

Faster quantum mixing for slowly evolving sequences of Markov chains

Davide Orsucci¹, Hans J. Briegel^{1,2}, and Vedran Dunjko^{1,3,4}

¹Institute for Theoretical Physics, University of Innsbruck, Technikerstraße 21a, 6020 Innsbruck, Austria

²Department of Philosophy, University of Konstanz, Fach 17, 78457 Konstanz, Germany

³Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, 85748 Garching, Germany

⁴LIACS, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands

October 31, 2018

Markov chain methods are remarkably successful in computational physics, machine learning, and combinatorial optimization. The cost of such methods often reduces to the mixing time, *i.e.*, the time required to reach the steady state of the Markov chain, which scales as δ^{-1} , the inverse of the spectral gap. It has long been conjectured that quantum computers offer nearly generic quadratic improvements for mixing problems. However, except in special cases, quantum algorithms achieve a run-time of $\mathcal{O}(\sqrt{\delta^{-1}}\sqrt{N})$, which introduces a costly dependence on the Markov chain size N , not present in the classical case. Here, we re-address the problem of mixing of Markov chains when these form a slowly evolving sequence. This setting is akin to the simulated annealing setting and is commonly encountered in physics, material sciences and machine learning. We provide a quantum memory-efficient algorithm with a run-time of $\mathcal{O}(\sqrt{\delta^{-1}}\sqrt[4]{N})$, neglecting logarithmic terms, which is an important improvement for large state spaces. Moreover, our algorithms output quantum encodings of distributions, which has advantages over classical outputs. Finally, we discuss the run-time bounds of mixing algorithms and show that, under certain assumptions, our algorithms are optimal.

1 Introduction

Markov chains (MCs) are central in computational approaches to physics [1], in computer science [2], and machine learning [3], and they form the crux of the ubiquitous Markov Chain Monte Carlo meth-

ods [4]. In MC-based approaches the underlying objective is to produce samples from the steady state, *i.e.*, the stationary distribution of a given MC. The MC is constructed so that this distribution encodes the solution of the problem at hand. The solution can then be reached by “mixing”, *i.e.*, by applying the MC transitions many times. For some problems, mixing processes constitute the fastest known classical solving algorithms, and play a vital role, *e.g.*, in the Metropolis-Hastings methods [5], periodic Gibbs sampling [6], and Glauber dynamics [7].

The fundamental parameter governing the time complexity of MC-based algorithms is thus the *mixing time*, that is, the number of steps required to attain stationarity. In most applications the MC is ergodic, *i.e.*, has a unique stationary distribution, and time-reversible, *i.e.*, satisfies detailed balance [8, 9]. The mixing time is tightly related to the spectral gap δ of the MC¹ and is bounded by $\Omega(\delta^{-1})$ [10].

Oftentimes direct mixing can be computationally prohibitive and thus heuristic methods, such as *simulated annealing* [11, 12], are employed. Here one constructs a sequence of Markov chains which, for instance, encode the Gibbs (thermal) distributions at gradually decreasing values of the temperature, where the target distribution is specified by the final MC, *i.e.*, the final temperature. Intuitively, this process increases efficiency by avoiding local minima, although the performance is typically not guaranteed. In simulated annealing, the neighbouring chains in the sequence are similar, in other words, the sequence is *slowly evolving*.

The emergence of quantum computation offers a new possibility to utilize quantum effects to achieve guaranteed mixing more rapidly. In particular, it has been conjectured that run-times of $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ ²

Davide Orsucci: davide.orsucci@uibk.ac.at

Hans J. Briegel: hans.briegel@uibk.ac.at

Vedran Dunjko: v.dunjko@liacs.leidenuniv.nl

¹The spectral gap is defined with $\delta = 1 - |\lambda_2|$, where λ_2 is the second largest eigenvalue (in absolute value) of the transition matrix of the time-reversible Markov chain.

²For expressing run-times we adopt the soft- \mathcal{O} notation ($\tilde{\mathcal{O}}$),




Figure 1: Schematic depiction of the scenarios of sequences of slowly evolving MCs. Panels (a,b) and (c,d) depict classical and quantum sampling tasks, respectively. Panels (a,c) and (b,d) respectively delineate finite and continuing sequences, the latter having step-wise outputs. This work is predominantly concerned with (b,d). Although panel (d) allows quantum states to be carried from one time-step to another, our algorithm actually works by forwarding just classical information, without sacrificing efficiency. That is, no quantum memory from one time-step to the next is required.

should be possible [13] for the mixing problem. Such quadratic speed-ups have been demonstrated for various special cases of MCs [13–18], mostly relying on quantum walk [19, 20] approaches. Quantum walks have also been utilized to speed-up simulated annealing [21–23], which often leads to the best run-times in practice. However, considering provable results for guaranteed mixing of general Markov chains, the best quantum algorithms achieve $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}}\sqrt{N})$, which falls short of the conjectured quadratic speed-up, as it introduces the dependence on the system size N . Avoiding the $\mathcal{O}(\sqrt{N})$ dependence seems to be challenging, which further motivates investigating the settings with relaxed constraints, *e.g.*, by restricting the MC family [13, 18, 24].

In this work, we obtain improved $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}}\sqrt{N})$ run-times for mixing problems not by restricting the Markov chain family, but rather by relying on additional context. In particular, we consider the settings where we are tasked to sequentially produce *independent samples* from a sequence of slowly evolving Markov chains. This setting is natural in statistical and quantum physics, *e.g.*, when studying phase boundaries, which requires many independent samples from near-by points in the parameter space [25]. Another motivation for studying this setting is in the context of machine learning (ML), appearing both in reinforcement learning [26] and in the training of generative models [27], as we discuss later in the paper.

Our setting is similar to simulated annealing in that it considers a sequence of pair-wise similar Markov-chains, and indeed our methods are similar to those in [21, 22]. However our setting brings about a key dis-

an extension of the \mathcal{O} notation where polylogarithmic multiplicative factors are neglected.

inction: in annealing the goal is to produce a sample from final MC, and the intermediary chains have only an auxiliary role; in our case the goal is to produce independent samples from each MC in the sequence. Further, in principle the sequence can be exponentially large, or having a length which is not a priori specified. This is schematically illustrated in Fig. 1.

2 Problem and methods

We now specify the setting more precisely and introduce the required notation. We consider finite-space MCs, of size N , where a distribution over the space is specified by a vector $\boldsymbol{\pi} := (\pi(1), \pi(2), \dots, \pi(N))^T$ of non-negative entries summing to one. The problem of sampling from this distribution corresponds to producing a single element $x \in \{1, 2, \dots, N\}$ according to $\boldsymbol{\pi}$. Due to the methods used, our algorithms will actually produce a quantum (or coherent) sample, *i.e.*, the quantum state $|\boldsymbol{\pi}\rangle := \sum_x \sqrt{\pi(x)}|x\rangle$, which is called the *coherent encoding* of $\boldsymbol{\pi}$. As we elaborate later, coherent encodings have substantial advantages over classical samples, although they are in general computationally more difficult to prepare.

In abstract terms our sampling problem can be formulated as follows. We consider an infinite sequence of ergodic time-reversible Markov chains $\{\text{MC}_t\}_{t=1}^{\infty}$. In our approach, we will at each time-step t generate a coherent sample, that is the state $|\boldsymbol{\pi}_t\rangle$, corresponding to the stationary distribution $\boldsymbol{\pi}_t$ of MC_t . The measurement of $|\boldsymbol{\pi}_t\rangle$ in the computational basis yields a classical sample from $\boldsymbol{\pi}_t$, thus preparation of $|\boldsymbol{\pi}_t\rangle$ allows for both classical and quantum sampling. We say that the sequence of MCs is slowly evolving if the stationary distributions of consecutive MCs in the sequence are sufficiently close, specifically, if at

every time step t we have $|\langle \pi_{t+1} | \pi_t \rangle|^2 \geq \eta$ for some constant $0 < \eta < 1$.

Our techniques rely on Szegedy-type quantum walks [28] to perform the above sequential sampling task with a $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}} \sqrt[4]{N})$ time complexity in the context of slowly evolving MCs. We thus briefly introduce the properties of the Szegedy constructions for the convenience of the reader and provide in App. A more background on MC theory.

2.1 Szegedy quantum walk

Each MC is specified by a stochastic matrix P which specifies the transition probabilities in a single step of the chain. For a given transition matrix P of an ergodic time-reversible MC one can construct the corresponding Szegedy quantum walk operator $W(P)$; this is a unitary operator having the crucial property that $|\pi\rangle$ is the unique $+1$ -eigenstate of $W(P)$, with all other eigenstates of $W(P)$ having an eigenphase which is at least quadratically larger than the spectral gap δ of P . In other words, if $|\theta\rangle$ is such that $W(P)|\theta\rangle = e^{i\theta}|\theta\rangle$, then $|\theta\rangle \in \mathcal{O}(\sqrt{\delta})$, see App. B for details and the construction.

These properties allow us to realize useful quantum subroutines with a run-time which is quadratically smaller than the classical mixing time. Namely, the Szegedy walk operator can be used in conjunction with the phase detection algorithm [29], a simple variant of phase estimation [30], to (approximately) distinguish $|\pi\rangle$ from all other eigenstates of $W(P)$. The run-time is in $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ and has only logarithmic dependence on the approximation error³ [29–32]. In turn, the capacity to identify $|\pi\rangle$ can be leveraged to implement an approximate projective measurement onto $|\pi\rangle\langle\pi|$, by measuring whether the quantum register containing the phase estimate is zero. Similarly, applying a Pauli- Z rotation onto the qubit that specifies whether the phase value is zero we obtain an approximation of the reflection operator $R(\pi) := \mathbb{I} - 2|\pi\rangle\langle\pi|$.

2.2 Amplitude amplification

This brings us to our key subroutine. Using the reflection $R(\pi)$ we use amplitude amplification [33, 34] to rotate an initial state $|\psi_{in}\rangle$ to an approximation of $|\pi\rangle$ in time $\tilde{\mathcal{O}}(\sqrt{\gamma^{-1}}\sqrt{\delta^{-1}})$, where $|\langle\psi_{in}|\pi\rangle|^2 \geq \gamma$. Equivalently, we can also use the fixed-point ampli-

³All algorithms we consider are approximate, but the dependence on the target error ε is at most $\log^2(\varepsilon^{-1})$ and thus always ignored in the $\tilde{\mathcal{O}}$ notation.

tude amplification algorithm⁴ of Yoder *et al.* [36], an algorithm that has the same quadratic speed-up as standard amplitude amplification. We refer the reader to App. C for further details on fixed-point amplitude amplification and to App. D for an analysis of how the runtime depends on the target precision.

With these key subroutines defined we can explain a straightforward algorithm that allows to prepare coherent encodings of $|\pi\rangle$ with $\mathcal{O}(\sqrt{\delta^{-1}}\sqrt{N})$ run-time. One simply utilizes amplitude amplification to rotate the uniform superposition $|\mathbf{u}\rangle := \frac{1}{\sqrt{N}}\sum_x|x\rangle$ to the target state $|\pi\rangle$. Since the target states are encodings of probability distributions, all amplitudes are real and non-negative and, thus, the fidelity always satisfies $\gamma > 1/N$. This bound is attained by distributions approaching a Kronecker delta.

We remark that amplitude amplification requires the ability to reflect both around the source state $|\mathbf{u}\rangle$ and the target state, *i.e.*, to implement both $R(\pi)$ and $R(\mathbf{u}) := \mathbb{I} - 2|\mathbf{u}\rangle\langle\mathbf{u}|$. In this work we restrict our attention to cases in which the preparation of $|\mathbf{u}\rangle$ is efficient and, therefore, also $R(\mathbf{u})$ can be easily implemented. Notice that preparing the uniform distribution is not always simple [17], *e.g.*, this happens when it is computationally difficult to decide if an element $x \in \mathbb{N}$ belongs to the space of the MC⁵. However, assuming that the space of the MC is $\{1, \dots, N\}$ the uniform distribution can be prepared with quantum circuits of depth $\mathcal{O}(\log(N))$ [37]; this is a case that finds application to quantum machine learning problems [38]. Even more simply, when $N = 2^n$ the uniform distribution is obtained from $|0\rangle$ via the Hadamard transform, that is, $|\mathbf{u}\rangle = H^{\otimes n}|0\rangle$. Consequently, our methods can be applied to spin $1/2$ systems, where the preparation of the uniform superposition of all the configuration of n spins is trivial, yet producing Gibbs distributions at low temperature for certain classical Hamiltonians is NP-hard [39].

3 Preparation from uniform distribution and from samples

To speed-up the basic algorithm described in the previous section, the idea is to eliminate the worst-case

⁴The fixed-point property means that the output state converges to the ideal target state for increasing run-times [35].

⁵As a concrete example, consider the following MC inspired by the Graph Isomorphism problem: the space of the MC consists of all graphs that can be obtained from an initial graph via permutation of the vertices of a given initial graph; and a transition in the MC is obtained by randomly selecting two vertices of a graph and swapping them. Then, deciding if a graph belongs to the space is equivalent to solving the Graph Isomorphism problem.

preparation scenario. Specifically, in the case when the distribution is highly clumped, one should attempt the preparation from an element having high probability in the target distribution $\boldsymbol{\pi}$. However, we still have to choose the candidate element to start from, which alone would lead to a $\Omega(\sqrt{N})$ run-time (by the optimality of Grover’s search [40]). We will first show that this issue can be resolved when one has access beforehand to a few classical samples from the target distribution. This seems to require that the solution we are looking for are already provided as input. But we will utilize the slowly evolving context to ensure such samples are available, and therefore samples for the subsequent step can be prepared without the necessity of back-tracking in the sequence.

To utilize these ideas we first show how to prepare the coherent encoding $|\boldsymbol{\pi}\rangle$ by choosing a suitable initial state $|\psi_{in}\rangle$ and then amplitude amplify $|\psi_{in}\rangle$ to obtain $|\boldsymbol{\pi}\rangle$. Specifically, the initial state is either the uniform distribution, $|\psi_{in}\rangle \equiv |\mathbf{u}\rangle$, or a classical sample x_j which is taken from a small set of classical samples $\vec{x} = \{x_1, \dots, x_c\}$ that are available beforehand, $|\psi_{in}\rangle \equiv |x_j\rangle$. We will call these subroutines `PrepareFromUniform` and `PrepareFromSamples`, respectively, and simply `Prepare` whenever the distinction is not relevant.

As mentioned, `PrepareFromUniform` is efficient in the extreme case where $\boldsymbol{\pi}$ is very close to being uniform, while the procedure can require up to $\mathcal{O}(\sqrt{N})$ operations in the opposite extreme case where $\boldsymbol{\pi}$ has support over only one element x ; this last case corresponds, in fact, to a standard Grover search for the element x . However, when most of the “weight” (probability) of $\boldsymbol{\pi}$ is concentrated on a few elements (which need not be nearby) these must have a large overlap with $|\boldsymbol{\pi}\rangle$, which is a sufficient condition to efficiently perform amplitude amplification; that is, running a search algorithm in reverse (un-searching) from one of these elements allows for a fast re-preparation of $|\boldsymbol{\pi}\rangle$ [24]. Then, a classical sample drawn from $\boldsymbol{\pi}$ will probably come from elements having large “weight” and thus the `PrepareFromSamples` subroutine will be efficient. The main idea of our algorithms is to discover which of the two `Prepare` algorithms is the most efficient and then use it for state preparation. The worst regime is for distributions which are neither too uniform nor too clumped, where both algorithms have a $\mathcal{O}(\sqrt[4]{N})$ time complexity.

Before continuing with the complete description of the `Prepare` subroutines, we remark that we use amplitude amplification starting from $|\psi_{in}\rangle$ to produce $|\boldsymbol{\pi}\rangle$ and thus we require the ability to perform reflections both around $|\psi_{in}\rangle$ and around $|\boldsymbol{\pi}\rangle$. Both choices for the initial state can be prepared efficiently: $|x_j\rangle$ is simply a classical state, while $|\mathbf{u}\rangle$ can be pre-

pared easily when the MC space is explicitly known. Thus, also reflectors around them can be efficiently implemented. A reflection around $|\boldsymbol{\pi}\rangle$ is instead approximated with $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ operations using Szegedy operator. Therefore, the total gate cost of `Prepare` is $\tilde{\mathcal{O}}(\sqrt{\gamma^{-1}}\sqrt{\delta^{-1}})$ where $|\langle\psi_{in}|\boldsymbol{\pi}\rangle|^2 \geq \gamma$. We are then left with the task of estimating this lower bound γ .

3.1 Preparing from uniform

This subroutine is the straightforward algorithm we mentioned earlier. Suppose for the moment that the value of $|\langle\mathbf{u}|\boldsymbol{\pi}\rangle|$ is known. Then we can amplitude amplify $|\mathbf{u}\rangle$ to $|\boldsymbol{\pi}\rangle$, operation having a gate cost which is proportional to

$$|\langle\mathbf{u}|\boldsymbol{\pi}\rangle|^{-1} = \frac{\sqrt{N}}{f(\boldsymbol{\pi})}, \quad (1)$$

in which we have introduced the notation:

$$f(\boldsymbol{\pi}) := \sum_{x=1}^N \sqrt{\pi(x)}. \quad (2)$$

By norm inequalities we get $1 \leq f(\boldsymbol{\pi}) \leq \sqrt{N}$, where the lower and upper bounds are saturated by a Kronecker delta and the uniform distribution, respectively.

If the value of $|\langle\mathbf{u}|\boldsymbol{\pi}\rangle|$ is not known we proceed as follows. We arbitrarily choose a value $\chi' \equiv \sqrt{N}/\chi$ as a tentative estimate of $\sqrt{N}/f(\boldsymbol{\pi})$ so that amplitude amplification produces an approximation of $|\boldsymbol{\pi}\rangle$ when $|\langle\mathbf{u}|\boldsymbol{\pi}\rangle| \geq \chi'^{-1}$ holds⁶. However, we do not know if the initial overlap is large enough and thus preparation of $|\boldsymbol{\pi}\rangle$ is not guaranteed. To amend this, we subsequently apply to the output of amplitude amplification a projective measurement onto $|\boldsymbol{\pi}\rangle\langle\boldsymbol{\pi}|$ (or onto the orthogonal complement) which, if successful, heralds the correct preparation of $|\boldsymbol{\pi}\rangle$. As we mentioned, this projective measurement can be implemented with Szegedy quantum walks with run-time $\tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ and therefore its run-time is independent from the initial overlap $|\langle\mathbf{u}|\boldsymbol{\pi}\rangle|$. This constitutes a *heralded preparation* of $|\boldsymbol{\pi}\rangle$ from $|\mathbf{u}\rangle$.

`PrepareFromUniform` is summarized in [Alg. 1](#), and in [App. E](#) further details and error analysis are provided. The run-time for preparing c copies is in $\tilde{\mathcal{O}}(c\chi'\sqrt{\delta^{-1}})$ with an exponentially decaying failure probability when $\chi' \geq \sqrt{N}/f(\boldsymbol{\pi})$.

⁶This can be easily enforced using the fixed-point version of amplitude amplification.

Algorithm 1 PrepareFromUniform

Output: a bit signalling success; in case of success, c quantum samples (*i.e.*, c copies of the state $|\boldsymbol{\pi}\rangle$).

Input: quantum access to the transition matrix P ; c , the number of copies to be produced; $\chi' \equiv \sqrt{N}/\chi$, a (tentative) estimate of $\sqrt{N}/f(\boldsymbol{\pi}) = |\langle \mathbf{u} | \boldsymbol{\pi} \rangle|^{-1}$.

Algorithm:

1. For $j = 1, \dots, 2c$:
Run a heralded preparation of $|\boldsymbol{\pi}\rangle$ from $|\mathbf{u}\rangle$ as described in the main text, assuming a initial overlap larger than $1/\chi'$.
 2. If at least c successful preparations have been heralded in step 1., output a bit signalling success, together with the quantum states obtained in c successful runs of heralded preparation of $|\boldsymbol{\pi}\rangle$. Else, return a bit signalling failure.
-

3.2 Preparing from samples

The second subroutine we will utilize, named `PrepareFromSamples`, requires extra inputs, namely a set of c samples $\vec{x} = \{x_1, \dots, x_c\}$ from the desired target distribution $\boldsymbol{\pi}$, and is based on amplitude amplification of $|x_j\rangle$ to $|\boldsymbol{\pi}\rangle$. Later we will show how these samples can be efficiently obtained in a slowly evolving sequence.

The run-time of amplitude amplification scales as $|\langle x_j | \boldsymbol{\pi} \rangle|^{-1} = 1/\sqrt{\pi(x_j)}$ assuming, for the moment being, that the value of $|\langle x_j | \boldsymbol{\pi} \rangle|$ is known. We thus introduce a random variable X distributed according to $\boldsymbol{\pi}$, *i.e.*, X takes a value x with probability $\pi(x)$. The run-time of amplitude amplification is also a random variable, proportional to $|\langle X | \boldsymbol{\pi} \rangle|^{-1} = \pi^{-1/2}(X)$. The average run-time scales as

$$\mathbb{E}_{\boldsymbol{\pi}}[\pi^{-1/2}(X)] = \sum_{x=1}^N \pi(x) \pi^{-1/2}(x) = f(\boldsymbol{\pi}). \quad (3)$$

Note that we have bounded only the *expected* run-time of our algorithm and not of a specific instance of the algorithm, *i.e.*, for a particular choice of x_j . The sampling procedure could return in fact a sample x_j for which the run-time factor $[\pi^{-1/2}(x_j)]$ is much larger than its average value.

However, this can be prevented by using a few samples, since the run-time when starting from a randomly sampled element x_j is, with constant probability, close to the average run-time. To formalize this we use Markov's inequality:

$$\Pr\{ \pi^{-1/2}(X) \geq a \mathbb{E}[\pi^{-1/2}(X)] \} \leq \frac{1}{a}. \quad (4)$$

We then consider the case $a = 2$ and proceed similarly as we did for `PrepareFromUniform`. Namely, we guess an estimate χ for $f(\boldsymbol{\pi})$ and amplitude amplify $|x_j\rangle$ to $|\boldsymbol{\pi}\rangle$ assuming that $|\langle x_j | \boldsymbol{\pi} \rangle| \geq 1/(2\chi)$ holds (that is, we use $\mathcal{O}(2\chi)$ reflections) and then repeat for all samples in \vec{x} . Subsequently we apply a projective measurement onto $|\boldsymbol{\pi}\rangle\langle\boldsymbol{\pi}|$ to herald the successful preparation of $|\boldsymbol{\pi}\rangle$.

Suppose that $\chi \geq f(\boldsymbol{\pi})$. Since the c samples are independent, the probability that `PrepareFromSamples` fails for all the samples in \vec{x} is then exponentially small in c (for instance, `PrepareFromSamples` can fail if $|\langle x_j | \boldsymbol{\pi} \rangle|^{-1} > 2f(\boldsymbol{\pi})$ holds for all x_j). A formal specification of `PrepareFromSamples` for preparing c new samples is given in [Alg. 2](#) and has run-time in $\tilde{\mathcal{O}}(c^2 \chi \sqrt{\delta^{-1}})$. The failure probability again goes down exponentially if $\chi \geq f(\boldsymbol{\pi})$. See [App. E](#) for details and error analysis.

Algorithm 2 PrepareFromSamples

Output: a bit signalling success; in case of success, c quantum samples (*i.e.*, c copies of the state $|\boldsymbol{\pi}\rangle$).

Input: quantum access to the transition matrix P ; $\vec{x} = \{x_1, \dots, x_c\}$, a set of c classical samples approximately drawn from $\boldsymbol{\pi}$ (correspondingly, we require to produce c new copies of $|\boldsymbol{\pi}\rangle$); χ , a (tentative) estimate of $f(\boldsymbol{\pi}) = \mathbb{E}[\pi^{-1/2}(X)]$.

Algorithm:

1. For $j = 1, \dots, c$:
For $2c$ times: run a heralded preparation of $|\boldsymbol{\pi}\rangle$ from $|\mathbf{u}\rangle$ as described in the main text, assuming a initial overlap larger than $1/(2\chi)$.
 2. If at least c successful preparations have been heralded in step 1., output a bit signalling success, together with the quantum states coming from c successful runs of heralded preparation of $|\boldsymbol{\pi}\rangle$. Else, return a bit signalling failure.
-

3.3 Combined algorithm

Now we will put together the two `Prepare` subroutines in a single combined algorithm. In the case in which the value $f(\boldsymbol{\pi})$ is known one simply runs whichever of the two `Prepare` algorithms is faster.

Summarizing [Eq. \(1\)](#) and [Eq. \(3\)](#) the run-time is then in $\tilde{\mathcal{O}}(c^2 C(\boldsymbol{\pi}) \sqrt{\delta^{-1}})$, where:

$$C(\boldsymbol{\pi}) := \min \left\{ \frac{\sqrt{N}}{f(\boldsymbol{\pi})}, f(\boldsymbol{\pi}) \right\} \leq \sqrt[4]{N}. \quad (5)$$

To deal with the situation when $f(\boldsymbol{\pi})$ is not known, we modify the algorithm in a manner similar to how

Grover’s search is adapted to work without an estimate on the number of marked numbers [40]. Essentially, one runs both preparation algorithms one after another, starting from $\chi = 1$ and $\chi' = 1$ for `PrepareFromSamples` and `PrepareFromUniform`, respectively; then, in each iteration the values of χ and χ' are set to twice larger values, terminating either when c copies of the coherent encoding are produced or when both χ and χ' exceed $2\sqrt[4]{N}$. Then the total number of reflections required scales as $C(\boldsymbol{\pi})$, which is $\sqrt[4]{N}$ in the worst case, and the global failure probability is, again, exponentially small.

4 Application in the context of slowly evolving sequences

The algorithm given in the preceding paragraphs can be implemented also for a stand-alone MC, *i.e.*, for a chain not coming from a slowly evolving sequence. However, it comes with the unrealistic requirement that c samples drawn from the stationary distribution of $\boldsymbol{\pi}$ are available beforehand: it seems that, paradoxically, the output of the algorithm is also required as input. Nonetheless, the result is non-trivial even for stand-alone MCs because of the following two observations. First, the initial classical samples can be re-used to prepare multiple coherent copies and this, in turn, allows us to prepare an arbitrary number of fresh independent samples, given only a small number (c) of seed examples. Second, our algorithm outputs a coherent encoding of the stationary distribution, allowing quantum information post-processing to be applied.

Going back to our primary objective, we now show how these initial samples can be made available in the context of slowly evolving sequences.

We proceed inductively. We suppose that at time step t we have at hand c samples from $\boldsymbol{\pi}_t$. This allows us to produce c copies of $|\boldsymbol{\pi}_t\rangle$ in time $\tilde{O}(c^2 C(\boldsymbol{\pi}_t) \sqrt{\delta_t^{-1}})$. Next, using a Szegedy quantum walk we can implement reflections both around $|\boldsymbol{\pi}_t\rangle$ and around $|\boldsymbol{\pi}_{t+1}\rangle$. This allows us to use amplitude amplification (or its fixed-point variant) to approximately map each of the c copies of $|\boldsymbol{\pi}_t\rangle$ to a copy of $|\boldsymbol{\pi}_{t+1}\rangle$. In turn, these copies of $|\boldsymbol{\pi}_{t+1}\rangle$ can be measured in the computational basis to obtain c samples from $\boldsymbol{\pi}_{t+1}$, allowing to proceed iteratively in the state preparation in the sequence.

By the slowly evolving assumption, the overlap between $|\boldsymbol{\pi}_t\rangle$ and $|\boldsymbol{\pi}_{t+1}\rangle$ is constant and amplitude amplification to $|\boldsymbol{\pi}_{t+1}\rangle$ for c copies in parallel has gate complexity in $\tilde{O}(c \sqrt{\delta_t'^{-1}})$, where $\delta_t' := \min\{\delta_t, \delta_{t+1}\}$. To simplify the statement of the result we can as-

sume that δ_t and δ_{t+1} are multiplicatively close, that is, $\frac{1}{\kappa} \delta_t \leq \delta_{t+1} \leq \kappa \delta_t$ for some constant $\kappa > 1$. Hence the time complexity of this final amplitude amplification is in $\tilde{O}(c \sqrt{\delta_t^{-1}})$, which is dominated by the run-time necessary to initially prepare $|\boldsymbol{\pi}_t\rangle^{\otimes c}$. Consequently, the overall run-time of this process is $\tilde{O}(c^2 C(\boldsymbol{\pi}) \sqrt{\delta^{-1}})$ per each MC, which is no worse than $\sqrt[4]{N}$, as advertised.

We highlight that the entire procedure only requires classical memory between consecutive time steps in the sequence, in the form of c classical samples stored in memory⁷. This is without loss of generality since, as we show later on, $\Omega(C(\boldsymbol{\pi}))$ reflections are needed even if one allows for quantum memory. Moreover, assuming that the classical memory is devoid of errors, this observation that no quantum memory is required shows that approximation errors do not accumulate in the slowly evolving sequence. In fact, the quantum algorithm performed at step $t+1$ does not receive any quantum state as input from step t , but only classical information. Of course, each step t still entails a finite failure probability, albeit exponentially small in c .

If quantum memory is available, the algorithm can be made slightly more efficient, in terms of how many samples (classical or quantum) are required. Instead of c classical samples, one need store only one coherent sample. The basis of this is a simple near-deterministic cloning algorithm producing two copies of $|\boldsymbol{\pi}\rangle$ from one, which may be of independent interest. Details of this quantum memory algorithm are provided in [App. F](#).

5 Application in quantum machine learning

We now consider a modification of the `Prepare` algorithm, which finds application, *e.g.*, in quantization of the reflective Projective Simulation (rPS) model. For the reader interested in quantum ML, details about the rPS can be found in [38]. Here it is sufficient to point out that the outputs in the rPS model are not samples from $\{\boldsymbol{\pi}_t\}_t$ but come from restricting to a subset of “marked elements” $\mathcal{M} \subseteq \{1, 2, \dots, N\}$ and, typically, the number of marked elements $M = |\mathcal{M}|$ is much smaller than N . That is, we want to sample from $\boldsymbol{\pi}^{\mathcal{M}}$, the (normalized) probability distribution obtained by restricting $\boldsymbol{\pi}$ to \mathcal{M} :

$$\boldsymbol{\pi}^{\mathcal{M}}(x) := \begin{cases} \frac{1}{\mu} \pi(x) & \text{if } x \in \mathcal{M} \\ 0 & \text{if } x \notin \mathcal{M} \end{cases}, \quad (6)$$

⁷Quantum access to both P_t and P_{t+1} is also assumed, which entails a factor of two increase in the required memory size.

where $\mu := \sum_{x \in \mathcal{M}} \pi(x)$ and therefore $|\pi^{\mathcal{M}}\rangle = \frac{1}{\sqrt{\mu}} \sum_{x \in \mathcal{M}} \sqrt{\pi(x)} |x\rangle$. The set of marked elements is specified by two black-box unitary maps, a *membership oracle* $\mathcal{P}_{\mathcal{M}}$ and by a *sparse oracle* $\mathcal{Q}_{\mathcal{M}}$. The former, given an element x , specifies whether $x \in \mathcal{M}$ or not; the latter is a quantum accessible memory that, upon input of a number $\nu \in \{1, \dots, M\}$, outputs the ν -th element of \mathcal{M} (*i.e.* $x_{\nu} \in \mathcal{M}$).

Accessing the oracle $\mathcal{P}_{\mathcal{M}}$ twice, one can implement a reflection over the subspace of marked elements. This allows to run amplitude amplification, as done, *e.g.*, in the context of rPS [38], to rotate $|\pi\rangle$ to $|\pi^{\mathcal{M}}\rangle$. This operation has a run-time of $\tilde{O}(\sqrt{\delta^{-1}}\sqrt{\mu^{-1}})$, yielding a quadratic improvement in both mixing time and hitting time with respect to classical methods. This can be done provided that an initial copy of $|\pi\rangle$ is available. We now explain how, in this context, the state $|\pi\rangle$ can sometimes be made available more cheaply.

We consider two modified algorithms for preparation of $|\pi\rangle$: these are equal to the algorithms specified before in Alg. 1 and Alg. 2, except for the choice of the initial states, which now are chosen to have support on the marked elements. Specifically, the initial state is either $|\psi_{in}\rangle = |\mathbf{u}^{\mathcal{M}}\rangle$ or $|\psi_{in}\rangle = |x_j\rangle$ for $x_j \in \vec{x}$, where now \vec{x} is a set of c samples drawn previously from $\pi^{\mathcal{M}}$. These modified state preparation algorithms require that the new input states $|\psi_{in}\rangle$ can be efficiently produced. This is obviously true for classical samples, while $|\mathbf{u}^{\mathcal{M}}\rangle$ can be prepared with one access to the $\mathcal{Q}_{\mathcal{M}}$ oracle, since $|\mathbf{u}^{\mathcal{M}}\rangle = \mathcal{Q}_{\mathcal{M}} \sum_{\nu=1}^M \frac{1}{\sqrt{M}} |\nu\rangle$. This is then sufficient to perform amplitude amplification of $|\psi_{in}\rangle$ to $|\pi\rangle$, for both choices of $|\psi_{in}\rangle$. Using similar reasoning as done previously, we see that the amplitude amplification has run-time scaling as

$$|\langle \mathbf{u}^{\mathcal{M}} | \pi \rangle|^{-1} = \sqrt{\mu^{-1}} \frac{\sqrt{M}}{f(\pi^{\mathcal{M}})}, \quad (7)$$

when trying to prepare $|\pi^{\mathcal{M}}\rangle$ from $|\mathbf{u}^{\mathcal{M}}\rangle$; and when performing amplitude amplification starting from the available samples \vec{x} the expected run-time is

$$\mathbb{E}_{\pi^{\mathcal{M}}} [|\langle X | \pi \rangle|^{-1}] = \sqrt{\mu^{-1}} f(\pi^{\mathcal{M}}). \quad (8)$$

Again, combining these two state preparation algorithms (which start from initial states having support on the marked subspace \mathcal{M}) into a single procedure we obtain a run-time in $\tilde{O}(c^2 C(\pi^{\mathcal{M}}) \sqrt{\mu^{-1}} \sqrt{\delta^{-1}})$ for preparing c copies of $|\pi\rangle$. The preparation from \mathcal{M} is then more efficient whenever $C(\pi^{\mathcal{M}}) \sqrt{\mu^{-1}} < C(\pi)$; notice that $C(\pi^{\mathcal{M}}) \leq \sqrt[4]{M}$, since $|\pi^{\mathcal{M}}\rangle$ has support on just M elements.

For the problem of sampling from marked elements, being in a slowly evolving sequence of MCs allows to

make the initial set of c samples available. In this case, we again use the state preparation for MC_t to prepare $|\pi_t\rangle^{\otimes c}$ and then map them to $|\pi_{t+1}\rangle^{\otimes c}$. The final step consists in running the algorithm of [38] (namely, an amplitude amplification of the marked subspace) in order to obtain c samples from $|\pi_{t+1}^{\mathcal{M}}\rangle$. This allows then to proceed inductively with sampling in the sequence. The final projection has $\tilde{O}(c \sqrt{\mu^{-1}} \sqrt{\delta^{-1}})$ gate complexity and thus is dominated by the cost of preparing $|\pi_{t+1}^{\mathcal{M}}\rangle^{\otimes c}$.

We notice, finally, that the method just presented can be directly used to produce new copies of $|\pi^{\mathcal{M}}\rangle$, thus directly solving the problem considered in [38]. It can be straightforwardly realized by running any of the algorithms for preparing $|\pi\rangle$ followed by amplitude amplification of the subspace of marked elements.

6 Optimality analysis

To begin with, notice that preparing coherent encodings $|\pi\rangle$ is in general a difficult task, even when sampling from π can be done efficiently. Consider a randomized algorithm that produces an outcome x with probability $\pi(x)$ which makes a number of binary random choices selecting a computational branch b_1 or b_2 with probabilities p or $1-p$. One can “purify” this algorithm to a unitary quantum circuit by substituting every random choice with a controlled dependence on a pure qubit prepared in the state $\sqrt{p}|b_1\rangle + \sqrt{1-p}|b_2\rangle$. The resulting output state then has the form $|\tilde{\pi}\rangle = \sum_x \sqrt{\pi(x)} |x\rangle |\phi(x)\rangle$ where $|\phi(x)\rangle$ contains residual information of all the choices. Starting from $|\tilde{\pi}\rangle$ one cannot directly obtain $|\pi\rangle$ since there is, in general, no efficient deterministic method that allows one to erase the information contained in the second register. In fact, the possibility to efficiently produce a coherent encoding $|\pi\rangle$ for all probability distributions which can be efficiently sampled would imply $\text{SZK} \subseteq \text{BQP}$ [41], that is, Statistical Zero Knowledge problems (including, *e.g.*, Graph Isomorphism) could be solved in quantum polynomial time. While the inclusion $\text{SZK} \subseteq \text{BQP}$ is not impossible, it is expected that specific structures of the problems have to be exploited (*e.g.*, graph-theoretic properties in Graph Isomorphism), while the methods based on MC mixing are oblivious to such problem structures. Consequently, it is highly unlikely that any quantum algorithm can prepare $|\pi\rangle$ encoding stationary distributions of time-reversible MCs in polylog(N) time, not even when the classical mixing process is fast (*i.e.*, when the MC mixes in polylog(N) time).

We prove in the App. G that our algorithm is

strictly optimal in the class of sampling algorithms which utilize oracle access to reflections about $|\pi\rangle$ (and do not use other properties of the transition matrix P of the MC) as is the case of many algorithms based on Szegedy quantum walk [42–46]. Specifically, we show that if we start from $c + 1$ copies of $|\pi\rangle$ and the goal is to obtain $c + 1$ classical samples from π (for some constant c) then $\Omega(\sqrt[4]{N})$ accesses to the reflection oracle are required (more tightly, we can prove a $\Omega(C(\pi))$ lower bound). Our proof relies on the “inner-product adversary” method developed in the context of quantum money [47], a so-called computational no-cloning theorem.

This $\Omega(\sqrt[4]{N})$ lower bound actually applies to any MC, also outside of the context of slowly evolving sequences of MCs. Any algorithm that uses quantum walks just to realize the reflection around $|\pi\rangle$, and then subsequently uses such reflections in a black-box fashion, cannot avoid an $\mathcal{O}(\sqrt[4]{N})$ dependence in its run-time. In particular, algorithms of this type cannot generically achieve the conjectured quadratic speed-up for sampling from stationary distributions of time-reversible MCs [13]. Hence, other techniques are needed.

We finally point out that, however, in our algorithms we have full access to the transition matrix P and, thus, we are not restricted to using reflections around $|\pi\rangle$. In particular, we can implement a classical random walk as well. If the MC is rapidly mixing then, by definition, the random walk allows to efficiently sample from π , while achieving the same goal having access only to reflections around $|\pi\rangle$ and some initial copies of $|\pi\rangle$ could take an exponentially longer time.

7 Discussion

We have presented quantum algorithms for generating samples from stationary distributions of a sequence of Markov chains which achieve a quadratic improvement over previous approaches that can guarantee the generation of the correct output, and work for all time-reversible chains. To achieve this improvement we do not assume special properties of the chain (except detailed balance) but rather we have considered settings where the chains come in a context, namely in a slowly evolving sequence. This result thus has application to all MCs where this framework is natural.

An important domain of application includes statistical physics and material science, where the slowly evolving context, and the need for independent samples, arise when studying phase transitions [25].

A second important family of applications occurs

in machine learning (ML), both in the reinforcement learning case [26] and in the context of generative models [27]. To briefly comment on this domain, as mentioned earlier in reinforcement learning settings [26] where the learner’s distribution over actions is specified by MCs, the MCs are sequentially updated as the system learns [38, 48–50]. The other facet involves the training of certain generative models (used, *e.g.*, for unsupervised learning), such as Boltzmann machines [51]. Here one encounters the need for producing samples from stationary distributions (*e.g.*, Gibbs states) which are themselves slowly modified as the model is updated [52, 53].

We remark that, in ML, the subsequent Markov chains in the sequence are generated according to a training algorithm which depends on the external outputs of previous Markov chains. Whenever this is the case, the methods developed for quantum-enhanced annealing methods become unsuitable, as they need to keep coherence through the protocol steps [22, 23].

We conclude observing that, as a feature of our protocol, at each time step we do not output just a classical sample from the target stationary distribution, but a coherent encoding of this distribution. This is not a guaranteed characteristic of quantum mixing protocols [13] and makes our approach suitable for combining with other quantum protocols which start from such a coherent encoding [29, 38, 44, 54].

Acknowledgements

The authors acknowledge support by the Austrian Science Fund (FWF) through the SFB FoQuS F4012, the Templeton World Charity Foundation grant TWCF0078/AB46, and the DK-ALM: W1259-N27. V.D. thanks G. D. Paparo for initial discussions, and acknowledges the support from the Alexander von Humboldt Foundation. D.O. thanks F. Guatieri for discussion. The authors also thank an anonymous reviewer for pointing out a mistake in a previous version of the work, and a method to correct it.

Author contributions

V.D. and H.J.B. wrote a preliminary version of the article. D.O. has worked on analysing and extending the algorithm as here presented and has written the current version of the article, all under the supervision of H.J.B. and V.D.

References

- [1] Newman, M. E. J. and Barkema, G. T., *Monte Carlo Methods in Statistical Physics*. Oxford University Press (1999).

- [2] Sinclair, A., *Algorithms for random generation and counting: a Markov chain approach*. Springer (1993).
- [3] Bellman, R., *A Markovian decision process*. *Journal of Mathematics and Mechanics* **6(5)**, 679–684 (1957).
- [4] Gilks, W. R., Richardson, S. and Spiegelhalter, D. *Markov chain Monte Carlo in practice*. CRC press (1995).
- [5] Hastings, W. K., *Monte Carlo sampling methods using Markov chains and their applications*. *Biometrika* **57(1)**, 97–109 (1970).
- [6] Geman, S. and Geman, D., *Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images*. *Readings in Computer Vision*, 564–584 (1987).
- [7] Martinelli, F., *Lectures on Glauber dynamics for discrete spin models*. *Lectures on probability theory and statistics*, Springer, 93–191 (1999).
- [8] Norris, J. R., *Markov chains*. Cambridge University Press (1998).
- [9] Levin, D. A. and Peres, Y., *Markov chains and mixing times*. American Mathematical Soc. (2017).
- [10] Aldous, D., László, L. and Winkler, P., *Mixing times for uniformly ergodic Markov chains*. *Stochastic Processes and their Applications* **71(2)**, 165–182 (1995).
- [11] Kirkpatrick, S., Gelatt, C. D. and Vecchi, M. P., *Optimization by simulated annealing*. *Science* **220(4598)**, 671–680 (1983).
- [12] Van Laarhoven, P. J., and Aarts, E. H., *Simulated annealing*. *Simulated annealing: Theory and applications* **37** (1987).
- [13] Richter, P. C., *Quantum speedup of classical mixing processes*. *Phys. Rev. A* **76**, 042306 (2007) [arXiv:0609204].
- [14] Nayak, A. and Vishwanath, A., *Quantum walk on the line*. arXiv:quant-ph/0010117 (2000).
- [15] Ambainis, A., Bach, E., Nayak, A., Vishwanath, A. and Watrous, J., *One-dimensional quantum walks*. *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, 37–49 (2001).
- [16] Aharonov, D., Ambainis, A., Kempe, J. and Vazirani, U., *Quantum walks on graphs*. *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, 50–59 (2001) [arXiv:0012090].
- [17] Richter, P. C., *Almost uniform sampling via quantum walks*. *New J. Phys.* **9(72)** (2007) [arXiv:0606202].
- [18] Dunjko, V. and Briegel, H. J., *Quantum mixing of Markov chains for special distributions*. *New J. Phys.* **17(7)**, 073004 (2015) [arXiv:1502.05511].
- [19] Kempe, J., *Quantum random walks - an introductory overview*. *Contemp. Phys.* **44(4)**, 307–327 (2003) [arXiv:0303081].
- [20] Reitzner, D., Nagaj, D. and Bužek, V., *Quantum Walks*. *Acta Phys. Slovaca* **61(6)**, 603–725 (2011) [arXiv:1207.7283].
- [21] Somma, R. D., Boixo, S., Barnum, H. and Knill, E., *Quantum simulations of classical annealing processes*. *Phys. Rev. Lett.* **101**, 130504 (2008) [arXiv:0804.1571].
- [22] Wocjan, P. and Abeyesinghe, A., *Speedup via quantum sampling*. *Phys. Rev. A* **78**, 042336 (2008) [arXiv:0804.4259].
- [23] Wocjan, P., Chiang, C., Nagaj, D. and Abeyesinghe, A., *Quantum algorithm for approximating partition functions*. *Phys. Rev. A* **80**, 022340 (2009) [arXiv:1405.2749].
- [24] Childs, A., *Quantum information processing in continuous time*. Ph. D. Thesis, Massachusetts Institute of Technology (2004).
- [25] Nishimori, H. and Ortiz, G., *Elements of phase transitions and critical phenomena*. OUP Oxford (2010).
- [26] Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. MIT Press, Cambridge Massachusetts (1998).
- [27] Bishop, C. M., *Pattern recognition and machine learning*. Springer-Verlag, New York (2016).
- [28] Szegedy, M., *Quantum speed-up of Markov chain based algorithms*. *45th Annual IEEE Symposium on Foundations of Computer Science*, 32–41(2004).
- [29] Magniez, F., Nayak, A., Roland, J. and Santha, M., *Search via quantum walk*. *SIAM Journal on Computing* **40(1)**, 142–164 (2011) [arXiv:0608026].
- [30] Kitaev, A. Y., *Quantum measurements and the Abelian Stabilizer Problem*. arXiv preprint quant-ph/9511026 (1995).
- [31] Svore, K. M., Hastings, M. B. and Freedman, M., *Faster Phase Estimation*. *Quantum Information & Computation* **14(3-4)**, 306–328 (2014) [arXiv:1304.0741].
- [32] Wiebe, N. and Granade, C. E., *Efficient Bayesian Phase Estimation* *Phys. Rev. Lett.* **117**, 010503 (2016) [arXiv:1508.00869].
- [33] Grover, L. K., *A fast quantum mechanical algorithm for database search*. *Proceedings of the 28th annual ACM Symposium on the Theory of Computing*, 212–219 (1996) [arXiv:9605043].
- [34] Brassard, G., Hoyer, P., Mosca, M. and Tapp, A., *Quantum Amplitude Amplification and Estimation*. *Contemporary Mathematics* **305**, 53–74 (2002) [arXiv:0005055].

- [35] Grover, L. K., *Fixed-Point Quantum Search*. *Phys. Rev. Lett.* **95**, 150501 (2005) [arXiv:0503205].
- [36] Yoder, T. J., Low, G. H. and Chuang, I. L., *Fixed-Point Quantum Search with an Optimal Number of Queries*. *Phys. Rev. Lett.* **113**, 210501 (2014)
- [37] Grover, L. and Rudolph, T., *Creating superpositions that correspond to efficiently integrable probability distributions*. *arXiv preprint quant-ph/0208112* (2002).
- [38] Paparo, G. D., Dunjko, V., Makmal, A., Matrin-Delgado, MA. and Briegel, H. J. *Quantum speedup for active learning agents*. *Phys. Rev. X* **4**, 031002 (2014) [arXiv:1401.4997].
- [39] Sly, A. *Computational transition at the uniqueness threshold*. 51st Annual IEEE Symposium on Foundations of Computer Science, 287–296 (2010) [arXiv:1005.5584].
- [40] Boyer, M., Brassard, G., Høyer, P. and Tapp, A., *Tight bounds on quantum searching*. *Progress of Physics* **46(4-5)**, 493–505 (1998) [arXiv:9605034].
- [41] Aharonov, D. and Ta-Shma, A., *Adiabatic Quantum State Generation and Statistical Zero Knowledge*. *Proceedings of the 35th annual ACM symposium on Theory of computing*, 20–29 (2003) [arXiv:0301023].
- [42] Ambainis, A., *Quantum walk algorithms for element distinctness*. *SIAM Journal on Computing* **37(1)**, 22–31 (2004) [arXiv:0311001].
- [43] Magniez, F., Santha, M. and Szegedy, M., *Quantum Algorithms for the Triangle Problem*. *SIAM Journal on Computing* **37(2)**, 413–424 (2007) [arXiv:0310134].
- [44] Krovi, H., Magniez, F., Ozols, M. and Roland, J., *Quantum walks can find a marked element on any graph*. *Algorithmica* **74(2)**, 851–907 (2016) [arXiv:1002.2419].
- [45] Temme, K., Osborne, T. J., Vollbrecht, K. G. H., Poulin, D. and Verstraete, F., *Quantum metropolis sampling*. *Nature* **471**, 87–90 (2011), [arXiv:0911.3635].
- [46] Yung, M.-H. and Aspuru-Guzik, A., *A quantum-quantum metropolis algorithm*. *Proceedings of the National Academy of Sciences* **109(3)**, 754–759 (2012) [arXiv:1011.1468].
- [47] Aaronson, S. and Christiano, P., *Quantum Money from Hidden Subspaces*. *Theory of Computing* **9(9)**, 349–401 (2013) [arXiv:1203.4740].
- [48] Briegel, H. J. and De las Cuevas, G., *Projective simulation for artificial intelligence*. *Sci. Rep.* **2**, 400 (2012).
- [49] Mautner, J., Makmal, A., Manzano, D., Tiersch, M. and Briegel, H. J., *Projective simulation for classical learning agents: a comprehensive investigation*. *New Generat. Comput.* **33(1)**, 69–114 (2015) [arXiv:1305.1578].
- [50] Dunjko, V. and Briegel, H. J., *Machine learning & artificial intelligence in the quantum domain: a review of recent progress*. *Reports on Progress in Physics* **81(7)**, 074001 (2018) [arXiv:1709.02779].
- [51] Fischer, A. and Christian, I., *An introduction to restricted Boltzmann machines*. *Iberoamerican Congress on Pattern Recognition*, 14–36 (2012).
- [52] Tieleman, T., *Training restricted Boltzmann machines using approximations to the likelihood gradient*. *Proceedings of the 25th international conference on Machine learning*, 1064–1071 (2008).
- [53] Wiebe, N., Kapoor, A. and Svore, K. M., *Quantum deep learning*. *Quantum Information & Computation* **16(7-8)**, 541–587 (2016) [arXiv:1412.3489].
- [54] Montanaro, A., *Quantum speedup of Monte Carlo methods*. *Proceedings of the Royal Society A* **471(2181)**, 0301 (2015) [arXiv:1504.06987].

A Markov chain notions

Here we review the fundamental notions of Markov chain theory and refer to [8, 9] for further details.

Transition matrices and probability distributions: We deal with discrete-time Markov chains having a finite number N of states. Therefore, to a MC is associated a left-stochastic matrix P (a matrix with non-negative entries which add up to one in every column) of size $N \times N$, and each entry $P_{x,y}$ specifies the transition probability from the state x to state y . Correspondingly, the non-negative (column) vector $\boldsymbol{\pi}$ denotes a probability distribution over the state space as

$$\begin{aligned} \boldsymbol{\pi} &= (\pi(1), \dots, \pi(N))^T \\ \text{with } \sum_{x=1}^N \pi(x) &= 1. \end{aligned} \quad (9)$$

A MC is then specified by a the transition matrix P and an initial probability distribution $\boldsymbol{\pi}_{in}$. We stick to the convention of left-stochastic matrices which act from the left on column vectors $\boldsymbol{\pi}$ representing probability distributions, that is $\boldsymbol{\pi}' = P\boldsymbol{\pi}$. This convention is not customary in the MC literature (where the usage of right-stochastic matrices prevails), but it matches the one adopted in the quantum information community. In particular, $P_{y,x}$ denotes the transition probability from the element x to the element y .

Ergodic MCs: A N -state MC is *irreducible* if it is possible from each state x to reach any other state y in a finite number of steps and with non-zero probability. The *period* of a state x is the largest positive integer such that any return to x can occur only at multiples of that integer. If the period of all states is 1, the MC is said to be *aperiodic*. If P is irreducible and aperiodic, then there exists a unique *stationary distribution* $\boldsymbol{\pi}$, such that:

$$P\boldsymbol{\pi} = \boldsymbol{\pi} \quad (10)$$

and, moreover, $\boldsymbol{\pi}$ has support over all the elements of the MC. This also implies that, under application of a sufficiently large number of steps any initial probability distribution $\tilde{\boldsymbol{\pi}}$ will converge to the unique stationary distribution, $\lim_{k \rightarrow \infty} P^k \tilde{\boldsymbol{\pi}} = \boldsymbol{\pi}$. This convergence process is called *mixing*, and since MCs mix if and only if they are irreducible and aperiodic, these are called *ergodic Markov chains*.

Time reversal: The *time reversal* \hat{P} of a Markov chain P having stationary distribution $\boldsymbol{\pi}$ is defined as:

$$\hat{P}_{y,x} := P_{x,y} \frac{\pi(y)}{\pi(x)} \quad (11)$$

and a MC is said to be time-reversible if it is equal to its time-reversed version, $P = \hat{P}$. Equivalently, a time-reversible MC is one that satisfies the *detailed balance equation*:

$$P_{y,x} \pi(x) = P_{x,y} \pi(y). \quad (12)$$

We can also write the time reversed MC in matrix form as $\hat{P} = D(\boldsymbol{\pi})P^T D(\boldsymbol{\pi})^{-1}$, where $D(\boldsymbol{\pi})$ is the diagonal matrix $D(\boldsymbol{\pi}) := \text{diag}(\pi(1), \dots, \pi(N))$. This implies that if P is time reversible, then its spectrum is real. In the following, we will always consider ergodic and time-reversible MCs.

Mixing times: Obviously, not all mixing process of ergodic MCs are equally fast. We use the *total variation distance*, defined as $d(\boldsymbol{\pi}, \boldsymbol{\pi}') := \frac{1}{2} \sum_x |\pi(x) - \pi'(x)|$ to assess the speed of mixing (the total variation distance exactly matches the trace distance in the quantum information context). We then define $d(k) := \max_{\boldsymbol{\sigma}} d(P^k \boldsymbol{\sigma}, \boldsymbol{\pi})$ as the distance in distributions between a sample drawn after k walk steps starting from any distribution $\boldsymbol{\sigma}$

and stationary distribution π of P . The *mixing time* $t_{\text{mix}}(\epsilon)$ then is defined as the smallest time necessary to bring any initial distribution within distance ϵ from the stationary distribution, $d(t_{\text{mix}}(\epsilon)) \leq \epsilon$. We then set $t_{\text{mix}} := t_{\text{mix}}(1/4)$. It can be shown then that the convergence of an ergodic MC is exponentially fast in terms of the mixing time, that is:

$$d(\ell t_{\text{mix}}) \leq 2^{-\ell}. \quad (13)$$

The mixing times often play the critical role in the computational complexity of MC-based algorithms. There are many techniques that can be employed for upper and lower bounding the mixing time, but one of the most useful characterizations is the following. Because of Perron-Frobenius theorem all eigenvalues of a left-stochastic matrix P are smaller or equal to 1 in modulus. If P is ergodic then its stationary distribution $|\pi\rangle$ is the only eigenvector of P having eigenvalue equal to +1. That is, all other eigenvectors have eigenvalues λ with $|\lambda| < 1$. Let $\sigma(P)$ be the spectrum of a time-reversible Markov chain P ; we define the *spectral gap* δ of P as:

$$\delta := 1 - \max_{\substack{\lambda \in \sigma(P): \\ \lambda \neq 1}} |\lambda| \quad (14)$$

i.e. the minimum of $1 - |\lambda|$ over the eigenvalues of P which differ from one. The spectral gap is a rather tight estimate for the inverse of the mixing time, since

$$\left(\frac{1}{\delta} - 1\right) \log\left(\frac{1}{2\epsilon}\right) \leq t_{\text{mix}}(\epsilon) \leq \frac{1}{\delta} \log\left(\frac{1}{\epsilon \pi_{\min}}\right) \quad (15)$$

holds for all time-reversible MCs (where π_{\min} is the smallest probability in π). In short we have $t_{\text{mix}} \in \tilde{\mathcal{O}}(1/\delta)$ and $1/\delta \in \tilde{\mathcal{O}}(t_{\text{mix}})$, giving asymptotic upper and lower bounds to the mixing time.

B Szegedy quantum walk

Here we review the basics of Szegedy quantum walks [28]. For further details see [29, 44] and references therein.

Szegedy walk operator: The Szegedy walk operator $W(P)$ can be implemented for any transition matrix P , and not only for those associated to ergodic and time-reversible MCs (but $W(P)$ has nice spectral properties only if P is ergodic and time reversible). The basic building block to define $W(P)$ is the *diffusion operator* U_P which acts on two quantum registers of N states and is (partially) defined as follows:

$$U_P |x\rangle_1 |0\rangle_2 := |x\rangle_1 \sum_{y=1}^N \sqrt{P_{x,y}} |y\rangle_2. \quad (16)$$

By measuring the second register in the computational basis a step of the classical random walk is obtained, hence U_P is a natural way of defining a quantum extension of the classical MC. When we say that we have quantum access to P , we specifically mean that we have access to a diffusion operator of the form (16). The diffusion operator U_P can be efficiently realized, for instance, when P is a sparse transition matrix. Then, we can define the *Szegedy walk operator* as the unitary

$$W(P) := \text{SWAP } U_P (\mathbb{I}_1 \otimes Z_2) U_P^\dagger, \quad (17)$$

where $Z_2 := 2|0\rangle\langle 0|_2 - \mathbb{I}$ and SWAP interchanges the first and second register. $W(P)$ acts non-trivially on the invariant subspace $A + B$, where $A := \text{span}\{U_P |x\rangle |0\rangle\}_x$ and $B := \text{span}\{\text{SWAP } U_P |x\rangle |0\rangle\}_x$.

Spectral properties of $W(P)$: When P is ergodic and time-reversible the space $A + B$ has dimension $2N - 1$ and the intersection $A \cap B$ contains only the state $U_P |\pi, 0\rangle = \text{SWAP } U_P |\pi, 0\rangle$, as one can verify using the detailed balance equation for P . The state $U_P |\pi, 0\rangle$ is the only +1-eigenstate of $W(P)\Pi_{A+B}$, where Π_{A+B} is a projector on the invariant subspace $A + B$. Moreover on the invariant subspace the other $2N - 2$ eigenvalues

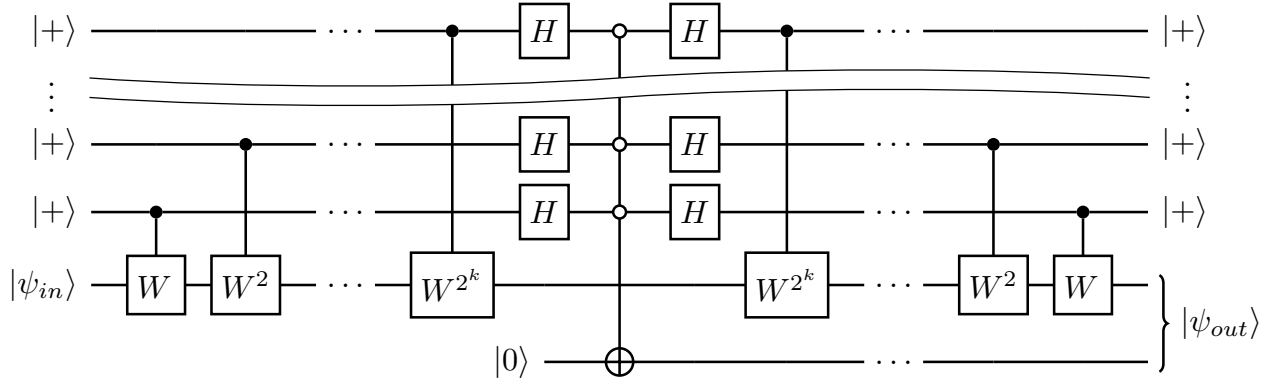


Figure 2: Phase detection algorithm applied to the Szegedy operator $W(P)$. The left part of the circuit implements the standard phase estimation algorithm, except for the fact that a final Hadamard transform is applied instead of an inverse quantum Fourier transform. Using the Hadamard transform is sufficient since we only require to discriminate the $+1$ eigenvector of $W(P)$ from eigenvectors having eigenvalue different from $+1$. The central multi-controlled Toffoli gate flips an ancilla qubit if and only if all control qubits are in $|0\rangle$. The right part of the circuit finally uncomputes the value contained in the ancillary registers.

of $W(P)$ are given by $\{e^{\pm i\theta_\ell}\}_{\ell \in [N]}$ where $\{\cos \theta_\ell\}_{\ell \in [N]}$ are the eigenvalues of P that are different from one (remember, the spectrum of P is real for time-reversible MCs). With this notation the phase gap (Δ) and spectral gap (δ) are given by

$$\begin{cases} \Delta := \min_{\ell} \{\theta_\ell\} \\ \delta := \min_{\ell} \{1 - |\cos \theta_\ell|\} \end{cases} \quad \text{with } \theta_\ell \in (0, \pi) \quad (18)$$

and therefore the phase gap is quadratically larger than the spectral gap, $\Delta \geq \sqrt{2\delta}$.

C Subroutines based on quantum walks

Here we show how to use Szegedy walk operator within the phase detection algorithm to implement projective measurements onto $|\pi\rangle$ and partial reflections around $|\pi\rangle$, as originally done in [29]; we also show how to use fixed-point amplitude amplification to deterministically map a given input state to $|\pi\rangle$.

Phase estimation and phase detection: The phase detection algorithm applied to the Szegedy walk operator $W(P)$ is illustrated in Fig. 2 and its action is as follows. Call $\{|\theta_\ell\rangle\}_\ell$ the eigenvectors of $W(P)$ having eigenvalue $e^{i\theta_\ell}$, $W(P)|\theta_\ell\rangle = e^{i\theta_\ell}|\theta_\ell\rangle$. In particular $|\theta_\ell = 0\rangle \equiv |\pi\rangle$. Then, given an input state of the form $|\psi_{in}\rangle = \sum_{\ell} \psi_{\theta_\ell} |\theta_\ell\rangle$ the phase detection algorithm outputs an approximation of the state

$$|\psi_{out}\rangle = \psi_0 |\pi\rangle |1\rangle + \sum_{\ell: \theta_\ell \neq 0} \psi_{\theta_\ell} |\theta_\ell\rangle |0\rangle. \quad (19)$$

That is, the second register contains a bit signalling whether the $|\theta_\ell\rangle$ is the $+1$ -eigenvector or not. The phase detection algorithm produces a state within trace distance ε from the state in Eq. (19) using $\mathcal{O}(\Delta^{-1} \log \varepsilon^{-1}) = \tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ oracle accesses to controlled- $W(P)$ and $\mathcal{O}(\Delta^{-1} \log \varepsilon^{-1}) = \tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$ extra gates. In particular, the cost has only a logarithmic dependence on the error, see *e.g.*, [29, 31, 32].

Projective measurement onto $|\pi\rangle$: The phase detection algorithm allows to directly implement the projective measurement given by the projectors $\{|\pi\rangle\langle\pi|, \mathbb{I} - |\pi\rangle\langle\pi|\}$, applied to any arbitrary input state $|\psi_{in}\rangle$: it is realized by measuring the second register of the state in Eq. (19) in the computational basis. The gate and oracle complexity of this projective measurement is thus the same of the phase detection algorithm,

$\tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$. The success probability of the measurement, applied on an input pure state $|\psi_{in}\rangle$, is approximately $|\langle \psi_{in} | \boldsymbol{\pi} \rangle|^2$. Moreover, the classical outcome of this projective measurement is a bit that signals whether the projection onto $|\boldsymbol{\pi}\rangle\langle \boldsymbol{\pi}|$ was successful or not. This allows, *e.g.*, to redo the preparation and measurement process until the algorithm succeeds in obtaining the target state $|\boldsymbol{\pi}\rangle$.

Partial reflections around $|\boldsymbol{\pi}\rangle$: Next, the phase detection algorithm can be used to approximately implement the *partial reflection*

$$R_\phi(\boldsymbol{\pi}) := e^{i\phi} |\boldsymbol{\pi}\rangle\langle \boldsymbol{\pi}| + (\mathbb{I} - |\boldsymbol{\pi}\rangle\langle \boldsymbol{\pi}|) \quad (20)$$

where ϕ is a tunable parameter. Notice that for $\phi = 180^\circ$ the partial reflection becomes a standard reflection around $|\boldsymbol{\pi}\rangle$, $R(\boldsymbol{\pi}) = \mathbb{I} - 2|\boldsymbol{\pi}\rangle\langle \boldsymbol{\pi}|$. A partial reflection can be implemented using once the circuit in Fig. 2 and once its inverse: with the first call a input state $|\psi_{in}\rangle$ is mapped to a state as in Eq. (19); then, a phase $e^{i\phi}$ is applied selectively on the ancilla qubit being in the $|1\rangle$ state; finally, the phase detection algorithm is run in reverse to uncompute the bit contained in the second register. In summary:

$$\begin{aligned} |\psi\rangle|0\rangle &\mapsto \psi_0 |\boldsymbol{\pi}\rangle|1\rangle + \sum_{\ell: \theta_\ell \neq 0} \psi_\ell |\theta_\ell\rangle|0\rangle \\ &\mapsto e^{i\phi} \psi_0 |\boldsymbol{\pi}\rangle|1\rangle + \sum_{\ell: \theta_\ell \neq 0} \psi_\ell |\theta_\ell\rangle|0\rangle \\ &\mapsto \left(e^{i\phi} \psi_0 |\boldsymbol{\pi}\rangle + \sum_{\ell: \theta_\ell \neq 0} \psi_\ell |\theta_\ell\rangle \right) |0\rangle. \end{aligned} \quad (21)$$

Notice that in the operation given above we apply in sequence a phase estimation and its inverse, and these two operations cancel out. In conclusion, the oracle and gate cost of approximating a partial reflection is also in $\mathcal{O}(\Delta^{-1} \log \varepsilon^{-1}) = \tilde{\mathcal{O}}(\sqrt{\delta^{-1}})$.

Fixed-point amplitude amplification: Partial reflections are fundamental for implementing *fixed-point amplitude amplification* (FPAA). This is a variant of amplitude amplification whereby the output state can get arbitrarily close to the ideal target state. In standard amplitude amplification usually one has the “soufflé problem” [35]: if the rotation in the amplitude amplification process is not stopped at the right moment, the fidelity with the target state starts decreasing again; moreover, even using the optimal number of reflections, only a constant fidelity between the outputs state and the ideal target is reached. In contrast, in FPAA one gets exponentially close to the target state $|\psi_{out}\rangle$, allowing an almost exact preparation of this state, with only a logarithmic dependence of run-time on the approximation error. Details follow.

A FPAA algorithm takes a single copy of a input state $|\psi_{in}\rangle$ and maps it to a state ε -close to $|\psi_{out}\rangle \equiv \Pi_{out}|\psi_{in}\rangle / \|\Pi_{out}|\psi_{in}\rangle\|$, where Π_{out} is a projector over a target subspace (after the process the input state $|\psi_{in}\rangle$ is no longer available). More precisely, in order to implement FPAA three ingredients are required:

1. a single copy of a input quantum state $|\psi_{in}\rangle$;
2. the ability of implementing *partial reflections* around the input state;
3. the ability of implementing *partial reflections* around the target subspace.

Specifically, these partial reflections are respectively given by $R_\phi(\psi_{in}) = e^{i\phi} |\psi_{in}\rangle\langle \psi_{in}| + (\mathbb{I} - |\psi_{in}\rangle\langle \psi_{in}|)$ and $R_{\phi'}(\Pi_{out}) = e^{i\phi'} \Pi_{out} + (\mathbb{I} - \Pi_{out})$, for arbitrary angles ϕ, ϕ' .

The FPAA algorithm of Yoder et. al. [36] can be implemented, provided that the conditions 1.-3. hold, and has both the quadratic speedup of Grover search and the fixed-point property. This algorithm is parametric, depending on two input parameters, $\varepsilon \in [0, 1]$ and $\gamma \in (0, 1)$: if $|\langle \psi_{out} | \psi_{in} \rangle| \geq \sqrt{\gamma}$ holds, then the output of the FPAA algorithm is a state with ε distance in trace norm from the ideal $|\psi_{out}\rangle$. The number of calls to $R_\phi(\psi)$ and $R_{\phi'}(\Pi_{out})$ is in $\mathcal{O}(\sqrt{\gamma^{-1}} \log \varepsilon^{-1})$.

Heralded state preparation: We finally show how to use FPAA followed by a projective measurement to implement (in an efficient way) a *heralded preparation* of $|\pi\rangle$. If we start from an initial state $|\psi_{in}\rangle$ and then apply the projective measurement $\{|\pi\rangle\langle\pi|, \mathbb{I} - |\pi\rangle\langle\pi|\}$ the success probability is $|\langle\pi|\psi_{in}\rangle|^2$; but the success probability of the measurement process can be increased by preceding the measurement by a round of amplitude amplification. On average, the procedure using amplitude amplification has a quadratically smaller run-time in producing a copy of $|\pi\rangle$.

More precisely, the heralded state preparation works as follows. We first run the optimal FPAA algorithm [36] using reflections around a initial state $|\psi_{in}\rangle$ and around the target state $|\pi\rangle$; the algorithm is run setting ε as the target approximation error and setting some value $\gamma > 0$ as overlap parameter. This means that, if the inequality $|\langle\pi|\psi_{in}\rangle|^2 \geq \gamma$ holds, then FPAA guarantees to output a state within ε distance from $|\pi\rangle$; however, if instead $|\langle\pi|\psi_{in}\rangle|^2 < \gamma$ holds, the output state can be arbitrarily far from $|\pi\rangle$. To ameliorate this issue, after the FPAA we apply a projective measurement onto $|\pi\rangle$ (or onto the orthogonal subspace). When the measurements succeeds, the preparation of (a approximation of) $|\pi\rangle$ is guaranteed, independently from the initial overlap $|\langle\pi|\psi_{in}\rangle|$. We also remind that this final measurement succeeds almost deterministically (with probability $1 - \varepsilon$) when $|\langle\pi|\psi_{in}\rangle|^2 \geq \gamma$ holds. The number of reflections needed in the heralded preparation of $|\pi\rangle$ is then $\mathcal{O}(\sqrt{\gamma^{-1}} \log(\varepsilon^{-1}))$ and the total run-time is in $\tilde{\mathcal{O}}(\sqrt{\gamma^{-1}} \log^2(\varepsilon^{-1}))$, as we will prove in the next Appendix.

D Analysis of imperfect reflection operators

Here we consider the propagation of errors when the partial reflection used within FPAA are approximate and how the run-time is affected. For this section only, we assume that the relevant parameters in the soft- \mathcal{O} notation are δ, γ and $\log \varepsilon^{-1}$; namely, we keep $\log \varepsilon^{-1}$ terms and discard $\log \log \varepsilon^{-1}$ dependencies.

We suppose that $\sqrt{\gamma}$ is (a lower bound to) the overlap between $|\psi_{in}\rangle$ and $|\pi\rangle$, and the final targeted error is ε . FPAA entails the use of $\mathcal{O}(\sqrt{\gamma^{-1}} \log \varepsilon^{-1})$ perfect reflections in order to achieve the desired accuracy goal. However, the same can be achieved with imperfect reflections, provided that each reflection has an error smaller of $\varepsilon/(\text{number of steps})$: by the triangle inequality the total accumulated error will be upper bounded by ε . That is, we need to implement a partial reflection with an accuracy

$$\varepsilon^R = \mathcal{O}\left(\frac{\varepsilon}{\sqrt{\gamma^{-1}} \log \varepsilon^{-1}}\right). \quad (22)$$

Hence the total gate cost of the heralded state preparation procedure (nesting approximate reflections within FPAA) is given by:

$$\begin{aligned} & \mathcal{O}\left(\sqrt{\gamma^{-1}} \log \varepsilon^{-1}\right) \times \mathcal{O}\left(\sqrt{\delta^{-1}} \log(1/\varepsilon^R)\right) \\ &= \mathcal{O}\left(\sqrt{\gamma^{-1}} \sqrt{\delta^{-1}} \log(\varepsilon^{-1}) \left[\log \sqrt{\gamma^{-1}} + \log \varepsilon^{-1} + \log \log \varepsilon^{-1}\right]\right) \\ &= \tilde{\mathcal{O}}\left(\sqrt{\gamma^{-1}} \sqrt{\delta^{-1}} \log^2(\varepsilon^{-1})\right). \end{aligned} \quad (23)$$

Thus, heralded preparation of $|\pi\rangle$ within ε final precision has a overall gate complexity scaling as $\log^2(\varepsilon^{-1})$. We also remark that errors do not propagate from one time step to the next in the slowly evolving sequence, since at each time step we freshly prepare new copies of $|\pi\rangle$. This is possible since we have access to projectors onto the required states, which allow to decrease approximation errors.

E Failure probabilities of preparation from uniform and from samples

We here show that the Prepare subroutines are not overly sensitive to small imperfections and that failure probabilities decrease exponentially with c , the number of classical samples carried over in each step of the slowly evolving sequence.

Preparation from uniform distribution: For PrepareFromUniform, as given in [Alg. 1](#), the analysis is simple, since the input state $|\mathbf{u}\rangle$ has no error. The algorithm succeeds when at least c out of $2c$ heralded preparations of $|\boldsymbol{\pi}\rangle$ starting from $|\mathbf{u}\rangle$ are successful. In the case in which $\chi' \geq \sqrt{N}/f(\boldsymbol{\pi})$ the global probability of failure is $2^{-\mathcal{O}(c)}$. In fact the $2c$ runs have independent outcomes; then, we can apply the Chernoff bound:

$$\Pr[\text{number of failures} \geq (1 + \delta) 2\varepsilon c] \leq \exp\left(-\frac{\delta^2}{2 + \delta} 2\varepsilon c\right) \quad (24)$$

where ε is an upper bound to the failure probability in the preparation of $|\boldsymbol{\pi}\rangle$ and $\delta > 0$ is a free parameter. Choosing $1 + \delta = 1/(2\varepsilon)$ we get:

$$\begin{aligned} \Pr[\text{number of failures} \geq c] &\leq \exp\left(-\frac{(1 - 2\varepsilon)^2}{1 + 2\varepsilon} c\right) \\ &\leq \exp(-0.9 c), \end{aligned} \quad (25)$$

where the second inequality holds for sufficiently small ε .

Preparation from samples: A similar analysis holds for PrepareFromSamples, as given in [Alg. 2](#). Notice that the input samples $\vec{x} = \{x_1, \dots, x_c\}$ are not (exactly) distributed with $\boldsymbol{\pi}$, but with a distribution $\tilde{\boldsymbol{\pi}}$ which is ε -close to $\boldsymbol{\pi}$, say, in total variation distance. Considering a random variable \tilde{X} distributed as $\tilde{\boldsymbol{\pi}}$ we have, for any $v > 0$:

$$\begin{aligned} \Pr\left[\pi^{-1/2}(\tilde{X}) \geq v\right] &= \sum_{x: \pi^{-1/2}(x) \geq v} \tilde{\boldsymbol{\pi}}(x) \\ &\leq \varepsilon + \sum_{x: \pi^{-1/2}(x) \geq v} \boldsymbol{\pi}(x) \\ &= \varepsilon + \Pr\left[\pi^{-1/2}(X) \geq v\right] \\ &\leq \varepsilon + \frac{\mathbb{E}[\pi^{-1/2}(X)]}{v}. \end{aligned} \quad (26)$$

In the first inequality we have applied the definition of total variation distance and in the second Markov's inequality. Thus, we have

$$\Pr\left[\pi^{-1/2}(\tilde{X}) \geq 2\mathbb{E}[\pi^{-1/2}(X)]\right] \leq \varepsilon + \frac{1}{2}. \quad (27)$$

Thus, with high probability at least one sample in $x_* \in \vec{x}$ satisfies $\pi^{-1/2}(x_*) < 2\mathbb{E}[\pi^{-1/2}(X)] = 2f(\boldsymbol{\pi})$; namely, this happens with probability at least:

$$1 - \left(\frac{1 + 2\varepsilon}{2}\right)^c = 1 - 2^{-\mathcal{O}(c)}. \quad (28)$$

Then, we consider a heralded state preparation of $|\boldsymbol{\pi}\rangle$ starting from $|x_*\rangle$ for $2c$ times. If $\chi \geq f(\boldsymbol{\pi})$ the analysis proceeds exactly as the one performed for PrepareFromUniform and, hence, with probability $1 - 2^{-\mathcal{O}(c)}$ at least c of these $2c$ runs will be successful in producing approximations of $|\boldsymbol{\pi}\rangle$. The global failure probability of PrepareFromSamples is thus in $2^{-\mathcal{O}(c)}$.

Combined algorithm: The algorithms PrepareFromUniform and PrepareFromSamples are used as subroutines of the combined state preparation algorithm, as specified in the main text. In this algorithm the values of χ and $\chi' = \sqrt{N}/\chi$ are doubled until they exceed $2\sqrt[4]{N}$, in which case either $\chi \geq f(\boldsymbol{\pi})$ or $\chi' \geq \sqrt[4]{N}/f(\boldsymbol{\pi})$ is satisfied: then, the algorithm has to succeed, except with probability $2^{-\mathcal{O}(c)}$.

F Quantum memory algorithm

Here we show that, if a long-term quantum memory is available, only one quantum sample (*i.e.*, $|\pi_t\rangle$) has to be stored in memory between consecutive steps in the slowly evolving sequence. That is, we assume that the quantum state $|\pi_t\rangle$ does not decohere during the time in which P_t is updated to P_{t+1} . Then, one can store a single copy $|\pi_t\rangle$ and employ it to (almost deterministically) prepare two copies of $|\pi_t\rangle$. One copy of $|\pi_t\rangle$ is provided as external output, while the other copy is rotated to $|\pi_{t+1}\rangle$ using fixed-point amplitude amplification and $|\pi_{t+1}\rangle$ is provided as input to the successive MC. Therefore, we only have to show how to implement this state duplication algorithm.

State duplication algorithm: The state duplication algorithm works as follows, assuming that $f(\pi)$ is known. If $f(\pi) \geq \sqrt[4]{N}$, then we use `PrepareFromUniform` and the second copy of $|\pi\rangle$ is produced *de novo* from the uniform distribution. Else ($f(\pi) < \sqrt[4]{N}$), we employ U_{PFS} , a coherent version of `PrepareFromSamples` as described in [Alg. 2](#), for the case $c = 1$. Notice that U_{PFS} is a quantum algorithm that tries to prepare $|\pi\rangle$ from a single classical sample x drawn from π ; hence it can be written as an isometry (that is, as a unitary operation, plus the ability to add ancillary quantum systems) acting on a register initialized in $|x\rangle$:

$$U_{\text{PFS}}|x\rangle = |x\rangle \left[\sqrt{p_{\text{succ}}(x)} |\pi\rangle |\text{ok}\rangle + \sqrt{1 - p_{\text{succ}}(x)} |\psi_x\rangle |\text{err}\rangle \right]. \quad (29)$$

Here the first register is the control (input) register, the third register contains a flag heralding the successful preparation of $|\pi\rangle$, and the second register either contains $|\pi\rangle$ or an arbitrary state $|\psi_x\rangle$ in case of failure. The algorithm U_{PFS} has a run-time proportional to $f(\pi)$ and has, averaging on x , a constant success probability, say larger than $1/2$: $\mathbb{E}_{\pi}[p_{\text{succ}}(X)] = \sum_x \pi(x) p_{\text{succ}}(x) \geq \frac{1}{2}$. We then consider the following quantum computation

$$\begin{aligned} |\pi\rangle &\mapsto \sum_x \sqrt{\pi(x)} |x\rangle |x\rangle \\ &\mapsto \sum_x \sqrt{\pi(x)} U_{\text{PFS}}|x\rangle U_{\text{PFS}}|x\rangle. \end{aligned} \quad (30)$$

The state in Eq. (30) can be rewritten as follows, after rearrangement of the quantum registers:

$$|\tilde{\pi}^{(2)}\rangle := \left[\sum_x \sqrt{\pi(x)} p_{\text{succ}}(x) |x\rangle^{\otimes 2} \right] |\pi\rangle^{\otimes 2} |\text{ok}'\rangle + \sqrt{1 - p'_{\text{succ}}} |\psi\rangle |\text{err}'\rangle. \quad (31)$$

Here the rightmost register is in $|\text{ok}'\rangle$ if both instances of U_{PFS} have raised a success flag and is in $|\text{err}'\rangle$ otherwise; also, the probability of two successful preparations of $|\pi\rangle$ is given by

$$p'_{\text{succ}} = \left\| \sum_x \sqrt{\pi(x)} p_{\text{succ}}(x) |x\rangle^{\otimes 2} \right\|^2 = \mathbb{E}_{\pi}[p_{\text{succ}}^2(X)] \geq \frac{1}{4}. \quad (32)$$

Hence upon measurement of the last register of $|\tilde{\pi}^{(2)}\rangle$, when we obtain as outcome $|\text{ok}'\rangle \equiv |\text{ok}\rangle|\text{ok}\rangle$ we also obtain two copies of $|\pi\rangle$. This happens with probability larger than $1/4$. The final step of the algorithm is to use FPAA to deterministically apply the projector $\Pi_{\text{ok}} := \mathbb{I} \otimes |\text{ok}'\rangle\langle\text{ok}'|$ to the state $|\tilde{\pi}^{(2)}\rangle$, thus deterministically recovering $|\pi\rangle^{\otimes 2}$ (together with an ancillary register in a separable quantum state, which can be discarded). Here FPAA can be implemented using $\tilde{\mathcal{O}}(1)$ accesses to U_{PFS} and thus has essentially the same run-time as the classical-memory `PrepareFromSamples` algorithm.

Further remarks: Notice that if we want to output $c \geq 3$ copies of $|\pi\rangle$, this can be obtained by applying the state duplication algorithm in sequence many times (which is more efficient than using a modification of Eq. (30) in which U_{PFS} is used c times in parallel). Finally, if the value of $f(\pi)$ is not known, one can simply revert to the classical-memory strategy, at the cost of carrying over c classical samples in order to have a $2^{-\mathcal{O}(c)}$ failure probability.

G Lower bound on the oracle cost of sampling

Here we prove a lower bound on the number of reflections around $|\pi\rangle$ that are needed to produce two classical samples drawn from π , starting from a single copy of $|\pi\rangle$. The lower bound also applies when $c + 1$ classical samples have to be produced from c copies of $|\pi\rangle$, thus we directly prove this more general case.

The state preparation algorithms presented in this work allow to prepare $|\pi\rangle$ using $\tilde{\mathcal{O}}(C(\pi)) \leq \mathcal{O}(\sqrt[4]{N})$ reflections around $|\pi\rangle$, using at most logarithmically many copies of $|\pi\rangle$. A result of Aaronson and Christiano [47] is that there exists a class of states with all positive real amplitudes (effectively, coherent encodings of probability distributions) which require, on average, $\Omega(\sqrt[4]{N})$ accesses to the reflection oracle in order to be duplicated. This already shows that our algorithms have essentially optimal worst-case performance.

We strengthen the result in two ways. First, we prove a $\Omega(C(\pi))$ lower bound in the number of oracle accesses needed, thus matching the $\mathcal{O}(C(\pi))$ oracle complexity attained by our algorithms, for all values of $C(\pi)$. Secondly, we show that the same lower bound holds also for classical sampling problems. Namely, suppose that we have c initial copies of $|\pi\rangle$ and the ability to implement controlled reflections around $|\pi\rangle$, while the goal is to obtain $c + 1$ classical samples distributed according to π . We show that in order to accomplish this task $\Omega(C(\pi)/c^2)$ controlled-reflection around $|\pi\rangle$ are required. This means that our algorithm is asymptotically optimal (up to polylogarithmic factors) in the number of queries to a reflection oracle.

The proof of this lower bound hinges upon the ‘‘inner-product adversary’’ method [47]. We can condense the results of Section 4.2 and Appendix B of [47] into the following theorem.

Theorem 1. *Suppose that we have access to reflection oracles U_ψ (and to its controlled version $c-U_\psi$) so that:*

$$\begin{aligned} U_\psi|\psi\rangle &= -|\psi\rangle \\ U_\psi|\eta\rangle &= +|\eta\rangle \quad \forall |\eta\rangle \text{ orthogonal to } |\psi\rangle. \end{aligned} \quad (33)$$

The states $|\psi\rangle$ come from a subset \mathcal{Z} of the entire Hilbert space \mathcal{H} . Moreover we require that on these states there is a symmetric binary relation $\mathcal{R} \subseteq \mathcal{Z} \times \mathcal{Z}$ such that

$$\forall |\psi\rangle \in \mathcal{Z} : (\psi, \psi) \notin \mathcal{R} \quad (34)$$

$$\forall |\psi\rangle \in \mathcal{Z}, \exists |\phi\rangle \in \mathcal{Z} : (\psi, \phi) \in \mathcal{R}. \quad (35)$$

Suppose, next, that for all $|\psi\rangle \in \mathcal{Z}$ and for all $|\eta\rangle \in \mathcal{H}$ that are orthogonal to $|\psi\rangle$ the following inequality holds

$$\mathbb{E}_{\substack{\phi \in \mathcal{Z}: \\ (\psi, \phi) \in \mathcal{R}}} \left[|\langle \eta | \phi \rangle|^2 \right] \leq \gamma \quad (36)$$

for some $\gamma \in \mathbb{R}^+$.

Then, consider a quantum circuit Q^ψ consisting of a fixed set of unitary operations Q^ that make oracle calls to $c-U_\psi$ (and similarly, Q^ϕ is obtained when Q^* calls $c-U_\phi$). Suppose that for all $(\psi, \phi) \in \mathcal{R}$ the quantum states $|\Psi_{in}^\psi\rangle$ and $|\Psi_{in}^\phi\rangle$ are two input states such that $|\langle \Psi_{in}^\phi | \Psi_{in}^\psi \rangle| \geq \alpha$, while the output states $|\Psi_{out}^\psi\rangle = Q^\psi |\Psi_{in}^\psi\rangle$ and $|\Psi_{out}^\phi\rangle = Q^\phi |\Psi_{in}^\phi\rangle$ have to satisfy $|\langle \Psi_{out}^\phi | \Psi_{out}^\psi \rangle| \leq \beta$. Then Q^* must make*

$$\Omega\left(\frac{\alpha - \beta}{\sqrt{\gamma}}\right) \quad (37)$$

accesses to a $c-U_\psi$ or $c-U_\phi$ to obtain these output states.

This theorem can be applied to our case as follows. The input state consists of c copies of $|\pi\rangle$, while the output state consists of $c + 1$ classical samples from π . Namely, we are considering a quantum circuit Q^π which aims at producing these $c + 1$ classical samples. Q^π consists of a sequence of CPTP maps making oracle calls to $c-U_\pi$, controlled reflections around $|\pi\rangle$; then, purifying the maps of Q^* to unitary operations we obtain a circuit Q^* which employs the same number of oracle calls to $c-U_\pi$. The output of Q^π then has the form:

$$Q^\pi(|\pi\rangle^{\otimes c} |0\rangle) = \sum_{x_1, \dots, x_{c+1}} \sqrt{\pi(x_1) \cdots \pi(x_{c+1})} |x_1, \dots, x_c\rangle |\phi(x_1, \dots, x_{c+1})\rangle, \quad (38)$$

where $|\phi(x_1, \dots, x_{c+1})\rangle$ is a state containing all the residual information. The output state in Eq. (38) also generalizes other tasks, *e.g.*, choosing $|\phi(x_1, \dots, x_{c+1})\rangle = |0\rangle$ corresponds to preparing $c + 1$ copies of $|\pi\rangle$.

Next, we consider a set \mathcal{Z} of quantum states which are coherent encodings of specific probability distributions; on these states we impose a relation \mathcal{R} which is suitable for computing a bound as in Eq. (36). In turn, using Eq. (37), this will provide a lower bound to the number of reflectors required by the quantum circuit \mathcal{Q}^* or, equivalently, by its purification Q^* .

Proposition 1. *We consider the set \mathcal{Z} of states of the form*

$$|\mathbf{u}_S\rangle := \frac{1}{\sqrt{K}} \sum_{x \in S} |x\rangle, \quad (39)$$

where $S \subseteq [N]$ is a subset containing a fixed number K of elements, with $K \leq N$. The state $|\mathbf{u}_S\rangle$ corresponds to the coherent encoding of the probability distribution \mathbf{u}_S . Notice that $f(\mathbf{u}_S) \equiv \sum_{x \in S} \frac{1}{\sqrt{K}} = \sqrt{K}$, hence the value of $f(\mathbf{u}_S)$ can take any value in the interval $[1, \sqrt{N}]$. Moreover, we say that two states $|\mathbf{u}_S\rangle, |\mathbf{u}_{S'}\rangle$ are in relation \mathcal{R} iff

$$\langle \mathbf{u}_S | \mathbf{u}_{S'} \rangle = a \iff |S \cap S'| = aK, \quad (40)$$

where $a \in (0, 1)$ is a constant.

Then, the inequality (36) can be expressed as follows: $\forall S \subseteq [N]$ with $|S| = K$, $\forall |\eta\rangle$ orthogonal to $|\mathbf{u}_S\rangle$

$$\mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[|\langle \eta | \mathbf{u}_{S'} \rangle|^2 \right] \leq \frac{a}{K} + 6(1-a)^2 \frac{K}{N} \equiv \gamma, \quad (41)$$

provided that $1/(1-a) \leq K \leq N/2$.

Proof. First, notice that $|\eta\rangle = \sum_x \eta_x |x\rangle$ is orthogonal to $|\mathbf{u}_S\rangle$, hence $\sum_{x \in S} \eta_x = 0$. Then we expand:

$$\begin{aligned} \mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[|\langle \eta | \mathbf{u}_{S'} \rangle|^2 \right] &= \mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[\left| \sum_{i \in S'} \frac{\eta_i}{\sqrt{K}} \right|^2 \right] \\ &= \frac{1}{K} \mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[\sum_{i \in S'} \sum_{j \in S'} \eta_i^* \eta_j \right] \end{aligned} \quad (42)$$

Next, we split the sum over elements in S' as sum of elements in $S' \cap S$ and elements in $S' \setminus S$:

$$\begin{aligned} (42) &= \frac{1}{K} \mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[\sum_{i \in S' \cap S} \sum_{j \in S' \cap S} \eta_i^* \eta_j + \sum_{i \in S' \setminus S} \sum_{j \in S' \setminus S} \eta_i^* \eta_j + \right. \\ &\quad \left. + \sum_{i \in S' \cap S} \sum_{j \in S' \setminus S} (\eta_i^* \eta_j + \eta_j^* \eta_i) \right] \end{aligned} \quad (43)$$

The first term in the sum in Eq. (43) evaluates to

$$\mathbb{E}_{\substack{I \subseteq S: \\ |I|=aK}} \left[\sum_{i \in I} \sum_{j \in I} \eta_i^* \eta_j \right] = a \sum_{i \in S} |\eta_i|^2 + \frac{aK(aK-1)}{K(K-1)} \sum_{\substack{i, j \in S \\ i \neq j}} \eta_i^* \eta_j; \quad (44)$$

the second term in the sum evaluates to (S^c is the complementary of S)

$$\mathbb{E}_{\substack{I \subseteq S^c: \\ |I|=(1-a)K}} \left[\sum_{i \in I} \sum_{j \in I} \eta_i^* \eta_j \right] = \frac{(1-a)K}{N-K} \sum_{i \in S^c} |\eta_i|^2 + \frac{(1-a)K[(1-a)K-1]}{(N-K)(N-K-1)} \sum_{\substack{i, j \in S^c \\ i \neq j}} \eta_i^* \eta_j; \quad (45)$$

while the third term evaluates to

$$\mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[\sum_{i \in S' \cap S} \sum_{j \in S' \setminus S} (\eta_i^* \eta_j + \eta_j^* \eta_i) \right] = \frac{aK}{K} \frac{(1-a)K}{N-K} \sum_{\substack{i \in S \\ j \in S^c}} (\eta_i^* \eta_j + \eta_j^* \eta_i). \quad (46)$$

Finally, using the equation $\sum_{x \in S} \eta_x = 0$, we get:

$$\begin{aligned}
\mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=aK}} \left[|\langle \eta | \mathbf{u}_{S'} \rangle|^2 \right] &= \frac{a}{K} \sum_{i \in S} |\eta_i|^2 + \frac{1-a}{N-K} \left(\sum_{i \in S^c} |\eta_i|^2 + \frac{(1-a)K-1}{N-K-1} \sum_{\substack{i,j \in S^c \\ i \neq j}} \eta_i^* \eta_j \right) \\
&\leq \frac{a}{K} + \frac{1-a}{N-K} \left(1 + \frac{(1-a)K-1}{N-K-1} (N-K) \right) \\
&\leq \frac{a}{K} + \frac{1-a}{N-K} \left(1 + 2(1-a)K \right) \\
&\leq \frac{a}{K} + \frac{1-a}{N-K} 3(1-a)K \\
&\leq \frac{a}{K} + 6(1-a)^2 \frac{K}{N}
\end{aligned} \tag{47}$$

for appropriate choices of K ; namely, we have used respectively $K \leq N-2$, $K \geq 1/(1-a)$ and $K \leq N/2$ for the last three inequalities in the derivation above. \square

Corollary 1. Consider the family of quantum states \mathcal{Z} and the relation \mathcal{R} given in [Prop. 1](#), setting the constant $a = 1 - 1/c$. Consider then a (purified) quantum circuit Q^* having access to controlled reflections around $|\mathbf{u}_S\rangle$. The input states to Q^* have the form $|\Psi_{in}^\psi\rangle = |\mathbf{u}_S\rangle^{\otimes c}$, $|\Psi_{in}^\phi\rangle = |\mathbf{u}_{S'}\rangle^{\otimes c}$. The output states $|\Psi_{out}^\psi\rangle$ and $|\Psi_{out}^\phi\rangle$ have the same form as the right hand side of [Eq. \(38\)](#) for $\boldsymbol{\pi} = \mathbf{u}_S$ and $\boldsymbol{\pi} = \mathbf{u}_{S'}$, respectively.

Then the circuit Q^* makes $\frac{1}{c^2} \Omega(C(\mathbf{u}_S))$ controlled reflections around $|\mathbf{u}_S\rangle$.

Proof. First, substituting $a = 1 - 1/c$ in [Eq. \(41\)](#) we obtain:

$$\mathbb{E}_{\substack{S': |S'|=K \\ |S \cap S'|=K-K/c}} \left[|\langle \eta | \mathbf{u}_{S'} \rangle|^2 \right] \leq \frac{1}{K} + 6c^2 \frac{K}{N} \equiv \gamma'. \tag{48}$$

Second, notice that for our choice of input and output states:

$$|\langle \Psi_{in}^\phi | \Psi_{in}^\psi \rangle| = a^c \tag{49}$$

$$|\langle \Psi_{out}^\phi | \Psi_{out}^\psi \rangle| \leq a^{c+1}. \tag{50}$$

Then, the application of [Thm. 1](#) for $\alpha = a^c$ and $\beta = a^{c+1}$ directly yields a $\Omega\left(\frac{a^c - a^{c+1}}{\sqrt{\gamma'}}\right)$ lower bound to the number of oracle access that are required by Q^* . Setting $a = 1 - 1/c$ this lower bound becomes $\Omega\left(\frac{1}{c\sqrt{\gamma'}}\right)$ and, finally, this can be further simplified to obtain the required result:

$$\begin{aligned}
\Omega\left(\frac{1}{c\sqrt{\gamma'}}\right) &\geq \Omega\left(\min\left\{\frac{1}{c^2}\sqrt{\frac{N}{K}}, \frac{1}{c}\sqrt{K}\right\}\right) \\
&\geq \frac{1}{c^2} \Omega\left(\min\left\{\frac{\sqrt{N}}{f(\mathbf{u}_S)}, f(\mathbf{u}_S)\right\}\right) \\
&= \frac{1}{c^2} \Omega(C(\mathbf{u}_S)).
\end{aligned} \tag{51}$$

\square

This Corollary shows that in general $\mathcal{O}(C(\boldsymbol{\pi}))$ reflections around $|\boldsymbol{\pi}\rangle$ are needed in order to obtain multiple samples from $\boldsymbol{\pi}$. Using the results of [Corollary 5.3](#) and [5.4](#) of [\[47\]](#) the same $\mathcal{O}(C(\boldsymbol{\pi}))$ lower bound applies for drawing approximate samples from $\boldsymbol{\pi}$, say, within constant approximation error. This is important since the distribution \mathbf{u}_S cannot be, strictly speaking, a stationary distribution of a irreducible MC, since it would need to have support over the entire set of N elements.