Simpler (Classical) and Faster (Quantum) Algorithms for Gibbs Partition Functions

Srinivasan Arunachalam[1], Vojtech Havlicek[1,2], Giacomo Nannicini[1], Kristan Temme[1], and Pawel Wocjan[1]

[1]IBM Quantum, IBM T.J. Watson Research Center
[2]School of Mathematics, University of Bristol

We present classical and quantum algorithms for approximating partition functions of classical Hamiltonians at a given temperature. Our work has two main contributions: first, we modify the classical algorithm of Štefankovič, Vempala and Vigoda (*J. ACM*, 56(3), 2009) to improve its sample complexity; second, we quantize this new algorithm, improving upon the previously fastest quantum algorithm for this problem, due to Harrow and Wei (SODA 2020). The conventional approach to estimating partition functions requires approximating the means of Gibbs distributions at a set of inverse temperatures that form the so-called cooling schedule. The length of the cooling schedule directly affects the complexity of the algorithm. Combining our improved version of the algorithm of Štefankovič, Vempala and Vigoda with the paired-product estimator of Huber (*Ann. Appl. Probab.*, 25(2), 2015), our new quantum algorithm uses a shorter cooling schedule than previously known. This length matches the optimal length conjectured by Štefankovič, Vempala and Vigoda. The quantum algorithm also achieves a quadratic advantage in the number of required quantum samples compared to the number of random samples drawn by the best classical algorithm, and its computational complexity has quadratically better dependence on the spectral gap of the Markov chains used to produce the quantum samples.

Note: a shorter version of this work previously appeared in the proceedings of the 2021 IEEE International Conference on Quantum Computing and Engineering (QCE).

Contents

1 Introduction

Markov Chain Monte Carlo (MCMC) method is a fundamental computational tool in statistics. It is a strategy for sampling from certain high-dimensional probability distributions, which can be used to estimate properties of systems that would be difficult to study otherwise. In the last two decades, MCMC methods found numerous applications in Bayesian inference [10, 8], counting problems [23], volume estimation of convex bodies [7], approximation of the permanent [16], estimating thermodynamic properties of systems [2] or in finance for simulating the performance and volatility of portfolios [9].

A major algorithmic task in the area of MCMC methods is the approximation of *partition functions*. In physics, the partition function describes the statistical properties of a physical system at a fixed inverse temperature, but a wide range of problems can be naturally cast as questions about partition functions. In machine learning for example, the partition function appears in the definition of probabilistic graphical models such as Markov Random Fields (e.g., restricted Boltzman machines [13]), and has been extensively studied before [5, 19, 14, 18]. Given the many of their applications, it is important to design efficient algorithms for their computation, and explore the potential of using faster techniques to compute them.

We study the advantages that quantum computers have over classical computers for the computation of partition functions. It is well known that quantum algorithms provide quadratic speedups over classical algorithms for a variety of tasks related to MCMC, such as amplitude amplification and estimation [3] and spectral gap amplification of Markov Chains [24]. Even though a general quantum speedup for MCMC methods has not been found, fast quantum algorithms for partition functions using simulated annealing have been proposed [28, 21, 12]. Here we study classical and quantum algorithms for computing partition functions; our main contribution is a unification of two *classical* algorithms, leading to a *quantum* algorithm that improves upon the best known (quantum or classical) approaches. Before we describe our main results, we first introduce some basic notation and the problem of computing Gibbs partition functions.

1.1 Estimating partition functions

Let Ω be a finite set and $H : \Omega \to \{0, \ldots, n\}$ be a function called the *Hamiltonian*[1]. The *Gibbs distribution* is defined as:

$$\mu_\beta(x) = \frac{1}{\mathsf{Z}(\beta)} \exp(-\beta H(x)) \quad \text{for every } x \in \Omega,$$

where β is the *inverse temperature* and $\mathsf{Z}(\beta) := \sum_{x \in \Omega} e^{-\beta H(x)}$ is the *partition function*. Given $\beta_{\max}$ and $\varepsilon > 0$, we aim to find $\hat{Z}$ such that $(1 - \varepsilon) \cdot \mathsf{Z}(\beta_{\max}) \leq \hat{Z} \leq (1 + \varepsilon) \cdot \mathsf{Z}(\beta_{\max})$ using the least amount of samples from μ_β for various choices of β. An approach for this task is based on the following telescoping product:

$$\mathsf{Z}(\beta_{\max}) = \mathsf{Z}(0) \cdot \frac{\mathsf{Z}(\beta_1)}{\mathsf{Z}(0)} \cdot \frac{\mathsf{Z}(\beta_2)}{\mathsf{Z}(\beta_1)} \cdots \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\ell-1})} = \mathsf{Z}(0) \cdot \prod_{i=0}^{\ell-1} \frac{\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\beta_i)}, \tag{1}$$

where $0 = \beta_0 < \beta_1 < \cdots < \beta_\ell = \beta_{\max}$ is a sequence of inverse temperatures called a *cooling schedule*. Eq. 1 approximates $\mathsf{Z}(\beta_{\max})$ if each factor $\mathsf{Z}(\beta_{i+1})/\mathsf{Z}(\beta_i)$ is estimated with sufficient precision. We fix $\beta_{\max} = \infty$ in the rest of the introduction for simplicity.

Two approaches for this approximation were introduced in Refs. [23, 14]. The algorithm of Štefankovič et al. [23] computes $\mathsf{Z}(\infty)$ in two steps: first, it produces a cooling schedule consisting of $O(\sqrt{\ln |\Omega|} \ln n \ln \ln |\Omega|)$ inverse temperatures. Defining $X_i = e^{(\beta_i - \beta_{i+1})H(x)}$, we have $\mathbb{E}_{x \sim \mu_{\beta_i}}[X_i] = \mathsf{Z}(\beta_{i+1})/\mathsf{Z}(\beta_i)$. The partition function can be approximated by sampling many $x \sim \mu_{\beta_i}$, empirically estimating the factors at subsequent temperatures and taking their product. This is called the *product estimator*. The resulting algorithm for partition function estimation has sample complexity $10^{10} \cdot \ln |\Omega| \cdot \varepsilon^{-2} \cdot (\ln n + \ln \ln |\Omega|)^5$. We call this algorithm SVV, for Štefankovič, Vempala and Vigoda.

SVV contains many subtle technical points that lead to some of the polylogarithmic factors and a large prefactor in its running time. Huber [14] and Kolmogorov [18] eliminated these hurdles using the so-called Tootsie Pop Algorithm (TPA) [15]. This algorithm uses a Poisson point process to build a cooling schedule with $\sim \ln |\Omega|$ temperatures. The product estimator would be inefficient in this context as it would pick up a quadratic overhead in the schedule length. To address this, Huber introduced the *paired-product estimator*. Given a cooling schedule, define $d_{i,i+1} = (\beta_{i+1} - \beta_i)/2$, and random variables V_i, W_i:

$$V_i = \exp\left(-d_{i,i+1} H(x_i)\right), \quad W_i = \exp\left(d_{i,i+1} H(x_{i+1})\right),$$

---

[1]Hamiltonians are usually defined as real-valued functions. In this paper we make the simplifying assumption that the range of the Hamiltonian is contained in the set of nonnegative integers. This assumption is common in the literature and can often be relaxed, see Kolmogorov [18] for a discussion.

where $x_i \sim \mu_{\beta_i}$. The telescoping product in Eq. (1) can be rewritten as

$$\mathsf{Z}(\infty) = \mathsf{Z}(0) \cdot \prod_{i=0}^{\ell-1} \frac{\mathbb{E}[W_i]}{\mathbb{E}[V_i]} = \mathsf{Z}(0) \cdot \frac{\prod_i \mathbb{E}[W_i]}{\prod_i \mathbb{E}[V_i]}. \qquad (2)$$

While in SVV each ratio $\mathbb{E}[W_i]/\mathbb{E}[V_i]$ is estimated separately, Huber estimates the entire numerator $\prod_i \mathbb{E}[W_i]$ and denominator $\prod_i \mathbb{E}[V_i]$, then takes their ratio. This, remarkably, leads to a better estimator: the overall complexity of TPA (using the improved analysis given by Kolmogorov [18]) is $O(\ln|\Omega| \cdot \varepsilon^{-2} \cdot \ln n)$.

1.2   Results

We give a quantum algorithm that estimates partition functions by combining elements of SVV and TPA. Assuming the ability to query an oracle that provides coherent encodings of the Gibbs distribution at arbitrary inverse temperature (*qsamples*), we show that the quantum algorithm requires fewer qsamples compared to the number of samples drawn by the best classical algorithm.

The above assumption is not restrictive, as qsamples can be generated from a classical Markov chain that has the corresponding Gibbs distribution as its limit distribution; the total running time of the algorithm (classical and quantum) must also account for the time – known as *mixing time* – taken to generate these samples. Quantum algorithms for approximating partition functions [28, 21, 12] do not perform the classical steps of the MCMC algorithm. Instead, they utilize unitary quantum walks [24] that are related to the classical Markov chain by a suitable embedding as the product of two reflection operators. Such quantum walks are constructed so that a natural quantum encoding of the Gibbs distribution is an eigenvector with eigenvalue one of the quantum walk unitary. This leads to a quadratic advantage of the quantum algorithm over classical, in terms of the mixing time.

1.2.1   Classical contribution

We propose a classical algorithm for the approximation of partition functions that simplifies [23], and almost matches the sample complexity of TPA, the best classical algorithm [18, 14]. It uses the paired-product estimator within SVV: this allows for a shorter cooling schedule and a simpler estimator.

This estimator becomes simpler (motivating the "simpler (classsical)" in the paper title), because the approximation of the partition function ratios in the cooling schedule algorithm does not depend on temperatures outside the interval on which these are estimated. This gets rid of the nested search in the SVV cooling schedule algorithm, a siginficant simplification. Moreover, unlike TPA, where these ratios were estimated using a Poisson Point Process, our algorithm is easier to quantize.

Ref. [23] conjectured that an optimal adaptive cooling schedule has length $\Theta(\sqrt{\ln|\Omega| \ln n})$ and proved a lower bound of $\Omega(\sqrt{\ln|\Omega|})$.[2] They give an algorithm that produces a schedule of length $O(\sqrt{\ln|\Omega|} \ln n \ln\ln|\Omega|)$. Our algorithm, albeit quantum, produces a schedule of length $O(\sqrt{\ln|\Omega|} \ln n)$, while achieving the same approximation precision.

---

[2] The lower bound of [23] is for the so-called B-Chebyshev schedules, however their proof applies to the *slowly-varying* schedules proposed in this paper, with straightforward modifications leading to a constant-factor difference.

1.2.2 Quantum contribution

Our main contribution is a quantum algorithm that improves the best known sample complexity and overall running time. In the quantum setting, the samples are the *coherent encodings* of the Gibbs distribution (qsamples):

$$|\mu_\beta\rangle = \sum_x \sqrt{\mu_\beta(x)}|x\rangle.$$

The quantum algorithm is based on SVV, but uses the paired-product estimator of [14]. We give a procedure, based on binary search, to obtain a cooling sequence of length $O(\sqrt{\ln |\Omega| \ln n})$, using the fact that we can efficiently compute inner products of qsamples of the form $|\langle\mu_{\beta_i}|\mu_{\beta_{i+1}}\rangle|$. The length of this cooling sequence matches the optimal length conjectured in [23].

1.2.3 A technical tool

To estimate the telescopic product (2), we give a quantum algorithm for the estimation of the expected value of random variables with bounded relative variance. Computing expectations of random variables is a fundamental task in statistics, applied probability, and related disciplines. Our contribution in this context is a simpler and faster quantum algorithm for computing the expectation of a random variable, given quantum samples. Recall that, in order to compute the Gibbs partition function, one needs to estimate ratios in the telescopic product (2). To do this, we give a simple quantum algorithm to estimate the expected value of random variables with bounded relative variance. In particular, for a distribution D, and given a random variable $V \sim D$ satisfying $\mathbb{E}[V^2]/\mathbb{E}[V]^2 \leq B$, we describe a quantum algorithm that uses $O(B)$ copies of $|\psi_D\rangle = \sum_x \sqrt{D(x)}|x\rangle$, and $\tilde{O}(\sqrt{B}/\varepsilon)$ reflections about $|\psi_D\rangle$, to obtain with high probability an ε-relative approximation of $\mathbb{E}[V]$; additionally, the algorithm restores one copy of $|\psi\rangle$. The proposed algorithm is based on the work of Montanaro [21] and improves upon its scaling from $O(B/\varepsilon)$ to $O(\sqrt{B}/\varepsilon)$. Our algorithm has essentially the same scaling as the algorithm of Hamoudi and Magniez [11], but the analysis of our algorithm is simpler.

1.2.4 Applications

Many computational problems can be encoded into evaluations of the partition function and we briefly sketch out two applications of our algorithm. Our work also applies to other problems, such as counting independent sets, matchings and Bayesian inference as discussed in [12, 21, 23]; we refer the interested reader to these works. As discussed earlier, our classical and quantum improvements over [12, 23] directly imply better algorithms for estimating partition functions of these Hamiltonians.

Ferromagnetic Ising Model The ferromagnetic Ising model on a graph $G = (V, E)$ can be defined as[3]

$$H(x) := \sum_{(i,j)\in E} \mathbf{1}_{[x_i \neq x_j]}, \ x \in \{0,1\}^{|V|}.$$

---

[3]The usual convention is to define the model as $H(x) = -\sum_{(i,j)\in E} x_i x_j$ for $x \in \{\pm 1\}^{|V|}$. Here we use a different convention to ensure that the Hamiltonian is nonnegative, which is an assumption of our classical algorithm and other similar algorithms in the literature (e.g., [23, Lemma 4.2] does not hold for Hamiltonians that change sign).

The associated Gibbs distribution can be sampled using a Markov Chain called Glauber dynamics. Convergence of Glauber dynamics is well studied and several criteria for rapid mixing have been found. Montanaro [21] quotes a result of Mossel and Sly [22] that proves that the chain mixes in time $O(|V| \ln |V|)$ on general graphs with finite degree for a sufficiently low value of β. As the mixing time upper-bounds the inverse spectral gap (see [22, Eq. 6]), it follows that for $|V| = n$, our classical algorithm achieves complexity $\tilde{O}(n^2 \cdot \varepsilon^{-2})$. While there may exist algorithms that perform significantly better, the complexity matches (up to polylogarithmic factors) the best classical upper bound that we are aware of [18]. The quantum algorithm in comparison achieves complexity of $\tilde{O}(n^{3/2} \cdot \varepsilon^{-1})$. Apart from being simpler than the algorithm derived in [12], our approach improves the complexity by polylogarithmic factors.

k-Colorings   The k-state Potts model on a graph $G = (V, E)$ has the Hamiltonian

$$H(x) = \sum_{(i,j) \in E} \mathbf{1}_{[x_i = x_j]}, \quad x \in \Omega := \{1, \ldots, k\}^{|V|},$$

with the corresponding partition function

$$\mathsf{Z}(\beta) = \sum_{x \in \Omega} \exp(-\beta H(x)) = \sum_{x \in \Omega} \prod_{(i,j) \in E} \exp\left(-\beta \mathbf{1}_{[x_i = x_j]}\right).$$

Observe that $\lim_{\beta \to \infty} \mathsf{Z}(\beta) := \mathsf{Z}(\infty) = |\{x : x_i \neq x_j, \ \forall (i,j) \in E\}|$. A k-coloring of a graph is a map $\phi : V \to \{1, 2, \ldots k\}$, such that $\phi(i) \neq \phi(j)$ for all $(i,j) \in E$. The partition function at $\beta = \infty$ therefore gives the number of proper k-colorings of G. Vigoda [26] gave a Markov chain with mixinng time $O(nk \log n)$ whenever $k > 11/6 \cdot d$, where d is the maximum degree of G. Our quantum algorithm has $\tilde{O}(n^{3/2} \cdot \varepsilon^{-1})$ complexity for this problem.

1.2.5 Comparison to prior works

Table 1 summarizes the sample complexity of the algorithms presented in this paper and compares them to the existing results in literature.

2 Classical algorithm

We first introduce our notation and describe the improved version of SVV, that is the basis for the quantum algorithm.

Let $H : \Omega \to \{0, \ldots, n\}$ be a classical Hamiltonian. For fixed $\beta_{\min}, \beta_{\max}$, we let $Q = \mathsf{Z}(\beta_{\max})/\mathsf{Z}(\beta_{\min})$. We often use the shorthand $q = \ln |\Omega|$ and assume that $\ln n \geq 1, \ln q \geq 1, |\Omega| \geq \ln n$ as in [23].[5] We assume that the partition function is non-zero so that the Hamiltonian has at least one state in its ground state (i.e. state with a zero energy). This implies that $|\Omega| \geq \mathsf{Z}(\beta) = \sum_{x \in \Omega} e^{-\beta H(x)} \geq 1$ for all $\beta \in [0, \infty)$. Often in this paper we assume that $\beta_{\max} \leq q$; this is not restrictive because after a cooling schedule reaches

---

[3]We remark that complexity of [1] holds for constant ε.

[5]This assumption is justified because if $|\Omega|$ is small, Q could be estimated using a schedule of length 1

	Algorithm	Sample Complexity
Bezáková et al. [1]	Classical	$O(\ln^2 |\Omega| \cdot (\ln n)^2)$ [4]
SVV [23]	Classical	$O(\ln |\Omega|(\ln\ln |\Omega| + \ln n)^5 \varepsilon^{-2})$
Huber [14]	Classical	$O(\ln |\Omega| \ln n \cdot (\ln\ln |\Omega| + \ln\ln n + \varepsilon^{-2}))$
Kolmogorov [18]	Classical	$O(\ln |\Omega| \ln n \cdot \varepsilon^{-2})$
Montanaro [21]	Hybrid	$O(\ln |\Omega| \cdot (\ln\ln |\Omega| + \ln n)^{5/2} \cdot \varepsilon^{-1})$
Harrow and Wei [12]	Quantum	$O(\ln |\Omega| \cdot (\ln\ln |\Omega| + \ln n)^{5/2} \cdot \varepsilon^{-1})$
This work	Classical	$O(\ln |\Omega| \ln^2 n \cdot \varepsilon^{-2})$
This work	Quantum	$O(\ln |\Omega| \ln n \cdot \varepsilon^{-1})$

Table 1: Comparison of the sample complexity of several algorithms to estimate the partition function. Montanaro's algorithm [21] uses classical SVV to generate a cooling schedule, and a quantum routine to estimate the ratio given a cooling schedule; here we report only the sample complexity of the quantum routine, for the total complexity one should add the complexity of classical SVV schedule generation. The stated complexity of our classical algorithm holds for $\varepsilon < (\ln |\Omega| + \ln n)^{-1}$; the stated complexity of our quantum algorithm holds for $\varepsilon < \left(\frac{\ln |\Omega| + \ln\ln n}{\ln |\Omega| + \ln n}\right)^2$. For a more precise statement of the complexities we refer the reader to Theorems 2.13 and 3.6.

the inverse temperature $q = \ln |\Omega|$, any larger inverse temperature can be reached in a single step.[6]

2.1 Product Estimator

We aim to estimate $Q = \mathsf{Z}(\beta_{\max})/\mathsf{Z}(\beta_{\min})$ given access to samples from the Gibbs distribution at arbitrary inverse temperatures. A naive estimator of Q can be constructed as follows:

- sample $x \sim \mu_{\beta_{\min}}$ from the Gibbs distribution at $\beta_{\min}$.

- compute $X = \exp[(\beta_{\min} - \beta_{\max})H(x)]$.

- the expectation of X is $\mathbb{E}[X] = \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\min})}$.

Define $\mathbb{S}[X] = \mathbb{E}[X^2]/\mathbb{E}[X]^2$. With a slight overload of terminology, will call $\mathbb{S}[X]$ the relative variance.[7] It is equal to:

$$\mathbb{S}[X] = \frac{\mathsf{Z}(2\beta_{\max} - \beta_{\min})\mathsf{Z}(\beta_{\min})}{\mathsf{Z}(\beta_{\max})^2}. \tag{3}$$

This is usually prohibitively large. To address this, Valleau and Card introduced the *product estimator* in Ref. [25] as a low-relative variance estimator of Q. It uses a *cooling*

---

[6]Indeed, we have

$$\mathsf{Z}(\ln |\Omega|) = \sum_{\substack{x \in \Omega: \\ H(x)=0}} e^0 + \sum_{\substack{x \in \Omega: \\ H(x)>0}} e^{-H(x)\cdot \ln |\Omega|} \leq \mathsf{Z}(\infty) + 1.$$

[7]It is related to the usual definition of relative variance by a constant.

schedule of inverse temperatures $\beta_{\min} = \beta_0 < \beta_1 < \cdots < \beta_\ell = \beta_{\max}$. For every step $0 \le i < \ell$:

- Sample $x \sim \mu_{\beta_i}$.

- Compute $X_i = \exp[(\beta_i - \beta_{i+1})H(x)]$.

- The expectation of X is $\mathbb{E}[X_i] = \mathsf{Z}(\beta_{i+1})/\mathsf{Z}(\beta_i)$.

Let X be the product random variable defined as $X = \prod_{i=1}^{\ell-1} X_i$. The random variables X_i, X_j are independent for $i \ne j$ and it follows that:

$$\mathbb{E}[X] = \prod_{i=0}^{\ell-1} \mathbb{E}[X_i] = \prod_{i=0}^{\ell-1} \frac{\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\beta_i)} = \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\min})}.$$

The relative variance of X_i for each $i \in \{0, \ldots, \ell - 1\}$ is:

$$\mathbb{S}[X_i] = \frac{\mathbb{E}[X_i^2]}{\mathbb{E}[X_i]^2} = \frac{\mathsf{Z}(2\beta_{i+1} - \beta_i)\mathsf{Z}(\beta_i)}{\mathsf{Z}(\beta_{i+1})^2}, \tag{4}$$

The product estimator reduces the relative variance $\mathbb{S}[X]$ in (3). However, the dependence of $\mathbb{S}[X_i]$ on the temperature $2\beta_{i+1} - \beta_i$, which is outside the interval $[\beta_i, \beta_{i+1}]$, leads to several complications in the analysis of SVV [23].

2.2 Paired-product Estimator

Our algorithm replaces the *product estimator* with the *paired-product estimator*. This estimator was introduced by Huber in [14]. For $i \in \{0, \ldots, \ell - 1\}$, define the midpoint temperature $\bar{\beta}_{i,i+1} = \frac{\beta_i + \beta_{i+1}}{2}$ and the semi-distance $d_{i,i+1} = \frac{\beta_{i+1} - \beta_i}{2}$. Then:

- For $i \in \{0, \ldots, \ell\}$, sample $x_i \sim \mu_{\beta_i}$.

- For $i \in \{0, \ldots, \ell - 1\}$, compute

$$V_i = \exp\big(-d_{i,i+1} H(x_i)\big), \quad W_i = \exp\big(d_{i,i+1} H(x_{i+1})\big); \tag{5}$$

observe that x_i is used for V_i and x_{i+1} for W_i.

- The expectations of V_i and W_i are:

$$\mathbb{E}[V_i] = \frac{\mathsf{Z}(\bar{\beta}_{i,i+1})}{\mathsf{Z}(\beta_i)}, \qquad\qquad \mathbb{E}[W_i] = \frac{\mathsf{Z}(\bar{\beta}_{i,i+1})}{\mathsf{Z}(\beta_{i+1})}; \tag{6}$$

observe that $\frac{\mathbb{E}[V_i]}{\mathbb{E}[W_i]} = \frac{\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\beta_i)}$, i.e., the pair V_i and W_i replaces X_i.

Define $V_i = \prod_{i=0}^{\ell-1} V_i$ and $W = \prod_{i=0}^{\ell-1} W_i$. Since the V_i are independent and so are the W_i, the ratio of $\mathbb{E}[V]$ and $\mathbb{E}[W]$ can be expressed as:

$$\frac{\mathbb{E}[V]}{\mathbb{E}[W]} = \prod_{i=0}^{\ell-1} \frac{\mathbb{E}[V_i]}{\mathbb{E}[W_i]} = \prod_{i=0}^{\ell-1} \frac{\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\beta_i)} = \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\min})}. \tag{7}$$

The advantage of using the pairs V_i and W_i becomes evident from their relative variances:

$$\mathbb{S}[V_i] = \frac{\mathbb{E}[V_i^2]}{\mathbb{E}[V_i]^2} = \frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\bar{\beta}_{i,i+1})^2} = \frac{\mathbb{E}[W_i^2]}{\mathbb{E}[W_i]^2} = \mathbb{S}[W_i], \tag{8}$$

Both $\mathbb{S}[V_i]$ and $\mathbb{S}[W_i]$ depend on the midpoint $\bar{\beta}_{i,i+1} \in [\beta_i, \beta_{i+1}]$ and it holds that $\mathbb{S}[W_i] = \mathbb{S}[V_i] < \mathbb{S}[X_i]$.[8]

---

[8]This follows from the fact that the relative variance is a monotone increasing function in β. A "quantum" proof of this fact is given in Fact 3.1.

2.3 Dyer and Frieze's bound on product estimation

A key technical ingredient in our classical algorithm is a result od Dyer and Frieze [6] that bounds the number of samples to estimate the expectation of a product random variable with relative error. We use this to estimate the expectation of the product of the random variables V_i and W_i.

Theorem 2.1 (Dyer and Frieze [6]). *Let $B > 0$, $\eta \in (0, 1)$. Assume that the independent random variables $X_1, \ldots, X_\ell$ satisfy $\mathbb{S}[X_i] \leq B$ for all $i \in [\ell]$. By taking $m = 2B\ell/(\eta\varepsilon^2)$ samples from X_i for every $i \in [\ell]$, we can obtain $\widehat{X}$ that satisfies*

$$\Pr\left[(1 - \varepsilon) \cdot \prod_i \mathbb{E}[X_i] \leq \widehat{X} \leq (1 + \varepsilon) \cdot \prod_i \mathbb{E}[X_i]\right] \geq 1 - \eta.$$

2.4 Cooling schedule

We give a short cooling schedule such that the resulting relative variances of the random variables V_i and W_i are bounded from above by a constant. This allows us to apply Theorem 2.1. To do this, we use a result due to Štefankovič et al. [23, Lemma 4.3], establishing the existence of a short cooling schedule that satisfies:

$$\frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\bar{\beta}_{i,i+1})^2} \leq e^2, \tag{9}$$

for all $i \in \{0, \ldots, \ell - 1\}$. This, together with Eq. (8), implies that the relative variances $\mathbb{S}[V_i]$ and $\mathbb{S}[W_i]$ of all paired-product estimators V_i and W_i are at most e^2. We rewrite Eq. (9) as

$$f\left(\frac{\beta_i + \beta_{i+1}}{2}\right) \geq \frac{f(\beta_i) + f(\beta_{i+1})}{2} - 1, \tag{10}$$

where $f(\beta) = \ln\big(\mathsf{Z}(\beta)\big)$. It is important that $f(\beta)$ is a strictly decreasing convex function, which is the case if the Hamiltonian is non-negative.

Theorem 2.2 (Perfectly-balanced schedule length [Analogous to [23]]). *] There exists a sequence $\beta_0 < \cdots < \beta_\ell$ with $\beta_0 = \beta_{\min}$ and $\beta_\ell = \beta_{\max}$ satisfying the condition in (10) with equality and having length ℓ bounded from above by:*

$$\ell \leq \sqrt{(f(\beta_{\min}) - f(\beta_{\max})) \cdot \frac{1}{2} \ln\left(\frac{f'(\beta_{\min})}{f'(\beta_{\max})}\right)}. \tag{11}$$

The proof of Theorem 2.2 can be found in Appendix A.

Corollary 2.3 (Upper bound on perfectly-balanced schedule length). *The length ℓ of a perfectly-balanced schedule, as in Theorem 2.2, satisfies $\ell \leq \sqrt{q \cdot \ln n}$.*

Proof. By Theorem 2.2, the length of the schedule satisfies Eq. (11). We can further upper bound ℓ in Eq. (11) as follows: observe that $f(\cdot) = \ln(\mathsf{Z}(\cdot))$, so we have

$$f(\beta_{\min}) - f(\beta_{\max}) = \ln(\mathsf{Z}(\beta_{\min})/\mathsf{Z}(\beta_{\max})) \leq \ln|\Omega| = q.$$

The above inequality holds because $|\Omega| \geq Z(\beta) \geq 1$ for all $\beta \geq 0$. Moreover,

$$f'(\beta) = \Big(\ln\big(\mathsf{Z}(\beta)\big)\Big)' = \frac{\mathsf{Z}'(\beta)}{\mathsf{Z}(\beta)} = -\sum_{x \in \Omega} \mu_\beta(x)H(x).$$

We have:

$$\frac{f'(\beta_{\min})}{f'(\beta_{\max})} = \frac{-f'(\beta_{\min})}{-f'(\beta_{\max})} \leq \frac{e-1}{e} \cdot n,$$

where we used the fact that $-f'(\beta_{\min}) \leq n$ (which follows from the equation above and the assumption that $H(x) \leq n$) and $-f'(\beta_{\max}) \geq -f'(\ln|\Omega|) \geq \frac{e}{e-1}$ (see Ref. [23, Eq. 35]). $\square$

Theorem 2.2 establishes existence of a perfectly-balanced cooling schedule.

Determining temperatures satisfying Eq. (10) as equality are challenging – we instead work with *well-balanced* cooling schedules, which satisfy:

$$c_1 \leq \frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\bar{\beta}_{i,i+1})^2} \leq c_2, \ \forall i \in \{0, \ldots, \ell-1\}, \tag{12}$$

for suitably chosen constants c_1, c_2. The upper bound c_2 ensures that the relative variances $\mathbb{S}[V_i]$ and $\mathbb{S}[W_i]$ are not too large, while the lower bound c_1 ensures that the temperatures increase rapidly enough, so that the schedule is short. We will also call a cooling schedule that satisfies:

$$\frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\bar{\beta}_{i,i+1})^2} \leq c_2, \ \forall i \in \{0, \ldots, \ell-1\}, \tag{13}$$

a c_2-*slowly varying* cooling schedule.

We describe an iterative algorithm that, given a cooling schedule ending at β_i, finds β_{i+1} so that the condition in Eq. (12) is satisfied. To achieve this, we estimate the ratios $\frac{\mathsf{Z}(\beta_i)}{\mathsf{Z}(\bar{\beta}_{i,i+1})}$ and $\frac{\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}(\bar{\beta}_{i,i+1})}$ and use their product as the estimate for the relative variances $\mathbb{S}[V_i]$ and $\mathbb{S}[W_i]$. Obtaining highly accurate estimates of these ratios is difficult: indeed, if one could compute $\frac{\mathsf{Z}(\beta_i)}{\mathsf{Z}(\beta_{i+1})}$ with high accuracy, then one could follow a cooling schedule, estimate the products $\frac{\mathsf{Z}(\beta_i)}{\mathsf{Z}(\beta_{i+1})}$ for all $i \in \{0, \ldots, \ell\}$, multiply them and solve the original problem of estimating $\mathsf{Z}(\beta_\ell)$. Fortunately, to decide if the well-balanced condition holds, it suffices to have "crude" estimates obtained by taking less samples. To prove this, we use several results from Ref. [23]. The following lemma bounds the number of large temperature increases in any cooling schedule by comparing it to a perfectly-balanced cooling schedule.

Corollary 2.4 (Analogous to [23, Corollary 4.4]). *Let $\beta_{\min} = \gamma_0 < \cdots < \gamma_m = \beta_{\max}$ be an arbitrary cooling schedule such that:*

$$\frac{\mathsf{Z}(\gamma_j)\mathsf{Z}(\gamma_{j+1})}{\mathsf{Z}\left(\frac{\gamma_j + \gamma_{j+1}}{2}\right)^2} \geq e^2, \tag{14}$$

for every $j = 0, \ldots, m-1$. Then, $m \leq \ell$, where ℓ is the length of a perfectly-balanced cooling schedule $\beta_{\min} = \beta_0 < \cdots < \beta_\ell = \beta_{\max}$ as in Theorem 2.2.

Proof. Theorem 2.2 upper-bounds the length of a perfectly-balanced cooling schedule, which has relative variances equal to e^2. We show that a cooling schedule $\gamma_0 < \cdots < \gamma_m$ with relative variances at least e^2 cannot be longer than the upper bound derived in Theorem 2.2. This requires showing that increasing the relative variance can only make the schedule shorter. Define:

$$g(x,y) = \ln \frac{\mathsf{Z}(x)\mathsf{Z}(y)}{\mathsf{Z}\left(\frac{x+y}{2}\right)^2} = f(x) + f(y) - 2f\left(\frac{x+y}{2}\right).$$

$g(x, y)$ is decreasing in x for $x < y$ because

$$\frac{\partial g(x, y)}{\partial x} = f'(x) - f'\left(\frac{x + y}{2}\right) < 0$$

due to strict convexity of f and $x < (x + y)/2$. Similarly one can show that $g(x, y)$ is increasing in y for $x < y$. We now show by induction that $\beta_j \leq \gamma_j$ for every j. This is trivial for $j = 0$. By assumption $g(\beta_j, \beta_{j+1}) = 2$ because the β schedule is perfectly-balanced, and $g(\gamma_j, \gamma_{j+1}) \geq 2$ from Eq. (14). Furthermore, $\beta_j \leq \gamma_j$ by the induction hypothesis. Assume to the contrary that $\beta_{j+1} > \gamma_{j+1}$. In this case, the monotonicity of g (in x,y individually) implies that $g(\beta_j, \beta_{j+1}) > g(\gamma_j, \gamma_{j+1})$. But this is a contradiction since $g(\beta_j, \beta_{j+1}) = 2$ and $g(\gamma_j, \gamma_{j+1}) > 2$. This concludes the induction step $\beta_{j+1} \leq \gamma_{j+1}$. $\square$

Theorem 2.2, Corollary 2.3 and Corollary 2.4 upper-bounds the length of a set of well-balanced cooling schedules with $c_1 = e^2$ by $\sqrt{q \ln n}$.

Definition 2.5. *Let $I = [b, c] \subseteq \{0, \dots, n\}$ and $h \in (0, 1)$. We say that I is h-heavy for $\beta > 0$ if*

$$\Pr_{X \sim \mu_\beta}[H(X) \in I] \geq h.$$

The quantity $\Pr_{X \sim \mu_\beta}[H(X) \in I]$ is also called the weight of I at inverse temperature β.

Lemma 2.6 ([23, Lemma 5.3]). *Let $I = [b, c] \subseteq \{0, \dots, n\}$, $h \in (0, 1)$. If I is h-heavy for two temperatures $\beta_1 < \beta_2$, then I is also h-heavy for all temperatures $(1 - \tau)\beta_1 + \tau\beta_2$, where $\tau \in [0, 1]$, i.e., the set of temperatures for which I is h-heavy is a (possibly empty) subinterval of $[\beta_{\min}, \beta_{\max}]$.*

Lemma 2.7 shows that if we find some interval $I = [b, c]$ that is both sufficiently narrow and is h-heavy for two nearby temperatures β_1 and β_2, we can use that interval to estimate the ratio $\mathsf{Z}(\beta_2)/\mathsf{Z}(\beta_1)$:

Lemma 2.7 ([23, Lemma 5.8]). *Let $I = [b, c] \subseteq \{0, \dots, n\}$, $\delta \in (0, 1]$ and $h \in (0, 1)$. Suppose that I is h-heavy for $\beta_1, \beta_2 > 0$ satisfying:*

$$|\beta_1 - \beta_2|(c - b) \leq 1. \tag{15}$$

For $k \in \{1, 2\}$, let $X_k \sim \mu_{\beta_k}$ and let Y_k be the indicator function for the event $[H(X_k) \in I]$. Let $s = \lceil (8/h) \cdot \ln(1/\delta) \rceil$ and U_k be the average of s independent samples from Y_k. Let

$$E\text{ST}(I, \beta_2, \beta_1) := \frac{U_1}{U_2} \exp(b(\beta_1 - \beta_2)). \tag{16}$$

Then, with probability at least $1 - 4\delta$, we have:

$$\frac{1}{4e} \frac{\mathsf{Z}(\beta_2)}{\mathsf{Z}(\beta_1)} \leq E\text{ST}(I, \beta_2, \beta_1) \leq 4e \frac{\mathsf{Z}(\beta_2)}{\mathsf{Z}(\beta_1)}.$$

Lemma 2.7 provides a way to compute $\mathsf{Z}(\beta_i)/\mathsf{Z}(\bar{\beta}_{i,i+1})$ and $\mathsf{Z}(\beta_{i+1})/\mathsf{Z}(\bar{\beta}_{i,i+1})$: find an interval which is h-heavy for both β_i and β_{i+1}, and use Eq. (16) to estimate the two ratios up to a relative error of $4e$. This is possible thanks to Lemma 2.6, ensuring that the midpoint $\bar{\beta}_{i,i+1}$ between β_i and β_{i+1} satisfies the assumptions of Lemma 2.7. To generate a short cooling schedule, given some β_i we need to determine some β_{i+1} such that: (i) we have an interval I which is h-heavy for both β_i and β_{i+1}; (ii) the well-balanced condition Eq. (12) holds. The following lemma is used to determine if a given interval is h-heavy at a given temperature.

Lemma 2.8 ([23, Lemma 5.5]). *Let $I = [b, c] \subseteq \{0, \dots, n\}$, $\beta > 0$ and $\delta, h \in (0, 1]$. Let $X \sim \mu_\beta$ and let Y be the indicator function for the event $[H(X) \in I]$. Let $s = \lceil (8/h) \ln(1/\delta) \rceil$ and U be the average of s independent samples from Y. Define*

$$\textsc{IsHeavy}(I, \beta) = \begin{cases} \text{true} & \text{if } U \geq 2h, \\ \text{false} & \text{if } U < 2h. \end{cases}$$

If I is not h-heavy at β we have $\Pr[\textsc{IsHeavy}(I, \beta) = \text{true}] \leq \delta$ (where the probability is over the randomness in sampling from μ_β), and if I is $4h$-heavy at β we have $\Pr[\textsc{IsHeavy}(I, \beta) = \text{false}] \leq \delta$.

Proof. Assume that I is $4h$-heavy. The expected number of samples that fall into I is $\geq 4hs$. By Chernoff bound I will receive $\leq 2hs$ samples with probability at most δ, i.e., $\Pr[U \leq 2h] \leq e^{-sh/8} \leq \delta$. Now assume that I is not h-heavy. The expected number of samples that fall into I is $\leq hs$. Using Chernoff bound, I will receive more than $2hs$ samples with probability at most δ, i.e., $\Pr[U \geq 2h] \leq e^{-sh/8} \leq \delta$. $\square$

Corollary 2.9. *Let $\beta > 0$, P be an arbitrary partition of $\{0, \dots, n\}$, and $h = \frac{1}{k|P|}$ for $k \geq 4$. Let F be a (possibly empty) subset of P such that no interval in F is $\frac{1}{|P|}$-heavy at β. Suppose we obtain $s = \lceil (8/h) \ln(1/\delta) \rceil$ samples from μ_β and select an interval in $P \setminus F$ that received the largest number of samples. With probability at least $1 - \delta|P|$, this interval is h-heavy at β.*

Proof. Observe that there exists a $kh = (1/|P|)$-heavy interval in $P \setminus F$: if this were not the case, then every interval in the partition P would have weight $< 1/|P|$ (since intervals in F also have weight $< 1/|P|$) and the overall weight would be < 1 which is a contradiction. Furthermore, $k \geq 4$, therefore there is at least one $4h$-heavy interval in $P \setminus F$. We now apply the union bound for the following event: every $4h$-heavy interval receives at least $2hs$ samples, and every interval that is not h-heavy receives less than $2hs$ samples. This event occurs with probability at least $1 - |P| \cdot \delta$, and conditioned on this event, the interval returned by the algorithm (i.e., the one that received the largest number of samples) is h-heavy. $\square$

We use the following special partition P in our schedule generation algorithm, that provides suitable candidate intervals I from which one can select an h-heavy interval:

- Set $P \leftarrow \emptyset$, $b \leftarrow 0$.

- Repeat until $b \geq n$: add the interval $\{b, \dots, b + \lfloor b/\sqrt{q} \rfloor\}$ to P; set $b \leftarrow b + \lfloor b/\sqrt{q} \rfloor + 1$.

- Return the set of intervals P.

Lemma 2.10 ([23, Lemma 5.1]). *P has size $|P| \leq 4\sqrt{q} \ln n$.*

The last ingredient is the binary search described in Algorithm 1.

Theorem 2.11. *There exists a procedure that, with probability at least $1 - \delta$, computes a $2 \cdot 10^5$-slowly-varying schedule with length at most $11\sqrt{q} \ln n$, and uses at most $5 \cdot 10^4 \cdot q \ln^2 n \cdot (\ln q + \ln n)^2 \ln(1/\delta)$ samples from the Gibbs distribution.*

Proof. See Algorithm 2.

---

[9]A *monotone predicate* $\mathcal{P}$ is a Boolean function defined on a totally ordered set with the following property: if $\mathcal{P}(x) = \text{true}$, then $\mathcal{P}(y) = \text{true}$ for all $y \leq x$ in the domain, i.e., $\mathcal{P}$ is monotone decreasing.

---

input : Monotone predicate $\mathcal{P}$,[9] interval $[a, b]$ such that $\mathcal{P}(a) = \mathsf{true}$, precision α.

output: b if $\mathcal{P}(b) = \mathsf{true}$, otherwise an x such that $\mathcal{P}(x) = \mathsf{true}$ and $\mathcal{P}(x + \alpha) = \mathsf{false}$

if $\mathcal{P}(b)$ **then**
 | **return** b
end
Set $\lambda \leftarrow a, \rho \leftarrow b$;
while $\rho - \lambda > \alpha$ **do**
 | **if** $\mathcal{P}(\frac{\lambda+\rho}{2})$ **then**
 | | $\lambda \leftarrow \frac{\lambda+\rho}{2}$
 | **else**
 | | $\rho \leftarrow \frac{\lambda+\rho}{2}$
 | **end**
end
return λ

Algorithm 1: Binary search subroutine, denoted $\textsc{BinarySearch}(\mathcal{P}, [a, b], \alpha)$.

---

input : Initial temperature β_0, partition P, largest temperature $\beta_{\max}$, probability δ.

output: Set of temperatures $\beta_0, \ldots, \beta_k = \beta_{\max}$.

Set $h \leftarrow \frac{1}{8|P|}$ throughout the algorithm and subroutine calls;
Set $k \leftarrow 0, F \leftarrow \emptyset$;
while $\beta_k < \beta_{\max}$ **do**
 | Apply the procedure of Corollary 2.9 with $\beta = \beta_k$, $h = \frac{1}{8|P|}$, and F as the set
 | of forbidden intervals; let its return value be the interval $I = \{b, \ldots, c\}$;
 | Set $L \leftarrow \min\{\beta_k + 1/(c - b), q\}$, with the convention $1/0 = \infty$;
 | Compute $L^* \leftarrow \textsc{BinarySearch}(\textsc{IsHeavy}(I, \cdot) = \mathsf{true}, [\beta_k, L], 1/2n)$;
 | Compute
 | $\beta^* \leftarrow \textsc{BinarySearch}(\textsc{Est}(I, \frac{\cdot+\beta_k}{2}, \beta_k)\textsc{Est}(I, \cdot, \frac{\cdot+\beta_k}{2}) \leq 1500, [\beta_k, L^*], 1/2n)$;
 | If $\beta^* = L^* < L$, set $F \leftarrow F \cup \{I\}$;
 | Set $\beta_{k+1} \leftarrow \beta^*, k \leftarrow k + 1$;
end
return $\beta_1, \ldots, \beta_k, \beta_{\max}$

Algorithm 2: Classical schedule generation procedure. The partition P is generated according to Lemma 2.10.

We prove correctness of this algorithm and analyze its complexity. Fix $h := 1/(8|P|)$. By Lemma 2.10, we have $|P| \leq 4\sqrt{q} \ln n$. Using Corollary 2.9 in the first line of the "while" loop, we can always find an interval I that is h-heavy for β_k and does not belong to the set of forbidden intervals F, provided that none of the forbidden intervals in F is $8h = (1/|P|)$-heavy at β_k. This will be shown below. The interval I depends on the current inverse temperature β_k. For now, we neglect the failure probability of the algorithm for determining I in Corollary 2.9 as well as the failure probabilities of the two binary searches in the third and fourth lines of the "while" loop and account for these at the end of the proof by a union bound argument. We will also show that the two predicates used inside the binary searches in the third and fourth lines are monotone and are satisfied at the left ends of their respective search intervals, as required for Algorithm 1.

To bound the total number of temperatures in the final schedule, we analyze three mutually exclusive cases that can arise in each iteration: (1) $\beta^* = L^* = L$, (2) $\beta^* = L^* < L$, (3) $\beta^* < L^*$. Observe that in cases (2) and (3) the binary searches return inverse temperatures that are not equal to the right endpoint of their respective search intervals. Therefore, if these inverse temperatures are increased by the precision $\alpha = 1/(2n)$, the predicates are no longer satisfied. This allows us to bound the number of steps in cases (2) and (3).

Case 1 In case (1), we move by setting $L^* = L$ and $\beta^* = L^*$. These moves are called "long moves" in [23] because the inverse temperature is increased by the maximally possible value according to the requirement (15) in Lemma 2.7. To bound the number of long moves, we use the following lemma.

Lemma 2.12. *The number of "long moves" in Algorithm 2, where we set $\beta^* = L^* = L$, is at most $6\sqrt{q} \ln n$.*

We defer the proof of this lemma to Appendix B as it is rather long and technical.

Case 2 We move by setting $L^* < L$ and $\beta^* = L^*$ and show that when such a move takes place, the interval I is not $8h = (1/|P|)$-heavy for any $\beta \geq L^*$. This ensures that we can correctly apply Corollary 2.9 in all subsequent iterations, which are executed with a set F of forbidden intervals that now includes I and only consider inverse temperatures $\beta \geq L^*$. The value L^* is obtained by applying $\text{BINARYSEARCH}(\text{ISHEAVY}(I, \cdot) = \text{true}, [\beta_k, L], 1/2n)$. Clearly, it holds that $\text{ISHEAVY}(I, \beta_k) = \text{true}$. Thus, the predicate used in the first binary search of Algorithm 2 is satisfied at the left endpoint β_k of the search interval. This predicate is monotone, which is implied by Lemma 2.6.

The fact $L^* < L$ means that the value returned by binary search is not equal to the right end of the search interval. Therefore, by the properties of the binary search procedure we have $\text{ISHEAVY}(I, L^*) = \text{true}$ and $\text{ISHEAVY}(I, \rho) = \text{false}$ for some ρ with $L^* < \rho \leq L^* + 1/(2n)$. By Lemma 2.8, I is $4h$-heavy for L^* but is not $4h$-heavy anymore for ρ.

We first prove that I is not $8h$-heavy at any inverse temperature in the interval $[L^*, \rho]$. To this end, we bound the weight of I at β', for any $\beta' \in [L^*, \rho]$, as follows:

$$\frac{1}{\mathsf{Z}(\beta')} \sum_{x:H(x)\in I} e^{-\beta' H(x)} \leq \frac{1}{\mathsf{Z}(\rho)} \sum_{x:H(x)\in I} e^{-\rho H(x)} e^{1/2} \leq 4h e^{1/2} < 8h.$$

---

[9]We use $(\cdot)$ as the argument for inline function definitions; e.g., the expression $\text{BINARYSEARCH}(\text{ISHEAVY}(I, \cdot) = \text{true}, [\beta_k, L], 1/2n)$ means that the value of the predicate for binary search, $\mathcal{P}$, at a point x is the expression $\text{ISHEAVY}(I, x) = \text{true}$, i.e. $\mathcal{P}(x) = 1 \Leftrightarrow \text{ISHEAVY}(I, x) = \text{true}$.

In the chain of inequalities above, we used the following facts: for the first inequality, $\mathsf{Z}(\beta') \geq \mathsf{Z}(\rho)$ because $\mathsf{Z}(\cdot)$ is non-increasing, and $\beta' \geq \rho - 1/(2n)$; for the second inequality we used $H(x) \leq n$; for the third inequality, we used the fact that the weight of I at ρ is at most $4h$ (recall Definition 2.5).

It remains to prove that I is not $8h$-heavy for any $\beta' > \rho$. Assume to the contrary that such a β' exists. This would automatically imply that I is $4h$-heavy at β'. Combined with Lemma 2.6, this would imply that I is $4h$-heavy for all temperatures inside the interval $[L^*, \beta']$. But this contradicts that I is not $4h$-heavy for ρ.

We can therefore apply Corollary 2.9 in subsequent iterations where I is forbidden. Because I is added to the set F of forbidden intervals, an iteration in which we set $\beta^* = L^*$ with $L^* < L$ can only take place at most once per interval I. Hence, the number of these iterations is at most $|P| \leq 4\sqrt{q}\ln n$.

Case 3 In case (3), we move by setting $\beta^* < L^*$. Notice that

$$\text{Est}(I, \frac{\beta_k + \beta_k}{2}, \beta_k)\text{Est}(I, \beta_k, \frac{\beta_k + \beta_k}{2}) \leq 16e^2 \leq 1500$$

by Lemma 2.7. Thus, the predicate used in the second binary search of Algorithm 2 is satisfied at the left endpoint β_k of the search interval. Moreover, this predicate is monotone, which follows from the discussion in the proof of Corollary 2.4. Hence, binary search with precision $1/2n$ determines a value λ such that there exists ρ satisfying:

$$\text{Est}\Big(I, \frac{\lambda + \beta_k}{2}, \beta_k\Big)\text{Est}\Big(I, \lambda, \frac{\lambda + \beta_k}{2}\Big) \leq 1500,$$
$$\text{Est}\Big(I, \frac{\rho + \beta_k}{2}, \beta_k\Big)\text{Est}\Big(I, \rho, \frac{\rho + \beta_k}{2}\Big) > 1500,$$
$$\rho - \lambda \leq 1/2n,$$
$$\lambda \geq \beta_k.$$

Using Lemma 2.7, we have:

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\lambda)}{\mathsf{Z}(\frac{\beta_k + \lambda}{2})^2} \leq 16e^2\text{Est}\Big(I, \frac{\lambda + \beta_k}{2}, \beta_k\Big)\text{Est}\Big(I, \lambda, \frac{\lambda + \beta_k}{2}\Big) \leq 2 \cdot 10^5. \tag{17}$$

Furthermore,

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\rho)}{\mathsf{Z}(\frac{\beta_k + \rho}{2})^2} \geq \frac{1}{16e^2}\text{Est}\Big(I, \frac{\rho + \beta_k}{2}, \beta_k\Big)\text{Est}\Big(I, \rho, \frac{\rho + \beta_k}{2}\Big) \geq \frac{1500}{16e^2}, \tag{18}$$

because the predicate is false at ρ. It remains to lower bound the LHS of Eq. (17). In order to bound the ratio, first observe that $\mathsf{Z}(\lambda) \geq \mathsf{Z}(\rho)$. Then, observe that for every $\varepsilon > 0$, the following inequality holds:

$$\mathsf{Z}(\beta) \geq \mathsf{Z}(\beta + \varepsilon) = \sum_x e^{-(\beta+\varepsilon)H(x)} \geq e^{-n\varepsilon} \cdot \mathsf{Z}(\beta).$$

It then follows that $\mathsf{Z}(\frac{\beta_k + \lambda}{2}) \leq \mathsf{Z}(\frac{\beta_k + \rho}{2})e^{1/4}$, because $\frac{\beta_k + \rho}{2} - \frac{\beta_k + \lambda}{2} \leq \frac{1}{4n}$. The LHS of Eq. (17) can be lower bounded by

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\lambda)}{\mathsf{Z}(\frac{\beta_k + \lambda}{2})^2} \geq \frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\rho)}{e^{1/2}\mathsf{Z}(\frac{\beta_k + \rho}{2})^2} \geq \frac{1500}{16e^{2.5}} \geq 10 \geq e^2, \tag{19}$$

where the first inequality used $\mathsf{Z}(\frac{\beta_k + \lambda}{2}) \leq \mathsf{Z}(\frac{\beta_k + \rho}{2})e^{1/4}$ and the second inequality used Eq. (18).

We now conclude the proof. We showed above that, in the iterations for which $\beta^* < L^*$, we select inverse temperatures that satisfy condition (19). The number of such inverse temperatures is at most $\sqrt{q \ln n}$ by Corollary 2.4. We have also shown that the number of iterations in which we set $\beta^* = L^* = L$ is at most $6\sqrt{q} \ln n$ using Lemma 2.12, and the number of iterations in which se set $\beta^* = L^* < L$ is at most $4\sqrt{q} \ln n$. Therefore, the length of the computed schedule is at most $6\sqrt{q} \ln n + 5\sqrt{q \ln n} \le 11\sqrt{q} \ln n$.

In each step we perform binary search (twice) with precision $1/(2n)$ over a domain that is contained in $[0, \beta_{\max}]$; the total number of binary search iterations per step is at most $\log(4n\beta_{\max}) \le 8(\ln \beta_{\max} + \ln n)$. Each binary search iteration requires at most $2s$ samples, where $s = (8/h) \cdot \ln(1/\delta') = 64|P| \ln(1/\delta')$ is given by Lemma 2.7. To ensure that the algorithm is successful with probability at least $1 - \delta$, we apply a union bound over $T = 88\sqrt{q} \ln n(\ln \beta_{\max} + \ln n)$ subroutine calls, and we choose $\delta' = \delta \cdot T^{-1}$ for each subroutine.

We obtain a schedule with length $11\sqrt{q} \ln n$ by taking at most $2sT \le 5 \cdot 10^4 q \ln^2 n(\ln q + \ln n) \ln(1/\delta')$ samples from the Gibbs distribution μ_β. Substituting the value for δ' and simplifying by using the assumptions discussed at the beginning of Section 2, plus the assumption $\ln n \ge 5 + \ln(\ln q + \ln n) + \ln \ln n$, the total number of samples is at most $5 \cdot 10^4 q \ln^2 n(\ln q + \ln n)^2 \ln(1/\delta)$. $\qquad\square$

2.5 Classical algorithm

Theorem 2.13. *Let $n \ge 1$, $\varepsilon \in (0, 1)$, and let $0 \le \beta_{\min} < \beta_{\max}$. Let $H : \Omega \to \{0, \ldots, n\}$ be a Hamiltonian, $Q = \mathsf{Z}(\beta_{\max})/\mathsf{Z}(\beta_{\min})$ and $q = \ln|\Omega|$. There exists a classical algorithm that satisfies the following: with probability at least $4/5$, the algorithm uses $\tilde{O}(q \cdot \varepsilon^{-2})$ Gibbs distribution samples and outputs $\widehat{Q}$ that approximates Q up to relative error ε*

Proof. The proof has the following structure:

1. Produce a cooling schedule schedule $\{\beta_1, \ldots, \beta_\ell\}$.

2. Bound relative variance of $\{W_i, V_i\}$.

3. Estimate $\widehat{W}_i \approx W_i, \widehat{V}_i \approx V_i$ and output $\prod_i \widehat{W}_i / \prod_i \widehat{V}_i$.

Step (1) uses Theorem 2.11 to find a slowly-varying sequence of inverse temperatures. To see that the variables V_i, W_i in step (2) have bounded variance, observe that:

$$
\begin{aligned}
\mathbb{S}[W_i] = \mathbb{S}[V_i] &= \frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}\left(\frac{\beta_i + \beta_{i+1}}{2}\right)^2} \\
&= \exp\left(-2f\left(\frac{\beta_i + \beta_{i+1}}{2}\right) + f(\beta_i) + f(\beta_{i+1})\right) \le 2 \cdot 10^5,
\end{aligned}
\tag{20}
$$

where the last inequality used that the schedule from Theorem 2.11 is $2 \cdot 10^5$-slowly-varying. Furthermore,

$$
\frac{\prod_i \mathbb{E}[V_i]}{\prod_i \mathbb{E}[W_i]} = \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\min})} = Q.
$$

From Theorem 2.1, obtaining $4 \cdot 10^5 \cdot \ell/(\eta\bar{\varepsilon}^2)$ samples from W_i and V_i suffices to produce $\widehat{W}, \widetilde{V}$, such that:

$$
\Pr\left[(1 - \bar{\varepsilon}) \cdot \prod_i \mathbb{E}[W_i] \le \widehat{W} \le (1 + \bar{\varepsilon}) \cdot \prod_i \mathbb{E}[W_i]\right] \ge 1 - \eta,
$$

for a given $\bar{\varepsilon} > 0$ (and similarly for $\hat{V}$). With probability at least $1 - 2\eta$, we have:

$$(1 - 3\bar{\varepsilon}) \cdot Q \leq (1 - \bar{\varepsilon})^2 \cdot Q = (1 - \bar{\varepsilon})^2 \cdot \frac{\prod_i \mathbb{E}[V_i]}{\prod_i \mathbb{E}[W_i]}$$

$$\leq \hat{V}/\hat{W} \leq (1 + \bar{\varepsilon})^2 \cdot \frac{\prod_i \mathbb{E}[V_i]}{\prod_i \mathbb{E}[W_i]} \leq (1 + 3\bar{\varepsilon}) \cdot Q$$

Choosing $\bar{\varepsilon} = \varepsilon/3$ shows that the proposed scheme correctly outputs a ε-approximation of Q. From Theorem 2.11, $5 \cdot 10^4 \cdot q \ln^2 n (\ln q + \ln n)^2 \ln(1/\delta)$ samples suffice to find a slowly-varying schedule $\{\beta_1, \ldots, \beta_\ell\}$, where δ is the maximum failure probability of the schedule generation algorithm. Choose $\delta = 1/10$. For the application of Theorem. 2.1, choose $\eta = 1/20$; the complexity of obtaining the estimates $\hat{W}, \hat{V}$ is

$$2\ell \cdot (4 \cdot 10^5 \cdot \ell/(\eta\bar{\varepsilon}^2)) \leq 5 \cdot 10^7 \ell^2/\varepsilon^2 \leq 7 \cdot 10^9 q \ln^2 n/\varepsilon^2,$$

where we used the expression for $\ell = 11\sqrt{q} \ln n$ given in Theorem 2.11. Hence, with probability at least $4/5$, our classical algorithm produces an ε-relative estimator of Q using

$$5 \cdot 10^5 \cdot q \ln^2 n \cdot (\ln q + \ln n)^2 + 7 \cdot 10^9 \cdot q \ln^2 n \cdot \varepsilon^{-2}$$

Gibbs distribution samples. $\qquad\square$

3 Quantum algorithm

3.1 Cooling schedule, quantumly

Our quantum algorithm is much simpler than the classical algorithm. For the schedule generation, the variances of the random variables $\mathbb{S}[V_i]$ and $\mathbb{S}[W_i]$ can be estimated with quantum amplitude estimation thanks to a simple relationship between the variances of these random variables and the overlap between the qsamples $|\mu_{\beta_i}\rangle$ and $|\mu_{\beta_{i+1}}\rangle$.

Fact 3.1 (Variance and overlap). *The overlap of Gibbs qsamples can be written as*

$$\left|\langle\mu_{\beta_i}|\mu_{\beta_{i+1}}\rangle\right|^2 = \left|\sum_{x \in \Omega} \sqrt{\frac{e^{-\beta_i H(x)}}{\mathsf{Z}(\beta_i)}} \cdot \sqrt{\frac{e^{-\beta_{i+1} H(x)}}{\mathsf{Z}(\beta_{i+1})}}\right|^2 \tag{21}$$

$$= \frac{\mathsf{Z}(\bar{\beta}_{i,i+1})^2}{\mathsf{Z}(\beta_i) \cdot \mathsf{Z}(\beta_{i+1})} = \frac{1}{\mathbb{S}[W_i]} = \frac{1}{\mathbb{S}[V_i]},$$

i.e., it equals the reciprocals of the variances $\mathbb{S}[V_i], \mathbb{S}[W_i]$.[10]

We again aim to find a cooling schedule that satisfies: $c_1 \leq (\mathsf{Z}(\beta_i) \cdot \mathsf{Z}(\beta_{i+1}))(\mathsf{Z}(\bar{\beta}_{i,i+1})^2) \leq c_2$, where c_1, c_2 are constants (see Eq. (12)). Using Fact 3.1, this can be written as $1/c_2 \leq |\langle\mu_{\beta_i}|\mu_{\beta_{i+1}}\rangle|^2 \leq 1/c_1$. We now apply the result below with $|\psi\rangle = |\mu_{\beta_i}\rangle$, $P = |\mu_{\beta_{i+1}}\rangle\langle\mu_{\beta_{i+1}}|$, and constant ε to estimate $|\langle\mu_{\beta_i}|\mu_{\beta_{i+1}}\rangle|^2$. The following is a corollary of [12, Thm. 6].

Corollary 3.2 (Non-destructive amplitude estimation). *Let* $|\psi\rangle$ *be an arbitrary quantum state and* P *an arbitrary projector. Let* $R_\psi = 2|\psi\rangle\langle\psi| - \mathbb{I}$. *Let* $\varepsilon \in (0,1)$. *Then, there is a quantum algorithm* $\mathcal{A}$ *that starts in the initial state* $|\psi\rangle$ *and with probability at least* $1 - \eta$ *outputs an estimate* $\hat{p}$ *of* $p = \langle\psi|P|\psi\rangle$ *with additive error* ε. *Additionally,* $\mathcal{A}$ *restores the state* $|\psi\rangle$ *with probability* $1 - \eta$ *and invokes the controlled reflection* R_ψ $O((1/\varepsilon) \cdot \ln(1/\eta))$ *many times.*

---

[10]It is now easy to see that, for $\mathbb{S}[X_i] = 1/|\langle\mu_{\beta_i}|\mu_{\gamma_{i,i+1}}\rangle|^2$ and $\mathbb{S}[V_i] = \mathbb{S}[W_i] = 1/|\langle\mu_{\beta_i}|\mu_{\beta_{i+1}}\rangle|^2$, the quantity $|\langle\mu_{\beta_i}|\mu_\beta\rangle|$ is monotonically decreasing in β (where $\beta \geq \beta_i$), hence $\mathbb{S}[V_i], \mathbb{S}[W_i] \leq \mathbb{S}[X_i]$.

Let $|\phi\rangle$ and $|\psi\rangle$ be two arbitrary quantum states. Define the projectors $P_\phi = |\phi\rangle\langle\phi|$ and $P_\phi^\perp = \mathbb{I} - P$ for the state $|\phi\rangle$ and similarly define P_ψ.

Fact 3.3. *Let $|\phi\rangle, |\psi\rangle$ be quantum states. Assume that the transition probability $a = |\langle\phi|\psi\rangle|^2$ is bounded from below by some constant. Then, starting with $|\phi\rangle$, we can prepare $|\psi\rangle$ with probability $1 - \eta$, by performing at most $O(\ln(1/\eta))$ measurements $\{P_\phi, P_\phi^\perp\}$ and $\{P_\psi, P_\psi^\perp\}$.*

Proof. The probability that we fail to prepare $|\psi\rangle$ after performing $2k + 1$ measurements is given by $(1 - a)(a^2 + (1 - a)^2)^k$. This result is easily established by observing that everything can be analyzed in the two dimensional space spanned by $P_\psi|\phi\rangle$ and $P_\psi^\perp|\phi\rangle$. Consider the two bases $\{|\phi\rangle, |\phi^\perp\rangle\}$ and $\{|\psi\rangle, |\psi^\perp\rangle\}$ for this subspace, where $\langle\phi|\phi^\perp\rangle = 0$ and $\langle\psi|\psi^\perp\rangle = 0$. The factor $(1 - a)$ corresponds to the transition $|\phi\rangle \to |\psi^\perp\rangle$, i.e., that the first measurement fails to prepare $|\psi\rangle$. The term a^2 corresponds to the transition $|\psi^\perp\rangle \to |\phi^\perp\rangle \to |\psi^\perp\rangle$ and the term $1 - a$ corresponds to the transition $|\psi^\perp\rangle \to |\phi\rangle \to |\psi^\perp\rangle$, i.e., that P_ψ measurement followed by the P_ϕ measurement fails to prepare $|\psi\rangle$. It is clear that we can make the failure probability smaller than η by choosing $k = O(\ln(1/\eta))$ provided that a is bounded from below by a constant. $\square$

We now describe a quantum algorithm to obtain a schedule that matches Theorem 2.2. The difference between the classical and quantum setting is mainly due to the classical estimator in Eq. (8), inaccurate compared to the quantum estimator. The quantum algorithm is more efficient than the classical algorithm, and improves the length quadratically in $\ln n$.

Theorem 3.4 (Quantum schedule generation)**.** *With probability $\geq 1 - \delta$, Algorithm 3 computes a 15-slowly-varying schedule with length at most $\sqrt{q \ln n}$, and uses $O\big(\sqrt{q \ln n}\big(\ln q + \ln n\big) \cdot \big(\ln q + \ln\ln n + \ln(1/\delta)\big)\big)$ reflections about Gibbs states at different inverse temperatures.*

---

input : Initial temperature β_0, largest temperature $\beta_{\max}$, probability δ, constant c_2.
output: Set of inverse temperatures $\beta_0, \ldots, \beta_k = \beta_{\max}$.
Set $k \leftarrow 0$;
while $\beta_k < \beta_{\max}$ **do**
    Define the function $f_o(\beta) := |\langle\mu_{\beta_k}|\mu_\beta\rangle|^2$, where the overlap is evaluated using Corollary 3.2 with additive error 0.005;
    Compute $\beta^* \leftarrow \textsc{BinarySearch}(f_o(\cdot) \geq 0.075, [\beta_k, q], 1/2n)$;
    Set $\beta_{k+1} \leftarrow \beta^*, k \leftarrow k + 1$;
end
return $\beta_1, \ldots, \beta_k, \beta_{\max}$

Algorithm 3: Quantum schedule generation procedure.

Proof. We first bound the number of iterations of the algorithm. Assume for now that all subroutines are successful, and we use a union bound at the end. Notice that $|\langle\mu_{\beta_k}|\mu_{\beta_k}\rangle|^2 = 1$ so the predicate used in binary search is satisfied at the left endpoint of the interval $[\beta_k, q]$. If the binary search (as given in Alg. 1) returns q, we are done. Otherwise, it determines

λ, ρ such that:

$$|\langle \mu_{\beta_k} | \mu_\lambda \rangle|^2 \geq 0.075 - 0.005,$$
$$|\langle \mu_{\beta_k} | \mu_\rho \rangle|^2 < 0.075 + 0.005,$$
$$\rho - \lambda \leq 1/2n, \ \lambda \geq \beta_k.$$

We have

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\lambda)}{\mathsf{Z}(\frac{\beta_k + \lambda}{2})^2} = \frac{1}{|\langle \mu_{\beta_k} | \mu_\lambda \rangle|^2} \leq \frac{1}{0.075 - 0.005} \leq 15,$$

which shows that the schedule is 15-slowly-varying. Also,

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\rho)}{\mathsf{Z}(\frac{\beta_k + \rho}{2})^2} = \frac{1}{|\langle \mu_{\beta_k} | \mu_\rho \rangle|^2} \geq \frac{1}{0.075 + 0.005} = 12.5.$$

It can be shown that $\mathsf{Z}(\lambda) \geq \mathsf{Z}(\rho)$, and that $\mathsf{Z}(\frac{\beta_k + \lambda}{2}) \leq \mathsf{Z}(\frac{\beta_k + \rho}{2})e^{1/4}$.

Using these facts, we can write:

$$\frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\lambda)}{\mathsf{Z}(\frac{\beta_k + \lambda}{2})^2} \geq \frac{\mathsf{Z}(\beta_k)\mathsf{Z}(\rho)}{e^{1/2}\mathsf{Z}(\frac{\beta_k + \rho}{2})^2} \geq e^{-1/2}\, 12.5 \geq e^2. \tag{22}$$

We see that the cooling schedule obtained by setting $\beta_{k+1} = \lambda$ in each iteration satisfies the condition in Corollary 2.4, which coincides with the bound in (22). This implies that the length of the resulting cooling schedule is bounded from above by the length of a perfectly balanced cooling schedule. The latter in turn is bounded from above by $\sqrt{q \ln n}$, which is established in Corollary 2.3. In each step we perform binary search with precision $1/2n$ over a domain that is contained in $[0, \beta_{\max}]$, which implies that the total number of binary search iterations per step is at most $\log(2n\beta_{\max}) \leq 2(\ln \beta_{\max} + \ln n)$. The total number of binary searches in all steps is therefore at most $2\sqrt{q \ln n}(\ln \beta_{\max} + \ln n)$.

Each binary search invokes amplitude estimation subroutine in Corollary 3.2 with additive error ε set to the constant 0.005. To ensure that the entire algorithm succeeds with probability at least $1 - \delta$, we choose the maximum probability of failure $\eta = \delta/(4\sqrt{q \ln n}(\ln \beta_{\max} + \ln n))$ for the amplitude estimation subroutine. By the union bound, the probability that at least one amplitude estimation subroutine fails is $\leq \delta/2$. Each amplitude estimation requires $O(\ln(1/\eta))$ many reflections since ε is a constant. Hence, the total number of reflections for all binary searches is $O(\sqrt{q \ln n}(\ln \beta_{\max} + \ln n) \ln(1/\eta))$, which, in terms of δ and using the assumption $\beta_{max} \leq q$ (see the proof of Theorem 2.11), is

$$O(\sqrt{q \ln n}(\ln q + \ln n)(\ln q + \ln \ln n + \ln(1/\delta))).$$

It remains to bound the time it takes to iteratively prepare the states $|\mu_{\beta_0}\rangle, \ldots, |\mu_{\beta_k}\rangle$ in the schedule. The overlap between all adjacent qsamples is large since $|\langle \mu_{\beta_k} | \mu_{\beta_{k+1}} \rangle|^2 \geq \frac{1}{15}$. Therefore, we can use the method described in Fact 3.3 to "jump" from $|\mu_{\beta_k}\rangle$ to $|\mu_{\beta_{k+1}}\rangle$. The necessary projective measurements can be implemented with the reflections around the Gibbs qsamples. Since there are at most $\sqrt{q \ln n}$ stages, we want the failure probability of any jump to be smaller than $(\delta/2)/\sqrt{q \ln n}$ to ensure that the probability that any jump fails is $\leq \delta/2$. The number of reflections across the Gibbs qsamples is $O(\sqrt{q \ln n}(\ln q + \ln n)(\ln q + \ln \ln n + \ln(1/\delta)))$. $\qquad \square$

Before we prove the main theorem in this section, we need the following result about the sample complexity of computing functions of random variables, given access to coherent qsamples of distributions. The proof of this theorem has been deferred to Appendix C.

Theorem 3.5 (Bounded relative variance). *Suppose a distribution $D : \Omega \to [0,1]$ and function $f : \Omega \to [0,\infty)$ satisfy B-bounded relative variance for $B > 1$. Then, there is a quantum algorithm $\mathcal{A}$ that: given $16B \ln(2/\eta) + 1$ copies of $|\psi_D\rangle$, with probability at least $1 - \eta$, $\mathcal{A}$ outputs an ε-relative estimate $\hat{\mu}$ of $\mu = \mathbb{E}x \sim D[f(x)]$. Additionally, $\mathcal{A}$ restores one copy of $|\psi_D\rangle$ and invokes the reflection R_D (i.e., the reflection across the state $\sum_{x \in \Omega} \sqrt{D(x)}|x\rangle$)*

$$O\Big(\sqrt{B}/\varepsilon \cdot \big(\ln(B/\varepsilon)\big)^{1.5} \cdot \ln\big(\ln(B/\varepsilon)/\eta\big)\Big)$$

many times.

The theorem above allows us to estimate the expectations $\mathbb{E}[W_k]$ and $\mathbb{E}[V_k]$ by setting D to be the Gibbs distribution μ_{β_k} and f to be $\exp(-\bar{\beta}_{k,k+1}H(x_k))$ for the random variable V_k or to $\exp(\bar{\beta}_{k,k+1}H(x_{k+1}))$ for the random variable W_k, as defined in Section 2.2. We can use that $B \leq 15$ because the cooling schedule $\beta_0, \ldots, \beta_\ell$ generated by the algorithm in Theorem 3.4 is 15-slowly varying. The error parameter ε and the failure probability η will have to be of the order of $1/\ell$.

3.2 Quantum algorithm

Theorem 3.6. *Let $n \geq 1$, $\delta, \varepsilon, \eta \in (0,1)$ and $0 \leq \beta_{\min} < \beta_{\max}$. Let $H : \Omega \to [0,n]$ be a classical Hamiltonian. Let $Q = \mathsf{Z}(\beta_{\max})/\mathsf{Z}(\beta_{\min})$ and $q = \ln Q$. There exists a quantum algorithm that uses $\tilde{O}\left(q \cdot \varepsilon^{-1}\right)$ many applications of reflection operators around Gibbs state (at different inverse temperatures) and with probability at least $4/5$, approximates Q up to to relative error ε, i.e., outputs $\widehat{Q}$ such that $(1 - \varepsilon) \cdot Q \leq \widehat{Q} \leq (1 + \varepsilon) \cdot Q$*

Proof. The proof of the quantum algorithm is very similar to the classical proof in Theorem 2.13. The main differences are in estimating the cooling schedule using quantum techniques and estimating the mean values using quantum techniques. Overall, the structure of the quantum algorithm is:

1. Compute a slowly varying cooling schedule of length ℓ as described in Theorem 3.4.

2. Use the quantum mean estimation algorithm to estimate the expectations $\mathbb{E}[W_i]$ and $\mathbb{E}[V_i]$ with relative error ε/ℓ using Theorem 3.5.

3. Multiply the estimates of $\mathbb{E}[W_i]$ to obtain an estimate of $\mathbb{E}[W]$ and $\mathbb{E}[V_i]$ to obtain an estimate of $\mathbb{E}[V]$ and output their ratio as the final estimate.

We argue the correctness of the algorithm first. Step (1) is clear. Theorem 3.4 satisfies the following: with probability $\geq 9/10$, the algorithm generates a sequence of $\ell = \sqrt{q \ln n}$ inverse temperatures $\beta_1, \ldots, \beta_\ell$ satisfying

$$\mathbb{S}[W_i] = \mathbb{S}[V_i] = \frac{\mathsf{Z}(\beta_i)\mathsf{Z}(\beta_{i+1})}{\mathsf{Z}\left(\frac{\beta_i + \beta_{i+1}}{2}\right)^2} \leq 15 \qquad \text{for every } i \in [\ell].$$

Additionally observe that $\frac{\prod_i \mathbb{E}[V_i]}{\prod_i \mathbb{E}[W_i]} = \frac{\mathsf{Z}(\beta_{\max})}{\mathsf{Z}(\beta_{\min})} = Q$. One can use Theorem 3.5 which uses $O(\ln \ell)$ samples of each $|\mu_{\beta_i}\rangle$ and invokes the reflection R_{β_i} $\tilde{O}(\ell/\varepsilon)$ many times and: with probability $\geq 1 - 1/(20\ell)$ produces a $\widehat{W}_i, \widehat{V}_i$ such that

$$1 - \varepsilon/(2\ell) \leq \frac{\hat{V}_i}{\mathbb{E}[V_i]} \leq 1 + \varepsilon/(2\ell) \quad \text{and}$$

$$1 - \varepsilon/(2\ell) \leq \frac{\hat{W}_i}{\mathbb{E}[W_i]} \leq 1 + \varepsilon/(2\ell).$$

We use ratios of the lower and upper bounds to bound the ratio $\hat{W}_i/\hat{V}_i$ from below and above and employ the union bound to obtain a bound on the success probability. We obtain that

$$(1 - \varepsilon/(2\ell))^2 \leq \frac{\hat{W}_i/\hat{V}_i}{\mathbb{E}[W_i]/\mathbb{E}[V_i]} \leq (1 + \varepsilon/(2\ell))^2. \tag{23}$$

hold with probability at least $1 - 1/(10\ell)$ for all i. This in turn implies that

$$(1 - \varepsilon/(2\ell))^{2\ell} \leq \frac{\prod_i(\hat{W}_i/\hat{V}_i)}{\prod_i(\mathbb{E}[W_i]/\mathbb{E}[V_i])} \leq (1 + \varepsilon/(2\ell))^{2\ell}. \tag{24}$$

holds with probability at least $1 - 1/10 - 1/10 \geq 4/5$ (the first $1/10$ is from the union bound over the $1 - 1/(10\ell)$ in satisfying Eq. (23) and the second $1/10$ comes because Theorem 3.4 was assumed to fail with probability $\leq 1/10$). Recall that $\prod_i(\mathbb{E}[W_i]/\mathbb{E}[V_i]) = Q$, where $Q = \mathsf{Z}(\beta_{\max})/\mathsf{Z}(\beta_{\min})$ is the desired ratio. We obtain:

$$(1 - 2\varepsilon) \cdot Q \leq (1 - \varepsilon/(2\ell))^{2\ell} \cdot Q \leq \prod_i(\hat{W}_i/\hat{V}_i) \leq$$

$$(1 + \varepsilon/(2\ell))^{2\ell} \cdot Q \leq e^{\varepsilon} \cdot Q \leq (1 + 2\varepsilon) \cdot Q,$$

where the first inequality used $(1 + x)^t \geq 1 + xt$ (for $x \geq -1$ and $t \geq 2$), the second and third inequality used Eq. (24), the fourth inequality used $(1 + x)^t \leq e^{xt}$ (for $x, t \geq 0$), last inequality used $e^x \leq 1 + 2x$ (for $x \in [0, 1]$). Hence our final output $\prod_i(\hat{W}_i/\hat{V}_i)$ is a (2ε)-relative estimator of Q.

It remains to analyze the complexity of this quantum algorithm. First, Theorem 3.4 uses

$$O\Big(\ell \cdot (\ln q + \ln n) \cdot \ln \ell\Big)$$

invocations of a reflection around the Gibbs state (note that we fix $\delta = O(1)$ when invoking this theorem).

Next, Theorem 3.5 uses $O(\ell \ln \ell)$ copies of the Gibbs states and invokes the reflection operator around Gibbs state (at different inverse temperatures) $\tilde{O}(\ell^2/\varepsilon)$ many times (we specify the poly-logarithmic factors below). We can use Fact 3.3 and similar arguments as at the end of the proof of Theorem 3.4 to analyze the complexity of preparing these copies. Observe that we cannot directly prepare a qsample $|\mu_{\beta_k}\rangle$ for an arbitrary k, but have to move successively through the sequence $|\mu_{\beta_0}\rangle, \ldots, |\mu_{\beta_k}\rangle$. In total, we need to make at most $O(\ell^2 \ln \ell)$ transitions between Gibbs qsamples of adjacent stages by performing reflections around these qsamples. To use the union bound, we need that the failure probability of any transition should be on the order of $\ell^2 \ln \ell$, which can be accomplished with at most $O(\ln \ell)$ many reflections per transition. Therefore, $O(\ell^2 \cdot (\ln \ell)^2)$ many reflections are necessary in total to prepare all the copies. Overall, using $\ell = \sqrt{q \ln n}$, our algorithm uses

$$O(q + \ell^2 \cdot (\ln \ell)^2) = O(q \ln n \cdot (\ln q + \ln n)^2)$$

walk steps to prepare these Gibbs state and

$$O\Big(q/\varepsilon \cdot \ln n \cdot \big(\ln(q/\varepsilon \cdot \ln n)\big)^2\Big)$$

invocations of the reflection around Gibbs state. The total number of quantum walk steps is

$$O\Big(q/\varepsilon \cdot \ln n \cdot \big(\ln(q/\varepsilon \cdot \ln n)\big)^2 + q \ln n \cdot (\ln q + \ln n)^2\Big).$$

$\square$

Finally, in order to translate the complexity stated in Theorem 3.6 into a running time bound, observe that each invocation of the reflection operator in the theorem involves $O(1/\sqrt{\Delta})$ Markov chain steps on top of the sample complexity above: this is a well-known result from [24, 20].

Corollary 3.7. *In the setting of Theorem 3.6, there exists a quantum algorithm that makes $\tilde{O}\left(q \cdot \varepsilon^{-1} \cdot \Delta^{-1/2}\right)$ steps of the quantum walk operator (each with spectral gap lower bounded by Δ) and with probability at least 4/5, approximates Q up to to relative error ε, i.e., outputs $\widehat{Q}$ such that $(1-\varepsilon) \cdot Q \leq \widehat{Q} \leq (1+\varepsilon) \cdot Q$.*

4 Conclusion

First, we have improved a well-known classical algorithm for estimating partition functions. Second, we have shown how to quantize this improved algorithm, thereby obtaining a quadratic speed-up with respect to the estimation precision and to the spectral gaps of the underlying Markov chains. In particular, we have obtained the best quantum algorithm for estimating partition functions, improving upon the state-of-the-art algorithm due to Harrow and Wei [12].

4.1 Open questions.

We conclude with a few concrete open questions. (1) Can we *fully* quantize the classical algorithms of [14] and Kolmogorov [18], in the process removing large prefactor in our algorithms? (2) Does there exist a classical algorithm for computing partition functions with schedule length $O(\sqrt{\ln |\Omega| \ln n})$, matching the conjecture by [23]? (3) Can we prove *any* lower bound in the standard model of computation for computing partition functions (either quantum or classical)? [11] (4) Can one remove the assumption that SVV or TPA applies only for non-negative Hamiltonians?

5 Acknowledgment

References

[1] Ivona Bezáková, Daniel Štefankovič, Vijay V Vazirani, and Eric Vigoda. Accelerating simulated annealing for the permanent and combinatorial counting problems. *SIAM Journal on Computing*, 37(5):1429–1454, 2008. `doi:10.1137/050644033`.

[2] Kurt Binder and Dieter W. Heermann. *Monte Carlo Simulation in Statistical Physics: An Introduction (Graduate Texts in Physics)*. Springer, 2019.

---

[11] The lower bound of [18] assumes a non-standard oracle access.

---

[3] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.

[4] Shouvanik Chakrabarti, Andrew M Childs, Shih-Han Hung, Tongyang Li, Chunhao Wang, and Xiaodi Wu. Quantum algorithm for estimating volumes of convex bodies. *arXiv:1908.03903*, 2019.

[5] Guillaume Desjardins, Yoshua Bengio, and Aaron C Courville. On tracking the partition function. In *Advances in neural information processing systems*, pages 2501–2509. Citeseer, 2011.

[6] Martin Dyer and Alan Frieze. Computing the volume of convex bodies: a case where randomness provably helps. *Probabilistic combinatorics and its applications*, 44:123–170, 1991.

[7] Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991. `doi:10.1145/102782.102783`.

[8] Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, pages 163–185, 1998. `doi:10.1214/ss/1028905934`.

[9] Paul Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.

[10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

[11] Yassine Hamoudi and Frédéric Magniez. Quantum Chebyshev's Inequality and Applications. In *46th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 132, pages 69:1–69:16, 2019. `doi:10.4230/LIPIcs.ICALP.2019.69`.

[12] Aram W Harrow and Annie Y Wei. Adaptive quantum simulated annealing for bayesian inference and estimating partition functions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 193–212, 2020. `doi:10.1137/1.9781611975994.12`.

[13] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006. `doi:10.1162/neco.2006.18.7.1527`.

[14] Mark Huber. Approximation algorithms for the normalizing constant of Gibbs distributions. *The Annals of Applied Probability*, 25(2):974–985, 2015. `doi:10.1214/14-AAP1015`.

[15] Mark Huber and Sarah Schott. Using TPA for Bayesian inference. *Bayesian Statistics*, 9:257–282, 2010. `doi:10.1093/acprof:oso/9780199694587.003.0009`.

[16] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal On Computing*, 18(6):1149–1178, 1989. `doi:10.1137/0218077`.

[17] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169 – 188, 1986. `doi:10.1016/0304-3975(86)90174-X`.

[18] Vladimir Kolmogorov. A faster approximation algorithm for the Gibbs partition function. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 228–249. PMLR, 2018.

[19] Qiang Liu and Alexander T. Ihler. Bounding the partition function using holder's inequality. In *Proceedings of the 28th International Conference on Machine Learning, ICML*, pages 849–856. Omnipress, 2011.

[20] F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *SIAM Journal on Computing*, 40(1):142–164, 2011. `doi:10.1137/090745854`.

[21] Ashley Montanaro. Quantum speedup of monte carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181):20150301, 2015.

[22] Elchanan Mossel and Allan Sly. Exact thresholds for ising–gibbs samplers on general graphs. *Ann. Probab.*, 41(1):294–328, 01 2013. `doi:10.1214/11-AOP737`.

[23] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *Journal of the ACM (JACM)*, 56(3):1–36, 2009. `doi:10.1145/1516512.1516520`.

[24] Mario Szegedy. Quantum speed-up of markov chain based algorithms. In *45th Annual IEEE symposium on foundations of computer science*, pages 32–41. IEEE, 2004.

[25] J. P. Valleau and D. N. Card. Monte carlo estimation of the free energy by multistage sampling. *The Journal of Chemical Physics*, 57(12):5457–5462, 1972.

[26] Eric Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41(3):1555–1569, 2000. `doi:10.1063/1.533196`.

[27] Marc Vuffray, Sidhant Misra, Andrey Lokhov, and Michael Chertkov. Interaction screening: Efficient and sample-optimal learning of ising models. In *Advances in Neural Information Processing Systems*, pages 2595–2603, 2016.

[28] Pawel Wocjan and Anura Abeyesinghe. Speedup via quantum sampling. *Physical Review A*, 78(4):042336, 2008. `doi:10.1103/PhysRevA.78.042336`.

A Proof of schedule length

We report the statement of Theorem 2.2, and detail its proof.

Theorem A.1 (Perfectly-balanced schedule length [23])**.** *There exists a sequence $\beta_0 < \cdots < \beta_\ell$ with $\beta_0 = \beta_{\min}$ and $\beta_\ell = \beta_{\max}$ satisfying the condition*

$$f\left(\frac{\beta_i + \beta_{i+1}}{2}\right) = \frac{f(\beta_i) + f(\beta_{i+1})}{2} - 1 \tag{25}$$

and having length ℓ bounded from above by

$$\ell \le \sqrt{(f(\beta_{\min}) - f(\beta_{\max})) \cdot \frac{1}{2}\ln\left(\frac{f'(\beta_{\min})}{f'(\beta_{\max})}\right)}.$$

Proof. Suppose we have already constructed the sequence up to β_i and let β_{i+1} be the largest value in $[\beta_i, \beta_{\max}]$ so that β_i and β_{i+1} satisfy Eq. (25). For notational simplicity, let

$$\bar{\beta}_{i,i+1} = \frac{\beta_i + \beta_{i+1}}{2}, \quad d_{i,i+1} = \frac{\beta_{i+1} - \beta_i}{2}, \quad K_i = f(\beta_i) - f(\beta_{i+1}).$$

If $\beta_{i+1} \geq \beta_{\max}$, then we are done constructing the well-balanced sequence. Otherwise, by the maximality of β_{i+1} with equality, we have

$$f(\bar{\beta}_{i,i+1}) = \frac{f(\beta_i) + f(\beta_{i+1})}{2} - 1.$$

Additionally since f is convex,[12] for every $a, b \in [\beta_{\min}, \beta_{\max}]$ satisfying $a < b$, we have

$$f'(a) \leq \frac{f(b) - f(a)}{b - a} \leq f'(b). \tag{26}$$

Setting $a = \beta_i$ and $b = \beta_{i+1}$ in Eq.(26), we have

$$f'(\beta_i) \leq \frac{f(\beta_{i+1}) - f(\beta_i)}{\beta_{i+1} - \beta_i},$$

which we rewrite as

$$-f'(\beta_i) \geq \frac{K_i}{2d_{i,i+1}}. \tag{27}$$

Setting $a = \bar{\beta}_{i,i+1}$ and $b = \beta_{i+1}$ in Eq. (26), we obtain

$$\frac{f(\beta_{i+1}) - f(\bar{\beta}_{i,i+1})}{\beta_{i+1} - \bar{\beta}_{i,i+1}} \leq f'(\beta_{i+1}).$$

Given β_i, β_{i+1} satisfy Eq. (25), we have

$$-f'(\beta_{i+1}) \leq \frac{f(\bar{\beta}_{i,i+1}) - f(\beta_{i+1})}{d_{i,i+1}} = \frac{f(\beta_i) - f(\beta_{i+1}) - 2}{2d_{i,i+1}} = \frac{K_i - 2}{2d_{i,i+1}}. \tag{28}$$

Putting together Eq. (27), (28), we have

$$\frac{f'(\beta_{i+1})}{f'(\beta_i)} = \frac{-f'(\beta_{i+1})}{-f'(\beta_i)} \leq \frac{K_i - 2}{K_i} = 1 - \frac{2}{K_i} \leq \exp\left(-\frac{2}{K_i}\right). \tag{29}$$

Taking the product of Eq. (29) for all $i \in \{0, \ldots, \ell - 1\}$ and rearranging, we obtain

$$\prod_{i=0}^{\ell-1} \exp\left(\frac{2}{K_i}\right) \leq \prod_{i=0}^{\ell-1} \frac{f'(\beta_i)}{f'(\beta_{i+1})} = \frac{f'(\beta_{\min})}{f'(\beta_{\max})}. \tag{30}$$

Taking the logarithm Eq. (30) and dividing both sides by 2, we obtain

$$\sum_{i=0}^{\ell-1} \frac{1}{K_i} = \frac{1}{2} \ln\left(\frac{f'(\beta_{\min})}{f'(\beta_{\max})}\right). \tag{31}$$

Summing K_i for all $i \in \{0, \ldots, \ell - 1\}$, we obtain

$$\sum_{i=0}^{\ell-1} K_i = f(\beta_{\min}) - f(\beta_{\max}). \tag{32}$$

Applying Cauchy-Schwarz inequality to Eq. (31), (32) we obtain

$$\ell^2 \leq \left(f(\beta_{\min}) - f(\beta_{\max})\right) \cdot \frac{1}{2} \ln\left(\frac{f'(\beta_{\min})}{f'(\beta_{\max})}\right).$$

This concludes the proof of the theorem. $\qquad\square$

---

[12]In fact it is well-known that the partition function of a Hamiltonian is *strongly convex*, not just convex (see [27] for the proof of strong convexity for classical Hamiltonians). Using this strong convexity property it is possible that one could potentially improve the upper bound on the schedule length; we leave it as an interesting open question.

B Bounding number of iterations in classical schedule generation

Here we prove Lemma 2.12 (restated below for convenience) that bounds number of "long moves" in case 1 of the proof of Theorem 2.11.

Lemma B.1. *The number of "long moves" in Algorithm 2, where we set $\beta^* = L^* = L$, is at most $6\sqrt{q}\ln n$.*

Proof. We follow the same scheme as the proof of [23, Lemma 5.14], with a few modifications to account for the differences between our Algorithm 2 and the schedule generation algorithm in [23]. Because of these modifications, we are able to obtain a tighter upper bound.

Recall that $I \subseteq \{0, \ldots, n\}$ and I is a contiguous interval, therefore the set of possible interval widths is a subset of $\{0, \ldots, n\}$. We first consider long moves that happen with interval width of at least 1; as discussed at the end of this proof, there can be at most one long move with interval width 0, therefore a bound on the number of long moves with interval width in $\{1, \ldots, n\}$ immediately implies a bound on the total number of long moves. Let x_k be the number of times that we perform a long move with an interval of width k, for $k = 1, \ldots, n$. Let $k_{\max} := \arg\max_k\{x_k : x_k > 0\}$, i.e., the maximum interval width k for which we perform at least one long move.

Let us consider all the long moves performed with an interval width belonging to the set $\{k_{\min}, \ldots, k_{\max}\}$, where we arbitrarily fix a choice $k_{\min} \in \{1, \ldots, k_{\max}\}$. We want to derive a lower bound on the inverse temperature β reached after $x_{k_{\min}} + x_{k_{\min}+1} + \cdots + x_{k_{\max}} = t+1$ such long moves. Since β is only increasing in the course of the algorithm, clearly a lower bound after t moves is also valid after $t+1$ moves. In particular, since we want to obtain the tightest possible lower bound, we consider β after performing y_k moves with interval width k, where y_k is defined as follows for $k \in \{k_{\min}, \ldots, k_{\max}\}$:

$$y_k = \begin{cases} x_k & \text{if } k < k_{\max} \\ x_k - 1 & \text{if } k = k_{\max}. \end{cases}$$

Notice that this is equivalent to undercounting the number of moves at interval width $k_{\max}$ by exactly 1. Since for interval width k, the inverse temperature β increases by exactly $1/k$ during a long move, after t long moves that satisfy the move count $y_{k_{\min}}, \ldots, y_{k_{\max}}$ we have:

$$\beta \geq \sum_{k=k_{\min}}^{k_{\max}} \frac{y_k}{k}. \tag{33}$$

We claim that in addition, β must satisfy:

$$\beta \leq \frac{q + \ln\frac{1}{h}}{k\sqrt{q}}. \tag{34}$$

We now prove this claim. First, notice that by construction of P, since we choose the width w of an interval $I = \{b, \ldots, b+w = c\}$ to be $w = \lfloor b/\sqrt{q}\rfloor$, we must have:

$$b \geq w\sqrt{q} \tag{35}$$

Next, notice that in order for a long move to happen starting from inverse temperature β with an interval I of width w, there must exist some inverse temperature $\beta' > \beta$ at which I is h-heavy (more specifically, I must be h-heavy at $\beta' = \beta + 1/w$, otherwise the long

move cannot take place). The weight of I at β' is $\frac{1}{\mathsf{Z}(\beta')} \sum_{x:H(x)\in I} e^{-\beta' H(x)}$. Therefore we must have:

$$h \leq \frac{1}{\mathsf{Z}(\beta')} \sum_{x:H(x)\in I} e^{-\beta' H(x)} \leq \frac{1}{\mathsf{Z}(\beta')} \sum_{x:H(x)\in I} e^{-\beta H(x)} \leq \sum_{x:H(x)\in I} e^{-\beta H(x)},$$

where for the second inequality we used $\beta' > \beta$, and for the last inequality we used $\mathsf{Z}(\beta') \geq 1$, which is true by assumption. We obtain the following chain of inequalities:

$$h \leq \sum_{x:H(x)\in I} e^{-\beta H(x)} \leq |\Omega| e^{-\beta k_{\min}\sqrt{q}},$$

where we used the facts that the width of I is between $k_{\min}$ and $k_{\max}$, and Eq. 35 (which, together, imply $H(x) \geq k_{\min}\sqrt{q}$ for the states considered in the summation). Taking the logarithm on both sides yields

$$\ln h \leq q - \beta k_{\min}\sqrt{q},$$

which immediately implies (34). Combining (33) and (34), we obtain:

$$\sum_{k=k_{\min}}^{k_{\max}} \frac{y_k}{k} \leq \frac{q + \ln\frac{1}{h}}{k\sqrt{q}}. \tag{36}$$

Now notice that $\sum_{k_{\min}=1}^{k_{\max}} \sum_{k=k_{\min}}^{k_{\max}} \frac{y_k}{k} = \sum_{k=1}^{k_{\max}} y_k$, because each term $\frac{y_k}{k}$ appears exactly k times in the double summation. Therefore, taking (36) and taking the sum for $k_{\min} = 1, \ldots, k_{\max}$ on both sides, we obtain:

$$\sum_{k=1}^{k_{\max}} y_k = \sum_{k_{\min}=1}^{k_{\max}} \sum_{k=k_{\min}}^{k_{\max}} \frac{y_k}{k} \leq \sum_{k=1}^{k_{\max}} \frac{q + \ln\frac{1}{h}}{k\sqrt{q}} \leq (1 + \ln n)\frac{q + \ln\frac{1}{h}}{\sqrt{q}}, \tag{37}$$

where the last inequality exploits the fact that $k_{\max} \leq n$ and the well-known inequality $\sum_{i=1}^{n} 1/i \leq 1 + \ln n$. Using the choice $h = \frac{1}{8|P|}$ in Algorithm 2 and Lemma 2.10, we can write $\ln\frac{1}{h} = \ln(32\sqrt{q}\ln n) \leq 2\ln q \ln\ln n$. Finally, using (37) and $(1 + \ln n) \leq 2\ln n$, we obtain:

$$\sum_{k=1}^{k_{\max}} y_k \leq (2\ln n)\frac{q + 2\ln q \ln\ln n}{\sqrt{q}} \leq 4\sqrt{q}\ln n.$$

The LHS of the last equation is, by definition, equal to the number of long moves that are performed with interval width in $\{1, \ldots, n\}$, minus one. Recall that in Algorithm 2 we can also have long moves with intervals of width 0, but there can be at most one such move because then we set $L = q \geq \beta_{\max}$. Hence, the total number of long moves in the algorithm is at most $4\sqrt{q}\ln n + 2 \leq 6\sqrt{q}\ln n$. $\qquad\square$

C   Proof of non-destructive amplitude estimation

We first introduce some definitions and notation which we use throughout.

Definition C.1. *Let Ω be a finite set, $D : \Omega \to [0,1]$ be a distribution, and $f : \Omega \to [0,\infty)$ be an arbitrary function. Define the mean*

$$\mu = \mathbb{E}_{x\sim D}[f(x)] = \sum_{x\in\Omega} D(x)f(x)$$

and the second moment

$$\phi = \mathbb{E}_{x \sim D}[f(x)^2] = \sum_{x \in \Omega} D(x)f(x)^2,$$

and the relative variance $\frac{\phi}{\mu^2}$. *We say* (D, f) *satisfy* B-*bounded relative variance if*

$$\frac{\phi}{\mu^2} \le B.$$

We say that $\hat{\mu}$ *is an* ε-*relative estimate of* μ *if* $|\mu - \hat{\mu}| \le \varepsilon\mu$.

Let $|\psi_D\rangle$ denote the *coherent encoding* of the distribution D, i.e,

$$|\psi_D\rangle = \sum_{x \in \Omega} \sqrt{D(x)}|x\rangle,$$

and R_D the *reflection* around $|\psi_D\rangle$, i.e.,

$$R_D = 2|\psi_D\rangle\langle\psi_D| - \mathbb{I}.$$

The results in this section are stated as general subroutines. Note that when we apply these techniques for the task of estimating partition functions, the probability distribution D corresponds to the Gibbs distribution μ_{β_i} at β_i, the function f corresponds to the function $\exp(-\beta H(\cdot))$.[13]

The goal of this section is to present a quantum algorithm for estimating μ with relative error ε given access to copies of the coherent encoding $|\psi_D\rangle$ and reflection operator R_D. We assume that the relative variance $\frac{\phi}{\mu^2}$ is bounded from above by $B > 1$. Naturally, the goal is to estimate μ using as few copies of $|\psi_D\rangle$ and invocations of the reflection operator R_D as possible. Naively, a quantum algorithm could simply use multiple copies of $|\psi_D\rangle$ to obtain $O(B/\varepsilon^2)$ samples x according to D, and estimate μ using the classical Chebyshev's inequality. However, quantumly one can obtain a quadratic improvement in $1/\varepsilon$, which we prove in this section.

Theorem C.2 (Bounded relative variance). *Suppose a distribution* $D : \Omega \to [0, 1]$ *and function* $f : \Omega \to [0, \infty)$ *satisfy* B-*bounded relative variance. Then, there is a quantum algorithm* $\mathcal{A}$ *that: given* $16B \ln(2/\eta) + 1$ *copies of* $|\psi_D\rangle$, *with probability at least* $1 - \eta$, $\mathcal{A}$ *outputs an* ε-*relative estimate* $\hat{\mu}$ *of* $\mu = \mathbb{E}_{x \sim D}[f(x)]$. *Additionally, with probability* $\ge 1 - \eta$, $\mathcal{A}$ *restores one copy of* $|\psi_D\rangle$. *Overall* $\mathcal{A}$ *invokes the reflection* R_D

$$O\left(\sqrt{B}/\varepsilon \cdot \left(\ln(B/\varepsilon)\right)^{1.5} \cdot \ln\left(\ln(B/\varepsilon)/\eta\right)\right)$$

many times.

The theorem above builds upon and improves the results due to [21, Algorithm 4 and Theorem 6]. Most importantly, the complexity of our algorithm grows only with $\sqrt{B}/\varepsilon$, whereas the algorithm in [21] grows with B/ε. Note that [11] already proved that a scaling with $\sqrt{B}/\varepsilon$ is possible and referred to this result as the quantum Chebyshev inequality. Our approach provides a different algorithm, with a simpler proof, that achieves essentially the same running time. We also note that [12] used the algorithm in [21] as a subroutine for estimating partition functions.

---

[13]To be specific, f will either correspond to $\exp(-\bar{\beta}_{i,i+1}H(x))$ for the random variable W_i or to $\exp(\bar{\beta}_{i,i+1}H(x))$ for the random variable V_i, where these quantities are defined in Section 2.2.

In order to prove this theorem, we will need a few theorems which we state first. We first state *non-destructive* amplitude estimation: a variant of amplitude estimation where the initial resource state is not destroyed in the quantum algorithm.[14]

Theorem C.3 (Non-destructive amplitude estimation [12, Theorem 6]). *Let $|\psi\rangle$ be an arbitrary quantum state and P an arbitrary projector. Let $R_\psi = 2|\psi\rangle\langle\psi| - I$. For every $t > 0$, there is a quantum algorithm $\mathcal{A}$ that starts in the initial state $|\psi\rangle$ and with probability at least $1 - \eta$, outputs an estimate $\hat{p}$ of $p = \langle\psi|P|\psi\rangle$ such that*

$$|\hat{p} - p| \leq 2\pi\frac{\sqrt{p(1-p)}}{t} + \frac{\pi^2}{t^2}.$$

Additionally, $\mathcal{A}$ restores $|\psi\rangle$ with probability $1 - \eta$. Overall $\mathcal{A}$ invokes the controlled reflection R_ψ $O(t \cdot \ln(1/\eta))$ many times.

An immediate corollary of this theorem is the following.

Corollary C.4. *Suppose a distribution $D : \Omega \to [0,1]$ and function $f : \Omega \to [0,1]$ satisfy B-bounded relative variance (for some $B \geq 1$). For every $t > 0$, there is a quantum algorithm $\mathcal{A}$ that, given access to a single copy of $|\psi_D\rangle$ and controlled reflection R_D, with probability $\geq 1 - \eta$ outputs an estimate $\hat{\mu}$ of μ such that*

$$|\mu - \hat{\mu}| \leq 2\pi\frac{\sqrt{\mu(1-\mu)}}{t} + \frac{\pi^2}{t^2}. \tag{38}$$

Additionally, $\mathcal{A}$ restores the state $|\psi_D\rangle$ with probability $1 - \eta$. Overall $\mathcal{A}$ invokes the controlled reflection R_D $O(t \cdot \ln(1/\eta))$ many times.

The algorithm in the corollary above is straightforward: first transform $|\psi_D\rangle$ into

$$|\psi_{D,f}\rangle = \sum_{x\in\Omega}\sqrt{D(x)}|x\rangle \otimes (\sqrt{f(x)}|1\rangle + \sqrt{1-f(x)}|0\rangle).$$

by performing a controlled rotation on an additional qubit. For a projector $P = \mathbb{I}\otimes|1\rangle\langle 1|$, observe that $\langle\psi_{D,f}|P|\psi_{D,f}\rangle = \sum_{x\in\Omega} D(x)f(x) = \mu$. So the algorithm can simply estimate μ using non-destructive amplitude estimation in Theorem C.3. The reflection around $|\psi_{D,f}\rangle$ can be realized with the help of the reflection R_D and the controlled qubit rotation.

The lemma below provides an important subroutine required for proving Theorem C.2. The lemma makes use of the non-destructive amplitude estimation routine. Our proof follows closely [21, Algorithm 2 and Lemma 4], but we choose a different way of partitioning Ω, which allows us to reduce the complexity of the algorithm.

Lemma C.5 (Bounded second moment). *Suppose a distribution $D : \Omega \to [0,1]$ and function $f : \Omega \to [0,\infty]$ satisfy B-bounded relative variance. Then, there is a quantum algorithm $\mathcal{A}$ that, given one copy of $|\psi_D\rangle$ and access to reflection operators R_D, with probability $\geq 1 - \eta$ outputs $\hat{\mu}$ such that*

$$|\mu - \hat{\mu}| \leq \varepsilon.$$

Additionally, $\mathcal{A}$ restores the initial state $|\psi_D\rangle$ with probability $\geq 1 - \eta$ and invokes the controlled reflection R_D $O\left(\sqrt{B}/\varepsilon \cdot \left(\ln(B/\varepsilon)\right)^{1.5} \cdot \ln\left(\ln(B/\varepsilon)/\eta\right)\right)$ many times.

---

[14]We remark that [4] also contains a procedure that performs non-destructive amplitude estimation.

Proof. Let k be a parameter to be determined later and consider the following sets

$$
\begin{aligned}
\Omega_0 &= \{x : 0 \le f(x) < 1\}, \\
\Omega_\ell &= \{x : 2^{\ell-1} \le f(x) < 2^\ell\} \quad \text{for} \quad \ell \in \{1, \dots, k\}, \\
\Omega_{k+1} &= \{x : 2^k \le f(x)\}.
\end{aligned}
$$

We write the expectation

$$
\mu = \sum_{\ell=0}^{k} 2^\ell \sum_{x \in \Omega_\ell} D(x) \frac{f(x)}{2^\ell} \;+\; \sum_{x \in \Omega_{k+1}} D(x) f(x). \tag{39}
$$

The second term in (39) can be bounded from above as follows

$$
\sum_{x \in \Omega_{k+1}} D(x) f(x) \le \frac{1}{2^k} \sum_{x \in \Omega} D(x) f(x)^2 \le \frac{B}{2^k}, \tag{40}
$$

where we used $2^k \le f(x)$ for $x \in \Omega_{k+1}$. We can "ignore" this term provided that k is sufficiently large so its contribution to μ becomes negligible.

Let us now focus on the double sum on the left in Eq.(39), which can be understood as a weighted sum of $k+1$ means. For each $\ell \in \{0, \dots, k\}$, we can estimate these means $\mu_\ell = \sum_{x \in \Omega_\ell} D(x) f(x) \cdot 2^{-\ell}$ with the help of Corollary C.4. To this end, we define the modified functions f_ℓ by setting $f_\ell(x) = f(x) \cdot 2^{-\ell}$ if $x \in \Omega_\ell$ and 0 otherwise. Let $\hat{\mu}_\ell$ be the estimates returned by amplitude estimation when applied to the probability distribution p and the functions f_ℓ. Our final estimate $\hat{\mu}$ of μ will then be

$$
\hat{\mu} = \sum_{\ell=0}^{k} 2^\ell \cdot \hat{\mu}_\ell,
$$

We now analyze the resulting estimation error. Let $t > 0$ be a parameter determined later and let $\eta' = \eta/(k+1)$. We use Corollary C.4 (with parameters t, η') to obtain $\hat{\mu}_\ell$ satisfying

$$
|\mu_\ell - \hat{\mu}_\ell| \le 2\pi \frac{\sqrt{\mu_\ell(1 - \mu_\ell)}}{t} + \frac{\pi^2}{t^2}. \tag{41}
$$

Note that the algorithm in Corollary C.4 succeeds with probability $1 - \eta'$, but since we invoke this corollary $k+1$ times for each $\ell \in \{0, \dots, k\}$, by union bound the success

probability is at least $1 - \eta' \cdot (k+1) = 1 - \eta$. We have

$$\sum_{\ell=0}^{k} 2^\ell |\mu_\ell - \hat{\mu}_\ell| \leq \sum_{\ell=0}^{k} 2^\ell \left(\frac{2\pi \sqrt{\mu_\ell(1-\mu_\ell)}}{t} + \frac{\pi^2}{t^2} \right)$$

$$\leq \sum_{\ell=1}^{k} 2^\ell \left(\frac{2\pi \sqrt{\mu_\ell}}{t} + \frac{\pi^2}{t^2} \right) + \frac{2\pi}{t} + \frac{\pi^2}{t^2}$$

$$= \sum_{\ell=1}^{k} \left(\frac{2\pi \sqrt{\sum_{x \in \Omega_\ell} D(x) 2^\ell f(x)}}{t} + \frac{2^\ell \pi^2}{t^2} \right) + \frac{2\pi}{t} + \frac{\pi^2}{t^2} \tag{42}$$

$$\leq \frac{2\pi}{t} \sum_{\ell=1}^{k} \sqrt{\sum_{x \in \Omega_\ell} D(x) f(x)^2} + \frac{2\pi}{t} + \frac{\pi^2 2^{k+1}}{t^2} \tag{43}$$

$$\leq \frac{2\pi}{t} \sqrt{k} \sqrt{\sum_{\ell=0}^{k} \sum_{x \in \Omega_\ell} D(x) f(x)^2} + \frac{2\pi}{t} + \frac{\pi^2 2^{k+1}}{t^2} \tag{44}$$

$$\leq \frac{2\pi \sqrt{k} \sqrt{B}}{t} + \frac{2\pi}{t} + \frac{\pi^2 2^{k+1}}{t^2}, \tag{45}$$

where the first inequality follows from Eq. (41), (42) used the definition of $\mu_\ell = \sum_{x \in \Omega_\ell} D(x) f(x) \cdot 2^{-\ell}$, Eq. (43) used $2^\ell \leq f(x)$ for $\ell \in \{1, \ldots, k\}$, Eq. (44) used Cauchy-Schwarz inequality and the final inequality used the relative variance upper bound (in the lemma statement).

Using this, we can bound $|\mu - \hat{\mu}|$ as follows:

$$|\mu - \hat{\mu}| \leq \sum_{\ell=0}^{k} 2^\ell |\mu_\ell - \hat{\mu}_\ell| + \sum_{x \in \Omega_{k+1}} D(x) f(x) \tag{46}$$

$$\leq \frac{2\pi \sqrt{k} \sqrt{B}}{t} + \frac{2\pi}{t} + \frac{\pi^2 2^{k+1}}{t^2} + \frac{B}{2^k}, \tag{47}$$

where we used the inequality in Eq. (45) and the last inequality used Eq. (40). We now set the parameters k and t as follows.

- Setting $k = \ln(2B/\varepsilon)$ ensures that the fourth error term $B/2^k$ in Eq. (46) is at most $\varepsilon/2$.

- Choosing $t \geq 4\pi\sqrt{B}/\varepsilon$ ensures that the third term $\pi^2 2^{k+1}/t^2$ in Eq. (46) is at most $\varepsilon/4$.

- Finally, choosing $t \geq 8\pi(\sqrt{B}\sqrt{\ln(2B/\varepsilon)} + 1)/\varepsilon$ ensures that the sum $2\pi\sqrt{k}\sqrt{B}/t + 2\pi/t$ of the first two terms in Eq. (46) is at most $\varepsilon/4$.

We see that $k = O(\ln(B/\varepsilon))$ and $t = O(\sqrt{B}/\varepsilon \cdot \sqrt{\ln(B/\varepsilon)})$. Overall, the number of times the reflection R_D needs to be invoked is proportional to

$$(k+1) \cdot t \cdot \ln\left((k+1)/\eta\right) = O\left(\sqrt{B}/\varepsilon \cdot \left(\ln(B/\varepsilon)\right)^{1.5} \cdot \ln\left(\ln(B/\varepsilon)/\eta\right)\right).$$

This proves the lemma statement. $\qquad\square$

We are now ready to prove the main theorem in this section.

Proof of Theorem C.2. The goal is to estimate $\mu = \sum_{x \in \Omega} D(x) f(x)$. We first use *classical* Chebyshev to obtain the following: take $16B$ qsamples $|\psi_D\rangle$, measure them in the computational basis to obtain $x \sim D$ and their mean produces a constant-factor estimate $\tilde{\mu}$ of μ such that

$$\mu/2 \leq \tilde{\mu} \leq 2\mu \tag{48}$$

with probability at least $3/4$. In order to boost this probability: it is well-known that we can decrease the failure probability to $1 - \eta/2$ simply by repeating the above process $\ln(2/\eta)$ many times and outputting the median of all the estimates.[15] The total number of required samples used here is $m = 16B \ln(2/\eta)$.

However recall that the goal is to obtain an ε-approximation of μ. Suppose we have a good estimate $\tilde{\mu}$ satisfying Eq. (48). Consider a rescaled function

$$f_{\mathrm{res}}(x) = f(x)/\tilde{\mu}.$$

Its mean μ_{res} and second moment ϕ_{res} satisfy

$$\mu_{\mathrm{res}} = \frac{\mu}{\tilde{\mu}} \quad \text{and} \quad \phi_{\mathrm{res}} = \frac{\phi}{\tilde{\mu}^2} \leq \frac{4\phi}{\mu^2} \leq 4B,$$

where we used the lower bound on $\tilde{\mu}$ in Eq. (48) and the assumption $\phi/\mu^2 \leq B$ in the theorem statement. We now invoke the subroutine in Lemma C.5 for the distribution D and function f_{res} (clearly (D, f_{res}) satisfies B-bounded variance): with probability $\geq 1 - \eta/2$, this subroutine produces an estimate $\hat{\mu}_{\mathrm{res}}$ satisfying

$$|\mu_{\mathrm{res}} - \hat{\mu}_{\mathrm{res}}| \leq \varepsilon/2,$$

which in particular also implies

$$|\mu - \tilde{\mu} \cdot \hat{\mu}_{\mathrm{res}}| = |\tilde{\mu} \cdot (\mu_{\mathrm{res}} - \hat{\mu}_{\mathrm{res}})| \leq \varepsilon/2 \cdot \tilde{\mu} \leq \varepsilon\mu$$

so $\hat{\mu} = \hat{\mu}_{\mathrm{res}}$ is the desired ε-relative estimate of μ. By a union bound (over classical Chebyshev step and the subroutine in Lemma C.5), the probability of obtaining this ε-estimate is $\geq 1 - \eta$. Overall the number of copies of $|\psi_D\rangle$ used is $16B \ln(2/\eta) + 1$ (the first term is because of classical Chebyshev inequality and the second term is from Lemma C.5). Moreover, note that since Lemma C.5 is non-destructive, one copy of $|\psi_D\rangle$ is restored in the process. Additionally, the algorithm in Lemma C.5 uses the reflection operator R_D

$$O\left(\sqrt{B}/\varepsilon \cdot \left(\ln(B/\varepsilon)\right)^{1.5} \cdot \ln\left(\ln(B/\varepsilon)/\eta\right)\right)$$

many times. $\qquad\square$

---

[15]For a proof of the powering lemma, we refer the interested reader to [17].