Here, the Π_i are projectors onto the subspace with energy λ_i. We assume that the spectrum of H is contained in the interval $[0, 1]$. For inverse temperature $\beta > 0$, the **Gibbs state** ρ_β is the state such that*

$$\rho_\beta \propto \exp(-\beta H). \tag{6}$$

*The **partition function** $\mathcal{Z}_\beta$ is the normalization factor given by*

$$\mathcal{Z}_\beta = \mathrm{Tr}\left(\exp(-\beta H)\right). \tag{7}$$

Our quantum algorithm makes it possible to approximately prepare thermal states ρ_β. It is based on the Davies generators defined below. The Davies generators describe dissipative dynamics that converge to thermal states.

Definition 3. *Let $\{S_\alpha\}_\alpha$ be a collection of Hermitian operators acting on $\mathcal{H}$. The **Davies generator** with respect to the system Hamiltonian H and the **coupling operators** S_α is the Lindbladian $\mathcal{L}$ in the Schrödinger picture given by*

$$\mathcal{L}(\sigma) = \sum_\omega G(\omega) \left(\sum_\alpha S_\alpha(\omega) \sigma S_\alpha(\omega)^\dagger - \frac{1}{2} \left(S_\alpha(\omega)^\dagger S_\alpha(\omega) \sigma + \sigma S_\alpha(\omega)^\dagger S_\alpha(\omega) \right) \right). \tag{8}$$

*The **jump operators** $S_\alpha(\omega)$ are enumerated by the **Bohr frequencies** ω of H and are obtained from the coupling operators S_α by*

$$S_\alpha(\omega) = \sum_{\substack{i,j \\ \lambda_i - \lambda_j = \omega}} \Pi_i S_\alpha \Pi_j. \tag{9}$$

*The **filter function** $G(\omega)$ is a real-valued function satisfying $G(\omega) = e^{\beta\omega} G(-\omega)$.*

The time evolution of quantum states is given by the quantum channels

$$\mathcal{T}_t = \exp(t\mathcal{L}) \tag{10}$$

for $t \geq 0$.

*We say that a quantum state ρ is a **fixed point** of the Davies generator $\mathcal{L}$ if*

$$\mathcal{L}(\rho) = 0. \tag{11}$$

It is known that if the fixed point of the Davies generator $\mathcal{L}$ is unique, then the Lindbladian time evolution is relaxing in the sense that

$$\lim_{t\to\infty} \mathcal{T}_t(\sigma) = \rho_\beta \tag{12}$$

for any initial state σ (see [Nig19] and the references therein). The converse clearly holds.

Definition 4. *We call the coupling operators S_α of a Davis generator $\mathcal{L}$* ***mixing*** *if the thermal state ρ_β is the unique fixed point of $\mathcal{L}$. In this case, t_{mix} denotes the* ***mixing time*** *of $\mathcal{L}$, that is, the smallest time t such that $\mathcal{T}_t(\sigma)$ is guaranteed to be sufficiently close to the desired thermal state ρ_β for any initial state σ.*

There are several sufficient conditions guaranteeing the uniqueness of the fixed point. For instance, it can be shown that the thermal state ρ_β is the unique fixed point of the Davis generator $\mathcal{L}$, if the matrix algebra generated by the coupling operators S_α is the full matrix algebra (this statement follows from [Spo77]; see also the discussion in [Nig19] for an overview of other sufficient conditions).

The Lindbladian $\mathcal{L}$ of the Davies generator can also be written in terms of a collection of jump operators $\{L_{\omega,\alpha}\}_{\omega,\alpha}$:

$$\mathcal{L}(\sigma) = \sum_{\omega,\alpha} L_{\omega,\alpha}\sigma L_{\omega,\alpha}^\dagger - \frac{1}{2}\left(L_{\omega,\alpha}^\dagger L_{\omega,\alpha}\sigma + \sigma L_{\omega,\alpha}^\dagger L_{\omega,\alpha}\right) \tag{13}$$

where $L_{\omega,\alpha} = \sqrt{G(\omega)}S_\alpha(\omega)$. Having brought the Davies generator into this form, we can leverage existing results for Lindblad simulation [CW17].

In particular, say we can implement a block encoding of the operator $\mathcal{O}_\mathcal{L}$ that implements the jump operators as follows:

$$\mathcal{O}_\mathcal{L}(|0\rangle \otimes |\psi\rangle) = \sum_{\omega,\alpha} |\omega,\alpha\rangle \otimes L_{\omega,\alpha}|\psi\rangle \tag{14}$$

Then, given access to $\mathcal{O}_\mathcal{L}$, we have technical tools to simulate the time evolution $e^{\mathcal{L}t}$. We use the simulation algorithm from [LW22], which simplifies the simulation algorithm of [CW17] and also generalizes their input models to block-encodings.

Proposition 5 (Lindblad simulation, adapted from [LW22]). *Say $\mathcal{L}$ is a Davies generator acting on k qubits with m many jump operators $L_{\omega,\alpha}$ with $\|L_{\omega,\alpha}\| \leq 1$, and say we are given access to oracles to the jump operators via the block encoding $\mathcal{O}_\mathcal{L}$ above. Then, for any $t \geq 0$ and any $\delta_{\text{sim}} > 0$, there exists a quantum algorithm that implements a quantum channel δ_{sim}-close in the diamond norm to $e^{t\mathcal{L}}$, making*

$$O\left(mt\,\frac{\log(mt/\delta_{\text{sim}})}{\log\log(mt/\delta_{\text{sim}})}\right) \tag{15}$$

uses of the block encoding of $\mathcal{O}_\mathcal{L}$, and

$$O\left(m^2 t\left(\frac{\log(mt/\delta_{\text{sim}})}{\log\log(mt/\delta_{\text{sim}})}\right)^2\right) \tag{16}$$

additional 1- and 2-qubit gates.

In our adaptation of the theorem above from [LW22], we have replaced a quantity $||\mathcal{L}_{\text{be}}||$ with a bound $O(m)$ which follows from the fact that $||L_{\alpha,\omega}|| \leq 1$.

So, the central task is to implement a block encoding of $\mathcal{O}_{\mathcal{L}}$ for the Davies generator, which essentially amounts to implementing block encodings of the $S_\alpha(\omega)$ operators. Once this is accomplished, we can simulate the evolution $e^{t_{\text{mix}}\mathcal{L}}$ on any input state, and obtain an approximation of the thermal state ρ_β.

2.2 Rounding promises

Unfortunately, it is not possible to block encode the jump operators $S_\alpha(\omega)$ of Davies generators without prior knowledge of the spectrum of H. Ideally, we would like to implement a unitary that performs the isometry:

$$\sum_i \Pi_i \otimes |\lambda_i\rangle\,, \tag{17}$$

that is, computes a binary representation of the energy λ_i into a new register. However, there are two unavoidable limitations. First, λ_i can only be estimated to precision ε using only $O(1/\varepsilon)$ many resources. So, unless we select ε to be less than the smallest gap between any pair of energies, which requires at least $O(\dim(H))$ many resources, we will be unable to distinguish certain energies. We find that, leveraging a trick from Poulin and Wocjan [PW09], this error can be dealt with formally.

However, the second limitation is more challenging to deal with. Rall [Ral21] observed that for any energy estimation algorithm there will exist particular energies λ_i such that a corresponding eigenstate $|\psi_i\rangle$ will produce a superposition of estimates:

$$|\psi_i\rangle \mapsto |\psi_i\rangle \otimes \left(\alpha\,|\tilde{\lambda}_i\rangle + \beta\,|\tilde{\lambda}'_i\rangle\right) \tag{18}$$

where α, β are complex coefficients and $\tilde{\lambda}_i, \tilde{\lambda}'_i$ are two estimates of the energy λ_i. Essentially, certain energies λ_i are indecisive about which direction they want to round, and end up being rounded up or down in superposition.

The superposition of rounding directions creates cross terms between the rounding results in the construction of approximate block encodings of the jump operators $S_\alpha(\omega)$, which cause errors in spectral norm of constant size no matter how high the precision of energy estimation, at least using SVT-techniques. A new approach is needed.

To remove the superposition of rounding errors, Rall [Ral21] introduced the notion of a 'rounding promise'. Observing that certain λ_i are indecisive about their rounding, the rounding promise simply asserts that these λ_i do not appear in H. This assumption on H is very unphysical. In our work, we introduce a new notion of a rounding promise that can be guaranteed without restricting the Hamiltonian, and makes a similar guarantee. The central idea is similar: certain energies λ_i are disallowed. But instead of being an assumption on the Hamiltonian itself, our rounding promises define subspaces of the Hilbert space.

Definition 6. *A **rounding promise** M is a collection of $s^{(M)}$ many intervals $[a_x^{(M)}, b_x^{(M)}] \subseteq [0,1]$, enumerated by the label $x \in \{0, 1, \ldots, s^{(M)} - 1\}$, such that $b_x^{(M)} < a_{x+1}^{(M)}$ for all x. We always use the convention that the first interval starts at 0 and the last one ends at 1. A **gap** of M is a connected open subinterval $(b_x^{(M)}, a_{x+1}^{(M)})$ that spans the gap between two adjacent*

intervals $[a_x^{(M)}, b_x^{(M)}]$ *and* $[a_{x+1}^{(M)}, b_{x+1}^{(M)}]$. *If we write* $\lambda \in M$, *we mean that* λ *is contained in the union of all the intervals.*

The ***promised eigenspace projector*** $P_x^{(M)}$ *is the projector onto the eigenspaces of* $H = \sum_i \lambda_i \Pi_i$ *whose eigenvalues lie in the interval* $[a_x^{(M)}, b_x^{(M)}]$, *that is:*

$$P_x^{(M)} := \sum_{\substack{i \\ \lambda_i \in [a_x^{(M)}, b_x^{(M)}]}} \Pi_i. \tag{19}$$

The ***promised subspace projector*** $P^{(M)}$ *is the projector onto the eigensubspaces of* H *whose eigenvalues lie in* M, *that is:*

$$P^{(M)} := \sum_{x=0}^{s^{(M)}-1} P_x^{(M)}. \tag{20}$$

The ***promised subspace*** $\mathcal{P}^{(M)}$ *is*

$$\mathcal{P}^{(M)} := \textit{image of } P^{(M)}. \tag{21}$$

Remark 7. *When the rounding promise* M *is fixed, we often omit the superscript* (M) *to abbreviate the notation. For instance, we write* $\mathcal{P}$ *instead of* $\mathcal{P}^{(M)}$ *for the promised subspace projector.*

Rather than restricting the Hamiltonian itself, we have defined a subspace $\mathcal{P}^{(M)}$ on which energy estimation can be performed without superpositions of rounding errors. Furthermore, the rounding promise also conveniently specifies the estimates themselves: estimating λ amounts to identification of the index x such that $\lambda \in [a_x^{(M)}, b_x^{(M)}]$. Indeed, we have the following result:

Proposition 8 (Energy estimation given a rounding promise [Ral21]**).** *Say* M *is a rounding promise, say* κ *is the length of the smallest gap. Suppose also that the number of intervals* $s^{(M)}$ *satisfies* $s^{(M)} \leq 2^{n+1}$, *so each label can be thought of as a bit string* $x \in \{0,1\}^{n+1}$.

Then, for any $\delta_{\text{est}} > 0$, *there exists a family of operators* $\tilde{P}_x$ *that approximate* P_x *in the sense that:*

$$\left\| \tilde{P}_x P^{(M)} - P_x \right\| \leq \delta_{\text{est}}. \tag{22}$$

In fact, $P_x, \tilde{P}_x$ *and* $P^{(M)}$ *commute. Moreover, say we have a block encoding of a Hamiltonian* H. *Then, for any* δ_{est}, *there exists a quantum circuit that implements an isometry* $\tilde{E}^{(M)}$ *satisfying:*

$$\tilde{E}^{(M)} = \sum_x |x\rangle \otimes G_{\tilde{P}_x} \tag{23}$$

where the $G_{\tilde{P}_x}$ *are isometries satisfying* $G_{\tilde{P}_x}^\dagger G_{\tilde{P}_x} = \tilde{P}_x$. *This circuit can be implemented using* $O(n^2 \kappa^{-1} \log\left(\delta_{\text{est}}^{-1}\right))$ *applications of the block encoding of* H *and 1- and 2-qubit gates.*

We prove this in Appendix B. The factor of $O(n^2)$ can be removed with some additional care using techniques from [Ral21], but we keep it here to simplify the algorithm. This factor corresponds to the $O(\log^2(\beta/\varepsilon))$ in the performance in Theorem 1.

3 Algorithm overview

While we cannot use Davies generators to exactly prepare the Gibbs state ρ_β on the entire space, we can prepare the promised Gibbs states $\rho_\beta^{(M)}$ which are restricted to a particular promised subspace $\mathcal{P}^{(M)}$.

Definition 9. *Let M be a rounding promise. For each x, let $m_x^{(M)}$ denote the midpoint of the interval $[a_x, b_x]$. For $\beta \geq 0$, the **promised Gibbs state** $\rho_\beta^{(M)}$ is the density matrix supported only on the promised subspace $\mathcal{P}^{(M)}$ such that*

$$\rho_\beta^{(M)} \propto \sum_x \exp\big(-\beta m_x^{(M)}\big)\, P_x^{(M)}. \tag{24}$$

*The **promised partition function** $\mathcal{Z}_\beta^{(M)}$ is the normalization factor*

$$\mathcal{Z}_\beta^{(M)} := \sum_x \exp\big(-\beta m_x^{(M)}\big)\, \mathrm{Tr}(P_x^{(M)}). \tag{25}$$

With this idea in place, we can give an informal high-level overview of our algorithm. Each of the major challenges in the algorithm's construction is treated in a section of this paper.

First, promised thermal states are in general not close to the ideal thermal state. However, we find that a suitable ensemble of promised thermal states is satisfactory.

Claim 10 (Section 4)**.** *The desired thermal state ρ_β for the Hamiltonian H can be approximated by an average of promised thermal states $\rho_\beta^{(M_j)}$ for suitably chosen rounding promises M_j, where $j \in \{0, \ldots, 2^r - 1\}$. In particular, let ρ^* be the ensemble average over the promised thermal states of the M_j. Then, if we perform energy estimation to n bits of precision:*

$$\|\rho^* - \rho_\beta\|_1 \leq \sqrt{\beta \cdot 2^{-n}} + 2 \cdot 2^{-r} \tag{26}$$

The main idea is that the rounding promises M_j have to be chosen such that each eigenvalue λ_i of the Hamiltonian H is contained in at least $2^r - 1$ rounding promises M_j. That way any individual eigenspace can be missing with probability at most 2^{-r}.

Second, we must restrict the dynamics of a Davies generator to a promised subspace. This 'promised Davies generator' will then be used to prepare the promised thermal states.

Definition 11 (Promised Davies generator). *Say $\mathcal{L}$ is a Davies generator on the full Hilbert space with coupling operators S_α. Say M is a rounding promise, and we are given an 'attenuation operator' $A^{(M)}$ that commutes with the Hamiltonian and satisfies:*

$$\left\| A^{(M)}(I - P^{(M)}) \right\| = 0. \tag{27}$$

*Then, the **promised Davies generator** $\mathcal{L}^{(M)}$ is a Lindblad operator defined by the jump operators $L_{\nu,\alpha}$ given by:*

$$L_{\nu,\alpha}^{(M)} = \sqrt{G(\nu)} S_\alpha^{(M)}(\nu). \tag{28}$$

Its Bohr frequencies ν are differences of the form $m_x^{(M)} - m_y^{(M)}$, where $m_x^{(M)}$ and $m_y^{(M)}$ are the midpoints of the intervals $[a_x, b_y]$ and $[a_y, b_y] \in M$, respectively. Its jump operators $S_\alpha^{(M)}(\nu)$ are given by

$$S_\alpha^{(M)}(\nu) = \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} P_x S_\alpha^{(M)} P_y. \tag{29}$$

Finally, the coupling operators $S_\alpha^{(M)}$ are given by

$$S_\alpha^{(M)} = A^{(M)} S_\alpha A^{(M)}. \tag{30}$$

In Section 4 we will construct the attenuation operator $A^{(M)}$ as well as show that the promised Davies generator has the following properties:

Claim 12 (Section 5)**.** *Say $\mathcal{L}$ is a Davies generator on the full Hilbert space with coupling operators S_α. Then, for any rounding promise M, there exists a* ***promised Davies generator*** *$\mathcal{L}^{(M)}$ with the following properties:*

- *If an input state σ is supported only on $\mathcal{P}^{(M)}$, the output state $e^{t\mathcal{L}^{(M)}}(\sigma)$ will be as well.*
- *The unique fixed point of $\mathcal{L}^{(M)}$ is the promised thermal state $\rho_\beta^{(M)}$.*
- *Numerical simulations indicate that the mixing time $\tilde{t}_{mix}$ of $\mathcal{L}^{(M)}$ is not too much slower than that of $\mathcal{L}$.*

The purpose of the attenuation operator $A^{(M)}$ is to ensure that the output of the Lindblad evolution according to $\mathcal{L}^{(M)}$ is confined to the promised subspace. Then, so long as time evolution starts in an initial state $\tilde{\sigma}^{(M)}$ that is approximately supported on the promised subspace $\mathcal{P}^{(M)}$, then the output state will be as well.

The attenuation operator can make the mixing time of the Lindblad evolution slower. First, even in the ideal case when $A^{(M)} = P^{(M)}$, the elimination of certain energy eigenspaces may result in slower mixing. Second, it is not possible to project onto the promised subspace perfectly, for similar reasons that it is not possible to estimate energies without superpositions of rounding errors. We solve this problem by attenuating some of the eigenspaces in the promised subspace as well. This may also result in slower mixing. We show via numerical study that neither of these effects make the mixing time too much worse in practice.

Remark 13 (**Well-rounded Hamiltonian**)**.** *Note that the promised Davies generator $\mathcal{L}^{(M)}$ is block diagonal with respect to the orthogonal decomposition $\mathcal{H} = \mathcal{P}^{(M)} \oplus \mathcal{P}^{(M)^\perp}$. It acts as 0 on $\mathcal{P}^{(M)^\perp}$, and acts on $\mathcal{P}^{(M)}$ as the Davies generator with respect to the promised Hamiltonian $H^{(M)} \coloneqq \sum_x m_x^{(M)} P_x^{(M)}$ and coupling operators $S_\alpha^{(M)}$ (where we view all operators as restricted to the promised subspace $\mathcal{P}^{(M)}$).*

Third and finally, we show how to construct an approximate block encoding $O_{\mathcal{L}^{(M)}}$ for each promised Davies generator. This lets us leverage Proposition 5 to simulate the Lindblad dynamics and prepare the promised thermal states.

Claim 14 (Section 6)**.** *Say we are given block encodings of some coupling operators S_α, and one of the rounding promises $M \in \{M_0, \ldots, M_{2^r-1}\}$. Suppose we perform energy estimation to n bits of precision. Then, there exists a quantum algorithm that implements time evolution for time t according to $\mathcal{L}^{(M)}$ to any precision $\delta_{\mathcal{L}} > 0$ in the diamond norm using:*

$$O\left(\gamma^{-1} \cdot n^2 2^{3n+r} t \cdot \mathrm{polylog}(t/\delta_{\mathcal{L}})\right) \tag{31}$$

invocations to the block encoding of H, where γ is an attenuation coefficient introduced in Section 5. We accomplish this by constructing and simulating time evolution according to an approximate promised Davies generator $\tilde{\mathcal{L}}^{(M)}$.

4 Averaging together promised thermal states

By selecting a rounding promise M restricting to a promised subspace $\mathcal{P}^{(M)}$, we have gained the ability to approximately prepare promised Gibbs states ${\rho_\beta}^{(M)}$. These Gibbs states have support only on $\mathcal{P}^{(M)}$, and thus may be far from the true thermal state ρ_β. In this section we show how to construct several rounding promises $M_j \in \{M_0, \ldots, M_{2^r-1}\}$, such that the ensemble average over the ${\rho_\beta}^{(M_j)}$'s is provably close to ρ_β.

Furthermore, to prepare an approximation of ${\rho_\beta}^{(M_j)}$, we require an initial state σ that is also approximately supported only on $\mathcal{P}^{(M_j)}$. This state is then fed as input to the promised Davies generator, whose dynamics are trivial outside of $\mathcal{P}^{(M_j)}$. We also show how to prepare these initial states in this section: it is achieved by measuring a POVM called the 'left-right POVM' on an arbitrary initial state.

The left-right POVM splits the computation into two branches, each corresponding to a family of left rounding promises $\bar{L}, L_j$ and right rounding promises $\bar{R}, R_j$. We use M as a symbol to denote either L or R depending on which branch we are in. An overview of the algorithm as a whole is as follows:

Step 1. Take an arbitrary initial state σ, and measure the left-right POVM defined in Section 4.1. Depending on the measurement outcome, the resulting state $\tilde{\sigma}^{(M)}$ will be approximately supported on $\mathcal{P}^{(\bar{M})}$, where $\bar{M} \in \{\bar{L}, \bar{R}\}$ is one of two 'fine-grained' rounding promises.

Step 2. Pick an index $j \in \{0, ..., 2^r - 1\}$ uniformly at random. This index determines a 'coarse-grained' rounding promise M_j, which is either L_j or R_j depending on the measurement outcome of the left-right POVM. These are defined in Section 4.2. Since the M_j are coarse-grainings of $\bar{M}$, the input state $\tilde{\sigma}^{(M)}$ is supported only on $\mathcal{P}^{(M_j)}$ for any j.

Step 3. Use the promised Davies generator to approximately prepare ${\rho_\beta}^{(M_j)}$. This is discussed in Sections 5 and 6.

Analysis. Let ρ^{*M} be the ensemble average over the index j that we selected in step 2 and used to prepare ${\rho_\beta}^{(M_j)}$. We show that for both $M \in \{L, R\}$, the density matrix ρ^{*M} is close in trace distance to the ideal thermal state ρ_β. We perform this analysis in Section 4.3.

This protocol is represented diagrammatically in Fig. 1.

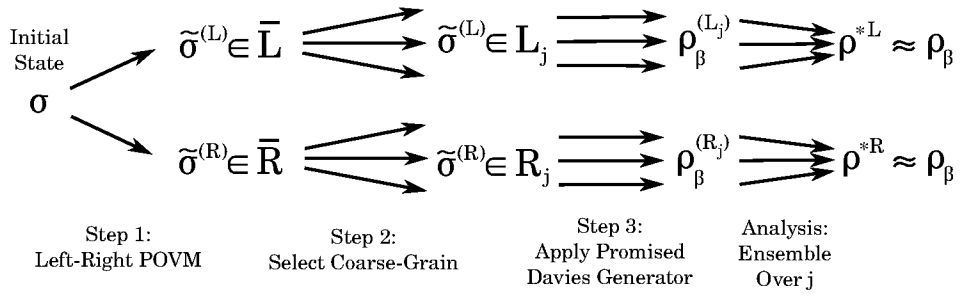

Figure 1: Sketch of a protocol that leverages promised Davies generators to prepare an approximation ρ^{*M} of an ideal thermal state ρ_β (for $M \in \{L, R\}$). In the figure above, $\rho \in \bar{M}$ is a shorthand for 'ρ is approximately supported entirely on $\mathcal{P}^{(\bar{M})}$', which we define more rigorously in Definition 15.

4.1 The Left-Right POVM

The Left-Right POVM has the purpose of producing a quantum state $\tilde{\sigma}^{(M)}$ which is guaranteed to be approximately supported on one of the two fine-grained rounding promises $\bar{M} \in \{\bar{L}, \bar{R}\}$ depending on the measurement outcome. In this section we show how to implement this POVM, and prove that the post-measurement state has the desired property. We define this notion rigorously now:

Definition 15. *Say M is a rounding promise and $\tilde{\sigma}$ a density matrix. We say $\tilde{\sigma}$ is δ_{sup}-**approximately supported on** M if*

$$\text{Tr}\left(\tilde{\sigma} P^{(M)}\right) \geq 1 - \delta_{\text{sup}}. \tag{32}$$

In other words, if we measure the projector $P^{(M)}$ in order to project into the promised subspace $\mathcal{P}^{(M)}$, we succeed with probability $\geq 1 - \delta_{\text{sup}}$.

The construction depends on two parameters: n and r. Here n is the number of bits of precision for energy estimation, and r determines the number of coarse-grained rounding promises, which is 2^r. The implementation of the left-right POVM as well as energy measurement using Proposition 8 will require $O(2^{n+r})$ applications of the block encoding of H. This is because all the rounding promises in this construction have a minimum gap size $\kappa = 2^{-n-r-2}$.

The quantities 2^{-n} and 2^{-r} correspond to two different errors on the final state. The quantity 2^{-n} corresponds to the accuracy of energy estimation, and 2^{-r} is the probability with which any particular energy is excluded from the ensemble. We will show in Theorem 24 that if output state is ρ^*, then we have the guarantee is $\|\rho^* - \rho_\beta\|_1 \leq \sqrt{\beta 2^{-n}} + 2 \cdot 2^{-r}$. In our final construction in Section 6 we will select $2^n \sim \beta/\varepsilon^2$ and $2^r \sim 1/\varepsilon$, achieving an accuracy of $\sim \varepsilon$.

As we present the construction, we recommend following along with Fig. 2. The main idea is that we would like to eliminate small regions of the spectrum via the Left-Right POVM, which is defined by an operator P_{LR} whose spectrum is sketched in Fig. 2. When P_{LR} has no support on an eigenstate, and we observe the POVM outcome corresponding to P_{LR}, then the output state has no support on that eigenstate. Consequently the output state satisfies the rounding promise $\bar{L}$. Similarly, the rounding promise $\bar{R}$ is defined by the eigenstates on which P_{LR} has eigenvalue 1. The key property of $\bar{L}$ and $\bar{R}$ that we require for the remainder of the construction is that there are 2^{n+r} many evenly spaced gaps in the spectrum. Later,

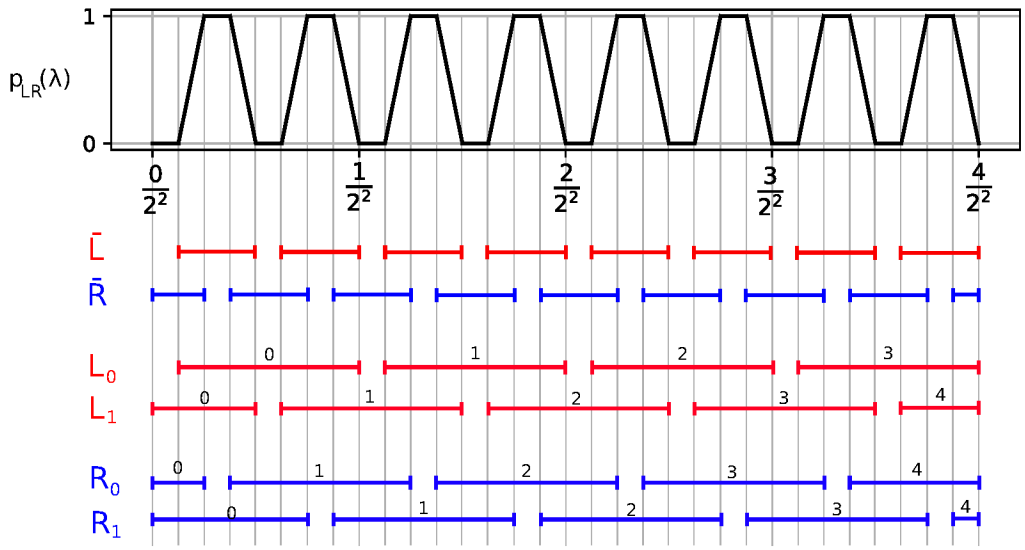

Figure 2: Sketch of the rounding promises from this section with $n = 2$, and $r = 1$. The diagram features a sketch of the function $p_{LR}(\lambda)$ that defines the projector P_{LR} from Lemma 17, a sketch of the fine-grained rounding promises from Definition 16, and a sketch of the coarse-grained rounding promises from Definition 21. The numbers above the intervals in the coarse-grained promises indicate the index x.

we will define coarse-grained rounding promises that close all but 2^n of these gaps at random, so that the probability of any individual energy being excluded by the coarse-grained promise is at most 2^{-r}.

Definition 16 (The fine-grained rounding promises). *For $n, r \in \mathbb{Z}_+$, let the rounding promises $\bar{L}$ and $\bar{R}$ be obtained by taking the full interval $[0,1]$ and removing some subintervals (a', b'). Specifically:*

$$\text{To construct } \bar{L}, \text{ delete } \left(\frac{k}{2^{n+r+2}}, \frac{k+1}{2^{n+r+2}},\right) \text{ for } k \in \{0, \ldots, 2^{n+r} - 1\}. \tag{33}$$

$$\text{To construct } \bar{R}, \text{ delete } \left(\frac{k+2}{2^{n+r+2}}, \frac{k+3}{2^{n+r+2}},\right) \text{ for } k \in \{0, \ldots, 2^{n+r} - 1\}. \tag{34}$$

Let the $[a_x, b_x]$ denote the resulting connected components. These closed intervals define the rounding promise in $\bar{L}$ and $\bar{R}$ respectively.

To prepare a state supported only on $\bar{L}$ or $\bar{R}$, we need to take our input state σ and remove all the support on the eigenspaces in the deleted regions above. To this end, we construct a block encoding of an operator P_{LR} that makes $P_{\mathrm{LR}}\sigma P_{\mathrm{LR}}$ be supported only on $\bar{L}$ and $(I - P_{\mathrm{LR}})\sigma(I - P_{\mathrm{LR}})$ be supported only on $\bar{R}$.

Lemma 17 (Left-right projection operator). *For any $n, r \in \mathbb{Z}_+$ and any $\delta_{\mathrm{sup}} > 0$, there exists a Hermitian operator P_{LR} that satisfies $\left\|P_{\mathrm{LR}}^2 \Pi_i\right\| \leq \delta_{\mathrm{sup}}/3$ for $\lambda_i \notin \bar{L}$, and satisfies $\left\|(I - P_{\mathrm{LR}}^2)\Pi_i\right\| \leq \delta_{\mathrm{sup}}/3$ for $\lambda_i \notin \bar{R}$. This operator has a block encoding that can be implemented using $O((n+r)2^{n+r} \log\left(\delta_{\mathrm{sup}}^{-1}\right))$ invocations of a block encoding of H.*

Proof. See Appendix B. We will ensure P_{LR} commutes with the Hamiltonian, and takes the form $P_{\mathrm{LR}} = \sum_i p_{\mathrm{LR}}(\lambda_i)\Pi_i$ for a function p_{LR} that is plotted in Fig. 2. □

It remains to show how to turn P_{LR} into a POVM, and how to implement that POVM using a quantum circuit. The starting point for the implementation is the block-measurement theorem from [Ral21]. This theorem takes a block encoding of an operator A satisfying $|A^2 - \Pi| \leq \delta$ for some projector Π, and uses it to implement a quantum channel $4\sqrt{2\delta}$-close in $\diamond$-norm to the isometry:

$$|1\rangle \otimes \Pi + |0\rangle \otimes (I - \Pi). \tag{35}$$

This does not quite suffice for us, since P_{LR} is not a projector. But, if we allow ourselves to introduce a constant amount of postselection, then we can implement a very similar operation that suffices for our purposes.

Lemma 18 (Postselective block-measurement). *Say A is a Hermitian matrix with a block encoding U_A. Then there exists a quantum circuit with postselection using one U_A and one $U_A^\dagger$ that implements the map $\sigma \to V_A \sigma V_A^\dagger / Tr(V_A \sigma V_A^\dagger)$, where:*

$$V_A \coloneqq |1\rangle \otimes A^2 + |0\rangle \otimes (I - A^2). \tag{36}$$

The postselection succeeds with probability at least $1/2$.

Proof. Say the block encoding U_A has k ancillae, that is $\left(\left\langle 0^k\right| \otimes I\right) U_A \left(\left|0^k\right\rangle \otimes I\right) = A$. We consider the following circuit:

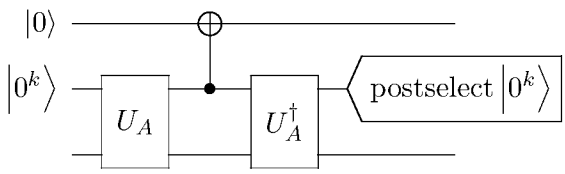

(37)

where the CNOT above denotes the operator $X \otimes |0^k\rangle\langle 0^k| + I \otimes (I - |0^k\rangle\langle 0^k|)$. When the postselection succeeds, this implements V_A as desired. It remains to bound the postselection probability, which is:

$$\mathrm{Tr}(V_A \sigma V_A^\dagger) = \mathrm{Tr}(A^2 \sigma A^2 + (I - A^2)\sigma(I - A^2)) \tag{38}$$

$$= \mathrm{Tr}(\sigma \cdot (A^4 + (I - A^2)^2)). \tag{39}$$

Since A is Hermitian and $x^4 + (1 - x^2)^2 \geq 1/2$, we have $A^4 + (I - A^2)^2 \geq I/2$. So we succeed with probability $\geq 1/2$. □

If we measure the top qubit after applying V_A, we will implement a POVM defined by the operator A.

All the tools are in place to define and implement the left-right POVM and demonstrate that it guarantees that the post-measurement state $\tilde{\sigma}^{(M)}$ has the desired property of being approximately supported on $\bar{M}$. One caveat is that the procedure only works if the probability of the observed outcome is bounded away from 0 by a constant. This is because the normalization of the state that occurs post-measurement may blow up the support outside the promised subspace. Fortunately, such a lower bound on the outcome probability is easily attained as we will see later.

Proposition 19 (Left-right POVM). *For any $n, r \in \mathbb{Z}_+$ and any $\delta_{\sup} > 0$, there exists a two-outcome POVM with the following property. We label the two outcomes 'L' and 'R', and for any input state σ let the output state be called $\tilde{\sigma}^{(M)}$ for $M \in \{L, R\}$. Say p_M is the probability of the M outcome, and suppose we obtain an outcome with $p_M \geq 1/3$. Then the output state $\tilde{\sigma}^{(M)}$ is $\delta_{\sup}$-approximately supported on $\bar{M}$.*

Furthermore, there exists a quantum circuit that implements the POVM with success probability $\geq 1/2$ using $O(2^{n+r})$ invocations of a block encoding of H.

Proof. The circuit is just an invocation of postselective block-measurement from Lemma 18 with the left-right projection operator P_{LR} from Lemma 17. After successfully applying the isometry $V_{P_{\mathrm{LR}}}$ we measure the label qubit and output 'L' in the 0 case and 'R' in the 1 case. This immediately gives the bound on the circuit complexity and success probability. It remains to show that the output state has the desired property.

The probabilities of the measurement outcomes are:

$$p_L := \mathrm{Tr}(P_{\mathrm{LR}}^2 \sigma P_{\mathrm{LR}}^2) \qquad \text{and} \qquad p_R := \mathrm{Tr}((I - P_{\mathrm{LR}}^2)\sigma(I - P_{\mathrm{LR}}^2). \tag{40}$$

The corresponding output states are:

$$\tilde{\sigma}^{(L)} := P_{\mathrm{LR}}^2 \sigma P_{\mathrm{LR}}^2 / p_L \qquad \text{and} \qquad \tilde{\sigma}^{(R)} := (1 - P_{\mathrm{LR}}^2)\sigma(I - P_{\mathrm{LR}}^2)/p_R. \tag{41}$$

Another way of writing the results of Lemma 17 is:

$$\left\| P_{\mathrm{LR}}^2 (I - P^{(\bar{L})}) \right\| \leq \delta_{\sup}/3 \text{ and } \left\| (I - P_{\mathrm{LR}}^2)(I - P^{(\bar{R})}) \right\| \leq \delta_{\sup}/3. \tag{42}$$

Having set the scene, we are ready the compute the probability of the output register being approximately supported on $\mathcal{P}^{(\bar{M})}$. We use the inequality $\mathrm{Tr}(XY) \leq \mathrm{Tr}(X)\mathrm{Tr}(Y)$ that holds for all positive semidefinite matrices X and Y. We obtain:

$$\begin{aligned}
\mathrm{Tr}(p_L \tilde{\sigma}^{(L)} \cdot (I - P^{(\bar{L})})) &= \mathrm{Tr}(P_{\mathrm{LR}}^2 \sigma P_{\mathrm{LR}}^2 \cdot (I - P^{(\bar{L})})) && (43)\\
&= \mathrm{Tr}(\sigma \cdot P_{\mathrm{LR}}^2 (I - P^{(\bar{L})}) P_{\mathrm{LR}}^2) && (44)\\
&\leq \mathrm{Tr}(\sigma) \cdot (\delta_{\sup}/3)^2 \leq \delta_{\sup}/3 \quad \text{and} && (45)\\
\mathrm{Tr}(p_R \tilde{\sigma}^{(R)} \cdot (I - P^{(\bar{R})})) &= \mathrm{Tr}((I - P_{\mathrm{LR}}^2)\sigma(I - P_{\mathrm{LR}}^2) \cdot (I - P^{(\bar{R})})) && (46)\\
&= \mathrm{Tr}(\sigma \cdot (I - P_{\mathrm{LR}}^2)(I - P^{(\bar{R})})(I - P_{\mathrm{LR}}^2)) && (47)\\
&\leq \mathrm{Tr}(\sigma) \cdot (\delta_{\sup}/3)^2 \leq \delta_{\sup}/3. && (48)
\end{aligned}$$

So for $M \in \{L, R\}$, we have:

$$p_M \cdot \mathrm{Tr}\left(\tilde{\sigma}^{(M)} P^{(\bar{M})}\right) \leq \delta_{\sup}/3. \tag{49}$$

Since we assumed $p_M \geq 1/3$, we have $\mathrm{Tr}\left(\tilde{\sigma}^{(M)} P^{(\bar{M})}\right) \leq \delta_{\sup}$ as desired.

□

Achieving $p_M \geq 1/3$ is very easy: we just repeat the protocol a couple of times and pick the outcome we saw the most frequently.

Corollary 20 (Selecting a high-probability outcome). *For any $\delta_{\text{fail}}, \delta_{\text{sup}} > 0$ there exists a procedure that produces a quantum state $\tilde{\sigma}^{(M)}$ that, for some random $\bar{M} \in \{\bar{L}, \bar{R}\}$, is δ_{sup}-approximately supported on $\bar{M}$ with failure probability at most δ_{fail}. Here, 'failure' refers to selecting an outcome $\bar{M}$ with $p_M < 1/3$. This procedure requires at most $20 \cdot \log\left(\delta_{\text{fail}}^{-1}\right)$ many re-preparations of the input state σ and implementations of the left-right POVM from Proposition 19.*

Proof. The procedure is as follows. Let N be the smallest odd number greater than $20\log(\delta_{\text{fail}})$. We repeat the left-right POVM from Proposition 19 a total of N times, preparing a new initial state σ each time. We let M be the most frequently observed outcome, and return the output state of any of our attempts that measured the M result.

We fail if $p_M < 1/3$. Let K be the number of times we observed the M outcome. Then, using the Chernoff-Hoeffding theorem, the probability that we observe $K/N > 1/2$ despite the fact that $p_M < 1/3$ is:

$$\Pr[K/N > 1/2] \leq \Pr[K/N > \mathbb{E}[K]/N + 1/6] \leq \exp(2N(1/6)^2) \leq \delta_{\text{fail}}. \tag{50}$$

□

4.2 Coarse-grained rounding promises

After having obtained a state $\tilde{\sigma}^{(M)}$ that (on average) is approximately supported on a fine-grained rounding promise $\bar{M}$, the next step of the algorithm is to select a coarse-grained rounding promise $M_j \in \{M_0, ..., M_{2^r-1}\}$. Because $\bar{M} \subset M_j$ for all M_j, we automatically have that $\tilde{\sigma}^{(M)}$ is approximately supported on each of the M_j.

The purpose of the coarse-grained rounding promises M_j is to ensure that the ensemble average ρ^{*M} of the promised thermal states ${\rho_\beta}^{(M_j)}$ is close to the ideal thermal state ρ_β. This would not be true if we just prepared the two states ${\rho_\beta}^{(\bar{M})}$ instead and considered their ensemble average, as we will explain now.

Recall that the main mechanism of the rounding promise is to eliminate certain energy eigenspaces. On the one hand, this enables energy estimation without superpositions of rounding errors. On the other hand, since certain energies are eliminated, an individual promised thermal state could never guaranteed be close in trace norm to ρ_β.

In an ensemble average ρ^{*M} over promised thermal states ${\rho_\beta}^{(M_j)}$ however, we can guarantee that the probability of each individual energy being eliminated is small ($\leq 2^{-r}$). We achieve this condition by selecting each of the M_j with probability $1/2^r$, and showing each energy is eliminated by at most one of the M_j's. We will show in the next section that this suffices to guarantee that $\|\rho^{*M} - \rho_\beta\| \leq \sqrt{\beta 2^{-n}} + 2 \cdot 2^{-r}$.

An ensemble over $\rho_\beta^{(\bar{M})}$ does not have this property: if an energy λ_i is eliminated by, say, $\bar{L}$, then the probability of the eigenspace being eliminated in the ensemble is p_L. If we leave the initial state arbitrary, then p_L could be as large as 1. But even if we set σ to the maximally mixed state, thereby ensuring $p_L = p_R = 1/2$, the probability of eliminating the energy is still constant, resulting in a constant-size error. To guarantee an error of at most ε, we require $O(\varepsilon^{-1})$ many rounding promises. We will show that they have width at most 2^{-n}, so the error from energy rounding will be at most $\sim 2^{-n}$.

We require one more property of the M_j's: recall from Definition 9 that the promised thermal state's eigenvalues are $\propto \exp\left(-\beta m_x^{(M_j)}\right)$, where m_x is the midpoint of the interval

$[a_x, b_x]$. This is because when we perform energy estimation according to Proposition 8, all of the energies in the interval $[a_x, b_x]$ are rounded to m_x. To avoid introducing too much error here, we must ensure that the intervals are not too wide.

The construction of the M_j is also visualized in Fig. 2. The intuition for the construction is as follows. We want to achieve $\bar{M} \subset M_j$, so we construct the M_j my merging gaps between the intervals of $\bar{M}$. As discussed above, the widths of the intervals in $\bar{M}$ are $\sim 2^{-n-r}$, and we want the widths of the intervals of M_j to be $\sim 2^{-n}$. This can be achieved by merging all but every 2^r'th gap. The promise M_j is defined by keeping every gap with index $= j$ modulo 2^r.

Definition 21 (**The coarse-grained rounding promises**). *For $n, r \in \mathbb{Z}_+$ and $M \in \{L, R\}$, let the rounding promises $M_j \in \{M_0, ..., M_{2^r-1}\}$ be constructed by taking the fine-grained rounding promise $\bar{M}$ from Definition 16 and merging some of the gaps.*

By merging the x'th gap, we mean the following transformation on a rounding promise: if the promise contains the intervals $[a_x, b_x]$ and $[a_{x+1}, b_{x+1}]$, then we take any intervals $[a', b_x], [a_{x+1}, b']$ for some a', b', and replace them with $[a', b']$. If $[a', b_x]$ is the final interval, then we replace it with $[a', 1]$. Finally, if $x = 0$, then we take the first interval $[a', b']$ and replace it with $[0, b']$.

Then, M_j is obtained by starting with $\bar{M}$ and merging all the x'th gaps where $x \neq j \bmod 2^{n+r}$.

Lemma 22 (**Properties of the coarse-grained rounding promises**). *First, if $\tilde{\sigma}^{(M)}$ is δ_{sup}-approximately supported on $\bar{M}$, then it is also approximately supported on all the M_j's.*

Second, say $\lambda \in [0, 1]$ is any energy. Then at most one of the M_j's does not contain it. Here is another way of saying this: let $1_{\lambda \in M_j}$ be 1 if $\lambda \in M_j$ and 0 otherwise. Then:

$$2^{-r} \sum_{j=0}^{2^r-1} 1_{\lambda \notin M_j} \leq 2^{-r}. \tag{51}$$

Finally, each interval $[a_x, b_x] \in M_j$ satisfies $b_x - a_x \leq 2^{-n}$.

Proof. For the first claim, since the gap merging procedure ensures that $\bar{M} \subset M_j$, we have that $P^{(\bar{M})} \leq P^{(M_j)}$. Consequently, following Definition 15:

$$\text{Tr}(\tilde{\sigma}^{(M)} P^{(M_j)}) \geq \text{Tr}(\tilde{\sigma}^{(M)} P^{(\bar{M})}) \geq 1 - \delta_{\text{sup}}. \tag{52}$$

For the second claim, we first consider what would happen if we merged *all* the gaps in $\bar{M}$: since the gap-merging procedure includes cases for the first and last interval, this would just result in the promise $\{[0, 1]\}$. So that means that for each M_j the only energies that are missing are those present in the gaps (b_x, a_{x+1}) where $x = j \bmod 2^{n+r}$. For every gap index x, there is a unique j for which this equation holds.

For every λ there are two cases. If we have $\lambda \in \bar{M}$, then since $\bar{M} \subset M_j$ we also have $\lambda \in M_j$ for all M_j. Otherwise, since $\lambda \notin \bar{M}$, λ must be contained in some gap of $\bar{M}$. By the above, that gap is contained in all but one of the M_j, so λ must be contained in all but one of the M_j's.

Finally, we must show that the intervals in the M_j are not too wide. The intervals in $\bar{M}$ all have width $3/4 \cdot 2^{-n-r}$, and are $1/4 \cdot 2^{-n-r}$ apart. If we close all but every 2^r'th gap, then the intervals will have width at most:

$$2^r \cdot 3/4 \cdot 2^{-n-r} + (2^r - 1) \cdot 1/4 \cdot 2^{-n-r} \leq (3/4 + 1/4) \cdot 2^{-n} \leq 2^{-n}. \tag{53}$$

□

4.3 Ensemble analysis

In this subsection we prove that the ensembles ρ^{*M} are close to the ideal thermal state ρ_β.

The error comes from two sources: from the accuracy of energy estimation (via n), and from number of rounding promises (via r). We first deal with the error from the accuracy of energy estimation by defining a new notion of an 'exact' promised thermal state $\hat{\rho}_\beta^{(M_j)}$.

Definition 23. *Let M be a rounding promise. Let the **exact promised partition function** $\hat{\mathcal{Z}}_\beta^{(M)}$ and **exact promised thermal state** $\hat{\rho}_\beta^{(M)}$ be defined similarly as their non-exact versions, just leaving λ_i intact rather than swapping them out with the midpoints m_x of the intervals containing them.*

$$\hat{\mathcal{Z}}_\beta^{(M)} := \sum_i 1_{\lambda_i \in M} \exp(-\beta\lambda_i), \quad \textit{and} \tag{54}$$

$$\hat{\rho}_\beta^{(M)} := \sum_i 1_{\lambda_i \in M} \Pi_i \exp(-\beta\lambda_i) / \hat{\mathcal{Z}}_\beta^{(M)}. \tag{55}$$

These exact promised thermal states are still restricted to the promised subspace, but unlike the regular promised thermal states ${\rho_\beta}^{(M_j)}$ their energies are unrounded and left intact. Thus, the difference between $\hat{\rho}_\beta^{(M_j)}$ and ${\rho_\beta}^{(M_j)}$ captures the error from energy estimation only. We will show that $\left\|\hat{\rho}_\beta^{(M_j)} - {\rho_\beta}^{(M_j)}\right\|_1 \leq \sqrt{\beta 2^{-n}}$. Then, we consider the ensemble $\hat{\rho}^{*M}$ of the states $\hat{\rho}_\beta^{(M_j)}$. It follows that $\left\|\hat{\rho}^{*M} - \rho^{*M}\right\|_1 \leq \sqrt{\beta 2^{-n}}$.

Second, we bound the error stemming from the elimination of certain energies in the rounding promises. In the previous subsection, we established that the probability that any particular energy λ is missing from M_j is at most $1/2^r$. Through careful consideration of the promised partition functions ${\mathcal{Z}_\beta}^{(M_j)}$, exact promised partition functions $\hat{\mathcal{Z}}_\beta^{(M_j)}$ and the ideal partition function $\mathcal{Z}_\beta$, we use this property to show that $\left\|\hat{\rho}^{*M} - \rho_\beta\right\| \leq 2 \cdot 2^{-r}$.

Combining these two facts yields the main theorem of this subsection:

Theorem 24 (Accuracy of the final ensemble). *Consider the following ensembles of density matrices:*

$$\rho^{*M} := \frac{1}{2^r} \sum_{j=0}^{2^r-1} \rho_\beta^{(M_j)} \quad \textit{for } M \in \{L, R\}. \tag{56}$$

Then, we have $\left\|\rho^{*M} - \rho_\beta\right\|_1 \leq \sqrt{\beta 2^{-n}} + 2 \cdot 2^{-r}$,

Proof. Combining Lemmas 27 and 28 with the triangle inequality, we obtain $\left\|\rho^{*M} - \rho_\beta\right\|_1 \leq \sqrt{\beta 2^{-n}} + 2 \cdot 2^{-r}$ for both $M \in \{L, R\}$. □

We start with the analysis of the rounding errors via the exact promised thermal states. To prove that these are close together, we leverage a fact from [PW09]:

Lemma 25 (Comparing thermal states via their Hamiltonians [PW09, Appendix C]**).** *Say H and H' are two Hamiltonians. Let $\rho_\beta := e^{-\beta H} / \operatorname{Tr}\left(e^{-\beta H}\right)$ and $\rho'_\beta := e^{-\beta H'} / \operatorname{Tr}\left(e^{-\beta H'}\right)$ be their Gibbs states, respectively. Then, if $F(\sigma, \sigma') = \left\|\sqrt{\sigma}\sqrt{\sigma'}\right\|_1$ is the fidelity:*

$$F(\rho, \rho') \geq \exp(-\beta |H - H'|). \tag{57}$$

This tool is convenient because a naive analysis of the trace distance yields an contribution from every eigenvalue, which accumulates into a total error proportional to the dimension. By considering the spectral distance in the Hamiltonian instead, we avoid the exponential blowup in the error. However, this trick is a bit tricky in our case since, when viewed as operators on the full Hilbert space, ${\rho_\beta}^{(M_j)}$ is not full rank. But states of the form $e^{-\beta H}/\mathrm{Tr}(e^{-\beta H})$ are always full rank whenever H is.

To remedy this issue, we recall an observation made in Remark 13, where we noticed that if we view ${\rho_\beta}^{(M_j)}$ as a density matrix over $\mathcal{P}^{(M_j)}$ only, then it can be written as $e^{-\beta H^{(M_j)}}/\mathrm{Tr}(e^{-\beta \bar{H}^{(M_j)}})$ where $H^{(M_j)}$ is the promised Hamiltonian. We will leverage this viewpoint where we restrict to $\mathcal{P}^{(M_j)}$ for Definition 26 and Lemma 25.

Definition 26. *Let M be a rounding promise. Let the **exact promised Hamiltonian** $\hat{H}^{(M)}$ be the part of H with support on the promised subspace $\mathcal{P}^{(M)}$. But unlike the non-exact promised Hamiltonian $H^{(M)}$, the eigenvalues are unchanged. It is given by:*

$$\hat{H}^{(M)} \coloneqq \sum_i \lambda_i \Pi_i. \tag{58}$$

Here, we view this operator as $\hat{H}^{(M)} \in \mathrm{L}(\mathcal{P}^{(M)})$.

Lemma 27 (Rounding errors from energy estimation). *For $M \in \{L, R\}$ consider the following ensemble over exact promised thermal states:*

$$\hat{\rho}^{*M} \coloneqq 2^{-r} \sum_j \hat{\rho}_\beta^{(M_j)}. \tag{59}$$

This density matrix satisfies $\left\|\hat{\rho}^{*M} - \rho^{*M}\right\|_1 \leq \sqrt{\beta 2^{-n}}$.

Proof. Recall from Lemma 22 that each of the intervals $[a_x, b_x]$ of the M_j satisfies $b_x - a_x \leq 2^{-n}$. Thus, if $\lambda_i \in [a_x, b_x]$, we have $|\lambda_i - m_x| \leq 2^{-n-1}$ where m_x is the midpoint of $[a_x, b_x]$. Consequently, we have $\left\|\hat{H}^{(M_j)} - H^{(M_j)}\right\| \leq 2^{-n-1}$.

Following Lemma 25, we have $F(\hat{\rho}_\beta^{(M_j)}, \rho_\beta^{(M_j)}) \geq e^{-\beta 2^{-n}}$, so, exploiting the fact that $\|\sigma - \sigma'\|_1 \leq \sqrt{1 - F(\sigma, \sigma')^2}$, we have:

$$\left\|\hat{\rho}_\beta^{(M_j)} - \rho_\beta^{(M_j)}\right\|_1 \leq \sqrt{1 - F(\hat{\rho}_\beta^{(M_j)}, \rho_\beta^{(M_j)})^2} \leq \sqrt{1 - e^{-2\beta 2^{-n-1}}} \leq \sqrt{\beta 2^{-n}}. \tag{60}$$

Finally, we average this error over the two ensembles:

$$\left\|\hat{\rho}^{*M} - \rho^{*M}\right\|_1 = \left\|2^{-r} \sum_j \left(\hat{\rho}_\beta^{(M_j)} - \rho_\beta^{(M_j)}\right)\right\|_1 \leq 2^{-r} \sum_j \left\|\hat{\rho}_\beta^{(M_j)} - \rho_\beta^{(M_j)}\right\|_1 \leq \sqrt{\beta 2^{-n}}. \tag{61}$$

□

We move on to the error stemming from the elimination of energies in the promises M_j.

Lemma 28 (The average of exact promised thermal states approximates the thermal state). *For $M \in \{L, R\}$, we have $\left\|\hat{\rho}^{*M} - \rho_\beta\right\|_1 \leq 2 \cdot 2^{-r}$.*

Proof. Letting $d_i := \mathrm{Tr}(\Pi_i)$ be the occupation numbers, we have:

$$\left\|\hat{\rho}^{*(M)} - \rho_\beta\right\|_1 = \sum_i d_i \left| 2^{-r} \sum_j 1_{\lambda_i \in M_j} \exp\left(-\beta\lambda_i\right) / \hat{\mathcal{Z}}_\beta^{(M_j)} - \exp(-\beta\lambda_i)/\mathcal{Z}_\beta \right| \tag{62}$$

$$\leq \sum_i d_i \left| 2^{-r} \sum_{\substack{j \\ \lambda_i \notin M_j}} (0 - \exp(-\beta\lambda_i)/\mathcal{Z}_\beta) \right| \tag{63}$$

$$+ \sum_i d_i \left| 2^{-r} \sum_{\substack{j \\ \lambda_i \in M_j}} \left(\exp\left(-\beta\lambda_i\right) / \hat{\mathcal{Z}}_\beta^{(M_j)} - \exp(-\beta\lambda_i)/\mathcal{Z}_\beta\right) \right|. \tag{64}$$

Here, for each λ_i, we split the rounding promises M_j into two categories: those not containing λ_i, for which $\hat{\rho}_\beta^{(M)}$ has eigenvalue 0, and those containing λ_i, for which $\hat{\rho}_\beta^{(M)}$ has eigenvalue $\exp\left(-\beta\lambda_i\right) / \hat{\mathcal{Z}}_\beta^{(M_j)}$. After applying the triangle inequality to these two groups, these correspond to the first and second terms in the above sum, respectively. We will show that both of these terms are bounded by 2^{-r}, so the overall bound is $2 \cdot 2^{-r}$ as desired.

The first term is bounded by:

$$\sum_i d_i \left| 2^{-r} \sum_{\substack{j \\ \lambda_i \notin M_j}} (0 - \exp(-\beta\lambda_j)/\mathcal{Z}_\beta) \right| \leq 2^{-r} \cdot \frac{1}{\mathcal{Z}_\beta} \sum_i d_i \exp(-\beta\lambda_i) = 2^{-r} \cdot \frac{\mathcal{Z}_\beta}{\mathcal{Z}_\beta} = 2^{-r}. \tag{65}$$

So all that is left is the second term. Observe that $\hat{\mathcal{Z}}_\beta^{(M_j)} \leq \mathcal{Z}_\beta$ which implies that $1/\hat{\mathcal{Z}}_\beta^{(M_j)} - 1/\mathcal{Z}_\beta \geq 0$. This allows us to remove the absolute value:

$$\sum_i d_i \left| 2^{-r} \sum_{\substack{j \\ \lambda_i \in M_j}} \left(\exp\left(-\beta\lambda_i\right) / \hat{\mathcal{Z}}_\beta^{(M_j)} - \exp(-\beta\lambda_i)/\mathcal{Z}_\beta\right) \right| \tag{66}$$

$$= \sum_i d_i \exp\left(-\beta\lambda_i\right) \cdot \left| 2^{-r} \sum_j 1_{\lambda_i \in M_j} \left(1/\hat{\mathcal{Z}}_\beta^{(M_j)} - 1/\mathcal{Z}_\beta\right) \right| \tag{67}$$

$$= \sum_i d_i \exp\left(-\beta\lambda_i\right) \cdot 2^{-r} \sum_j 1_{\lambda_i \in M_j} \left(1/\hat{\mathcal{Z}}_\beta^{(M_j)} - 1/\mathcal{Z}_\beta\right). \tag{68}$$

Now we can bound the two terms corresponding to the difference $1/\hat{\mathcal{Z}}_\beta^{(M_j)} - 1/\mathcal{Z}_\beta$ indi-

vidually. The term with $1/\hat{\mathcal{Z}}_\beta^{(M_j)}$ is is just 1:

$$\sum_i d_i \exp\left(-\beta\lambda_i\right) \cdot 2^{-r} \sum_j 1_{\lambda_i \in M_j} \left(1/\mathcal{Z}_\beta^{(M_j)}\right) \tag{69}$$

$$= 2^{-r} \sum_j \left[\sum_i d_i \exp\left(-\beta\lambda_i\right) 1_{\lambda_i \in M_j}\right] \left(1/\mathcal{Z}_\beta^{(M_j)}\right) \tag{70}$$

$$= 2^{-r} \sum_j \mathcal{Z}_\beta^{(M_j)} \left(1/\mathcal{Z}_\beta^{(M_j)}\right) = 1. \tag{71}$$

The term with $-1/\mathcal{Z}_\beta$ can be bounded by leveraging that for any λ_i, at most one $M^{(j)}$ doesn't contain it, as shown in Lemma 22. That is, $2^{-r}\sum_j 1_{\lambda_i \in M_j} \geq 1 - 2^{-r}$.

$$-\sum_i d_i \exp\left(-\beta\lambda_i\right) \cdot 2^{-r} \sum_j 1_{\lambda_i \in M_j} \left(1/\mathcal{Z}_\beta\right) \tag{72}$$

$$\leq -\sum_i d_i \exp\left(-\beta\lambda_i\right) \cdot \left(1 - 2^{-r}\right) \left(1/\mathcal{Z}_\beta\right) \tag{73}$$

$$= -\mathcal{Z}_\beta \cdot \left(1 - 2^{-r}\right) \left(1/\mathcal{Z}_\beta\right) = -(1 - 2^{-r}). \tag{74}$$

So a bound on the second term is $1 - (1 - 2^{-r}) = 2^{-r}$. So, each of the two terms is at most 2^{-r}, so the total error is at most $2 \cdot 2^{-r}$. □

5 Lindblad dynamics on the promised subspace

In the previous section we required the following capability: to take a state $\sigma^{(M)}$ that is supported on a rounding promise M, and to transform it to the promised thermal state ${\rho_\beta}^{(M)}$. Once we have this capability, we can prepare good approximations to the true thermal state ρ_β.

We attain this capability by running a promised Davies generator $\mathcal{L}^{(M)}$ on the system, which we defined in Definition 11. We constructed $\mathcal{L}^{(M)}$ by starting with a regular Davies generator $\mathcal{L}$ and truncating it so that its dynamics are restricted to $\mathcal{P}$. The basic idea is to consider the coupling operators S_α of the original Davies generator, and to conjugate them with an attenuation operator $A^{(M)}$ to obtain:

$$S_\alpha^{(M)} \coloneqq A^{(M)} S_\alpha A^{(M)} \tag{75}$$

Then, $\mathcal{L}^{(M)}$ is defined the exact same way as $\mathcal{L}$, just with ${S_\alpha}^{(M)}$ instead of S_α. The other difference is that $\mathcal{L}$ is defined with respect to the ideal Hamiltonian $H = \sum_i \lambda_i \Pi_i$, and $\mathcal{L}^{(M)}$ with respect to the promised Hamiltonian $H^{(M)} = \sum_x m_x \Pi_x$.

In this section we detail the construction of the attenuation operator $A^{(M)}$ and how that the resulting promised Davies generator has some desirable properties. Ideally, we would like to have $A^{(M)} = P^{(M)}$, but due to limitations of singular value transformation, we cannot achieve this. This results in two challenges.

The first challenge is *leakage*: we cannot perfectly eliminate the transfer of probability mass from $\mathcal{P}$ to $\mathcal{P}^\perp$; we can only suppress it by a factor δ_{leak}. Fortunately, this error is both exponentially suppressible, and its ramifications on the final output error can be quantified

formally. We give a construction of an approximate attenuation operator $\tilde{A}^{(M)}$ that has support at most δ_{leak} on $\mathcal{P}^{\perp}$. The operator $A^{(M)}$ itself has no support on $\mathcal{P}^{\perp}$ at all, ensuring that $\mathcal{L}^{(M)}$ does not leak. The leakage error can be dealt with by studying its approximate implementation $\tilde{\mathcal{L}}^{(M)}$ which we construct in Section 6.

The second challenge, which is more significant, is *attenuation*. Since $A^{(M)}$ is implemented via singular value transformation, it takes the form $A^{(M)} = \sum_i a^{(M)}(\lambda_i)\Pi_i$ for some continuous function $a^{(M)} : [0,1] \to [0,1]$. We have already achieved that $a^{(M)}(\lambda) = 0$ for $\lambda \notin M$. Ideally, we would like to also have $a^{(M)}(\lambda) = 1$ for $\lambda \in M$, but this means that $a^{(M)}$ has a discontinuity at the boundaries of M. The function $a^{(M)}$ must be continuous, so we need to smoothly interpolate from 0 to 1 near the edges of the intervals of M. Formally, we can define a truncated rounding promise M^T whose intervals are a little bit narrower than those of M, and then ensure $a^{(M)}(\lambda) = 1$ for $\lambda \in M^T$. However, this choice means that the operator also attenuates certain eigenvectors of H whose eigenvalues are close to the edges of the intervals of M. This attenuation might slow the rate of convergence of $\mathcal{L}^{(M)}$.

Having highlighted the trade-offs of the construction, we formally define the attenuation operator $A^{(M)}$ and its approximation $\tilde{A}^{(M)}$.

Lemma 29 (**Attenuation operators**). *Consider any rounding promise M with minimum gap width κ whose intervals are also at least κ wide. Then, for any leakage error $\delta_{\text{leak}} > 0$, and attenuation factor γ there exists an* ***attenuation operator*** *$A^{(M)}$ and an* ***approximate attenuation operator*** *$\tilde{A}^{(M)}$. They both commute with the Hamiltonian, and take the form:*

$$A^{(M)} = \sum_i a^{(M)}(\lambda_i)\Pi_i \qquad \text{and} \qquad \tilde{A}^{(M)} = \sum_i \tilde{a}^{(M)}(\lambda_i)\Pi_i \tag{76}$$

for some functions $a^{(M)}, \tilde{a}^{(M)} : [0,1] \to [0,1]$. Let M^T be a $\kappa\gamma$-truncated rounding promise where if $[a_x, b_x] \in M$, then $[a_x + \kappa\gamma, b_x - \kappa\gamma] \in M^T$. Then, $a^{(M)}, \tilde{a}^{(M)} : [0,1] \to [0,1]$ satisfy:

$$\text{if } \lambda \notin M \text{ then } a^{(M)}(\lambda) = 0 \text{ and } \tilde{a}^{(M)}(\lambda) \leq \delta_{\text{leak}}, \tag{77}$$

$$\text{if } \lambda \in M^T \text{ then } a^{(M)}(\lambda) = 1 \text{ and } \tilde{a}^{(M)}(\lambda) \geq 1 - \delta_{\text{leak}}, \tag{78}$$

$$\text{otherwise } a^{(M)}(\lambda) = \tilde{a}^{(M)}(\lambda). \tag{79}$$

Consequently, we have $\|A^{(M)} - \tilde{A}^{(M)}\| \leq \delta_{\text{leak}}$, and $\|A^{(M)}(I - P^{(M)})\| = 0$ as required by Definition 11. Furthermore, there exists a block encoding of $\tilde{A}^{(M)}$ with circuit complexity $O\left(\kappa^{-1}\gamma^{-1}\log\left(s^{(M)}\delta_{\text{leak}}^{-1}\right)\right)$.

The proof is in Appendix B. We will show how to deal with leakage errors in Section 6. This section is dedicated to assessing the impact of truncation and attenuation on the dynamics of $\mathcal{L}^{(M)}$. Our goal is to show that $\mathcal{L}^{(M)}$ is mixing (recall Definition 4) and that the mixing time t_{mix} is not too much slower than that of $\mathcal{L}$.

On a positive note, we emphasize that truncation to $\mathcal{P}$ as well as attenuation cannot cause $\mathcal{L}^{(M)}$ to fail to converge onto the promised thermal state unless the coupling operators S_α are extremely contrived. For example, one way $\mathcal{L}$ could be mixing while $\mathcal{L}^{(M)}$ could fail to be so is if the S_α encode a particular random walk over the energy eigenspaces. Specifically, if the random walk is over a bipartite graph where one half of the eigenspaces are in $\mathcal{P}$ and the other half are in $\mathcal{P}^{\perp}$, then a Davies generator based on S_α may be mixing, while $S_\alpha^{(M)}$

will not be. But construction of such a coupling operator requires careful knowledge over the energy eigenspaces. The chances of accidentally selecting such an S_α in practice are slim.

The main problem is that the projection of S_α to ${S_\alpha}^{(M)}$ will slow the rate of mixing, and that attenuation only makes this problem worse. We assess the severity of this problem by performing numerical simulations on a particular physical system: the transverse field Ising model. Let (n_1, n_2) be the width and height of 2d-grid of qubits, and let v be the strength of the transverse field. Let $\vec{i}$ be indices in an $n_1 \times n_2$ grid, let $X_{\vec{i}}$ and $Z_{\vec{i}}$ denote Pauli X and Z on the qubit at grid position $\vec{i}$, and let $\vec{i} \sim \vec{j}$ denote that the two grid positions are adjacent. Then, the Hamiltonian we consider in this section is:

$$H = \sum_{\vec{i} \sim \vec{j}} Z_{\vec{i}} Z_{\vec{j}} + v \sum_{\vec{i}} X_{\vec{i}}. \tag{80}$$

We study the dynamics of Davies generators defined with respect to this Hamiltonian, as well as the coupling operators $X_{\vec{i}}$ and $Z_{\vec{i}}$ for all grid positions $\vec{i}$. To assess the mixing time t_{mix} we compute the spectral gap Δ of the Lindbladian.

In Fig. 3 we plot spectral gap as a function of the attenuation factor γ for various parameters. In all our experiments, we verified that the ideal Davies generator $\mathcal{L}$ had the ideal thermal state ρ_β as its unique fixed point, and that the promised Davies generators $\mathcal{L}^{(M_j)}$ had the promised thermal state $\rho_\beta^{(M_j)}$ as their unique fixed point. We analyze their exact versions with $A^{(M)}$ rather than $\tilde{A}^{(M)}$ and P_x rather than $\tilde{P}_x$.

We find that in the ideal case of $\gamma = 0$, the ideal Davies generator and the promised Davies generators have about the same convergence rate - the promised Davies generators are a little slower. So while there is some slowdown in mixing time from projection, the slowdown is largely due to attenuation. We find that setting reasonably small values of γ achieves convergence rates similar to the $\gamma = 0$ case. Furthermore, this slowdown only appears to be significant for certain rounding promises M_j. One interpretation of our technique in this paper is that we trade mixing time for a provable accuracy guarantee on the final thermal state. So, some amount of slowdown is acceptable in exchange for the rigor.

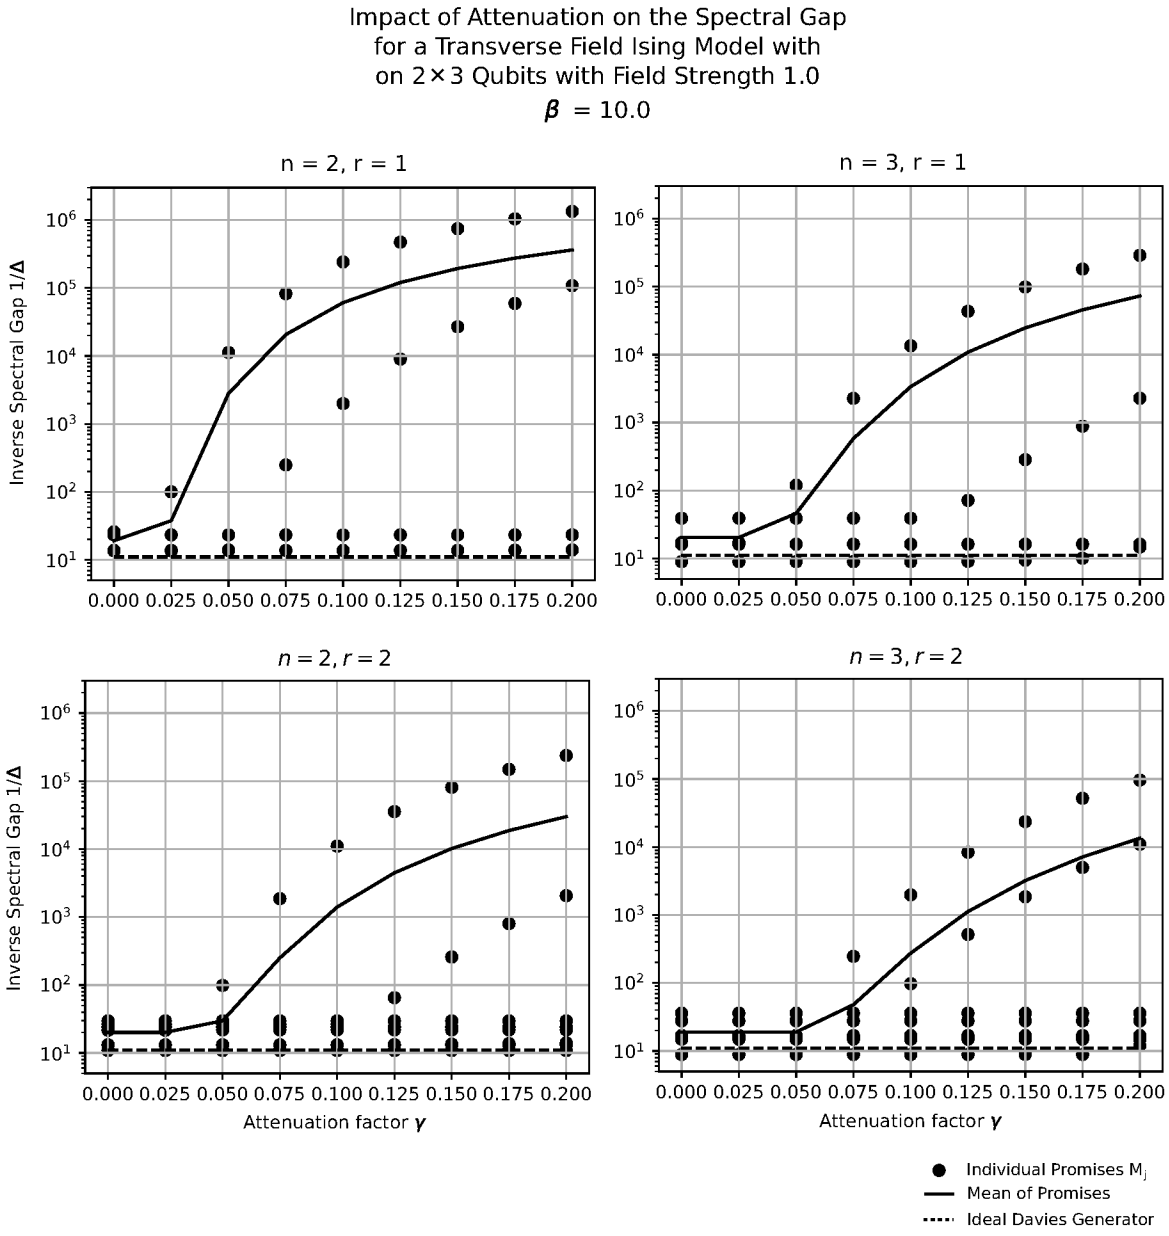

Figure 3: Analysis of the spectral gap of the Davies generators of the transverse field Ising model from Eq. (80). We set $\beta = 10.0$ throughout. We see that the convergence rates of the promised Davies generators are a little slower than the ideal Davies generator as $\gamma \to 0$. While the convergence rate slows down significantly with large γ, selecting reasonably small values of γ achieves a convergence time just as fast as the unphysical case of $\gamma = 0$. There is also significant variation among the promises M_j. The software that produced these data is available at: https://github.com/qiskit-community/promised-davies-generator.

6 Implementing Lindblad dynamics

The goal of this section is to show how to simulate the evolution of the promised Davies generator for time t given a block encoding of its jump operators. This block encoding can be implemented by energy estimation. This construction also involves the coupling operator projected onto the promised subspace. Both energy estimation and the projection are not perfect, and they lead to errors in the block-encoded jump operators. We begin with a lemma that shows that a small perturbation of the jump operator will not change the Lindbladian by much. Then we give the detailed construction and prove the main theorem.

For simplicity, in this section we assume that there is only one coupling operator so that we can drop the index α in for the jump operators $L_{\nu,\alpha}$ and the coupling operators S_α, and simply write L_ν for the jump operators of a promised Davies generator (Definition 11), and write S for the coupling operator. Our analysis can be easily generalized to the case of more coupling operators: the number of jump operators just grows by the factor of the number of the coupling operators.

Furthermore, it is reasonable to suppose that $\|S\| \leq 1$. First, this is generally the case in practice: the coupling operator S is usually selected to be a unitary transformation that scrambles the energy eigenbasis. Even if S were a block encoded operator, it would require a unitary implementation, so the only way to achieve $\|S\| \geq 1$ would be via some virtual scale factor. Second, in the analysis of Hamiltonian simulation e^{iHt}, it is usually assumed that $\|H\| \leq 1$, since we can always absorb a rescaling of the Hamiltonian into the time t. The exact same argument applies for the simulation of Davies generators.

To begin with, we show that the distance between Lindbladians in terms of the diamond norm can be bounded by the distances of their jump operators in terms of the spectral norm.

Lemma 30 (**Approximating Lindblad evolution via jump operators**). *Say $L_\nu \in \mathrm{L}(\mathcal{H}^{(M)})$ are jump operators defining a Lindbladian $\mathcal{L}$, and say there are m many of them. Say some $\tilde{L}_\nu \in \mathrm{L}(\mathcal{H})$ similarly define a Lindbladian $\tilde{\mathcal{L}}$, and say these satisfy:*

$$\|\tilde{L}_\nu - L_\nu\| \leq \delta_L. \tag{81}$$

Say we also have $\|L_\nu\| \leq 1$. Then, for any $t > 0$:

$$\|\tilde{\mathcal{L}} - \mathcal{L}\|_\diamond \leq 4m\delta_L + 2\delta_L^2 \tag{82}$$

Proof. First of all, by Eq. (81) and the triangle inequality, we have

$$\|\tilde{L}_\nu\| \leq \|L_\nu\| + \delta_L. \tag{83}$$

Now, recall that

$$\mathcal{L}(\rho) = \sum_\nu L_\nu \rho L_\nu^\dagger - \frac{1}{2}\left(L_\nu^\dagger L_\nu \rho + \rho L_\nu^\dagger L_\nu\right), \quad \text{and} \tag{84}$$

$$\tilde{\mathcal{L}}(\rho) = \sum_\nu \tilde{L}_\nu \rho \tilde{L}_\nu - \frac{1}{2}\left(\tilde{L}_\nu^\dagger \tilde{L}_\nu \rho + \rho \tilde{L}_\nu^\dagger \tilde{L}_\nu\right). \tag{85}$$

Let ρ be the matrix that achieves the induced trace norm of $\|\mathcal{L} - \tilde{\mathcal{L}}\|_1$, i.e.,

$$\|\mathcal{L} - \tilde{\mathcal{L}}\|_1 = \|\mathcal{L}(\rho) - \tilde{\mathcal{L}}(\rho)\|_1. \tag{86}$$

We have

$$\|\mathcal{L}-\tilde{\mathcal{L}}\|_1 \leq \sum_\nu \left\|L_\nu\rho L_\nu^\dagger - \tilde{L}_\nu\rho\tilde{L}_\nu^\dagger\right\|_1 + \frac{1}{2}\left\|\tilde{L}_\nu^\dagger\tilde{L}_\nu\rho - L_\nu^\dagger L_\nu\rho\right\|_1 + \frac{1}{2}\left\|\rho\tilde{L}_\nu^\dagger\tilde{L}_\nu - \rho L_\nu^\dagger L_\nu\right\|_1 \tag{87}$$

$$\leq \sum_\nu \|L_\nu\|\delta_L + \|\tilde{L}_\nu\|\,\delta_L + \|\tilde{L}_\nu\tilde{L}_\nu - L_\nu^\dagger L_\nu\| \tag{88}$$

$$\leq \sum_\nu 2(\|L_\nu\|\delta_L + \|\tilde{L}_\nu\|\,\delta_L) \tag{89}$$

$$\leq 4m\delta_L + 2\delta_L^2, \tag{90}$$

where we have used the fact that $\|ABC\|_1 \leq \|A\|\|B\|_1\|C\|$ for matrices A, B and C.

To extend this bound to the diamond norm, we consider the identity map $\mathcal{I} : \mathrm{L}(\mathcal{H}^{(M)}) \to \mathrm{L}(\mathcal{H}^{(M)})$. We have

$$\mathcal{L}\otimes\mathcal{I}(\rho) = \sum_\nu (L_\nu\otimes I)\rho(L_\nu^\dagger\otimes I) - \frac{1}{2}\left((L_\nu^\dagger L_\nu\otimes I)\rho + \rho(L_\nu^\dagger L_\nu\otimes I)\right), \quad \text{and} \tag{91}$$

$$\tilde{\mathcal{L}}\otimes\mathcal{I}(\rho) = \sum_\nu (\tilde{L}_\nu\otimes I)\rho(\tilde{L}_\nu^\dagger\otimes I) - \frac{1}{2}\left((\tilde{L}_\nu^\dagger\tilde{L}_\nu\otimes I)\rho + \rho(\tilde{L}_\nu^\dagger\tilde{L}_\nu\otimes I)\right). \tag{92}$$

Then it is easy to obtain that

$$\|\mathcal{L}-\tilde{\mathcal{L}}\|_\diamond = \|\mathcal{L}\otimes\mathcal{I} - \tilde{\mathcal{L}}\otimes\mathcal{I}\|_1 \leq 4m\delta_L + 2\delta_L^2, \tag{93}$$

using the fact that $\|A\otimes I\| = \|A\|$ for any matrix A. □

The Lindbladian we try to simulate is the promised Davies generator $\mathcal{L}^{(M)}$, which involves the promised coupling operator $S^{(M)} = A^{(M)}SA^{(M)}$. However, in our construction, we implement an approximate attenuation operator $\tilde{S}^{(M)} = \tilde{A}^{(M)}S\tilde{A}^{(M)}$ (see Section 5). Based on the distance between $A^{(M)}$ and $\tilde{A}^{(M)}$ given by Lemma 29, we prove a bound on the distance between $S^{(M)}$ and $\tilde{S}^{(M)}$ as follows.

Lemma 31. *Given a rounding promise M, let $S^{(M)}$ and $\tilde{S}^{(M)}$ be as defined in Section 5. For any leakage error $\delta_{\mathrm{leak}} > 0$, it holds that*

$$\left\|S^{(M)} - \tilde{S}^{(M)}\right\| \leq 2\delta_{\mathrm{leak}}. \tag{94}$$

Proof. Recall that $S^{(M)} = A^{(M)}SA^{(M)}$, and $\tilde{S}^{(M)} = \tilde{A}^{(M)}S\tilde{A}^{((M)}$. By Lemma 29, we have

$$\left\|A^{(M)} - \tilde{A}^{(M)}\right\| \leq \delta_{\mathrm{leak}}. \tag{95}$$

To prove the desired bound, we leverage $\left\|A^{(M)}\right\| \leq 1$ and $\left\|\tilde{A}^{(M)}\right\| \leq 1$ to show:

$$\left\|S^{(M)} - \tilde{S}^{(M)}\right\| = \left\|A^{(M)}SA^{(M)} - \tilde{A}^{(M)}S\tilde{A}^{(M)}\right\| \tag{96}$$

$$= \left\|A^{(M)}SA^{(M)} - A^{(M)}S\tilde{A}^{(M)} + A^{(M)}S\tilde{A}^{(M)} - \tilde{A}^{(M)}S\tilde{A}^{(M)}\right\| \tag{97}$$

$$\leq \delta_{\mathrm{leak}} + \delta_{\mathrm{leak}} = 2\delta_{\mathrm{leak}}. \tag{98}$$

Note that we have assumed that $\|S\| \leq 1$ as in the beginning of this section. □

Now, we have all the tools to prove the main theorem and present our quantum algorithm for simulating the Davies generator.

Theorem 32 (Implementation of the Davies generator given a rounding promise). *For some rounding promise M, say $\mathcal{L}^{(M)}$ is an promised Davies generator as defined in Definition 11. Then, for any $\delta_{\mathcal{L}}, t > 0$, there exists a quantum algorithm implementing a channel $\delta_{\mathcal{L}}$-close in diamond norm to $e^{t\mathcal{L}^{(M)}}$. If κ is the minimum gap of M, M has $s^{(M)}$ many intervals, and γ is the desired attenuation factor for the $\tilde{S}^{(M)}$, then this algorithm uses*

$$O\left((s^{(M)})^2 t \frac{\log\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}{\log\log\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}\right) \tag{99}$$

queries to the block encoding of S,

$$O\left(t \cdot \kappa^{-1} \cdot \gamma^{-1} \cdot (s^{(M)})^2 \log\left(s^{(M)}\right) \cdot \frac{\log^2\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}{\log\log\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}\right) \tag{100}$$

queries to the block encoding of H, and

$$O\left((s^{(M)})^4 t \left(\frac{\log\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}{\log\log\left(s^{(M)}t/\delta_{\mathcal{L}}\right)}\right)^2\right) \tag{101}$$

additional 1- and 2-qubit gates.

Proof. We aim to invoke Proposition 5, so we need an oracle $\mathcal{O}_{\tilde{\mathcal{L}}}$. This lets us approximately apply the map $e^{t\tilde{\mathcal{L}}}$. We then show that the $\tilde{L}_\nu$ defining $\tilde{\mathcal{L}}$ are close to the L_ν defining $\mathcal{L}^{(M)}$, allowing us to invoke Lemma 30.

Recall the approximate energy estimation isometry $\tilde{E}^{(M)}$ from Proposition 8, which extracts estimates according to $\tilde{P}_x$ satisfying $|\tilde{P}_x P^{(M)} - P_x| \leq \delta_{\text{est}}$ via isometries $G_{\tilde{P}_x}$ satisfying $G_{\tilde{P}_x}^\dagger G_{\tilde{P}_x} = \tilde{P}_x$. We define the block-encoded operator:

$$V_\pm^{(M)} :=$$

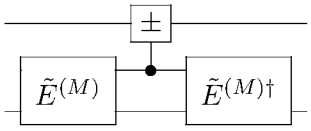

(102)

where the controlled ± refers to adding/subtracting $m_x{}^{(M)}$ to/from the value of the target register (recall Remark 13, the control register has the value of x), and the kets on the right side of the circuit denote postselection onto that state. Next, obtain a block encoding of:

$$G := \sum_\nu \sqrt{G(\nu)} \otimes |\nu\rangle\langle\nu| \,. \tag{103}$$

Finally, let $\mathcal{O}_{\tilde{\mathcal{L}}}$ be given by:

$$\mathcal{O}_{\tilde{\mathcal{L}}} :=$$

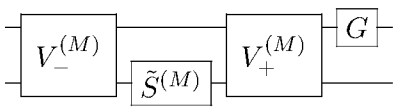

(104)

and hence $\mathcal{O}_{\tilde{\mathcal{L}}}$ implements some $\tilde{L}_\nu$ in the following manner:

$$\mathcal{O}_{\tilde{\mathcal{L}}}(|0\rangle \otimes |\psi\rangle) = \sum_\nu |\nu\rangle \otimes \tilde{L}_\nu |\psi\rangle . \tag{105}$$

Let's find an expression for the $\tilde{L}_\nu$. We begin by observing that:

$$V_\pm^{(M)} = \sum_\nu \sum_x |\nu \pm m_j{}^{(M)}\rangle\langle\nu| \otimes G^\dagger_{\tilde{P}_x} G_{\tilde{P}_x} = \sum_\nu \sum_x |\nu \pm m_j{}^{(M)}\rangle\langle\nu| \otimes \tilde{P}_x. \tag{106}$$

That way:

$$\mathcal{O}_{\tilde{\mathcal{L}}}(|0\rangle \otimes |\psi\rangle) = \sum_{x,y} \sqrt{G(m_x^{(M)} - m_y^{(M)})} \left| m_x^{(M)} - m_y^{(M)} \right\rangle \otimes \tilde{P}_x \tilde{S}^{(M)} \tilde{P}_y |\psi\rangle \tag{107}$$

$$= \sum_\nu \sqrt{G(\nu)} |\nu\rangle \otimes \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x \tilde{S}^{(M)} \tilde{P}_y |\psi\rangle . \tag{108}$$

We see that:

$$\tilde{L}_\nu = \sqrt{G(\nu)} \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x \tilde{S}^{(M)} \tilde{P}_y. \tag{109}$$

Observing $|G(\nu)| \leq 1$ and recalling the definition of $S(\omega)$, all that is left to do is to leverage Proposition 8 and Lemma 31 to bound the distance between $\tilde{L}_\nu$ and L_ν.

First observe that $\|\tilde{P}_x\| \leq 1$ which follows from the fact that $\tilde{P}_x$ is obtained from singular value transformation. By Proposition 5, we have

$$\left\| \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} P_x S^{(M)} P_y - \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x P^{(M)} S^{(M)} P^{(M)} \tilde{P}_y \right\| \tag{110}$$

$$= \left\| \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} P_x S^{(M)} P_y - \tilde{P}_x P^{(M)} S^{(M)} P_y + \tilde{P}_x P^{(M)} S^{(M)} P_y - \tilde{P}_x P^{(M)} S^{(M)} P^{(M)} \tilde{P}_y \right\| \tag{111}$$

$$\leq 2(s^{(M)})^2 \delta_{\text{est}}. \tag{112}$$

Note that $P^{(M)} S^{(M)} P^{(M)} = S^{(M)}$. The above inequality implies that

$$\left\| \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} P_x S^{(M)} P_y - \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x S^{(M)} \tilde{P}_y \right\| \leq 2(s^{(M)})^2 \delta_{\text{est}}. \tag{113}$$

Now, using Lemma 31, we obtain

$$\left\| \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x S^{(M)} \tilde{P}_y - \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x \tilde{S}^{(M)} \tilde{P}_y \right\| \tag{114}$$

$$= \left\| \sum_{\substack{x,y \\ m_x^{(M)} - m_y^{(M)} = \nu}} \tilde{P}_x (S^{(M)} - \tilde{S}^{(M)}) \tilde{P}_y \right\| \tag{115}$$

$$\leq (s^{(M)})^2 \delta_{\mathrm{leak}}. \tag{116}$$

Combining Eqs. (113) and (116), we have

$$\left\| L_\nu - \tilde{L}_\nu \right\| \leq (s^{(M)})^2 (2\delta_{\mathrm{est}} + \delta_{\mathrm{leak}}). \tag{117}$$

So we have the desired property with $(s^{(M)})^2(\delta_{\mathrm{leak}} + 2\delta_{\mathrm{est}})$. Letting $\delta_L := (s^{(M)})^2(\delta_{\mathrm{leak}} + 2\delta_{\mathrm{est}})$ and invoking Lemma 30 with $m = (s^{(M)})^2$, we have

$$\|e^{\mathcal{L}t} - e^{\tilde{\mathcal{L}}t}\|_\diamond \leq t \|\mathcal{L} - \tilde{\mathcal{L}}\|_\diamond \tag{118}$$

$$\leq t \cdot \left[4(s^{(M)})^2 (\delta_{\mathrm{leak}} + 2\delta_{\mathrm{est}}) + 2(\delta_{\mathrm{leak}} + 2\delta_{\mathrm{est}})^2 \right] \tag{119}$$

$$= 16 (s^{(M)})^2 t (\delta_{\mathrm{leak}} + 2\delta_{\mathrm{est}}), \tag{120}$$

where the first inequality follows from the observation that for all integers $k \geq 0$,

$$\|e^{\mathcal{L}t} - e^{\tilde{\mathcal{L}}t}\|_\diamond = \|(e^{\mathcal{L}t/k})^k - (e^{\tilde{\mathcal{L}}t/k})^k\|_\diamond \leq k \|e^{\mathcal{L}t/k} - e^{\tilde{\mathcal{L}}t/k}\|_\diamond \leq \|\mathcal{L} - \tilde{\mathcal{L}}\|_\diamond + O(t/k). \tag{121}$$

Here the first inequality follows from the subadditivity of the diamond norm [Wat18, Proposition 3.48], and the last inequality is due to Taylor expansion.

We assume S is given as a block encoding with normalizing constant 1 since $\|S\| = 1$. To use Proposition 5, it suffices to set

$$\delta_{\mathrm{leak}} = O\left(\frac{\delta_{\mathcal{L}}}{(s^{(M)})^4 t} \right), \quad \text{and } \delta_{\mathrm{est}} = O\left(\frac{\delta_{\mathcal{L}}}{(s^{(M)})^4 t} \right) \tag{122}$$

to $\delta_{\mathcal{L}}$-approximately simulate $e^{\mathcal{L}^{(M)} t}$ by simulating $e^{\tilde{\mathcal{L}}^{(M)} t}$. Observe that $\|L_\nu\| \leq 1$. Let k be the number of system qubits. The simulation algorithm costs

$$O\left((s^{(M)})^2 t \, \frac{\log\left(s^{(M)} t / \delta_{\mathcal{L}}\right)}{\log\log\left(s^{(M)} t / \delta_{\mathcal{L}}\right)} \right) \tag{123}$$

queries to $\mathcal{O}_{\tilde{\mathcal{L}}}$ and

$$O\left((s^{(M)})^4 t \left(\frac{\log\left(s^{(M)} t / \delta_{\mathcal{L}}\right)}{\log\log\left(s^{(M)} t / \delta_{\mathcal{L}}\right)} \right)^2 \right) \tag{124}$$

additional 1- and 2-qubit gates. Note that the number of queries to $O_{\tilde{\mathcal{L}}}$ is also the number of innovations to S.

To implement $\mathcal{O}_{\tilde{\mathcal{L}}}$, we invoke Proposition 8 with precision parameter δ_{est} and Lemma 29 with precision parameter δ_{leak}. Each application of $\mathcal{O}_{\tilde{\mathcal{L}}}$ has the following number of invocations of U_H:

$$O\left(\kappa^{-1}\log\left(s^{(M)}\right)^2\log\left(\delta_{\text{est}}^{-1}\right)+\kappa^{-1}\gamma^{-1}\log\left(s^{(M)}\delta_{\text{leak}}^{-1}\right)\right) \tag{125}$$

$$=O\left(\kappa^{-1}\log\left(s^{(M)}\right)^2\log\left(s^{(M)}t\delta_{\mathcal{L}}^{-1}\right)+\kappa^{-1}\gamma^{-1}\log\left(s^{(M)}\right)\log\left(s^{(M)}t\delta_{\mathcal{L}}^{-1}\right)\right) \tag{126}$$

$$=O\left(\kappa^{-1}\gamma^{-1}\log\left(s^{(M)}\right)^2\cdot\log\left(s^{(M)}t\delta_{\mathcal{L}}^{-1}\right)\right). \tag{127}$$

□

To achieve the final number of queries to the block encoding of H, we observe that the rounding promises M_j have $\kappa^{-1}\in O(2^{n+r})$ and $s^{(M)}\in O(2^n)$. Plugging these in, and accounting for the cost of the left-right POVM from Lemma 17 we obtain:

$$O\left(\gamma^{-1}\cdot n^2 2^{3n+r}t\cdot\text{polylog}(t/\delta_{\mathcal{L}})\right). \tag{128}$$

This establishes Claim 14. Now we recall Theorem 24 which states that the final output state ρ^{*M} satisfies:

$$\left\|\rho^{*M}-\rho_\beta\right\|_1\leq\sqrt{\beta 2^{-n}}+2\cdot 2^{-r}. \tag{129}$$

To achieve $\|\rho^{*M}-\rho_\beta\|_1\leq\varepsilon$, we select $n=\log_2(\beta(\varepsilon/2)^{-2})$ and $r=\log_2(4/\varepsilon)$. Plugging these into the query complexity, we get:

$$O\left(\gamma^{-1}\cdot t\cdot\beta^3\varepsilon^{-7}\cdot\text{polylog}(t/\delta_{\mathcal{L}})\cdot\log^2(\beta/\varepsilon)\right). \tag{130}$$

This establishes Theorem 1.

7 Some open questions

Our result achieved a time complexity that scales linearly in the mixing time t_{mix}. However, the performance with respect to the inverse temperature β and accuracy ε is $\tilde{O}(\beta^3\varepsilon^{-7})$ which has plenty of room for improvement. One potential path that may yield an accuracy of $\tilde{O}(\beta\varepsilon^{-2})$ is to remove the linear dependence on the number of jump operators in the Lindblad simulation algorithm from Proposition 5. Proposition 5 demands block encodings of the individual jump operators L_j, but the oracle $\mathcal{O}_{\mathcal{L}}$ we prepare may actually be much more powerful than this. To see this, first note that if we are given access to the isometry $\sum_j|j\rangle\otimes K_j$, then implementing the channel is trivial as we already have the Stinespring dilation. In our algorithm, we implemented an oracle in the form of $\sum_j|j\rangle\otimes L_j$, which is close to the Stinespring dilation we want because for the infinitesimal approximation channel in [CW17], all but one Kraus operators are proportional to L_j. That one special Kraus operator involves all the L_j's. Does there exist any special treatment of this special Kraus operator so that we can leverage the special structure of the oracle $\sum_j|j\rangle\otimes L_j$ to get rid of the $O(m)$ dependence?[1]

[1]As remarked earlier, after the first version of this manuscript was made public, recent work [CKBG23] resolved this question in the affirmative.

A central goal in our work is to attain a rigorous bound on the accuracy of the final output state. To this end, we assume that we are given a lower bound on the mixing time t_{mix}, so that we know for how long the dynamics of the Davies generator must be simulated to achieve a high-accuracy output state. But this assumption is rarely the case in practice. Furthermore, the attenuation discussed in Section 5 may slow the mixing time and exasperate this problem. Is there a technique that can detect if the Davies generator has been run for long enough? One approach might be to purify the dynamics of the Lindblad simulation, and then use amplitude estimation to compare the resulting pure output state to an ideal thermal state purification. A reflection operator around a purification of the thermal state could be obtained via the techniques from [WT23].

The only piece of our method that eludes rigorous mathematical treatment is the impact of attenuation on mixing time $\tilde{t}_{\text{mix}}$. Certainly the dependence of γ^{-1} in the circuit complexity the synthesis attenuation operators $A^{(M)}$ is optimal, due to lower bounds on approximation of threshold functions with polynomials. But perhaps given additional knowledge about the Hamiltonian, there may be other approaches for selecting coupling operators S_α that do not leak out of a given promised subspace.

8 Acknowledgements

The authors thank Kristan Temme and Chi-Fang Chen for helpful discussions, as well as anonymous reviewers for their helpful comments. CW was supported by a seed grant from the Institute of Computational and Data Science (ICDS) and a National Science Foundation grant CCF-2238766 (CAREER).

A Glossary

This manuscript features many quantities and mathematical symbols, so give a brief description of some of these with references to the relevant part of the text.

Hamiltonian. H is decomposed into eigenvalues λ_i and eigenspace projectors Π_i via $\sum_i \lambda_i \Pi_i$. Our goal is the prepare the thermal state ρ_β at inverse temperature β. See Definition 2.

Rounding promise. $M \subset [0, 1]$ defines a promised subspace $\mathcal{P}^{(M)}$ and a projector onto that subspace $P^{(M)}$. M consists of intervals $[a_x, b_x]$ with corresponding promised eigenspace projectors $P_x^{(M)}$. See Definition 6. We can perform energy estimation with respect to approximate promised eigenspace projectors $\tilde{P}_x^{(M)}$ using Proposition 8. We suppress the superscript (M) when a promise is clear from context.

Lindbladian. $\mathcal{L}$ is a superoperator defining a continuous time quantum Markov process describing open system dynamics. It is defined by the jump operators $L_{\omega,\alpha}$. See Eq. (13).

Davies Generator. A particular choice of jump operators $L_{\omega,\alpha} = \sqrt{G(\omega)} S_\alpha(\omega)$ that yields thermalizing dynamics. ω is among the Bohr frequencies of H: the set of pairwise energy differences. The filter function $G(\omega)$ biases the dynamics towards certain energy differences, and the coupling operators S_α 'scramble' the Hilbert space. See Definition 3.

Initial state. Our simulation of thermalizing dynamics starts in an arbitrary initial state σ. After measuring the left-right POVM defined by P_{LR}, we obtain an initial state $\tilde{\sigma}^{(M)}$ satisfying a rounding promise $M \in \{\bar{L}, \bar{R}\}$. See Section 4.1.

Attenuation. We require coupling operators $S^{(M)}$ that do not 'leak' out of the promised subspace $\mathcal{P}^{(M)}$. We achieve this by sandwiching them between attenuation operators $A^{(M)}$ with the same support as $P^{(M)}$, but also some eigenvalues in $\mathcal{P}^{(M)}$ are attenuated to be less than 1 which slows the mixing time. The attenuation coefficient γ controls the number of these eigenvalues, but an approximate implementation $\tilde{A}^{(M)}$ of $A^{(M)}$ requires circuit complexity $O(\gamma^{-1})$. See Section 5.

Promised thermal states. When the Hamiltonian H is truncated onto $\mathcal{P}^{(M)}$ we obtain the well-rounded Hamiltonian $H^{(M)}$ with eigenvalues $m_x^{(M)}$ and eigenspace projectors $P_x^{(M)}$, see Remark 13. On $\mathcal{P}^{(M)}$, it has the promised thermal state $\rho_\beta^{(M)}$ and promised partition function $\mathcal{Z}_\beta^{(M)}$, see Definition 9. In section Section 4.3, it is convenient to define 'exact' versions of these $\hat{\rho}_\beta^{(M)}, \hat{\mathcal{Z}}_\beta^{(M)}$ that are still truncated to $\mathcal{P}^{(M)}$ but are based on the eigenvalues λ_i of H.

Mixing time. We assume that the ideal Davies generator $\mathcal{L}$ with jump operators $L_{\omega,\alpha} = \sqrt{G(\omega)} S_\alpha(\omega)$ requires the mixing time t_{mix} in order to transform an arbitrary input state into something close to ρ_β. Due to attenuation, the promised Davies generator $\mathcal{L}^{(M)}$ with jump operators $L_{\omega,\alpha} = \sqrt{G(\omega)} S_\alpha^{(M)}(\omega)$ will have a slower mixing time $\tilde{t}_{\mathrm{mix}}$. See Definition 11.

B Polynomial construction

Throughout the paper, we required the construction of block encodings. Proposition 8 gave operators P_x that indicate the eigenspaces of a particular interval $[a_x, b_x]$ of a rounding promise M, and a unitary $U^{(M)}$ that computes the binary expansion $|x\rangle$ accordingly. Lemma 17 gave an operator P_{LR} can be used to force either the $\bar{L}$ or $\bar{R}$ rounding promise via POVM, by making P_{LR} small in $\mathcal{P}^{(\bar{L})}$ and large in $\mathcal{P}^{(\bar{R})}$. Finally, Lemma 29 gave an attenuation operator $A^{(M)}$ that vanishes outside of $\mathcal{P}^{(M)}$, but is simultaneously as large as possible within $\mathcal{P}^{(M)}$.

These operators P_x, P_{LR}, and $A^{(M)}$ have a lot in common: they all commute with the Hamiltonian, and can hence be seen as $\sum_i f(\lambda_i)\Pi_i$ for some function f. Furthermore, the requirements on f are always that $f(\lambda) = 0$ or $f(\lambda) = 1$ for λ in certain regions. We construct all of these operators through singular value transformation, which lets us construct such operators via polynomial approximations of f.

The requirements on the operators ensure that there are always gaps between the regions of λ where $f(\lambda) = 0$ or $f(\lambda) = 1$. This is essential for polynomial approximation of f, since polynomials are always continuous. The degree of the polynomial scales with the reciprocal of the width of the gap. Our starting point for the polynomial construction is a highly accurate polynomial approximation of a step function.

Lemma 33 (Polynomial approximation of a step function ([LC17, Appendix A])**.** *For any $\kappa, \delta > 0$, there exists an odd polynomial $\Theta_{\kappa,\delta}$ of degree $O(\kappa^{-1}\log(\delta^{-1}))$ such that*

$\forall \lambda \in [-2,2]$ *we have* $0 \le \Theta_{\kappa,\delta}(\lambda) \le 1$ *and:*

$$\text{if } \lambda \le -\kappa/2, \text{ then } \Theta_{\kappa,\delta}(x) < \delta; \tag{131}$$

$$\text{if } \kappa/2 \ge \lambda, \text{ then } 1-\delta \le \Theta_{\kappa,\delta}(\lambda). \tag{132}$$

Given a polynomial approximation of a single transition from $f(\lambda) = 0$ to $f(\lambda) = 1$, we can approximate an arbitrary sequence of transitions by shifting and adding multiple such polynomials together. Above, we guaranteed that the approximation of the step function holds for $\lambda \in [-2,2]$, so that we can shift the polynomial by up to 1 and still have good approximation on the interval $[-1,1]$.

Lemma 34 (Construction of approximate projection polynomials). *Say we have a collection of intervals* $\{[a_x, b_x]\}$ *in* $[0,1]$ *with* $a_x < b_x$ *and* $b_x < a_{x+1}$*, and each interval is labeled with a bit* $c_x \in \{0,1\}$*. Say furthermore that each of the intervals is at least* κ *far apart, that is,* $b_x + \kappa < a_{x+1}$*.*

Then, for any $\delta > 0$ *there exists a polynomial* $P(\lambda)$ *such that:*

$$\forall x, \text{ if } \forall \lambda \in [a_x, b_x], \text{ then } |P(\lambda) - c_x| \le \delta. \tag{133}$$

Furthermore, the polynomial satisfies $0 \le P(\lambda) \le 1$*. Say there are* K *many 'flips', that is, indexes* x *where* $c_x \ne c_{x+1}$*. Then the polynomial is of degree* $O(\kappa^{-1}\log(K\delta^{-1}))$*.*

Proof. We will construct $P(\lambda)$ by adding together several $\Theta_{\kappa_x,\delta/K}$ for various values of κ_x, all of which satisfy $\kappa_x \ge \kappa$. Since adding several polynomials doesn't change the degree, the resulting polynomial $P(\lambda)$ has degree $O(\kappa^{-1}\log(K\delta^{-1}))$.

There will be one $\Theta_{\kappa_x,\delta/K}$ for each flip, that is, each pair of intervals with $c_x \ne c_{x+1}$. Let $t_x \coloneqq (b_x + a_{x+1})/2$ be the midpoint between two intervals, and let $\kappa_x = a_{x+1} - b_x \ge \kappa$ be the distance between them. We construct:

$$P(\lambda) \coloneqq c_1 + \sum_{\substack{x \\ c_x \ne c_{x+1}}} \Theta_{\kappa_x,\delta/K}(\lambda - t_x) \cdot \begin{cases} +1 & \text{if } c_x < c_{x+1} \\ -1 & \text{if } c_x > c_{x+1} \end{cases}. \tag{134}$$

It remains to prove that $P(\lambda)$ satisfies the desired property. By Lemma 33 and the choice of t_x, κ_x it is guaranteed that $\Theta_{\kappa_x,\delta/K}(\lambda - t_x)$, is always either $\le \delta/K$ or $\ge 1 - \delta/K$ outside of the region $[b_x, a_{i+x}]$. So, for any interval $[a_x, b_x]$ we have that any of the $\Theta_{\kappa_x,\delta/K}(\lambda - t_x)$ is within δ/K of 0 or 1. Thus, for the purposes of analyzing $P(\lambda)$, let us pretend for the rest of the proof that they are *exactly* 0 or 1. In doing so, we will be wrong by at most $K \cdot \delta/K = \delta$.

Let us consider any particular interval $[a_x, b_x]$. We want to argue that for any λ in this interval, $P(\lambda) = c_i$. We achieve this through induction in K. The base case is very easy: if $K = 0$, then all the c_x are equal, so we can just set $P(\lambda) = c_1$.

Now, consider only the first $K-1$ switches, and let $P'(\lambda)$ be the polynomial only from these. If c_x is the last x such that $c_x \ne c_{x+1}$, then $P'(\lambda)$ satisfies $\lambda \in [a_y, b_y] \implies P'(\lambda) = c_y$ for $y \le x$, and $\lambda \in [a_y, b_y] \implies P'(\lambda) = 1 - c_y$ for $y > x$. The construction of $P(\lambda)$ depends on if $c_{x+1} = 0$ or $c_{x+1} = 1$. If $c_{x+1} = 0$, then $\lambda \in [a_{x+1}, b_{x+1}] \implies P'(\lambda) = 1$, so we need to subtract 1 for $\lambda > b_x$, which is achieved by subtracting $\Theta_{\kappa_x,\delta/K}(\lambda - t_x)$ from $P'(\lambda)$. Otherwise, if $c_{x+1} = 1$, then $\lambda \in [a_{x+1}, b_{x+1}] \implies P'(\lambda) = 0$, so we need to add 1 for $\lambda > b_x$, which is

achieved by adding $\Theta_{\kappa_x,\delta/K}(\lambda - t_x)$. The conditions $c_{x+1} = 0$ and $c_{x+1} = 1$ are equivalent to $c_x > c_{x+1}$ and $c_x < c_{x+1}$ respectively. Consequently:

$$P(x) = P'(x) + \Theta_{\kappa_x,\delta/K}(x - t_x) \cdot \begin{cases} +1 & \text{if } c_x < c_{x+1} \\ -1 & \text{if } c_x > c_{x+1} \end{cases}. \tag{135}$$

□

It remains to invoke singular value transformation in order to construct operators with this polynomial as their spectrum. The polynomial we have constructed has mixed parity, which is acceptable because we perform singular value transformation on a Hermitian operator. The method from [GSLW19] splits the polynomial into its even and odd parts and combines them together via a linear combination of block encodings. This introduces an extra factor of $\frac{1}{2}$. Since we only care about making eigenvalues either close to 0 or 1, we can use a simple version of oblivious amplitude amplification via the Chebyshev polynomial T_3 to remove this extra factor.

Lemma 35 (Construction of projectors via singular value transformation). *Say H is a Hermitian matrix with $-1 \leq H \leq 1$, and we are given a block encoding U_H of H. Say H has the eigendecomposition $H = \sum_i \lambda_i \Pi_i$. Say $P(\lambda)$ is a degree-d polynomial that, for some δ and for certain regions $\{[a_x, b_x]\}$ satisfies $\lambda \in [a_x, b_x] \implies |P(\lambda) - c_x| \leq \delta$ for $c_x \in \{0, 1\}$, and furthermore satisfies $0 \leq P(\lambda) \leq 1$ for all $\lambda \in [-1, 1]$.*

Then, there exists a quantum circuit $U_{\bar{P}(H)}$ which is a block encoding of a matrix $\bar{P}(H)$ defined by:

$$\bar{P}(H) := \sum_i \bar{P}(\lambda_i)\Pi_i. \tag{136}$$

where $\bar{P}$ is a function that satisfies $\lambda \in [a_x, b_x] \implies |\bar{P}(\lambda) - c_x| \leq 2\delta$. This circuit makes d many uses of controlled-U_H, and has circuit complexity $O(d)$ overall.

Proof. Our starting point is Theorem 56 of [GSLW19], which allows us to construct a block encoding of

$$\frac{1}{2}P(H) := \sum_i \frac{1}{2}P(\lambda_i)\Pi_i. \tag{137}$$

To get rid of the factor of $\frac{1}{2}$, let $T_3(\lambda)$ be the third Chebyshev polynomial of the first kind, and use Corollary 18 of [GSLW19] to construct a block encoding of:

$$-T_3\left(\frac{1}{2}P(H)\right) := \sum_i -T_3\left(\frac{1}{2}P(\lambda_i)\right)\Pi_i. \tag{138}$$

Observe that $-T_3(0) = 0$ and $-T_3(\frac{1}{2}) = 1$, and also that $-T_3(\delta/2) \leq 2\delta$ and $-T_3((1-\delta)/2) \geq 1 - 2\delta$. So if we let $\bar{P}(\lambda) := -T_3\left(\frac{1}{2}P(\lambda)\right)$ then $-T_3\left(\frac{1}{2}P(H)\right)$ is the $\bar{P}(H)$ we wanted to construct. Since $P(\lambda)$ has degree d, and $-T_3(\lambda)$ has degree $O(1)$, the total circuit complexity and number of invocations to controlled-U_H is $O(d)$ overall. □

Now we are ready to synthesize the operators we need in order to prove Proposition 8, Lemma 17, and Lemma 29.

We start with the energy estimation result from Proposition 8, which is adapted from [Ral21]. This work gives a construction for energy estimation that attempts to minimize the constant factors in polynomial degree as well as ancilla count. We do not concern ourselves with these, and just give an asymptotic bound that also has an extra factor of n in the complexity. A similar method also appears in [MRTC21].

Proof of Proposition 8. Say U_A is a block encoding with k ancilla qubits of a Hermitian operator A. For $y \in \{0,1\}^k$ we define $A_y \coloneqq (\langle y| \otimes I) U_A (|0^k\rangle \otimes I)$, so that:

$$U_A \left(|0^k\rangle \otimes I \right) = \sum_y |y\rangle \otimes A_y. \tag{139}$$

Then, let $G_A \coloneqq |0^k\rangle \otimes A$ and $\bar{G}_A \coloneqq \sum_{y \in \{0,1\}^k \setminus \{0^k\}} |y\rangle \otimes A_y$ and apply a generalized Toffoli gate to synthesize the isometry:

$$\left(|1\rangle \otimes |0^k\rangle\langle 0^k| + |0\rangle \otimes \sum_{y \in \{0,1\}^k \setminus \{0^k\}} |y\rangle\langle y| \right) U_A \left(|0^k\rangle \otimes I \right) \tag{140}$$

$$= |1\rangle \otimes G_A + |0\rangle \otimes \bar{G}_A. \tag{141}$$

Clearly $G_A^\dagger G_A = A^2$. We also have $\bar{G}_A^\dagger \bar{G}_A = \sum_{y \in \{0,1\}^k \setminus \{0^k\}} A_y^\dagger A_y = I - A^2$ following from the unitarity of U_A.

Our goal is to compute $|x\rangle$, where x is the index of the interval $[a_x, b_x]$ containing λ. We will repeatedly use the above isometry to compute $|x\rangle$ one bit at a time. To do so, we make use of Lemmas 34 and 35 to synthesize δ_j-accurate projectors for each bit of x — call them $\tilde{P}^{(j)}$ for the approximate projector for the j'th bit.

If we apply the above construction involving U_A above for each of $\tilde{P}^{(j)}$, we construct an isometry $\tilde{E}^{(M)}$ that achieves $\tilde{E}^{(M)} = \sum_x |x\rangle \otimes G_{\tilde{P}_x}$ with the desired properties. We observe that:

$$G_{\tilde{P}_x} \coloneqq \prod_{j=1}^{n+1} \begin{cases} G_{\tilde{P}^{(j)}} \text{ if } x_j = 1 \\ \bar{G}_{\tilde{P}^{(j)}} \text{ if } x_j = 0 \end{cases}. \tag{142}$$

Since the $\tilde{P}^{(j)}$ commute, we can evaluate:

$$\tilde{P}_x \coloneqq G_{\tilde{P}_x}^\dagger G_{\tilde{P}_x} = \prod_{j=1}^{n+1} \begin{cases} G_{\tilde{P}^{(j)}}^\dagger G_{\tilde{P}^{(j)}} \text{ if } x_j = 1 \\ \bar{G}_{\tilde{P}^{(j)}}^\dagger \bar{G}_{\tilde{P}^{(j)}} \text{ if } x_j = 0 \end{cases} = \prod_{j=1}^{n+1} \begin{cases} (\tilde{P}^{(j)})^2 \text{ if } x_j = 1 \\ I - (\tilde{P}^{(j)})^2 \text{ if } x_j = 0 \end{cases}. \tag{143}$$

It remains to show that $\left\| \tilde{P}_x P^{(M)} - P_x \right\| \leq \delta_{\text{est}}$. The j-th bit has 2^j many flips. We take $\delta_j \coloneqq \delta_{\text{est}} / 2 \cdot 2^{-j}$ such that $\left\| \tilde{P}^{(j)} P^{(M)} - P^{(j)} \right\| \leq \delta_j$, with $P^{(j)}$ being the ideal projector onto the j'th bit. That way the total error in spectral norm is at most $\delta_{\text{est}} \cdot \sum_{j=1}^{n+1} 2^{-j} \leq \delta_{\text{est}}$. Each projector has complexity $O\left(\kappa^{-1} \log\left(2^j \delta_j^{-1}\right)\right)$. The total complexity is:

$$\kappa^{-1} \sum_{j=1}^{n+1} \log\left(2^j \delta_j^{-1}\right) = \kappa^{-1} \sum_{j=1}^{n+1} \log\left(4^j \cdot 2\delta_{\text{est}}^{-1}\right) = O\left(\kappa^{-1} n^2 \log\left(\delta_{\text{est}}^{-1}\right)\right). \tag{144}$$

□

Next, we construct the operator that underlies the left-right POVM. To follow this argument, we refer again to Fig. 2 which gives a sketch of $\bar{L}, \bar{R}$ as well as the function approximated by Lemma 34.

Proof of Lemma 17. The operator P_{LR} is constructed using Lemmas 34 and 35: all we need to do is select the intervals $[a_x, b_x]$ and the labels c_x. Recall that the goal was to ensure that $\|P_{\mathrm{LR}}^2 \Pi_i\| \leq \delta_{\mathrm{sup}}$ for $\lambda_i \notin \bar{L}$ and $\|(I - P_{\mathrm{LR}}^2)\Pi_i\| \leq \delta_{\mathrm{sup}}$ for $\lambda_i \notin \bar{R}$. Recall Fig. 2. The intervals of interest are the *gaps* of $\bar{L}$ and $\bar{R}$, for which we set $c_x = 0$ and $c_x = 1$ respectively. If we invoke Lemma 34 with precision $\delta = \delta_{\mathrm{sup}}/6$, then for $\lambda_i \notin \bar{L}$ we have $\|P_{\mathrm{LR}}^2 \Pi_i\| \leq \delta^2 \leq \delta_{\mathrm{sup}}/3$ and for $\lambda_i \notin \bar{R}$ we have $\|(I - P_{\mathrm{LR}}^2)\Pi_i\| \leq 1 - (1-\delta)^2 \leq \delta_{\mathrm{sup}}/3$ as desired.

From the construction of $\bar{L}$ and $\bar{R}$ in Definition 16, we see that each of the intervals are exactly 2^{-n-r-2} apart, and there are 2^{n+r+1} of them. Since the intervals alternate with $c_x = 0$ and $c_x = 1$, there are $O(2^{n+r})$ many switches. So, the implementation requires $(n+r)2^{n+r} \log\left(\delta_{\mathrm{sup}}^{-1}\right)$ many invocations of the block encoding of H. □

Finally, we construct the attenuation operator $A^{(M)}$ and its approximation $\tilde{A}^{(M)}$. We demand that $A^{(M)}$ vanishes outside of the rounding promise M, and is close to 1 in inside of a truncated rounding promise M^T. For all the other eigenspaces, the approximation $\tilde{A}^{(M)}$ and $A^{(M)}$ agree. This means that we can use the construction of $\tilde{A}^{(M)}$ to define what $A^{(M)}$ should be outside these regions.

Proof of Lemma 29. As usual, we use Lemmas 34 and 35 to construct the block encoding of $\tilde{A}^{(M)}$. This construction results in a function $\tilde{a}^{(M)}$ which lets us define $a^{(M)}$ and hence $A^{(M)}$ via the requirements in Lemma 29.

We need to ensure that $\tilde{a}^{(M)}(\lambda) \leq \delta_{\mathrm{leak}}$ for $\lambda \notin M$, and that $\tilde{a}^{(M)}(\lambda) \geq 1 - \delta_{\mathrm{leak}}$ for $\lambda \in M^T$. This immediately shows what intervals to use for Lemma 29: the intervals with $c_x = 0$ are the gaps of M, and the intervals with $c_x = 1$ are the intervals of M^T.

There are 2 'flips' per interval of M, so there are $O(s^{(M)})$ flips total. Since M^T was obtained by taking M and shrinking the intervals by $\kappa\gamma$ on each side, the gap width for the purposes of Lemma 29 is $\kappa\gamma$. Accordingly, the polynomial degree and hence the query complexity are $O\left(\kappa^{-1}\gamma^{-1} \log\left(s^{(M)}\delta^{-1}\right)\right)$.

□

C The Approximate Lindbladian

We describe here a method that does not try to force any rounding promises to make energy estimation unambiguous. Rather this method works directly with approximate jump operators that arise from imperfect energy estimation of the Hamiltonian H on the entire Hilbert space $\mathcal{H}$.

We consider a general energy estimation unitary $V^{(\pm)}$ that acts as

$$V^{(\pm)} = \sum_i \Pi_i \otimes \sum_x f(\lambda_i, \tilde{\lambda}_x) \, |z\rangle\langle z \pm \tilde{\lambda}_x| \,, \tag{145}$$

where λ_i's are the energies of H and Π_i's the projectors onto the corresponding eigensubspaces, and $\tilde{\lambda}_x$'s denote the possible outcomes produced by the energy estimation method.

The summation over x means that any energy estimation method necessarily produces superpositions of energy estimates. Typically, the magnitude of $f(\lambda_i, \tilde{\lambda}_x)$ decreases with increasing distance between the true energy λ_i and the approximate energies $\tilde{\lambda}_x$.

For instance, when energy estimation is based on phase estimation the unitary $U = \exp(iH)$, the estimates are simply n-bit binary fractions $\tilde{\lambda}_x = x/2^n$ and the amplitudes $f(\lambda_i, \tilde{\lambda}_x)$ are given by

$$f(\lambda_i, \tilde{\lambda}_x) = \frac{1}{2^n} \frac{\exp\left(\pi i\, 2^n(\lambda_i - \tilde{\lambda}_x)\right)}{\exp\left(\pi i\, (\lambda_i - \tilde{\lambda}_x)\right)}. \tag{146}$$

The spread of $f(\lambda_i, \tilde{\lambda}_x)$ can be made narrower with the help of median amplification.

We now construct an oracle $\mathcal{O}_{\tilde{\mathcal{L}}}$ encoding jump operators of an approximate Lindbladian $\tilde{\mathcal{L}}$ using a general estimation unitary $V^{(\pm)}$ as in Eq. (145). We consider the case that there is only one coupling operator S. $\mathcal{O}_{\tilde{\mathcal{L}}}$ is given by the circuit:

$\mathcal{O}_{\tilde{\mathcal{L}}} :=$ 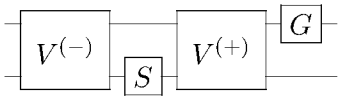

(147)

where G denotes a block encoding of:

$$G := \sum_{\nu} \sqrt{G(\nu)} \otimes |\nu\rangle\langle\nu| \,. \tag{148}$$

To describe the jump operators that are encoded by the above circuit, we define the approximate Bohr frequencies ν to be differences of the form $\tilde{\lambda}_x - \tilde{\lambda}_y$. The approximate jump operators $\tilde{L}(\nu)$ are given by

$$\tilde{L}(\nu) = \sqrt{G(\nu)}\, \tilde{S}(\nu), \tag{149}$$

where

$$\tilde{S}(\nu) = \sum_{\substack{x,y \\ \tilde{\lambda}_x - \tilde{\lambda}_y = \nu}} \sum_{i,j} f(\tilde{\lambda}_x, \lambda_i)\overline{f(\tilde{\lambda}_x, \lambda_j)} \Pi_i S \Pi_j \tag{150}$$

$$= \sum_{\substack{x,y \\ \tilde{\lambda}_x - \tilde{\lambda}_y = \nu}} A(x) S A(y)^\dagger. \tag{151}$$

Unfortunately, the operators

$$A(x) = \sum_i f(\tilde{\lambda}_x, \lambda_i) \Pi_i \tag{152}$$

that "sandwich" the coupling operator S are not projectors as in the case for Davies generators. Thus, it is not possible to interpret the operators $\tilde{L}(\nu)$ as jump operators of a Davies generator with respect to some Hamiltonian that is close to the original Hamiltonian. Therefore, it is much more difficult to determine the fixed point of the corresponding Lindbladian $\tilde{\mathcal{L}}$, which is given by

$$\tilde{\mathcal{L}}(\rho) = \sum_{\nu} G(\nu) \left[\tilde{S}(\nu)\rho\tilde{S}(\nu)^\dagger + \frac{1}{2}\left(\tilde{S}(\nu)^\dagger\tilde{S}(\nu)\rho + \rho\tilde{S}(\nu)^\dagger\tilde{S}(\nu)\right)\right]. \tag{153}$$

This 'approximate Davies generator' $\tilde{\mathcal{L}}$ is challenging to analyze. It is intuitive that this Davies generator's steady state must be somewhat close to the ideal thermal state: the error due to the finite precision of energy estimation can be dealt with in the same way as for the promised Davies generators, and the 'rounding errors' stemming from energies located at $x/2^n + 1/2^{n+1}$ also only shift the accuracy of the estimate slightly.

To our knowledge, no method exists for rigorously proving a bound on the distance between this Lindblad operator's steady state and the true thermal state. A potential candidate for a proof technique via 'approximate detailed balance' appeared in [TOV+11, CB21], where it was leveraged to prove the accuracy of the quantum Metropolis algorithm. But it is not clear how to translate this technique to Davies generators.

How severe are these rounding errors? In practice, it may be the case that the steady state of approximate Davies generator is close to the ideal thermal state. Here, we present some evidence that it may be difficult to prove such a claim without an assumption on the Hamiltonian. We construct an 'adversarial' Hamiltonian that places its eigenvalues exactly at the locations where the rounding errors are maximized. Specifically, its energies are located at:

$$\lambda_i = \frac{i + \alpha/2}{2^n} \tag{154}$$

where $0 \leq i < 2^n$ is an integer and $\alpha \in [0, 1]$ is an 'adversariality parameter'. When $\alpha = 0$, then energy estimation is perfect and the operators $A(x)$ are projectors exactly. But as α increases, the eigenvalues are shifted so that they are rounded up or down with increasing entropy. This Hamiltonian is specifically designed to capture these rounding errors alone: the energies are spaced out evenly with exactly the precision of the energy estimation protocol.

Fig. 4 shows the error in preparing the ideal thermal state when using the approximate Davies generator with this adversarial Hamiltonian. We see that at $\alpha = 0$ the method computes the thermal state exactly. But as α increases we see significant errors. If we amplify the accuracy of the energy measurement using median amplification, then the errors appear only for larger α. But no amount of amplification can remove the error for $\alpha = 1$.

Since the approximate Davies generator is significantly simpler and may be less costly to implement than our scheme of random promised Davies generators, a proof technique establishing rigorous accuracy bounds on this Davies generator may result in a substantially improved algorithm for thermal state preparation.

References

[Alh22] Álvaro M Alhambra. Quantum many-body systems in thermal equilibrium. *arXiv:2204.08349*, 2022.

[BCGW21] Sergey Bravyi, Anirban Chowdhury, David Gosset, and Pawel Wocjan. On the complexity of quantum partition functions. *arXiv:2110.15466*, 2021.

[BKL+19] 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 *46th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 132, page 27, 2019.

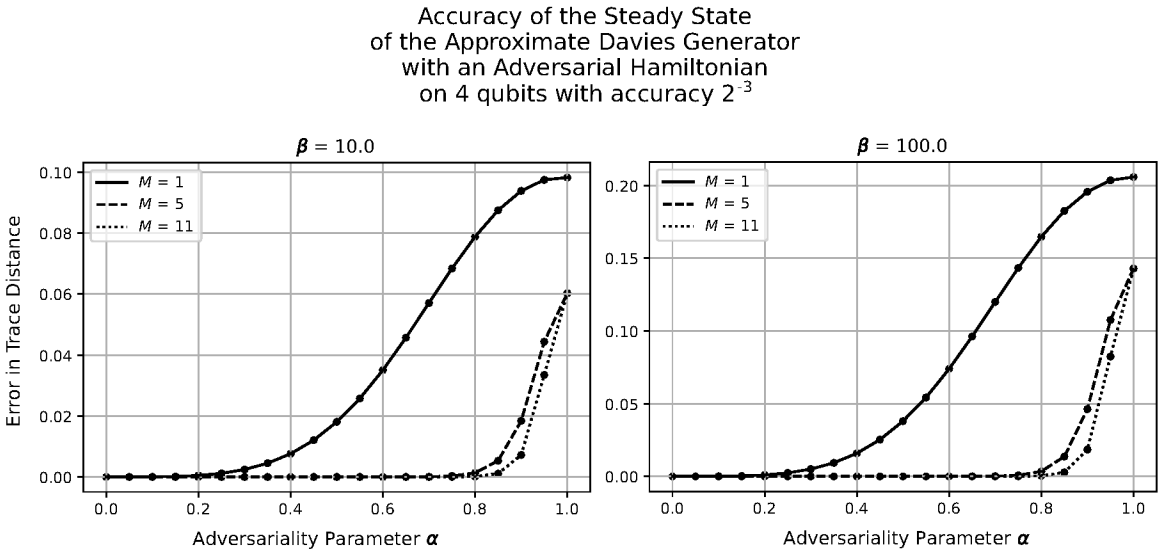

Figure 4: Accuracy of thermal state preparation using the approximate Davies generator in Eq. (153) with an adversarial Hamiltonian with eigenvalues as in Eq. (154). We perform energy estimation to 3 bits of precision on a 4 qubit system. Energy estimation is performed with median amplification, where M denotes the number of estimates over which the median is performed. We observe significant errors for large α. When M is increased, the large errors appear only for the largest α. The software that produced these data is available at: `https://github.com/qiskit-community/promised-davies-generator`.

[BP02] Heinz-Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 2002.

[CB21] Chi-Fang Chen and Fernando GSL Brandão. Fast thermalization from the eigenstate thermalization hypothesis. *arXiv:2112.07646*, 2021.

[CKBG23] Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation. *arXiv:2303.18224*, 2023.

[CS17] Anirban Narayan Chowdhury and Rolando D Somma. Quantum algorithms for Gibbs sampling and hitting-time estimation. *Quantum Information & Computation*, 17(1-2):41–64, 2017.

[CW17] Richard Cleve and Chunhao Wang. Efficient quantum algorithms for simulating Lindblad evolution. *arXiv:1612.09512 Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, 2017.

[Dav76] Edward Brian Davies. *Quantum Theory of Open Systems*. Academic Press, 1976.

[Dav79] Edward Brian Davies. Generators of dynamical semigroups. *Journal of Functional Analysis*, 34(3):421–432, 1979.

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan 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 2019)*, pages 193–204, 2019.

[HMS+22] Zoe Holmes, Gopikrishnan Muraleedharan, Rolando D. Somma, Yigit Subasi, and Burak Şahinoğlu. Quantum algorithms from fluctuation theorems: Thermal-state preparation. *Quantum*, 6:825, October 2022.

[KW17] Mária Kieferová and Nathan Wiebe. Tomography and generative training with quantum Boltzmann machines. *Physical Review A*, 96(6):062327, 2017.

[LC17] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by uniform spectral amplification. *arXiv:1707.05391*, 2017.

[Lin76] Goran Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, 1976.

[LW22] Xiantao Li and Chunhao Wang. Simulating Markovian open quantum systems using higher-order series expansion. *2212.02051*, 2022.

[MRTC21] John M Martyn, Zane M Rossi, Andrew K Tan, and Isaac L Chuang. Grand unification of quantum algorithms. *PRX Quantum*, 2(4):040203, 2021.

[Nig19] Davide Nigro. On the uniqueness of the steady-state solution of the Lindblad–gorini–Kossakowski–Sudarshan equation. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(4):043202, 2019.

[PW09] David Poulin and Pawel Wocjan. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters*, 103(22):220502, 2009.

[Ral21] Patrick Rall. Faster coherent quantum algorithms for phase, energy, and amplitude estimation. *Quantum*, 5:566, 2021.

[SC22] S. Slezak and E. Crosson. Eigenstate thermalization and quantum Metropolis sampling, 2022. Presentation at QIP 2022. `https://youtu.be/by4rvu7RMtY`.

[Spo77] Herbert Spohn. An algebraic condition for the approach to equilibrium of an open n-level system. *Letters in Mathematical Physics*, 2(1):33–38, 1977.

[TOV+11] Kristan Temme, Tobias J Osborne, Karl G Vollbrecht, David Poulin, and Frank Verstraete. Quantum Metropolis sampling. *Nature*, 471(7336):87–90, 2011.

[vAG19] 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 2019)*, 2019.

[VAGGdW20] Joran Van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020.

[Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018.

[WT23] Pawel Wocjan and Kristan Temme. Szegedy walk unitaries for quantum maps. *Commun. Math. Phys.*, 2023.