



Quantum Algorithm for Estimating Volumes of Convex Bodies

SHOUVANIK CHAKRABARTI and ANDREW M. CHILDS, University of Maryland, USA

SHIH-HAN HUNG, University of Texas at Austin, USA

TONGYANG LI, Peking University, China

CHUNHAO WANG, Pennsylvania State University, USA

XIAODI WU, University of Maryland, USA

Estimating the volume of a convex body is a central problem in convex geometry and can be viewed as a continuous version of counting. We present a quantum algorithm that estimates the volume of an n -dimensional convex body within multiplicative error ϵ using $\tilde{O}(n^3 + n^{2.5}/\epsilon)$ queries to a membership oracle and $\tilde{O}(n^5 + n^{4.5}/\epsilon)$ additional arithmetic operations. For comparison, the best known classical algorithm uses $\tilde{O}(n^{3.5} + n^3/\epsilon^2)$ queries and $\tilde{O}(n^{5.5} + n^5/\epsilon^2)$ additional arithmetic operations. To the best of our knowledge, this is the first quantum speedup for volume estimation. Our algorithm is based on a refined framework for speeding up simulated annealing algorithms that might be of independent interest. This framework applies in the setting of “Chebyshev cooling,” where the solution is expressed as a telescoping product of ratios, each having bounded variance. We develop several novel techniques when implementing our framework, including a theory of continuous-space quantum walks with rigorous bounds on discretization error. To complement our quantum algorithms, we also prove that volume estimation requires $\Omega(\sqrt{n} + 1/\epsilon)$ quantum membership queries, which rules out the possibility of exponential quantum speedup in n and shows optimality of our algorithm in $1/\epsilon$ up to poly-logarithmic factors.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**; **Quantum computation theory**; **Computational geometry**;

Additional Key Words and Phrases: Volume estimation, quantum algorithms, quantum query complexity, simulated annealing, Markov chain Monte Carlo methods

S. C., A. M. C., T. L., and X. W. were supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams program. A. M. C. also received support from the Army Research Office (MURI Award No. W911NF-16-1-0349), the Canadian Institute for Advanced Research, and the National Science Foundation (Grant No. CCF-1813814). S. H. H. received support from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Testbed Pathfinder program under Award No. DE-SC0019040. T. L. also received support from an IBM Ph.D. Fellowship and an NSF QISE-NET Triplet Award (Grant No. DMR-1747426). C. W. was supported by Scott Aaronson’s Vannevar Bush Faculty Fellowship from the U.S. Department of Defense. X. W. also received support from the National Science Foundation (Grants No. CCF-1755800 and No. CCF-1816695).

Authors’ addresses: S. Chakrabarti, A. M. Childs, and X. Wu, Department of Computer Science, Institute for Advanced Computer Studies, Joint Center for Quantum Information and Computer Science, University of Maryland, USA; emails: shouv@cs.umd.edu, amchilds@umd.edu, xwu@cs.umd.edu; S.-H. Hung, Department of Computer Science, University of Texas at Austin, USA; email: shung@cs.utexas.edu; T. Li, Center on Frontiers of Computing Studies and School of Computer Science, Peking University, China; email: tongyangli@pku.edu.cn; C. Wang, Pennsylvania State University, USA; email: cwang@psu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2643-6817/2023/05-ART20 \$15.00

<https://doi.org/10.1145/3588579>

ACM Reference format:

Shouvani Chakrabarti, Andrew M. Childs, Shih-Han Hung, Tongyang Li, Chunhao Wang, and Xiaodi Wu. 2023. Quantum Algorithm for Estimating Volumes of Convex Bodies. *ACM Trans. Quantum Comput.* 4, 3, Article 20 (May 2023), 60 pages.
<https://doi.org/10.1145/3588579>

1 INTRODUCTION

Estimating the volume of a convex body is a central challenge in theoretical computer science. Volume estimation is a basic problem in convex geometry and can be viewed as a continuous version of counting. Furthermore, algorithms for a generalization of volume estimation—namely, log-concave sampling—can be directly used to perform convex optimization, and hence can be widely applied to problems in statistics, machine learning, operations research, and so on. See the survey [72] for a more comprehensive introduction.

Volume estimation is a notoriously difficult problem. References [9, 26] proved that any *deterministic algorithm* that approximates the volume of an n -dimensional convex body within a factor of $n^{o(n)}$ necessarily makes exponentially many queries to a membership oracle for the convex body. Furthermore, References [25, 39, 40] showed that estimating the volume exactly (deterministically) is #P-hard, even for explicitly described polytopes.

Surprisingly, volumes of convex bodies can be approximated efficiently by *randomized algorithms*. Dyer, Frieze, and Kannan [24] gave the first polynomial-time randomized algorithm for estimating the volume of a convex body in $\mathbb{R}^n$. They present an iterative algorithm that constructs a sequence of convex bodies. The volume of the convex body of interest can be written as the telescoping product of the ratios of the volumes of consecutive convex bodies, and these ratios are estimated by **Markov chain Monte Carlo (MCMC)** methods via random walks inside these convex bodies. The algorithm in Reference [24] has complexity¹ $\tilde{O}(n^{23})$ with multiplicative error $\epsilon = \Theta(1)$ (the definition of multiplicative error is given in Equation (1.4)). Subsequent work [7, 23, 38, 48–50, 52] improved the design of the iterative framework and the choice of the random walks. The state-of-the-art algorithm for estimating the volume of a general convex body [36] uses $\tilde{O}(n^{3.5})$ queries to the membership oracle for the convex body (see “Formulation” below) and $\tilde{O}(n^{5.5})$ additional arithmetic operations.

It is natural to ask whether quantum computers can solve volume estimation even faster than classical randomized algorithms. Although there are frameworks with potential quantum speedup for simulated annealing algorithms in general, with volume estimation as a possible application [75], we are not aware of any previous quantum speedup for volume estimation. There are several reasons to develop such a result. First, quantum algorithms for volume estimation can be seen as performing a continuous version of quantum counting [12, 13], a key algorithmic technique with wide applications in quantum computing. Second, quantum algorithms for volume estimation can exploit quantum MCMC methods (e.g., References [57, 63, 74]), and a successful quantum volume estimation algorithm may illuminate the application of quantum MCMC methods in other scenarios. Third, there has been recent progress on quantum algorithms for convex optimization [5, 14], so it is natural to study the closely related task of estimating volumes of convex bodies.

Formulation. Given a convex set $K \subset \mathbb{R}^n$, we consider the problem of estimating its volume:

$$\text{Vol}(K) := \int_{x \in K} dx. \quad (1.1)$$

¹Throughout the article, $\tilde{O}$ omits factors in $\text{poly}(\log R/r, \log 1/\epsilon, \log n)$ where R and r are defined in Equation (1.2).

To get a basic sense about the location of K , we assume that it contains the origin. Furthermore, we assume that we are given inner and outer bounds on K , namely,

$$B_2^n(0, r) \subseteq K \subseteq B_2^n(0, R), \quad (1.2)$$

where $B_2^n(x, l)$ is the ball of radius l in ℓ_2 -norm centered at $x \in \mathbb{R}^n$. Denote $D := R/r$.

We consider the very general setting where the convex body K is only specified by an oracle. In particular, we have a *membership oracle* for K that determines whether a given $x \in \mathbb{R}^n$ belongs to K . The efficiency of volume estimation is then measured by the number of queries to the membership oracle (i.e., the *query complexity*) and the total number of other arithmetic operations. Note that the membership oracle is commonly used in convex optimization research (see, for example, Reference [29]). This model is not only general but also of practical interest. For instance, when K is a bounded convex polytope, the membership oracle can be efficiently implemented by checking if all its linear constraints are satisfied (see Reference [45]).

In the quantum setting, the membership oracle is a unitary operator O_K . Specifically, we have

$$O_K|x, 0\rangle = |x, \delta[x \in K]\rangle \quad \forall x \in \mathbb{R}^n, \quad (1.3)$$

where $\delta[P]$ is 1 if P is true and 0 if P is false.² In other words, we allow coherent superpositions of queries to the membership oracle. If the classical membership oracle can be implemented by an explicit classical circuit, then the corresponding quantum membership oracle can be implemented by a quantum circuit of about the same size. Therefore, the quantum query model provides a useful framework for understanding the quantum complexity of volume estimation.

1.1 Contributions

Our main result is a quantum algorithm for estimating volumes of convex bodies:

THEOREM 1.1 (MAIN THEOREM). *Let $K \subset \mathbb{R}^n$ be a convex set with $B_2^n(0, r) \subseteq K \subseteq B_2^n(0, R)$. Assume $0 < \epsilon < 1/2$. Then there is a quantum algorithm that returns a value $\text{Vol}(K)$ satisfying*

$$\frac{1}{1+\epsilon} \text{Vol}(K) \leq \widetilde{\text{Vol}}(K) \leq (1+\epsilon) \text{Vol}(K), \quad (1.4)$$

with probability at least $2/3$ using $\tilde{O}(n^3 + n^{2.5}/\epsilon)$ quantum queries to the membership oracle O_K (defined in Equation (1.3)) and $\tilde{O}(n^5 + n^{4.5}/\epsilon)$ additional arithmetic operations.

Note that arithmetic operations (e.g., addition, subtraction, multiplication, and division) can be in principle implemented by a universal set of quantum gates using the Solovay-Kitaev Theorem [22] up to a small overhead. In our quantum algorithm, the number of arithmetic operations is dominated by n -dimensional matrix-vector products computed in superposition for rounding the convex body (see Section 4.4).

To the best of our knowledge, this is the first quantum algorithm that achieves quantum speedup for this fundamental problem, compared to the classical state-of-the-art algorithm [21, 53] that uses $\tilde{O}(n^{3.5} + n^3/\epsilon^2)$ classical queries and $\tilde{O}(n^{5.5} + n^5/\epsilon^2)$ additional arithmetic operations.³ Furthermore, our quantum algorithm not only achieves a quantum speedup in query complexity but also in the number of arithmetic operations for executing the algorithm. This differs from previous quantum

²Here x can be approximated just as in the classical algorithms, such as with fixed-point numbers. Our algorithmic approach is robust under discretization (see Section 5), and our quantum lower bound holds even when x is stored with arbitrary precision (Section 6). We mostly assume for convenience that O_K operates on $x \in \mathbb{R}^n$, since this neither presents a serious obstacle nor conveys significant power.

³This is achieved by applying Reference [36] to preprocess the convex body to be well-rounded (i.e., $R/r = O(\sqrt{n})$) using $\tilde{O}(n^{3.5})$ queries and then applying Reference [21] using $\tilde{O}(n^3/\epsilon^2)$ queries to estimate the volume of the (well-rounded) convex body. The number of additional arithmetic operations has an overhead of $O(n^2)$ due to the affine transformation.

Table 1. Summary of Complexities of Volume Estimation, Where n Is the Dimension of the Convex Body, ϵ Is the Multiplicative Precision of Volume Estimation, and C_{MEM} Is the Cost of Applying the Membership Oracle Once

	Classical bounds	Quantum bounds (this article)
Query complexity	$\tilde{O}(n^{3.5} + n^3/\epsilon^2)$ [21, 36], $\tilde{\Omega}(n^2)$ [61]	$\tilde{O}(n^3 + n^{2.5}/\epsilon)$, $\Omega(\sqrt{n} + 1/\epsilon)$
Total complexity	$\tilde{O}((n^2 + C_{\text{MEM}}) \cdot (n^{3.5} + n^3/\epsilon^2))$ [21, 36]	$\tilde{O}((n^2 + C_{\text{MEM}}) \cdot (n^3 + n^{2.5}/\epsilon))$

Total complexity refers to the cost of the of queries plus the number of additional arithmetic operations.

algorithms for convex optimization [5, 14] where only the query complexity is improved, but the gate complexity is the same as that of the classical state-of-the-art algorithm [41, 42].

However, we prove in Section 6.1 that volume estimation with $\epsilon = \Theta(1)$ requires $\Omega(\sqrt{n})$ quantum queries to the membership oracle, ruling out the possibility of achieving superpolynomial quantum speedup for volume estimation. Classically, the best-known lower bound on the query complexity of volume estimation is $\tilde{\Omega}(n^2)$ due to Rademacher and Vempala [61], but there are technical difficulties to lift it to a quantum lower bound (see Section 1.2.3). For the dependence on $1/\epsilon$, we establish a quantum query lower bound of $\Omega(1/\epsilon)$, and the same argument shows a classical query lower bound of $\Omega(1/\epsilon^2)$ (see Section 6.2). As a result, our quantum algorithm in Theorem 1.1 achieves a provable quadratic quantum speedup in $1/\epsilon$ and is optimal in $1/\epsilon$ up to poly-logarithmic factors.

Technically, we refine a framework for achieving quantum speedups of simulated annealing algorithms, which might be of independent interest. Our framework applies to MCMC algorithms with cooling schedules that ensure each ratio in a telescoping product has bounded variance, an approach known as *Chebyshev cooling*. Furthermore, we propose several novel techniques when implementing this framework, including a theory of continuous-space quantum walks with rigorous bounds on discretization error, a quantum algorithm for nondestructive mean estimation, and a quantum algorithm with interlaced rounding and volume estimation of convex bodies (as described further in Section 1.2 below). In principle, our techniques apply not only to the integral of the identity function (as in Theorem 1.1), but could also be applied to any log-concave function defined on a convex body, following the approach in Reference [51].

We summarize our main results in Table 1.

1.2 Techniques

We now summarize the key technical aspects of our work.

1.2.1 Classical Volume Estimation Framework.

Volume estimation by simulated annealing. The volume of a convex body K can be estimated using simulated annealing. Consider the value

$$Z(a) := \int_K e^{-a\|x\|_2} dx, \quad (1.5)$$

where $\|x\|_2 := \sqrt{x_1^2 + \dots + x_n^2}$ is the ℓ_2 -norm of x . On the one hand, $Z(0) = \text{Vol}(K)$; on the other hand, because $e^{-\|x\|_2}$ decays exponentially fast with $\|x\|_2$, taking a large enough a ensures that the vast majority of $Z(a)$ concentrates near 0, so it can be well approximated by integrating on a small ball centered at 0. Therefore, a natural strategy is to consider a sequence $a_0 > a_1 > \dots > a_m$ with a_0 sufficiently large and a_m close to 0. We consider a simulated annealing algorithm that iteratively

changes a_i to a_{i+1} and estimates $\text{Vol}(K)$ by the telescoping product

$$\text{Vol}(K) \approx Z(a_m) = Z(a_0) \prod_{i=0}^{m-1} \frac{Z(a_{i+1})}{Z(a_i)}. \quad (1.6)$$

In the i th step, a random walk is used to sample the distribution over K with density proportional to $e^{-a_i \|x\|_2}$. Denote one such sample by X_i , and let $V_i := e^{(a_i - a_{i+1}) \|X_i\|_2}$. Then, we have

$$\mathbb{E}[V_i] = \int_K e^{(a_i - a_{i+1}) \|x\|_2} \frac{e^{-a_i \|x\|_2}}{Z(a_i)} dx = \int_K \frac{e^{-a_{i+1} \|x\|_2}}{Z(a_i)} dx = \frac{Z(a_{i+1})}{Z(a_i)}. \quad (1.7)$$

Therefore, each ratio $\frac{Z(a_{i+1})}{Z(a_i)}$ can be estimated by taking i.i.d. samples X_i , computing the corresponding V_i s, and taking their average.

We can analyze this volume estimation algorithm by considering its behavior at three levels:

- (1) High level: The algorithm follows the simulated annealing framework described above, where the volume is estimated by a telescoping product as in Equation (1.6).
- (2) Middle level: The number of i.i.d. samples used to estimate $\mathbb{E}[V_i]$ (a ratio in the telescoping product given by Equation (1.7)) is small. Intuitively, the annealing schedule should be slow enough that V_i has small variance.
- (3) Low level: The random walk converges fast so that we can take each i.i.d. sample of V_i efficiently.

Classical volume estimation algorithm. Our approach follows the classical volume estimation algorithm in Reference [53] (see also Section 4.1). At the high level, it is a simulated annealing algorithm that considers a cylinder $[0, 2R/r] \times K$ in $\mathbb{R}^{n+1}$ and a cone $C := \{x \in \mathbb{R}^{n+1} : x_0 \geq 0, \sum_{i=1}^n x_i^2 \leq x_0^2\}$. We denote the intersection of the $(n+1)$ -dimensional cylinder and $(n+1)$ -dimensional cone as an alternative convex body K' . This is called a *pencil construction*. It shares the same intuition as above, but replaces the integral Equation (1.5) by $Z(a) = \int_{K'} e^{-ax_0} dx$, because it is easier to calculate while can be directly used to estimate $\text{Vol}(K)$ when $a \approx 0$ by a standard Monte Carlo approach (see Lemma 4.1).

Without loss of generality, assume that $r = 1$. Lovász and Vempala [53] proved that if we take the sequence $a_0 > \dots > a_m$ where $a_0 = 2n$, $a_{i+1} = (1 - \frac{1}{\sqrt{n}})a_i$, and $m = \tilde{O}(\sqrt{n})$, then $Z(a_0) \approx \int_C e^{-a_0 x_0} dx$ and

$$\text{Var}[V_i^2] = O(1) \cdot \mathbb{E}[V_i]^2, \quad \forall i \in [m], \quad (1.8)$$

i.e., the variance of V_i is bounded by a constant multiple of the square of its expectation, which is small when a_i and a_{i+1} are close to each other. Such a simulated annealing schedule is known as *Chebyshev cooling* (see also Section 4.3.3). This establishes the middle-level requirement of the simulated annealing framework. Furthermore, Reference [53] proves that the product of the average of $\tilde{O}(\sqrt{n}/\epsilon^2)$ i.i.d. samples of V_i for all $i \in [m]$ gives an estimate of $\text{Vol}(K')$ within multiplicative error ϵ with high success probability.

At the low level, Reference [53] uses a *hit-and-run walk* to sample X_i . In this walk, starting from a point p , we uniformly sample a line ℓ through p and move to a random point along the chord $\ell \cap K$ with density proportional to e^{-ax_0} (see Section 2.4 for details). Reference [52] analyzes the convergence of the hit-and-run walk, proving that it converges to the distribution over K with density proportional to e^{-ax_0} within $\tilde{O}(n^3)$ steps, assuming that K is *well-rounded* (i.e., $R/r = O(\sqrt{n})$).

Finally, Reference [53] constructs an affine transformation that transforms a general K to be well-rounded with $\tilde{O}(n^4)$ classical queries to its membership oracle, hence removing the

constraint of the previous steps that K be well-rounded. Because the affine transformation is an n -dimensional matrix-vector product, this introduces an overhead of $O(n^2)$ in the number of arithmetic operations.

Overall, the algorithm has $\tilde{O}(\sqrt{n})$ iterations, where each iteration takes $\tilde{O}(\sqrt{n}/\epsilon^2)$ i.i.d. samples, and each sample takes $\tilde{O}(n^3)$ steps of the hit-and-run walk. In total, the query complexity is

$$\tilde{O}(\sqrt{n}) \cdot \tilde{O}(\sqrt{n}/\epsilon^2) \cdot \tilde{O}(n^3) = \tilde{O}(n^4/\epsilon^2). \quad (1.9)$$

The number of additional arithmetic operations is $\tilde{O}(n^4/\epsilon^2) \cdot O(n^2) = \tilde{O}(n^6/\epsilon^2)$ due to the affine transformation for rounding the convex body.

1.2.2 Quantum Algorithm for Volume Estimation. It is natural to consider a quantum algorithm for volume estimation following the classical framework in Section 1.2.1. A naive attempt might be to develop a quantum walk that achieves a generic quadratic speedup of the mixing time. However, this is unfortunately difficult to achieve in general. Quantum walks are unitary processes that do not converge to stationary distributions in the classical sense. As a result, alternative and indirect quantum analogues of mixing properties of Markov chains have been proposed and studied (see Section 1.3.2 for more detail). None of these methods provide a direct replacement for classical mixing, and we cannot directly apply them in our context.

Instead, we adapt one of the frameworks proposed in Reference [74]. To give a quantum speedup for volume estimation by this method, we address the following additional technical challenges:

- **Quantum walks in continuous space:** Quantum walks have mainly been studied in discrete spaces [55, 70], and we need to understand how to define a quantum counterpart of the hit-and-run walk.
- **Quantum mean estimation:** Quantum counting [12] is a general tool for estimating a probability $p \in [0, 1]$ with quadratic quantum speedup compared to classical sampling. However, estimating the mean of an unbounded random variable with a quantum version of Chebyshev concentration requires more advanced tools.
- **Rounding:** Classically, rounding a general convex body takes $\tilde{O}(n^4)$ queries [53], more expensive than volume estimation of a well-rounded body using $\tilde{O}(n^3/\epsilon^2)$ queries [21]. To achieve an overall quantum speedup, we also need to give a fast quantum algorithm for rounding convex bodies.
- **Error analysis of the quantum hit-and-run walk:** We must bound the error incurred when implementing the quantum walk on a digital quantum computer with finite precision. Existing classical error analyses (e.g., Reference [28]) do not automatically cover the quantum case.

We develop several novel techniques to resolve these issues:

Theory of continuous-space quantum walks (Section 3). Our first technical contribution is to develop a quantum implementation of the low-level framework, i.e., to replace the classical hit-and-run walk by a *quantum hit-and-run walk*. However, although quantum walks in discrete spaces have been well studied (see, for example, Reference [55, 70]), we are not aware of comparable results that can be used to analyze spectral properties and mixing times of quantum walks in continuous space. Here, we describe a framework for continuous-space quantum walks that can be instantiated to give a quantum version of the hit-and-run walk. In particular, we formally define such walks and analyze their spectral properties, generalizing Szegedy's theory [70] to continuous spaces (Section 3.1). We also show a direct correspondence between the stationary distribution of a classical walk and a certain eigenvector of the corresponding quantum walk (Section 3.2).

It is, in principle, possible to directly analyze a discretized version of the hit-and-run walk, and we carefully consider the effects of discretization in Section 5. However, it is more convenient to perform the bulk of the analysis in continuous space. Geometric random walks such as the hit-and-run walk are highly non-local and there is a non-zero probability of transition between any two points in K . Therefore, simple discrete walks such as the graph walk on a grid do not simulate it well, and a true discretization corresponds to a complete graph with a complicated weight structure. Furthermore, the essential properties of the walk, including its roundedness, conductance, spectral gap, and convergence, rely on the convexity of K , which has no simple analog in the discrete setting. A completely discrete analysis would require us to track complicated, but ultimately insignificant, corrections to each of these parameters, significantly obfuscating the main ideas. For these reasons, we make the choice to present the main algorithm in the continuous setting and defer the analysis of discretization issues to Section 5. We note that this choice is consistent with classical literature on convex geometry, and most of the classical algorithms referenced in this work are analyzed purely in the continuous setting.

Quantum volume estimation algorithm via simulated annealing (Section 4.2). Having described a quantum hit-and-run walk, the next step is to understand the high-level simulated annealing framework. As mentioned above, it is nontrivial to directly prepare stationary states of quantum walks. In this article, we follow a quantum MCMC framework proposed by Reference [74] that can prepare stationary states of quantum walks by simulated annealing (see Section 2.2). In this framework, we have a sequence of slowly varying Markov chains, and the stationary state of the initial Markov chain can be efficiently prepared. In each iteration, we apply fixed-point amplitude amplification (more specifically, the $\pi/3$ -amplitude amplification) of the quantum walk operator [32] due to Grover to transform the current stationary state to the next one; compared to classical slowly varying Markov chains, the convergence rate of such quantum procedure is *quadratically better in spectral gap*.

Our **main technical contribution** is to show how to adapt the Chebyshev cooling schedule in Reference [53] to the quantum MCMC framework in Reference [74] using our quantum hit-and-run walk. The conductance lower bound together with the classical $\tilde{O}(n^3)$ mixing time imply that we can perform one step of $\pi/3$ -amplitude amplification using $\tilde{O}(n^{1.5})$ queries to O_K . Furthermore, the inner product between consecutive stationary states is a constant. These two facts ensure that the stationary state in each iteration can be prepared with $\tilde{O}(n^{1.5})$ queries to the membership oracle O_K . The total number of iterations is still $\tilde{O}(\sqrt{n})$, as in the classical case.

Quantum algorithm for nondestructive mean estimation (Section 4.3.3). In the next step, we consider how to estimate each ratio in the telescoping product at the middle level. The basic tool is quantum counting [12], which estimates a probability $p \in [0, 1]$ with error ϵ and high success probability using $O(1/\epsilon)$ quantum queries, a quadratic speedup compared to the classical complexity $O(1/\epsilon^2)$. However, in our case, we need to estimate the expectation of a random variable with bounded variance. We use the “quantum Chebyshev inequality” developed in Reference [33], which truncates the random variable with reasonable upper and lower bounds and then reduces to quantum counting; see Section 2.3.⁴ Compared to the classical counterpart, it achieves quadratic speedup in the dependencies on both variance and multiplicative error.

There is an additional technical difficulty in quantum simulated annealing: classically, it is implicitly assumed that in the $(i + 1)$ st iteration, we have samples to the stationary distribution in the

⁴A related technique is the quantum Monte Carlo method of Montanaro [57]. Here, we use Reference [33] for two reasons: first, it has the advantage of handling multiplicative instead of additive errors, which is appropriate for estimating the telescoping ratios. Second, its quantum algorithm is based on amplitude estimation and hence can readily be made nondestructive, as discussed below.

i th iteration. Applying existing quantum mean estimation techniques to the quantum stationary state in the i th iteration would ruin that state and make it hard to use in the subsequent $(i + 1)$ st iteration. To resolve this issue, we estimate the mean *nondestructively* in the quantum Chebyshev inequality while keeping its quadratic speedup in the error dependence using a nondestructive amplitude estimation technique developed in Reference [34]. Nondestructive mean estimation relies on the following observation: applying amplitude estimation on a state $|\psi\rangle$ results with high probability in the measurement collapsing to one of two states $|\psi_+\rangle, |\psi_-\rangle$ with constant overlap with ψ . The algorithm repeatedly projects these states onto $|\psi\rangle$: if the projection is successful then the state is restored, otherwise amplitude estimation can be performed again to obtain $|\psi_+\rangle, |\psi_-\rangle$ and the projection can be repeated. Due to the constant overlap, $\text{poly}(\log(\delta^{-1}))$ repetitions suffice to ensure that at least one of the projections succeeds with probability δ . It remains to implement the required projection efficiently: We show how this can be accomplished using quantum walk operators corresponding to the Markov Chains in the MCMC framework; see Section 4.3.4.

In our quantum volume estimation algorithm, we apply the quantum Chebyshev inequality under the same compute-uncompute procedure. This gives a quadratic speedup in ϵ^{-1} when estimating the $\mathbb{E}[V_i]$ in Equation (1.7), so that $\tilde{O}(\sqrt{n}/\epsilon)$ copies of the stationary state suffice (see Lemma 4.3).

Quantum algorithm for volume estimation with interlaced rounding (Section 4.4). The stationary states of the quantum hit-and-run walk can be prepared with $\tilde{O}(n^{1.5})$ queries to O_K only when the corresponding density functions are *well-rounded*, i.e., every level set with probability μ contains a ball of radius μr and the variance of the density is bounded by R^2 , where $R/r = O(\sqrt{n})$. (When the density function is uniform in K , this definition of well-roundedness reduces to that in Footnote 3. The definition of level sets is the same as in Reference [53].) It remains to show how to ensure that the convex body is well-rounded.

We follow a classical framework for directly rounding logconcave densities [51]. The rounding is interlaced with the volume estimation algorithm, so that in each iteration of the simulated annealing framework, we use some of the samples to calculate an affine transformation that makes the next stationary state well-rounded. This ensures that the quantum hit-and-run walk continues to take only $\tilde{O}(n^{1.5})$ queries for each sample. Our algorithm maintains $\tilde{O}(n)$ extra quantum states for rounding, and the quantum hit-and-run walk is used to transform them from one stationary distribution to the next. In each iteration, we use a nondestructive measurement to sample the required affine transformation. With $\tilde{O}(\sqrt{n})$ iterations this results in an additional $\tilde{O}(\sqrt{n}) \cdot \tilde{O}(n) \cdot \tilde{O}(n^{1.5}) = \tilde{O}(n^3)$ cost for rounding.

We also show that this framework can be used as a preprocessing step that puts the convex body itself in well-rounded position (i.e., $B_2(0, r) \subseteq K \subseteq B_2(0, R)$ with $R/r = O(\sqrt{n})$) using $\tilde{O}(n^3)$ quantum queries. Putting a convex body in well-rounded position implies that several random walks used in simulated annealing algorithms (including the hit-and-run walk) mix fast without the need for further rounding. Therefore, as an alternative, we could preprocess the convex body to be well-rounded and then apply the simulated annealing algorithm to obtain a volume estimation algorithm that uses $\tilde{O}(n^3 + n^{2.5}/\epsilon)$ quantum queries.

Error analysis of discretized hit-and-run walks (Section 5). Although we defined quantum hit-and-run walks abstractly in Section 3, implementing a continuous-space quantum walk on a digital quantum computers will lead to discretization error, and the error analysis of classical walks in a discrete space approximating $\mathbb{R}^n$ (such as Reference [28]) does not automatically apply to the quantum counterpart. To ensure that discretization errors do not affect a realistic implementation of our algorithm, in Section 5, we propose a *discretized hit-and-run walk* and provide rigorous bounds on the discretization error.




Fig. 1. The structure of our quantum volume estimation algorithm. The four purple frames represent the four novel techniques that we propose, the yellow frame represents the known technique from Reference [32], and the green frame at the center represents our quantum algorithm.

Summary. Our quantum volume estimation algorithm can be summarized as follows.

- (1) High level: The quantum algorithm follows a simulated annealing framework using a quantum MCMC method [74], where the volume is estimated by a telescoping product (as in Equation (1.6)); the number of iterations is $\tilde{O}(\sqrt{n})$.
- (2) Middle level: We estimate the $\mathbb{E}[V_i]$ in Equation (1.7), a ratio in the telescoping product, using the nondestructive version of the quantum Chebyshev inequality [33]. This takes $\tilde{O}(\sqrt{n}/\epsilon)$ implementations of the quantum hit-and-run walk operators.
- (3) Low level: If the convex body K is well-rounded (i.e., $R/r = O(\sqrt{n})$), then each quantum hit-and-run walk operator can be implemented using $\tilde{O}(n^{1.5})$ queries to the membership oracle O_K in Equation (1.3).

Finally, we give a quantum algorithm that interlaces rounding and volume estimation of the convex body, using an additional $\tilde{O}(n^{2.5})$ quantum queries to O_K in each iteration. Because the affine transformation is an n -dimensional matrix-vector product, it introduces an overhead of $O(n^2)$ in the number of arithmetic operations (just as in the classical rounding algorithm).

Overall, our quantum volume estimation algorithm has $\tilde{O}(\sqrt{n})$ iterations. Each iteration implements $\tilde{O}(\sqrt{n}/\epsilon)$ quantum hit-and-run walks, and each quantum hit-and-run walk uses $\tilde{O}(n^{1.5})$ queries; there is also a cost of $\tilde{O}(n^{2.5})$ for rounding. Thus, the quantum query complexity is

$$\tilde{O}(\sqrt{n}) \cdot \left(\tilde{O}(\sqrt{n}/\epsilon) \cdot \tilde{O}(n^{1.5}) + \tilde{O}(n^{2.5}) \right) = \tilde{O}(n^3 + n^{2.5}/\epsilon). \quad (1.10)$$

The number of additional arithmetic operations is $\tilde{O}(n^3 + n^{2.5}/\epsilon) \cdot O(n^2) = \tilde{O}(n^5 + n^{4.5}/\epsilon)$ due to the affine transformations for interlaced rounding of the convex body.

Figure 1 summarizes our techniques. The volume estimation and interlaced rounding algorithms are given as Algorithms 3 and 4, respectively, in Section 4.

1.2.3 Quantum Lower Bounds (Section 6). The classical state-of-the-art query lower bound for volume estimation is a $\tilde{\Omega}(n^2)$ bound for n -dimensional parallelpipeds [61]. Their main tool is the dispersion of random determinants (of the matrix that characterizes the parallelpiped) and dispersion of the distance of a random point from a convex body. However, there is no such a counterpart in the quantum setting.

Nevertheless, we prove that volume estimation requires $\Omega(\sqrt{n})$ quantum queries to the membership oracle, ruling out the possibility of exponential quantum speedup (see Theorem 6.1). We establish this by a reduction to search: for a hyper-rectangle $K = \times_{i=1}^n [0, 2^{s_i}]$ specified by a binary string $s = (s_1, \dots, s_n) \in \{0, 1\}^n$ with $|s| = 0$ or 1, we prove that a membership query to K can be

simulated by a query to s . Thus, since $\text{Vol}(K) = 2$ if and only if $|s| = 1$, the $\Omega(\sqrt{n})$ quantum lower bound on search [10] applies to volume estimation.

In addition, we prove that volume estimation requires $\Omega(1/\epsilon)$ quantum queries (see Theorem 6.2), which means that our quantum algorithm is optimal in $1/\epsilon$ up to poly-logarithmic factors. The idea is to construct a convex body whose volume estimation reduces to the Hamming distance problem with known tight quantum query complexity [59]. To be more specific, we consider the n -dimensional unit hypercube and attach “hyperpyramids” to its faces, such that its central axis passes through the center of the hypercube. We show that adding or deleting any hyperpyramid of volume $1/2n$ does not influence the convexity of the convex body, and calculating the volume of the body reveals the Hamming weight of a binary string that encodes the presence or absence of the hyperpyramids.

1.3 Related Work

While our article gives the first quantum algorithm for volume estimation, classical volume estimation algorithms have been well-studied, as we review in Section 1.3.1. Our quantum algorithm builds upon quantum analogs of Markov chain Monte Carlo methods that we review in Section 1.3.2.

1.3.1 Classical Volume Estimation Algorithms. There is a rich literature on classical algorithms for estimating volumes of convex bodies (e.g., see the surveys in References [44, 72]). The general approach is to consider a sequence of random walks inside the convex body K whose stationary distributions converge quickly to the uniform distribution on K . Applying simulated annealing to this sequence of walks (as in Section 1.2), the volume of K can be approximated by a telescoping product.

The first polynomial-time algorithm for volume estimation was given by Reference [24]. It uses a *grid walk* in which the convex body K is approximated by a grid mesh K_{grid} of spacing δ (i.e., K_{grid} contains the points in K whose coordinates are integer multiples of δ). The walk proceeds as follows:

- (1) Pick a grid point y uniformly at random from the neighbors of the current point x .
- (2) If $y \in K_{\text{grid}}$, then go to y ; else, stay at x .

Dyer, Frieze, and Vempala [24] proved that for a properly chosen δ , the grid walk converges to the uniform distribution on K_{grid} in $\tilde{O}(n^{23})$ steps, and that $\delta^n |K_{\text{grid}}|$ is a good approximation of $\text{Vol}(K)$ (in the sense of Equation (1.4)). Subsequently, more refined analysis of the grid walk improved its cost to $\tilde{O}(n^8)$ [7, 23, 49]. However, this is still inefficient in practice.

Intuitively, the grid walk converges slowly, because each step only moves locally in K . Subsequent work improved the complexity by considering other types of random walk. These improvements mainly use two types of walk: the *hit-and-run walk* and the *ball walk*. In this article, we use the hit-and-run walk (see also Section 2.4). It is initialized at the original point, and then the walk proceeds as follows:

- (1) Pick a uniformly distributed random line ℓ through the current point p .
- (2) Move to a uniformly random point along the chord $\ell \cap K$.

Smith [66] proved that the stationary distribution of the hit-and-run walk is the uniform distribution on K . Regarding the convergence of the hit-and-run walk, Reference [48] showed that it mixes in $\tilde{O}(n^3)$ steps from a warm start (see Definition 4.1) after appropriate preprocessing, and Reference [52] subsequently proved that the hit-and-run walk mixes rapidly from any interior starting point (see also Theorem 2.4). Under the simulated annealing framework, the hit-and-run

Table 2. Summary of Classical Methods for Estimating the Volume of a Convex Body $K \subset \mathbb{R}^n$ When $\epsilon = \Theta(1)$, Where R, r Are the Radii of the Balls Centered at the Origin That Contain and Are Contained by the Convex Body, Respectively

Method	State-of-the-art query complexity	Restriction on the convex body
Grid walk	$\tilde{O}(n^8)$ [23]	General ($R/r = \text{poly}(n)$)
Hit-and-run walk	$\tilde{O}(n^4)$ [51, 53]	General ($R/r = \text{poly}(n)$)
Ball walk	$\tilde{O}(n^3)$ [20, 21]	Well-rounded ($R/r = O(\sqrt{n})$)

walk gives the state-of-the-art volume estimation algorithm with query complexity $\tilde{O}(n^4)$ [51, 53]. Our quantum volume estimation algorithm can be viewed as a quantization of this classical hit-and-run algorithm.

Given a radius parameter δ , the ball walk is defined as follows:

- (1) Pick a uniformly random point y from the ball of radius δ centered at the current point x .
- (2) If $y \in K$, then go to y ; else, stay at x .

Lovász and Simonovits [50] proved that the ball walk mixes in $\tilde{O}(n^6)$ steps. Kannan et al. [38] subsequently improved the mixing time to $\tilde{O}(n^3)$ starting from a warm start, giving a total query complexity of $\tilde{O}(n^5)$ for the volume estimation problem.

The analysis of the ball walk relies on a central conjecture in convex geometry, the **Kannan-Lovász-Simonovits (KLS)** conjecture (see Reference [44]). The KLS conjecture states that the Cheeger constant of any log-concave density is achieved to within a universal, dimension-independent constant factor by a hyperplane-induced subset. Here the Cheeger constant ψ is the minimum ratio between the measure of the boundary of a subset and the measure of the subset or its complement, whichever is smaller. In the special case when the convex body is well-rounded (i.e., $R/r = O(\sqrt{n})$), Cousins and Vempala [20, 21] proved the KLS conjecture for Gaussian distributions. Thus, they established a volume estimation algorithm with query complexity $\tilde{O}(n^3)$ in the well-rounded case. The state-of-the-art algorithm for rounding convex bodies has complexity $\tilde{O}(n^3\psi^2)$ [36], and in that work the authors used the best known upper bound $\psi = O(n^{1/4})$ [43] at that time. Combined with Reference [21], this results in a classical algorithm for n -dimensional volume estimation with multiplicative error ϵ using $\tilde{O}(n^{3.5} + n^3/\epsilon^2)$ queries.

More recently, a breakthrough on the KLS conjecture by Chen [16] showed that $\psi = O(d^{o(1)})$. Chen's result improves the convex body rounding bound $\tilde{O}(n^3\psi^2)$ in Reference [36] to $\tilde{O}(n^{3+o(1)})$, resulting in a classical algorithm for volume estimation with query complexity $\tilde{O}(n^{3+o(1)} + n^3/\epsilon^2)$ (again using [21]). This breakthrough appeared after the initial version of our article was completed, but straightforwardly improves our analysis. In particular, both Theorem 2.7 of Reference [36] and Theorem 2.5 below come from the observation that the walks converge in $\tilde{O}(n^2\psi^2)$ steps (see also Reference [38]), which means that in Lemma 4.8 the number of calls to the quantum walk operators in each iteration is essentially $\tilde{O}(n\psi)$. Therefore, Chen's result improves Lemma 4.8 by a factor of $n^{0.5-o(1)}$, implying that our quantum algorithm has query complexity $\tilde{O}(n^{2.5+o(1)} + n^{2+o(1)}/\epsilon)$, which is still better than the classical counterpart.

Table 2 summarizes classical algorithms for volume estimation.

1.3.2 Quantum Markov Chain Monte Carlo Methods. The performance of MCMC methods is determined by the rate of convergence to their stationary distributions (i.e., the mixing time). Suppose we have a reversible, ergodic Markov chain with unique stationary distribution π . Let π_k

denote the distribution obtained by applying the Markov chain for k steps from some arbitrary initial state. It is well-known (see, for example, Reference [46]) that $O(\frac{1}{\Delta} \log(1/(\epsilon \min_x \pi(x))))$ steps suffice to ensure $\|\pi_k - \pi\|_1 \leq \epsilon$, where Δ is the spectral gap of the Markov chain.

Many authors have studied quantum analogs of Markov chains (in both continuous [27] and discrete [1, 4, 70] time) and their mixing properties. While a quantum walk is a unitary process and hence does not converge to a stationary distribution, one can define notions of quantum mixing time by choosing the number of steps at random or by adding decoherence [1, 3, 4, 15, 17, 62, 63], and compare them to the classical mixing time. Note that distribution sampled by such a process may or may not be the same as the stationary distribution π of the corresponding classical Markov process, depending on the structure of the process and the notion of mixing. It is also natural to ask how efficiently we can prepare a quantum state close to $|\pi\rangle := \sum_x \sqrt{\pi_x} |x\rangle$, which can be viewed as a “quantum sample” from π . However, it is unclear how to do this efficiently in general, even in cases where a corresponding classical Markov process mixes quickly; in particular, a generic quantum algorithm for this task could be used to solve graph isomorphism [2, Section 8.4].

It is also possible to achieve quantum speedup of MCMC methods by not demanding speedup of the mixing time of each separate Markov chain, but only for the procedure as a whole. In particular, MCMC methods are often implemented by simulated annealing algorithms where the final output is a telescoping product of values at different temperatures. From this perspective, Somma et al. [11, 67, 68] used quantum walks to accelerate classical simulated annealing processes by exploiting the quantum Zeno effect, using measurements implemented by phase estimation of the quantum walk operators of these Markov chains. References [71, 77] also introduced how to implement Metropolis sampling on quantum computers.

Our quantum volume estimation algorithm is most closely related to work of Wocjan and Abeyesinghe [74], which achieves complexity $\tilde{O}(1/\sqrt{\Delta})$ for preparing the final stationary distribution of a sequence of slowly varying Markov chains, where Δ is the minimum of their spectral gaps. Their quantum algorithm transits between the stationary states of consecutive Markov chains by $\pi/3$ -amplitude amplification [32], which is implemented by amplitude estimation with $\tilde{O}(1/\sqrt{\Delta})$ implementations of the quantum walk operators of these Markov chains (see Section 2.2 for more details).

Our simulated annealing procedure preserves the slowly varying property, so we adopt the framework of Reference [74] in our algorithm for volume estimation (see Section 4.3.2). We develop several novel techniques (described in Section 1.2) that allow us to implement the steps of this framework efficiently. Note that the slowly varying property also facilitates other frameworks that give efficient adiabatic [2] or circuit-based [60] quantum algorithms for generating quantum samples of the stationary state.

Previous work has mainly applied these quantum simulated annealing algorithms to estimating partition functions of discrete systems. Given an inverse temperature $\beta > 0$ and a classical Hamiltonian $H: \Omega \rightarrow \mathbb{R}$, where Ω is a finite space, the goal is to estimate the partition function

$$Z(\beta) := \sum_{x \in \Omega} e^{-\beta H(x)} \quad (1.11)$$

within multiplicative error $\epsilon > 0$. Wocjan et al. [75] gave a quantum algorithm that achieves quadratic quantum speedup with respect to both mixing time and accuracy.

The classical algorithm that Reference [75] quantizes uses $\tilde{O}(\log |\Omega|)$ annealing steps to ensure that each ratio $Z(\beta_{i+1})/Z(\beta_i)$ is bounded. In fact, it is possible to relax this requirement and use a cooling schedule with only $\tilde{O}(\sqrt{\log |\Omega|})$ steps such that the variance of each ratio is bounded, so its mean can be well-approximated by Chebyshev’s inequality; this is exactly the Chebyshev cooling technique [69] introduced in Section 1.2 (see also Section 4.3.3). Montanaro [57] improved

upon [75] using Chebyshev cooling. Harrow and Wei [34] further quadratically improved the spectral gap dependence of the estimation of the partition function. More recently, Arunachalam et al. [19] further simplifies the Chebyshev cooling schedule and proposed faster quantum algorithms for estimating Gibbs partition functions, and Cornelissen and Hamoudi [19] proposed a sublinear-time quantum algorithm for approximating partition functions.

1.4 Open Questions

This work leaves several natural open questions for future investigation. In particular:

- Can we improve the complexity of our quantum volume estimation algorithm? Possible improvements might result from designing a shorter simulated annealing schedule, giving better analysis of the conductance of the hit-and-run walk, or even using other types of walks. In particular, an important step of our quantum algorithm is convex body rounding, where we quantize the classical algorithm by Lovasz and Vempala [53] with query complexity $\tilde{O}(n^4)$, but the state-of-the-art result in Reference [36] has query complexity $\tilde{O}(n^{3.5})$. It is natural to ask whether their algorithm fits into the interlaced structure of Algorithm 4 and achieves even better quantum query/gate complexities. Another natural direction is to combine our approach with recent quantum algorithms using shorter annealing schedules [6, 8, 19, 34]. One might also improve volume estimation using quantum algorithms for logconcave sampling, such as those recently proposed in Reference [18].
- Our algorithm directly quantizes a classical framework that has many moving pieces: the pencil construction, rounding the convex body, geometric random walks, and Monte Carlo estimation of the telescoping ratios. A natural question is whether quantum algorithms with better complexity can be obtained by using a simpler framework, or by combining some of these parts. For example, our treatment differs from that of Lovasz and Vempala [53] in that the rounding of the convex body is interlaced with the simulated annealing. This allows us to obtain a $\tilde{O}(n^3)$ overhead from rounding, whereas a direct quantization of Reference [53] would result in a $\tilde{O}(n^{3.5})$ overhead (which is not better than the best classical algorithm).
- Can we prove better quantum query lower bounds on volume estimation? Note that classically there is an $\tilde{\Omega}(n^2)$ query lower bound [61].
- Can we give faster quantum algorithms for volume estimation in some special circumstances? For instance, volume estimation of well-rounded convex bodies only takes $\tilde{O}(n^3)$ classical queries [21] (see also Section 1.3.1), and the volume of polytopes with m faces can be estimated with only $\tilde{O}(mn^{2/3})$ classical queries [45]. Specifically, it is natural to ask whether the ball walk in Reference [21] or the **Riemannian Hamiltonian Monte Carlo (RHMC)** method in Reference [45] can be implemented by continuous-space quantum walks (and their discretizations).
- Can we apply our simulated annealing framework to solve other problems? As a concrete example, it may be of interest to check whether our framework can recover the results of Reference [34] on estimating partition functions in counting problems.

Organization. We review necessary background in Section 2. We describe the theory of continuous-space quantum walks in Section 3. In Section 4, we first review the classical state-of-the-art volume estimation algorithm in Section 4.1, and then we give our quantum algorithm for estimating volumes of well-rounded convex bodies in Section 4.2. The proofs of our quantum algorithms are given in Section 4.3, and the quantum algorithm for rounding convex bodies is given in Section 4.4. The details of our discretized hit-and-run walk are given in Section 5, and we conclude with our quantum lower bound on volume estimation in Section 6.

2 PRELIMINARIES

We summarize necessary tools used in this article as follows.

2.1 Classical and Quantum Walks

A Markov chain over a finite state space Ω is a sequence of random variables $X_0, X_1, \dots$ such that for each $i \in \mathbb{N}$, the probability of transition to the next state $y \in \Omega$,

$$\Pr[X_{i+1} = y \mid X_i = x, X_{i-1} = x_{i-1}, \dots, X_0 = x_0] = \Pr[X_{i+1} = y \mid X_i = x] =: p_{x \rightarrow y},$$

only depends on the present state $x \in \Omega$. The Markov chain can be represented by the transition probabilities $p_{x \rightarrow y}$ satisfying $\sum_y p_{x \rightarrow y} = 1$. For each $i \in \mathbb{N}$, we denote by π_i the distribution over Ω with density $\pi_i(x) = \Pr[X_i = x]$ for all $x \in \Omega$. A *stationary distribution* π satisfies $\sum_{x \in \Omega} p_{x \rightarrow y} \pi(x) = \pi(y)$. A Markov chain is *reversible* if it has a stationary distribution π such that $\pi(x)p_{x \rightarrow y} = \pi(y)p_{y \rightarrow x}$ for all $x, y \in \Omega$. The *conductance* of a reversible Markov chain is defined as

$$\Phi := \inf_{S \subseteq \Omega} \frac{\sum_{x \in S} \sum_{y \in \Omega \setminus S} \pi(x) p_{x \rightarrow y}}{\min\{\sum_{x \in S} \pi(x), \sum_{x \in \Omega \setminus S} \pi(x)\}}. \quad (2.1)$$

Intuitively, we partition Ω into two sets, S and $\Omega \setminus S$. The conductance characterizes the infimum of the ratio between the probability of transitioning from S to $\Omega \setminus S$ under the stationary distribution π and the smaller of the probability mass of S and $\Omega \setminus S$ under π .

The theory of discrete-time quantum walks has also been well developed. Given a classical reversible Markov chain on Ω with transition probability p , we define a unitary operator U_p on $\mathbb{C}^{|\Omega|} \otimes \mathbb{C}^{|\Omega|}$ such that

$$U_p |x\rangle |0\rangle = |x\rangle |p_x\rangle, \text{ where } |p_x\rangle := \sum_{y \in \Omega} \sqrt{p_{x \rightarrow y}} |y\rangle. \quad (2.2)$$

The quantum walk is then defined as

$$W_p := S(2U_p(I_\Omega \otimes |0\rangle\langle 0|)U_p^\dagger - I_\Omega \otimes I_\Omega), \quad (2.3)$$

where I_Ω is the identity map on $\mathbb{C}^{|\Omega|}$ and $S := \sum_{x, y \in \Omega} |x, y\rangle\langle y, x| = S^\dagger$ is the swap gate on $\mathbb{C}^{|\Omega|} \otimes \mathbb{C}^{|\Omega|}$. This is the standard definition of discrete-time quantum walks due to Szegedy [70].

To understand the quantum walk, it is essential to analyze the spectrum of W_p . First, observing that $W_p = S(2\Pi - I)$ where $\Pi = U_p(I_\Omega \otimes |0\rangle\langle 0|)U_p^\dagger = \sum_{x \in \Omega} |x\rangle\langle x| \otimes |p_x\rangle\langle p_x|$ projects onto the span of the states $|x\rangle \otimes |p_x\rangle$, we consider the eigenvector $|\lambda\rangle$ of $\Pi S \Pi$ with eigenvalue λ . We have $\Pi S \Pi = \sum_{x, y \in \Omega} D_{xy} |x\rangle\langle y| \otimes |p_x\rangle\langle p_y|$, where $D_{xy} := \sqrt{p_{x \rightarrow y} p_{y \rightarrow x}}$. Since $W_p |\lambda\rangle = S |\lambda\rangle$ and $W_p S |\lambda\rangle = 2\lambda S |\lambda\rangle - |\lambda\rangle$, the subspace $\text{span}\{|\lambda\rangle, S |\lambda\rangle\}$ is invariant under W_p . The eigenvalues of W_p within this subspace are $\lambda \pm i\sqrt{1 - \lambda^2} = e^{\pm i \arccos \lambda}$. For more details, see Reference [70].

The phase gap $\arccos \lambda \geq \sqrt{2(1 - \lambda)} \geq \sqrt{2}\delta$, where δ is the spectral gap of D . Therefore, applying phase estimation using $O(1/\sqrt{\delta})$ calls to W_p suffices to distinguish the state corresponding to the stationary distribution of the classical Markov chain from the other eigenvectors.

2.2 Quantum Speedup of MCMC Sampling Via Simulated Annealing

Consider a Markov chain with spectral gap Δ and stationary distribution π . Classically, it takes $\Theta(\frac{1}{\Delta} \log(1/\epsilon \pi_{\min}))$ steps to sample from a distribution $\tilde{\pi}$ such that $\|\tilde{\pi} - \pi\| \leq \epsilon$, where $\pi_{\min} := \min_i \pi_i$. Quantumly, Reference [74] proved the following result about a sequence of slowly varying Markov chains:

THEOREM 2.1 ([74, THEOREM 2]). *Let $p_1, \dots, p_r$ be the transition probabilities of r Markov chains with stationary distributions $\pi_1, \dots, \pi_r$, spectral gaps $\delta_1, \dots, \delta_r$, and quantum walk operators*

$W_1, \dots, W_r$, respectively; let $\Delta := \min\{\delta_1, \dots, \delta_r\}$. Assume that $|\langle \pi_i | \pi_{i+1} \rangle|^2 \geq p$ for some $0 < p < 1$ and all $i \in [r-1]$, and assume that we can efficiently prepare the state $|\pi_1\rangle$ (where each $|\pi_i\rangle$ is a quantum sample defined as in Section 1.3.2). Then, for any $0 < \epsilon < 1$, there is a quantum algorithm that produces a quantum state $|\tilde{\pi}_r\rangle$ such that $\| |\tilde{\pi}_r\rangle - |\pi_r\rangle \| \leq \epsilon$, using $\tilde{O}(r/(p\sqrt{\Delta}))$ steps of the quantum walk operators $W_1, \dots, W_r$, where the $\tilde{O}$ omits poly-logarithmic terms in r , $1/\epsilon$, and $1/p\sqrt{\Delta}$.⁵

Their quantum algorithm produces the states $|\pi_1\rangle, \dots, |\pi_r\rangle$ sequentially, and can do so rapidly if consecutive states have significant overlap and the walks mix rapidly. Intuitively, this is achieved by amplitude amplification. However, to avoid overshooting, the article uses a variant of standard amplitude amplification, known as $\pi/3$ -amplitude amplification [32], that we now review.

Given two states $|\psi\rangle$ and $|\phi\rangle$, we let $\Pi_\psi := |\psi\rangle\langle\psi|$, $\Pi_\psi^\perp := I - \Pi_\psi$, $\Pi_\phi := |\phi\rangle\langle\phi|$, and $\Pi_\phi^\perp := I - \Pi_\phi$. Define the unitaries

$$R_\psi := \omega \Pi_\psi + \Pi_\psi^\perp, \quad R_\phi := \omega \Pi_\phi + \Pi_\phi^\perp, \quad \text{where } \omega = e^{i\frac{\pi}{3}}. \quad (2.4)$$

Given $|\langle\psi|\phi\rangle|^2 \geq p$, it can be shown that $|\langle\phi|R_\psi R_\phi|\psi\rangle|^2 \geq 1 - (1-p)^3$. Recursively, one can establish the following:

LEMMA 2.1 ([74, LEMMA 1]). *Let $|\psi\rangle$ and $|\phi\rangle$ be two quantum states with $|\langle\psi|\phi\rangle|^2 \geq p$ for some $0 < p \leq 1$. Define the unitaries R_ψ, R_ϕ as in Equation (2.4) and the unitaries U_m recursively as follows:*

$$U_0 = I, \quad U_{m+1} = U_m R_\psi U_m^\dagger R_\phi U_m. \quad (2.5)$$

Then, we have

$$|\langle\phi|U_m|\psi\rangle|^2 \geq 1 - (1-p)^{3^m}, \quad (2.6)$$

and the unitaries in $\{R_\psi, R_\psi^\dagger, R_\phi, R_\phi^\dagger\}$ are used at most 3^m times in U_m .

Taking $m = \lceil \log_3(\ln(1/\epsilon)/p) \rceil$, the inner product between $|\phi\rangle$ and $U_m|\psi\rangle$ in Equation (2.6) is at least $1 - \epsilon$, and we use $3^m = O(\log(1/\epsilon)/p)$ unitaries from the set $\{R_\psi, R_\psi^\dagger, R_\phi, R_\phi^\dagger\}$.

To establish Theorem 2.1 by Lemma 2.1, it remains to construct the unitaries $R_i := \omega |\pi_i\rangle\langle\pi_i| + (I - |\pi_i\rangle\langle\pi_i|)$. In Reference [74], this is achieved by phase estimation of the quantum walk operator W_i with precision $\sqrt{\Delta}/2$. Recall that if a classical Markov chain has spectral gap δ , then the corresponding quantum walk operator has phase gap of at least $2\sqrt{\delta}$ (see Section 2.1). Therefore, phase estimation with precision $\sqrt{\Delta}/2$ suffices to distinguish between $|\pi_i\rangle$ and other eigenvectors of W_i . As a result, we can take

$$R_i = \text{PhaseEst}(W_i)^\dagger \left(I \otimes \left(\omega |0\rangle\langle 0| + (I - |0\rangle\langle 0|) \right) \right) \text{PhaseEst}(W_i). \quad (2.7)$$

2.3 Quantum Chebyshev Inequality

Assume we are given a unitary U such that

$$U|0\rangle|0\rangle = \sqrt{p}|0\rangle|\phi\rangle + |0^\perp\rangle, \quad (2.8)$$

where $|\phi\rangle$ is a normalized pure state and $\langle 0| \otimes I |0^\perp\rangle = 0$. If we measure the output state, then we get 0 in the first register with probability p ; by the Chernoff bound, it takes $\Theta(1/\epsilon^2)$ samples to estimate p within ϵ with high success probability. However, there is a more efficient quantum

⁵Note that this is quadratically worse in $1/p$ than the Grover's algorithm [31] with complexity $O(1/\sqrt{p})$. This is because we use a simple fixed-point quantum search algorithm, i.e., $\pi/3$ -amplitude amplification [32], that does not require knowing p in advance. Notice that there exist other fixed-point quantum search algorithms that preserve the $O(1/\sqrt{p})$ speedup (e.g., Reference [76] and Reference [73, Chapter 6]), but in our quantum algorithm, the simpler algorithm suffices as $p = \Theta(1)$ (see Lemma 4.2).




Fig. 2. The quantum circuit for amplitude estimation.

algorithm, called *amplitude estimation* [12], that estimates the value of p using only $O(1/\epsilon)$ calls to U :

THEOREM 2.2 ([12, THEOREM 12]). *Given U satisfying Equation (2.8), the amplitude estimation algorithm in Figure 2 outputs an angle $\tilde{\theta}_p \in [-\pi, \pi]$ such that $\tilde{p} := \sin^2(\tilde{\theta}_p)$ satisfies*

$$|\tilde{p} - p| \leq \frac{2\pi\sqrt{p(1-p)}}{M} + \frac{\pi^2}{M^2}, \quad (2.9)$$

with success probability at least $8/\pi^2$, using M calls to U and $U^\dagger$.

Here QFT denotes the quantum Fourier transform over $\mathbb{Z}_M$ and $Q := -US_0U^\dagger S_1$ where S_0 and S_1 are reflections about $|0\rangle$ and the target state, respectively, following the pattern of Grover search. The controlled- Q gate denotes the operation $\sum_{j=0}^{M-1} |j\rangle\langle j| \otimes Q^j$. In fact, it was shown in the proof of Reference [12, Theorem 12] that the state after applying the circuit in Figure 2 is

$$\frac{e^{i\theta_p}}{\sqrt{2}} |\tilde{\theta}_p\rangle |\Psi_+\rangle - \frac{e^{-i\theta_p}}{\sqrt{2}} |-\tilde{\theta}_p\rangle |\Psi_-\rangle, \quad (2.10)$$

where $\theta_p \in [0, \pi]$ such that $p = \sin^2(\theta_p)$, and $\tilde{\theta}_p \in [0, \pi]$ is a random variable such that $\tilde{p} = \sin^2(\tilde{\theta}_p)$, and $|\Psi_\pm\rangle$ are two eigenvectors of Q . Measuring the first register either gives $\tilde{\theta}_p$ or $-\tilde{\theta}_p$ with probability 1/2, but since $\sin^2(\tilde{\theta}_p) = \sin^2(-\tilde{\theta}_p) = \tilde{p}$, this does not influence the success of Theorem 2.2.

In Equation (2.9), if we take $M = \lceil 2\pi(\frac{2\sqrt{p}}{\epsilon} + \frac{1}{\sqrt{\epsilon}}) \rceil = O(1/\epsilon)$, then we get

$$|\tilde{p} - p| \leq \frac{2\pi\sqrt{p(1-p)}}{2\pi} \epsilon + \frac{\pi^2}{4\pi^2} \epsilon^2 \leq \frac{\epsilon}{2} + \frac{\epsilon}{4} \leq \epsilon. \quad (2.11)$$

Furthermore, the success probability $8/\pi^2$ can be boosted to $1 - \nu$ by executing the algorithm $\Theta(\log 1/\nu)$ times and taking the median of the estimates.

Amplitude estimation can be generalized from estimating a single probability $p \in [0, 1]$ to estimating the expectation of a random variable. Assume that U is a unitary acting on $\mathbb{C}^S \otimes \mathbb{C}^{|\Omega|}$ such that

$$U|0\rangle|0\rangle = \sum_{x \in \Omega} \sqrt{p_x} |\psi_x\rangle |x\rangle, \quad (2.12)$$

where $S \in \mathbb{N}$ and $\{|\psi_x\rangle : x \in \Omega\}$ are unit vectors in $\mathbb{C}^S$. Let

$$\mu_U := \sum_{x \in \Omega} p_x x, \quad \sigma_U^2 := \sum_{x \in \Omega} p_x (x - \mu_U)^2 \quad (2.13)$$

denote the expectation and variance of the random variable, respectively. Several quantum algorithms have given speedups for estimating μ_U . Specifically, Reference [57] showed how to estimate

μ_U within additive error ϵ by $\tilde{O}(\sigma_U/\epsilon)$ calls to U and $U^\dagger$. Given an upper bound H and a lower bound $L > 0$ on the random variable, Reference [47] showed how to estimate μ_U with multiplicative error ϵ using $\tilde{O}(\sigma_U/\epsilon\mu_U \cdot H/L)$ calls to U and $U^\dagger$. More recently, Reference [33] mutually generalized these results and proposed a significantly better quantum algorithm:

THEOREM 2.3 ([33, THEOREM 3.5]). *There is a quantum algorithm that, given a quantum sampler U as in Equation (2.12), an integer Δ_U , a value $H > 0$, and two reals $\epsilon, \delta \in (0, 1)$, outputs an estimate $\tilde{\mu}_U$. If $\Delta_U \geq \sqrt{\sigma_U^2 + \mu_U^2}/\mu_U$ and $H > \mu_U$, then $|\tilde{\mu}_U - \mu_U| \leq \epsilon\mu_U$ with probability at least $1 - \delta$, and the algorithm uses $\tilde{O}(\Delta_U/\epsilon \cdot \log^3(H/\mu_U) \log(1/\delta))$ calls to U and $U^\dagger$.*

The quantum algorithm works as follows. First, assume $\Omega \subseteq [L, H]$ for given real numbers $L, H \geq 0$, there is a basic estimation algorithm (denoted BasicEst) that estimates $H^{-1}\mu_U$ up to ϵ -multiplicative error:

ALGORITHM 1: BasicEst: the basic estimation algorithm.

Input: A quantum sampler U acting on $\mathbb{C}^S \otimes \mathbb{C}^{|\Omega|}$, interval $[L, H]$, precision parameter $\epsilon \in (0, 1)$, failure parameter $\delta \in (0, 1)$.

Output: ϵ -multiplicative approximation of $H^{-1}\mu_U$.

1 Use controlled rotation to implement a unitary $R_{L,H}$ acting on $\mathbb{C}^{|\Omega|} \otimes \mathbb{C}^2$ such that for all $x \in \Omega$,

$$R_{L,H}|x\rangle|0\rangle = \begin{cases} |x\rangle(\sqrt{1 - \frac{x}{H}}|0\rangle + \sqrt{\frac{x}{H}}|1\rangle) & \text{if } L \leq x < H; \\ |x\rangle|0\rangle & \text{otherwise} \end{cases};$$

2 Let $V = (I_S \otimes R_{L,H})(U \otimes I_2)$ and $\Pi = I_S \otimes I_\Omega \otimes |1\rangle\langle 1|$;

3 **for** $i = 1, \dots, \Theta(\log(1/\delta))$ **do**

4 Compute $\tilde{p}_i$ by Theorem 2.2 with $U \leftarrow V$, $S_1 \leftarrow 2\Pi - I$, and $M \leftarrow \Theta(1/(\epsilon\sqrt{H^{-1}\mu_U}))$;

5 **Return** $\tilde{p} = \text{median}\{\tilde{p}_1, \dots, \tilde{p}_{\Theta(\log(1/\delta))}\}$.

However, usually the bounds L and H are not explicitly given. In this case, Reference [33] considered the truncated mean $\mu_{<b}$ defined by replacing the outcomes larger than b with 0. The article then runs Algorithm 1 (BasicEst) to estimate $\mu_{<b}/b$. A crucial observation is that $\sqrt{b/\mu_{<b}}$ is smaller than Δ_U for large values of b , and it becomes larger than Δ_U when $b \approx \mu_U \Delta_U^2$. As a result, by repeatedly running BasicEst with Δ_U quantum samples, and applying $O(\log(H/L))$ steps of a binary search on the values of b , the first non-zero value is obtained when $b/\Delta_U^2 \approx \mu_U$. In Reference [33], more precise truncation means are used to improve the precision of the result to $\tilde{O}(1/\epsilon)$ and remove the dependence on L .

Note that the quantum algorithm for Theorem 2.3 only relies on BasicEst. This is crucial when we estimate the mean of our simulated annealing algorithm in different iterations *nondestructively* (see Section 4.2 for more details).

2.4 Hit-and-run Walk

As introduced in Section 1.3.1, there are various random walks that mix fast in a convex body K , such as the grid walk [24] and the ball walk [21, 50]. In this article, we mainly use the *hit-and-run walk* [48, 52, 66]. It is defined as follows. Initially, the walk starts at the original point 0 (the centers of the balls with radius r and R in Equation (1.2)). In each iteration:

- (1) Pick a uniformly distributed random line ℓ through the current point p .
- (2) Move to a uniform random point along the chord $\ell \cap K$.

For any two points $p, q \in K$, we let $\ell(p, q)$ denote the length of the chord in K through p and q . Then the transition probability of the hit-and-run walk is determined by the following lemma:

LEMMA 2.2 ([48, LEMMA 3]). *At a step of the hit-and-run walk algorithm with current point u , the density function of the distribution of the next point $x \in K$ is*

$$p_u(x) = \frac{2}{nv_n} \cdot \frac{1}{\ell(u, x)|x - u|^{n-1}}, \quad (2.14)$$

where $v_n := \pi^{\frac{n}{2}}/\Gamma(1 + \frac{n}{2})$ is the volume of the n -dimensional unit ball. In other words, the probability that the next point is in a (measurable) set $A \subseteq K$ is

$$P_u(A) = \int_A \frac{2}{nv_n} \cdot \frac{1}{\ell(u, x)|x - u|^{n-1}} dx. \quad (2.15)$$

In general, we can also define a hit-and-run walk with a given density. Let f be a density function in $\mathbb{R}^n$. For any points $u, v \in \mathbb{R}^n$, we let

$$\mu_f(u, v) := \int_0^1 f((1-t)u + tv) dt. \quad (2.16)$$

For any line ℓ , let ℓ^+ and ℓ^- be the endpoints of the chord $\ell \cap K$ (with $+$ and $-$ assigned arbitrarily). The density f specifies the following hit-and-run walk:

- (1) Pick a uniformly distributed random line ℓ through the current point p .
- (2) Move to a random point x along the chord $\ell \cap K$ with density $\frac{f(x)}{\mu_f(\ell^-, \ell^+)}$.

Let π_K denote the uniform distribution over K . Smith [66] proved that the stationary distribution of the hit-and-run walk with uniform density is π_K . Furthermore, Lovász and Vempala [52] proved that the hit-and-run walk mixes rapidly from any initial distribution:

THEOREM 2.4 ([52, THEOREM 1.1]). *Let K be a convex body that satisfies Equation (1.2): $B_2(0, r) \subseteq K \subseteq B_2(0, R)$. Let σ be a starting distribution and let $\sigma^{(m)}$ be the distribution of the current point after m steps of the hit-and-run walk in K . Let $\epsilon > 0$, and suppose that the density function $d\sigma/d\pi_K$ is upper bounded by M except on a set S with $\sigma(S) \leq \epsilon/2$. Then for any*

$$m > 10^{10} \frac{n^2 R^2}{r^2} \ln \frac{M}{\epsilon}, \quad (2.17)$$

the total variation distance between $\sigma^{(m)}$ and π_K is less than ϵ .

Theorem 2.4 can also be generalized to exponential distributions on K , using the definition of level sets:

Definition 2.1. Given a probability density f supported on K , the level set of f of probability p is defined as the set

$$L_p := \{x : f(x) \geq f_p\} \quad (2.18)$$

for a value $f_p \in \mathbb{R}$ such that $\int_{x \in L_p} f(x) dx = p$.

THEOREM 2.5 ([52, THEOREM 1.3]). *Let $K \subset \mathbb{R}^n$ be a convex body and let f be a density supported on K that is proportional to $e^{-a^T x}$ for some vector $a \in \mathbb{R}^n$. Assume that the level set of f of probability $1/8$ contains a ball of radius r , and $\mathbb{E}_f(|x - z_f|^2) \leq R^2$, where z_f is the centroid of f . Let σ be a starting distribution and let σ^m be the distribution for the current point after m steps of the hit-and-run walk*

applied to f . Let $\epsilon > 0$, and suppose that the density function $\frac{d\sigma}{d\pi_f}$ is upper bounded by M except on a set S with $\sigma(S) \leq \frac{\epsilon}{2}$. Then for

$$m > 10^{30} \frac{n^2 R^2}{r^2} \ln^5 \frac{MnR}{r\epsilon},$$

the total variation distance between σ^m and π_f is less than ϵ .

Roughly speaking, the proofs of Theorem 2.4 and Theorem 2.5 have two steps. First, for any random walk on a continuous domain Ω with transition probability p and stationary distribution π , we define its conductance (which generalizes the discrete case in Equation (2.1)) as

$$\Phi := \inf_{S \subseteq \Omega} \frac{\int_S \int_{\Omega/S} dx dy \pi_x p_{x \rightarrow y}}{\min\{\int_S dx \pi_x, \int_{\Omega/S} dx \pi_x\}}. \quad (2.19)$$

It is well-known that the mixing time of this random walk is proportional to $1/\Phi^2$ up to logarithmic factors. This is captured by the following proposition:

PROPOSITION 2.1 ([50, COROLLARY 1.5]). *Let $M := \sup_{S \subseteq \Omega} \frac{\sigma(S)}{\pi(S)}$ where σ is the initial distribution. Then for every $S \subseteq \Omega$,*

$$|\sigma^{(k)}(S) - \pi(S)| \leq \sqrt{M} \left(1 - \frac{1}{2} \Phi^2\right)^k. \quad (2.20)$$

Furthermore, the conductance in Proposition 2.1 can be relaxed to that of sets with a fixed small probability p :

PROPOSITION 2.2 (ADAPTED FROM REFERENCE [50, COROLLARY 1.6]). *Let $M := \sup_{S \subseteq \Omega} \frac{\sigma(S)}{\pi(S)}$. If the conductance for all connected, measurable set $A \subseteq \Omega$ such that $\pi(A) = p \leq 1/2$ is at least Φ_p , then for all $S \subseteq \Omega$, we have*

$$|\sigma^{(k)}(S) - \pi(S)| \leq 2Mp + 2M \left(1 - \frac{1}{2} \Phi_p^2\right)^k. \quad (2.21)$$

Intuitively, conductance can directly characterize the convergence of the random walk to its stationary distribution. Specifically, the second term on the right-hand side of Equation (2.21) goes to 0 when $k = \Omega(1/\Phi_p^2)$.

Second, Reference [52] proved a lower bound on the conductance of the hit-and-run walk with exponential density:

PROPOSITION 2.3 ([52, THEOREM 6.9]). *Let f be a density in $\mathbb{R}^n$ proportional to $e^{-a^T x}$ whose support is a convex body K of diameter d . Assume that $B_2(0, r) \subseteq K$. Then for any subset S with $\pi_f(S) = p \leq 1/2$, the conductance of the hit-and-run walk satisfies*

$$\phi(S) \geq \frac{r}{10^{13} n d \ln(nd/rp)}. \quad (2.22)$$

Proposition 2.1 and Proposition 2.3 imply Theorem 2.4 and Theorem 2.5, respectively; complete proofs are given in Reference [52].

For the conductance of the hit-and-run walk with a uniform distribution, Reference [52] established a stronger lower bound that is independent of p :

PROPOSITION 2.4 ([52, THEOREM 4.2]). *Assume that K has diameter d and contains a unit ball. Then the conductance of the hit-and-run in K with uniform distribution is at least $\frac{1}{2^{24} n d}$.*

3 THEORY OF CONTINUOUS-SPACE QUANTUM WALKS

In this section, we develop the theory of continuous-space, discrete-time quantum walks.

Specifically, we generalize the discrete-time quantum walk of Szegedy [70] to continuous space. Let $n \in \mathbb{N}$ and suppose Ω is a continuous subset of $\mathbb{R}^n$ (we say that Ω is continuous if for any $x, y \in \Omega$ there is a continuous function $f_{x,y}: [0, 1] \rightarrow \Omega$ such that $f_{x,y}(0) = x$ and $f_{x,y}(1) = y$). A probability transition density p on Ω is a continuous function $p: \Omega \times \Omega \rightarrow [0, +\infty)$ such that

$$\int_{\Omega} dy p(x, y) = 1 \quad \forall x \in \Omega. \quad (3.1)$$

We also write $p_{x \rightarrow y} := p(x, y)$ for the transition density from x to y . Together, Ω and p specify a continuous-space Markov chain that we denote (Ω, p) throughout the article.

For background on the mathematical foundations of quantum mechanics over continuous state spaces, see Reference [65, Chapter 1]. In this section, we treat quantum states as square integrable functions $f: \Omega \rightarrow \mathbb{R}$ in $L^2(\Omega)$ if $\int_{\Omega} dx |f(x)|^2 < \infty$. The inner product $\langle \cdot, \cdot \rangle$ on $L^2(\Omega)$ is defined by

$$\langle f, g \rangle := \int_{\Omega} dx f(x)g(x) \quad \forall f, g \in L^2(\Omega). \quad (3.2)$$

Note that by the Cauchy-Schwarz inequality, we have

$$|\langle f, g \rangle|^2 \leq \left(\int_{\Omega} dx |f(x)|^2 \right) \left(\int_{\Omega} dx |g(x)|^2 \right) < \infty; \quad (3.3)$$

the norm of an $f \in L^2(\Omega)$ is subsequently defined as $\|f\| := \sqrt{\langle f, f \rangle}$. The pure states in Ω correspond to functions in the set

$$\text{St}(\Omega) := \left\{ f: \Omega \rightarrow \mathbb{R} \mid \int_{\Omega} dx |f(x)|^2 = 1 \right\}. \quad (3.4)$$

The computational basis elements $|y\rangle$ for all $y \in \Omega$ correspond to Dirac delta functions $\delta(x - y)$ centered at y , where $\delta(x) = 0$ for all $x \neq 0$, and $\int_{\mathbb{R}^n} \delta(x) dx = 1$. Delta functions are not members of the Hilbert space $L^2(\Omega)$; however, we interpret them in the following sense: For any $y \in \text{int}_{\epsilon}(\Omega)$, we consider a normalized Gaussian with width ϵ , given by $\delta_{y,\epsilon}(x) \propto \frac{1}{(2\pi\epsilon^2)^{n/2}} e^{-\|x-y\|^2/2\epsilon^2}$ for $x \in \Omega$ and 0 for $x \notin \Omega$. It is clear that $\delta_{y,\epsilon} \in L^2(\Omega)$ and its behavior approaches that of the delta function. In the remainder of the section statements such as $A = B$ are to be interpreted as $\lim_{\epsilon \rightarrow 0} |A_{\epsilon} - B_{\epsilon}| = 0$ where $A_{\epsilon}, B_{\epsilon}$ are obtained from A, B by replacing every occurrence of a computational basis vector $|y\rangle$ by $\delta_{y,\epsilon}$.⁶ Integrals over $L^2(\Omega)$ are computed pointwise. It can be verified that

$$\int_{\Omega} dx |x\rangle\langle x| = I \quad (3.5)$$

and

$$\langle x|x'\rangle = \delta(x - x'), \quad \forall x, x' \in \Omega. \quad (3.6)$$

⁶As in most treatments of continuous quantum mechanics, we shall not be fully rigorous with respect to operations such as interchanging orders of limits. We have two reasons to believe that pathological cases do not occur. First, the states that occur during the algorithm correspond to probability distributions that are obtained during classical geometric random walks. Second, we show in Section 5.1 that our algorithm can also be executed discretely; we work in the continuous setting only for ease of analysis.

3.1 Continuous-space Quantum Walk

Given a transition density function p , the quantum walk is characterized by the following states:

$$|\phi_x\rangle := |x\rangle \otimes \int_{\Omega} dy \sqrt{p_{x \rightarrow y}} |y\rangle \quad \forall x \in \mathbb{R}^n. \quad (3.7)$$

Since $p_{x \rightarrow y}$ is a normalized probability density function, $|\phi_x\rangle \in L^2(\Omega)$. Now, denote

$$U := \int_{\Omega} dx |\phi_x\rangle \langle x| \otimes \langle 0|, \quad \Pi := \int_{\Omega} dx |\phi_x\rangle \langle \phi_x|, \quad S := \int_{\Omega} \int_{\Omega} dx dy |x, y\rangle \langle y, x|. \quad (3.8)$$

Notice that U is an isometry on $L^2(\Omega) \otimes |0\rangle$, Π is the projection onto $\text{span}\{|\phi_x\rangle\}_{x \in \mathbb{R}^n}$, because

$$\Pi^2 = \int_{\Omega} \int_{\Omega} dx dx' |\phi_x\rangle \langle \phi_x| \phi_{x'} \rangle \langle \phi_{x'}| = \int_{\Omega} \int_{\Omega} dx dx' \delta(x - x') |\phi_x\rangle \langle \phi_{x'}| = \Pi, \quad (3.9)$$

and S is the swap operator for the two registers. A single step of the quantum walk is defined as the unitary operator

$$W := S(2\Pi - I). \quad (3.10)$$

The first main result of this subsection is the following theorem:

THEOREM 3.1. *Let*

$$D := \int_{\Omega} \int_{\Omega} dx dy \sqrt{p_{x \rightarrow y} p_{y \rightarrow x}} |x\rangle \langle y| \quad (3.11)$$

denote the discriminant operator of p . Let Λ be the set of eigenvalues of D , so that $D = \int_{\Lambda} d\lambda \lambda |\lambda\rangle \langle \lambda|$. Then the eigenvalues of the quantum walk operator W in Equation (3.10) are ± 1 and $\lambda \pm i\sqrt{1 - \lambda^2}$ for all $\lambda \in \Lambda$.

Notice that in Equation (3.11), D is real and symmetric, and thus its spectral decomposition exists. We defer the proof of Theorem 3.1 to Appendix A.

3.2 Stationary Distribution

Classically, the density $\pi = (\pi_x)_{x \in \Omega}$ corresponding to the stationary distribution of a Markov chain (Ω, p) satisfies

$$\int_{\Omega} dx \pi_x = 1; \quad \int_{\Omega} dy p_{y \rightarrow x} \pi_y = \pi_x \quad \forall x \in \Omega. \quad (3.12)$$

In other words, we can naturally define a transition operator as

$$P := \int_{\Omega} \int_{\Omega} dx dy p_{y \rightarrow x} |x\rangle \langle y|, \quad (3.13)$$

and the stationary density π satisfies $P\pi = \pi$. The Markov chain (Ω, p) is *reversible* if there exists a classical density $\sigma = (\sigma_x)_{x \in \Omega}$ such that

$$p_{y \rightarrow x} \sigma_y = p_{x \rightarrow y} \sigma_x \quad \forall x, y \in \Omega. \quad (3.14)$$

(This is called the *detailed balance condition*.) Notice that for all $x \in \Omega$,

$$\int_{\Omega} dy p_{y \rightarrow x} \sigma_y = \int_{\Omega} dy p_{x \rightarrow y} \sigma_x = \sigma_x \int_{\Omega} dy p_{x \rightarrow y} = \sigma_x; \quad (3.15)$$

therefore, we must have $P\sigma = \sigma$, i.e., σ is a stationary density of P . In this article, we focus on Markov chains (Ω, p) that are reversible and have a unique stationary distribution (i.e., $\sigma = \pi$). Such assumptions are natural for Markov chains in practice, including the Metropolis-Hastings algorithm, simple random walks on graphs, and so on.

We point out that if π is the classical stationary density of a reversible Markov chain (Ω, p) , then

$$|\pi_W\rangle := \int_{\Omega} dx \sqrt{\pi_x} |\phi_x\rangle \quad (3.16)$$

is the unique eigenvalue-1 eigenstate of the quantum walk operator W restricted to the subspace $\text{span}_{\lambda \in \Lambda} \{T|\lambda\rangle, ST|\lambda\rangle\}$. First, a simple calculation shows that

$$W|\pi_W\rangle = S(2\Pi - I)|\pi_W\rangle \quad (3.17)$$

$$= S|\pi_W\rangle \quad (3.18)$$

$$= \left(\int_{\Omega} \int_{\Omega} dx dy |x, y\rangle \langle y, x| \right) \left(\int_{\Omega} \int_{\Omega} dx dy \sqrt{\pi_y p_{y \rightarrow x}} |y, x\rangle \right) \quad (3.19)$$

$$= \int_{\Omega} \int_{\Omega} dx dy \sqrt{\pi_y p_{y \rightarrow x}} |x, y\rangle \quad (3.20)$$

$$= \int_{\Omega} \int_{\Omega} dx dy \sqrt{\pi_x p_{x \rightarrow y}} |x\rangle |y\rangle \quad (3.21)$$

$$= \int_{\Omega} dx \sqrt{\pi_x} |x\rangle \left(\int_{\Omega} dy \sqrt{p_{x \rightarrow y}} |y\rangle \right) \quad (3.22)$$

$$= \int_{\Omega} dx \sqrt{\pi_x} |\phi_x\rangle \quad (3.23)$$

$$= |\pi_W\rangle, \quad (3.24)$$

where Equation (3.18) follows from $|\pi_W\rangle \in \text{span}_{x \in \Omega} \{|\phi_x\rangle\}$, Equation (3.19) follows from the definition of S in Equation (3.8), Equation (3.21) follows from Equation (3.14), and Equation (3.21) follows from the definition of $|\phi_x\rangle$ in Equation (3.7). Thus, $|\pi_W\rangle$ is an eigenvector of W with eigenvalue 1. However, since (Ω, p) is reversible, P is similar to D : If we denote $D_{\pi} := \int_{\Omega} dx \sqrt{\pi_x} |x\rangle \langle x|$, then

$$\begin{aligned} D_{\pi} D D_{\pi}^{-1} &= \left(\int_{\Omega} dx \sqrt{\pi_x} |x\rangle \langle x| \right) \left(\int_{\Omega} \int_{\Omega} dx dy \sqrt{p_{x \rightarrow y} p_{y \rightarrow x}} |x\rangle \langle y| \right) \left(\int_{\Omega} dy \sqrt{\pi_y^{-1}} |y\rangle \langle y| \right) \\ &= \int_{\Omega} \int_{\Omega} dx dy \sqrt{\pi_x \pi_y^{-1} p_{x \rightarrow y} p_{y \rightarrow x}} |x\rangle \langle y| \end{aligned} \quad (3.25)$$

$$= \int_{\Omega} \int_{\Omega} dx dy p_{y \rightarrow x} |x\rangle \langle y| \quad (\text{by Equation (3.14)}) \quad (3.26)$$

$$= P. \quad (3.27)$$

As a result, D and P have the same set of eigenvalues. Furthermore, Lemma A.1 implies that all eigenvalues of P have norm at most 1, and the proof of Theorem 3.1 shows that $|\pi_W\rangle$ is the unique eigenvector with this eigenvalue within $\text{span}_{\lambda \in \Lambda} \{T|\lambda\rangle, ST|\lambda\rangle\}$.

The state

$$|\pi\rangle := \int_{\Omega} dx \sqrt{\pi_x} |x\rangle \quad (3.28)$$

represents a quantum sample from the density π ; in particular, measuring $|\pi\rangle$ in the computational basis gives a classical sample from π . Furthermore, the unitary operator in Equation (3.8) satisfies

$$U^{\dagger} |\pi_W\rangle = \left(\int_{\Omega} dx |x\rangle \langle 0| \langle \phi_x| \right) \left(\int_{\Omega} dx \sqrt{\pi_x} |\phi_x\rangle \right) = \int_{\Omega} dx \sqrt{\pi_x} |x\rangle \langle 0| = |\pi\rangle \langle 0|, \quad (3.29)$$

so we have $U|\pi\rangle \langle 0| = |\pi_W\rangle$.

4 QUANTUM SPEEDUP FOR VOLUME ESTIMATION

In this section, we present and analyze our quantum algorithm for volume estimation. First, we review the classical state-of-the-art volume estimation algorithm in Section 4.1. We then describe our quantum algorithm for estimating the volume of well-rounded convex bodies (i.e., $R/r = O(\sqrt{n})$) with query complexity $\tilde{O}(n^{2.5}/\epsilon)$ in Section 4.2, with detailed proofs given in Section 4.3. Finally, we remove the well-rounded condition by giving a quantum algorithm with interlaced rounding and volume estimation with additional cost $\tilde{O}(n^{2.5})$ in each iteration in Section 4.4.

4.1 Review of Classical Algorithms for Volume Estimation

The state-of-the-art classical algorithm for volume estimation uses $\tilde{O}(n^4 + n^3/\epsilon^2)$ classical queries, where $\tilde{O}(n^4)$ queries are used to construct the affine transformation that makes convex body well-rounded [53] and $\tilde{O}(n^3/\epsilon^2)$ queries are used to estimate the volume of the well-rounded convex body (after the affine transformation) [21].

We now review the algorithm of Reference [53] for estimating volumes of well-rounded convex bodies. This algorithm estimates the volume of a convex body obtained by the following *pencil construction*. Define the cone

$$C := \left\{ x \in \mathbb{R}^{n+1} : x_0 \geq 0, \sum_{i=1}^n x_i^2 \leq x_0^2 \right\}, \quad (4.1)$$

where x_j denotes the j th coordinate of x . Let K' be the intersection of the cone C and a cylinder $[0, 2D] \times K$, i.e.,

$$K' := ([0, 2D] \times K) \cap C \quad (4.2)$$

(recall $D = R/r$). Without loss of generality, we renormalize to $r = 1$, so that $B_2(0, 1) \subseteq K \subseteq B_2(0, D)$. Since $D \text{Vol}(K) \leq \text{Vol}(K') \leq 2D \text{Vol}(K)$, with the knowledge of $\text{Vol}(K')$, we can estimate $\text{Vol}(K)$ with multiplicative error ϵ by generating $O(1/\epsilon^2)$ sample points from the uniform distribution on $[0, 2D] \times K$ and then counting how many of them fall into K' . Such an approximation succeeds with high probability by a Chernoff-type argument (see Section 4.3.1 for a formal proof).

Lovász and Vempala [53] considers simulated annealing under the pencil construction. For any $a > 0$, define

$$Z(a) := \int_{K'} e^{-ax_0} dx. \quad (4.3)$$

It can be shown that for any $a \leq \epsilon/D$,

$$(1 - \epsilon) \text{Vol}(K') \leq Z(a) \leq \text{Vol}(K'). \quad (4.4)$$

However, it is shown in Reference [53, Section 2.3] that for any $a \geq 2n$ and $\epsilon > (3/4)^n$,

$$(1 - \epsilon) \int_C e^{-ax_0} dx \leq Z(a) \leq \int_C e^{-ax_0} dx. \quad (4.5)$$

This suggests using a simulated annealing procedure for estimating $\text{Vol}(K')$. Specifically, if we select a sequence $a_0 > a_1 > \dots > a_m$ for which $a_0 = 2n$ and $a_m \leq \epsilon/D$, then we can estimate $\text{Vol}(K')$ by

$$Z(a_m) = Z(a_0) \prod_{i=0}^{m-1} \frac{Z(a_{i+1})}{Z(a_i)}. \quad (4.6)$$

(Note that this procedure uses an increasing sequence of temperatures $1/a_i$, unlike standard simulated annealing in which temperature is decreased.)

Let π_i be the probability distribution over K' with density proportional to $e^{-a_i x_0}$, i.e., $d\pi_i(x) = \frac{e^{-a_i x_0}}{Z(a_i)} dx$. Let X_i be a random sample from π_i , and let $(X_i)_0$ be its zeroth coordinate (x_0 in Equation (4.1)); define $V_i := e^{(a_i - a_{i+1})(X_i)_0}$. We have

$$\mathbb{E}_{\pi_i}[V_i] = \int_{K'} e^{(a_i - a_{i+1})x_0} d\pi_i(x) = \int_{K'} e^{(a_i - a_{i+1})x_0} \frac{e^{-a_i x_0}}{Z(a_i)} dx = \frac{Z(a_{i+1})}{Z(a_i)}. \quad (4.7)$$

Furthermore, if the simulated annealing schedule satisfies $a_{i+1} \geq (1 - \frac{1}{\sqrt{n}})a_i$, then V_i satisfies (see Reference [53, Lemma 4.1])

$$\frac{\mathbb{E}_{\pi_i}[V_i^2]}{\mathbb{E}_{\pi_i}[V_i]^2} \leq \left(\frac{a_{i+1}^2}{a_i(2a_{i+1} - a_i)} \right)^{n+1} < 8 \quad \forall i \in [m], \quad (4.8)$$

i.e., the variance of V_i is bounded by a constant multiple of the square of its expectation. Thus, this simulated annealing procedure constitutes *Chebyshev cooling* (see also Section 4.3.3), ensuring its correctness (see Proposition 4.1).

Algorithm 2 presents this approach in detail. Note that sampling from π_0 in Line 2 is straightforward: select a random positive real number x_0 from the distribution with density e^{-2nx} and choose a uniformly random point $(v_1, \dots, v_n)$ from the unit ball. If $X = (x_0, x_0 v_1, \dots, x_0 v_n) \notin K'$, then try again; else return X . Equation (4.5) ensures that we succeed with probability at least $1 - \epsilon$ for each sample.

ALGORITHM 2: Volume estimation of well-rounded K with $\tilde{O}(n^4/\epsilon^2)$ classical queries [53].

Input: Membership oracle O_K of K ; R such that $B_2(0, 1) \subseteq K \subseteq B_2(0, R)$; $R = O(\sqrt{n})$, i.e., K is well-rounded.

Output: ϵ -multiplicative approximation of $\text{Vol}(K)$.

- 1 Set $m = 2\lceil \sqrt{n} \ln(n/\epsilon) \rceil$, $k = \frac{512}{\epsilon^2} \sqrt{n} \ln(n/\epsilon)$, $\delta = \epsilon^2 n^{-10}$, and $a_i = 2n(1 - \frac{1}{\sqrt{n}})^i$ for $i \in [m]$;
 - 2 Take k samples $X_0^{(1)}, \dots, X_0^{(k)}$ from π_0 ;
 - 3 **for** $i \in [m]$ **do**
 - 4 Take k samples from π_i with error parameter δ and starting points $X_{i-1}^{(1)}, \dots, X_{i-1}^{(k)}$, giving points $X_i^{(1)}, \dots, X_i^{(k)}$;
 - 5 Compute $V_i = \frac{1}{k} \sum_{j=1}^k e^{(a_i - a_{i+1})(X_i^{(j)})_0}$;
 - 6 Return $n!v_n(2n)^{-(n+1)}V_1 \cdots V_m$ as the estimate of the volume of K' , where $v_n := \pi^{\frac{n}{2}}/\Gamma(1 + \frac{n}{2})$ is the volume of the n -dimensional unit ball.
-

4.2 Quantum Algorithm for Volume Estimation

As introduced in Section 1.2, our quantum algorithm has four main improvements that contribute to the quantum speedup of Algorithm 2:

- (1) We replace the classical hit-and-run walk in Section 2.4 by a quantum hit-and-run walk, defined using the framework of Section 3. Classically, the hit-and-run walk mixes in $\tilde{O}(n^3)$ steps in a well-rounded convex body given a warm start (see Theorem 2.4). Quantumly, we can use the quantum hit-and-run walk operator to prepare its stationary state given a warm start state using only $\tilde{O}(n^{1.5})$ queries to the membership oracle for the well-rounded convex body.
- (2) We replace the simulated annealing framework in Algorithm 2 by the quantum MCMC framework described in Section 2.2. Classically, we sample from π_i in the i th iteration by running the classical hit-and-run walk starting from the samples taken in the $(i - 1)$ st iteration. Quantumly, we prepare the quantum sample $|\pi_i\rangle$ in the i th iteration by applying

$\pi/3$ -amplitude amplification to a quantum sample produced in the $(i - 1)$ st iteration, where the unitaries in the $\pi/3$ -amplitude amplification are implemented by phase estimation of the quantum hit-and-run walk operators as in Equation (2.7).

- (3) We use the quantum Chebyshev inequality (see Section 2.3) to give a quadratic quantum speedup in ϵ^{-1} when taking the average $e^{(a_i - a_{i+1})(\bar{X}_i)_0}$ in Line 5 of Algorithm 2. However, we must be cautious, because the resulting points $X_i^{(1)}, \dots, X_i^{(k)}$ in Line 4 follow the distribution π_i , which varies in different iterations of simulated annealing. Instead, our quantum algorithm must be *nondestructive*: it must still have a copy of $|\pi_i\rangle$ after estimating the average $e^{(a_i - a_{i+1})(\bar{X}_i)_0}$, so that we can map this state to $|\pi_{i+1}\rangle$ by $\pi/3$ -amplitude amplification for the next iteration. This is achieved in Section 4.3.3.
- (4) In Section 4.4, we show how the densities can be transformed to be well-rounded by an affine transformation at each stage of the algorithm. This is to ensure that the hit-and-run walk mixes fast assuming the densities π_i to be sampled from are well-rounded (see Theorem 2.5). The high-level idea is to sample points from density π_i and compute an affine transformation S_{i+1} that rounds π_i and the next density π_{i+1} (see Lemma 4.11). To sample these points, we use $\pi/3$ -amplitude amplification to map the states corresponding to the uniform distributions for one stage to those for the next. The affine transformation can be computed coherently using nondestructive mean estimation, with $\tilde{O}(n^{2.5})$ quantum queries in each iteration.

Algorithm 3 is our quantum volume estimation algorithm that satisfies our main theorem:

THEOREM 1.1 (MAIN THEOREM). *Let $K \subset \mathbb{R}^n$ be a convex set with $B_2^n(0, r) \subseteq K \subseteq B_2^n(0, R)$. Assume $0 < \epsilon < 1/2$. Then there is a quantum algorithm that returns a value $\text{Vol}(K)$ satisfying*

$$\frac{1}{1 + \epsilon} \text{Vol}(K) \leq \widetilde{\text{Vol}}(K) \leq (1 + \epsilon) \text{Vol}(K), \quad (1.4)$$

with probability at least $2/3$ using $\tilde{O}(n^3 + n^{2.5}/\epsilon)$ quantum queries to the membership oracle O_K (defined in Equation (1.3)) and $\tilde{O}(n^5 + n^{4.5}/\epsilon)$ additional arithmetic operations.

The proof of Theorem 1.1 is organized as follows. We first assume that in each iteration, S_{i+1} puts π_{i+1} in isotropic position, i.e., the densities are promised to be well-rounded. The rest of this subsection presents an overview of the proof of Theorem 1.1 (including a quantum circuit in Figure 3), and proofs details are given in Section 4.3. In Section 4.4, we show how the well-roundedness be maintained at an additional cost of $\tilde{O}(n^{2.5})$ quantum queries in each iteration.

Following the discussion in Section 1.2, our proof can be described at three levels:

High level (the simulated annealing framework). In Section 4.3.1, we show how to estimate $\text{Vol}(K)$ given an estimate of the volume of the pencil construction, $\text{Vol}(K')$:

LEMMA 4.1. *If we have access to $\widetilde{\text{Vol}}(K')$ such that*

$$\frac{1}{1 + \epsilon/2} \text{Vol}(K') \leq \widetilde{\text{Vol}}(K') \leq (1 + \epsilon/2) \text{Vol}(K'), \quad (4.9)$$

with probability at least 0.7 , then we can return a value $\widetilde{\text{Vol}}(K)$ such that

$$\frac{1}{1 + \epsilon} \text{Vol}(K) \leq \widetilde{\text{Vol}}(K) \leq (1 + \epsilon) \text{Vol}(K) \quad (4.10)$$

holds with probability at least $2/3$, using $\tilde{O}(n^{2.5}/\epsilon)$ quantum queries to the membership oracle O_K .

In Section 4.3.2, we prove that the inner product between stationary states of consecutive simulated annealing steps is at least a constant:

ALGORITHM 3: Volume estimation of convex K with $\tilde{O}(n^3 + n^{2.5}/\epsilon)$ quantum queries.

Input: Membership oracle O_K for K ; $R = O(\sqrt{n})$ such that $B_2(0, 1) \subseteq K \subseteq B_2(0, R)$.

Output: ϵ -multiplicative approximation of $\text{Vol}(K)$.

- 1 Set $m = \Theta(\sqrt{n} \log(n/\epsilon))$ to be the number of iterations of simulated annealing and $a_i = 2n(1 - \frac{1}{\sqrt{n}})^i$ for $i \in [m]$. Let π_i be the probability distribution over K' with density $\pi_i(x)$ proportional to $e^{-a_i x_0}$ (here x_0 is the zeroth coordinate of x , as in Equation (4.1));
Set error parameters $\delta, \epsilon' = \Theta(\epsilon/m^2), \epsilon_1 = \epsilon/2m$; let $k = \tilde{O}(\sqrt{n}/\epsilon)$ be the number of copies of stationary states for applying the quantum Chebyshev inequality; let $l = \tilde{O}(n)$ be the number of copies of stationary states needed to obtain the affine transformation S_i ;
 - Prepare $k + l$ (approximate) copies of $|\pi_0\rangle$, denoted $|\tilde{\pi}_0^{(1)}\rangle, \dots, |\tilde{\pi}_0^{(k+l)}\rangle$ (Lemma 4.4);
 - 2 **for** $i \in [m]$ **do**
 - 3 Use the quantum Chebyshev inequality on the k copies of the state $|\tilde{\pi}_{i-1}\rangle$ with parameters ϵ_1, δ to estimate the expectation $\mathbb{E}_{\pi_i}[V_i]$ (in Equation (4.7)) as $\tilde{V}_i$ (Lemma 4.9 and Figure 4). The post-measurement states are denoted $|\hat{\pi}_{i-1}^{(1)}\rangle, \dots, |\hat{\pi}_{i-1}^{(k)}\rangle$;
 - 4 Use the l copies of the state $|\pi_{i-1}\rangle$ to nondestructively obtain the affine transformation S_i that rounds π_{i-1} and π_i (Section 4.4). The post-measurement states are denoted $|\hat{\pi}_{i-1}^{(k+1)}\rangle, \dots, |\hat{\pi}_{i-1}^{(k+l)}\rangle$;
 - 5 Apply $\pi/3$ -amplitude amplification with error ϵ' (Section 2.2 and Lemma 4.8) and affine transformation S_i to map $|S_i \hat{\pi}_{i-1}^{(1)}\rangle, \dots, |S_i \hat{\pi}_{i-1}^{(k+l)}\rangle$ to $|S_i \tilde{\pi}_i^{(1)}\rangle, \dots, |S_i \tilde{\pi}_i^{(k+l)}\rangle$, using the quantum hit-and-run walk;
 - 6 Invert S_i to get $k + l$ (approximate) copies of the stationary distribution $|\pi_i\rangle$ for use in the next iteration;
 - 7 Compute an estimate $\widetilde{\text{Vol}(K')} = n!v_n(2n)^{-(n+1)}\tilde{V}_1 \dots \tilde{V}_m$ of the volume of K' , where v_n is the volume of the n -dimensional unit ball;
 - 8 Use $\widetilde{\text{Vol}(K')}$ to estimate the volume of K as $\widetilde{\text{Vol}(K)}$ (Section 4.3.1).
-

LEMMA 4.2. *Let $|\pi_i\rangle$ be the stationary distribution state of the quantum walk W_i for $i \in [m]$ defined in Equation (3.28). For $n \geq 2$, we have $\langle \pi_i | \pi_{i+1} \rangle > 1/3$ for $i \in [m-1]$.*

This allows $\pi/3$ -amplitude amplification to transform the stationary state of one Markov chain to that of the next. The total number of iterations of $\pi/3$ -amplitude amplification is thus $\tilde{O}(\sqrt{n})$, which equals to the number of iterations of the classical volume estimation algorithm by Reference [53].

Middle level (each telescoping ratio). In Section 4.3.3, we describe how we apply the quantum Chebyshev inequality (Theorem 2.3) to the Chebyshev cooling schedule.

LEMMA 4.3. *Given $\tilde{O}(\log(1/\delta)/\epsilon)$ copies of the quantum states $|\pi_{i-1}\rangle$, there exists a quantum algorithm that outputs an estimate of $\mathbb{E}_{\pi_i}[V_i]$ (in Equation (4.7)) with relative error less than ϵ with probability at least $1 - O(\delta)$ using $\tilde{O}(C \log(1/\delta)/\epsilon)$ oracle calls, where C oracle calls are required to implement a sampler for $|\pi_i\rangle$. Moreover, this quantum algorithm is nondestructive, i.e., the initial copies of quantum states $|\pi_{i-1}\rangle$ are restored after the computation with probability at least $1 - O(\delta)$.*

Because the relative error in each iteration for estimating the volume via Chebyshev cooling is $\Theta(\epsilon/m) = \tilde{O}(\epsilon/\sqrt{n})$, Lemma 4.3 implies that $O(\log(1/\delta)/(\epsilon/\sqrt{n})) = \tilde{O}(\sqrt{n}/\epsilon)$ copies of the stationary state suffice for the simulated annealing framework.⁷

⁷Notice that in the quantum Chebyshev inequality by Hamoudi and Magniez, they did not distinguish the number of copies of the initial state from the number of quantum samples [33, Theorem 5.3]. In fact, their proof uses only $O(\log(1/\delta))$ copies




Fig. 3. The quantum circuit for Algorithm 3 (assuming well-roundedness). Here $U_{CB,i}$ is the circuit of the quantum Chebyshev inequality (Theorem 2.3) in the i th iteration and $U_{i,l}$ is $\pi/3$ -amplitude amplification from $|\pi_i\rangle$ to $|\pi_{i+1}\rangle$.

Low level (the quantum hit-and-run walk). In Section 4.3.4, we give a careful analysis of the errors coming from the quantum Chebyshev inequality as well as the $\pi/3$ -amplitude amplification:

LEMMA 4.4. For $\epsilon_1 < 1$, given $\tilde{O}(\log(1/\delta)/\epsilon_1)$ copies of a state $|\tilde{\pi}_{i-1}\rangle$ such that $\| |\tilde{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| \leq \epsilon_1$, there exists a quantum procedure (using $\pi/3$ -amplitude amplification and the quantum Chebyshev inequality) that outputs a $\tilde{V}_i$ such that $|\tilde{V}_i - \mathbb{E}_{\pi_i}[V_i]| \leq \epsilon_1 \mathbb{E}_{\pi_i}[V_i]$ (where $\mathbb{E}_{\pi_i}[V_i]$ is defined in Equation (4.7)) with success probability $1 - \delta^4$ using $\tilde{O}(n^{3/2} \log(1/\delta)/\epsilon_1 + n^{3/2} \log(1/\epsilon'))$ calls to the membership oracle and returns $\tilde{O}(\log(1/\delta)/\epsilon_1)$ copies of final states $|\tilde{\pi}_i\rangle$ such that $\| |\tilde{\pi}_i\rangle - |\pi_i\rangle \| = O(\epsilon_1 + \delta + \epsilon')$.

Having the four lemmas above from all the three levels, we establish Theorem 1.1 as follows.

PROOF OF THEOREM 1.1. We prove the correctness and analyze the cost separately.

Correctness. By Lemma 4.1, it suffices to compute the volume of the pencil construction K' to relative error $\epsilon/2$ with probability at least 0.7 to compute the volume of the well-rounded convex body K . This is computed as a telescoping sum of $m = O(\sqrt{n} \log n/\epsilon)$ products of the form $Z(a_{i+1})/Z(a_i)$. The random variable V_i is an unbiased estimator for $Z(a_{i+1})/Z(a_i)$, i.e., $\mathbb{E}_{\pi_i}[V_i] = Z(a_{i+1})/Z(a_i)$. Consider applying Lemma 4.4 m times with sufficiently small $\delta, \epsilon' \leq \epsilon/12m^2$ and $\epsilon_1 = \epsilon/3m$.

of the initial state $|\pi_{i-1}\rangle$ in Lemma 4.3, which reduces the total number of copies of the stationary states in the simulated annealing framework to $O(\log(1/\delta))$. Nevertheless, this does not change the total query and time complexities of our quantum algorithms, because the total number of calls to the quantum sampler remains the same.

This will promise that $\epsilon_1 + \delta + \epsilon' \leq \epsilon/2m$. At each iteration i , we have a state $|\tilde{\pi}_{i-1}\rangle$ such that $\| |\tilde{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| \leq O(\epsilon/4m)$. Thus, each telescoping sum can be computed with a relative error of $\epsilon/2m$, resulting in a relative error of less than $\epsilon/2$ for the final volume. The probability of success for each iteration is at least $1 - \delta^4 = 1 - \Theta(\epsilon^4/4m^8)$. Thus, the probability of success for the whole algorithm is at least $1 - \Theta(\epsilon^4/4m^7) = 1 - \tilde{O}(\epsilon^{11}/n^{3.5})$, which is greater than 0.7 for a large enough n .

Cost. Ignoring the cost of obtaining the affine transformation to round the logconcave densities to be sampled (assuming that all the relevant densities are well rounded), we have from Lemma 4.4, the number of calls to the membership oracle in each iteration of Algorithm 3 is $\tilde{O}(n^{3/2} \log(1/\delta)/\epsilon_1 + n^{3/2} \log(1/\epsilon')) = \tilde{O}(n^2/\epsilon)$. The total number of oracle calls is thus $\tilde{O}(n^{2.5}/\epsilon)$. The argument for correctness above applies for well-rounded logconcave densities. This is ensured by maintaining $\tilde{\Theta}(n)$ states that are used to round the densities in each iteration (Algorithm 4). By Proposition 4.4, this procedure uses $\tilde{O}(n^{2.5})$ calls to the membership oracle in each iteration, resulting in a final query complexity of $\tilde{O}(n^3 + n^{2.5}/\epsilon)$. Since the affine transformation is an n -dimensional matrix-vector product, the number of additional arithmetic operations is hence $O(n^2) \cdot \tilde{O}(n^3 + n^{2.5}/\epsilon) = \tilde{O}(n^5 + n^{4.5}/\epsilon)$. $\square$

4.3 Proofs of Lemmas in Section 4.2

We now prove the lemmas in Section 4.2 that establish Theorem 1.1.

4.3.1 From the Pencil Construction to the Original Convex Body. Here, we prove

LEMMA 4.1. *If we have access to $\widetilde{\text{Vol}}(K')$ such that*

$$\frac{1}{1 + \epsilon/2} \text{Vol}(K') \leq \widetilde{\text{Vol}}(K') \leq (1 + \epsilon/2) \text{Vol}(K'), \quad (4.9)$$

with probability at least 0.7, then we can return a value $\widetilde{\text{Vol}}(K)$ such that

$$\frac{1}{1 + \epsilon} \text{Vol}(K) \leq \widetilde{\text{Vol}}(K) \leq (1 + \epsilon) \text{Vol}(K) \quad (4.10)$$

holds with probability at least 2/3, using $\tilde{O}(n^{2.5}/\epsilon)$ quantum queries to the membership oracle O_K .

PROOF. We follow the same notation in Section 4.1; i.e., without loss of generality, we assume that $r = 1$ and denote $D = R/r = R$. Since R and r are both given, D is a known value. The pencil construction is

$$K' := ([0, 2D] \times K) \cap \left\{ x \in \mathbb{R}^{n+1} : x_0 \geq 0, \sum_{i=1}^n x_i^2 \leq x_0^2 \right\}. \quad (4.11)$$

By the definition of D , for any $(x_1, \dots, x_n) \in K$, we have $\sum_{i=1}^n x_i^2 \leq D^2$, so $[D, 2D] \times K \subseteq K'$. This implies that $D \text{Vol}(K) \leq \text{Vol}(K') \leq 2D \text{Vol}(K)$. In other words, letting $\xi_K := \frac{\text{Vol}(K')}{2D \text{Vol}(K)}$, we have $0.5 \leq \xi_K \leq 1$.

Classically, we consider a Monte Carlo approach to approximating $\text{Vol}(K)$: We take k (approximately) uniform samples $x_1, \dots, x_k$ from $[0, 2D] \times K$, and if k' of them are in K' , then we return $\frac{k'}{k} \cdot \text{Vol}(K')$. For each $i \in [k]$, $\delta[x_i \in K']$ is a Bernoulli random variable with expectation $\xi_K = \Theta(1)$. Any Bernoulli random variable has variance $O(1)$. Therefore, by Chebyshev's inequality, taking $k = O(1/\epsilon^2)$ suffices to ensure that

$$\Pr \left[\left| \frac{k'}{k} - \xi_K \right| \leq \frac{\epsilon \xi_K}{2} \right] \geq 0.99. \quad (4.12)$$

Quantumly, we adopt the same Monte Carlo approach, but we implement two steps using quantum techniques:

- We take an approximately uniform sample from $K' = [0, 2D] \times K$ via the quantum hit-and-run walk. To obtain a quantum stationary state, we use a similar idea as in Reference [24] to construct a sequence of $m = \lceil n \log_2(2D) \rceil$ convex bodies. Let $\hat{K}_0 := B_2^{n+1}(0, 1)$ and $\hat{K}_i := 2^{i/n} B_2^{n+1}(0, 1) \cap K'$ for $i \in [m]$. As the length of the pencil is $2D$, $\hat{K}_m = K'$. The state $|\pi_0\rangle$ corresponding to $\hat{K}_0$ is easy to prepare. It is straightforward to verify that $\langle \pi_i | \pi_{i+1} \rangle \geq c$ for some constant c , as $\text{Vol}(\hat{K}_{i+1}) \leq 2 \text{Vol}(\hat{K}_i)$. To utilize the quantum speedup for MCMC framework (Theorem 2.1), it remains to lower bound the phase gap of the quantum walk operator for $\hat{K}_i$. It can be shown that the mixing property of the hit-and-run walk in Theorem 2.5 implies that the phase gap of the quantum walk operator is $\tilde{\Omega}(n^{-1.5})$; see the proof of Lemma 4.8. Thus, by Theorem 2.1, $|\pi_m\rangle$ can be prepared using $\tilde{O}(n) \cdot \tilde{O}(n^{1.5}) = \tilde{O}(n^{2.5})$ quantum queries to O_K .
- We estimate ξ_K with multiplicative error $\epsilon/2$ using the quantum Chebyshev inequality (Theorem 2.3) instead of its classical counterpart. This means that $O(1/\epsilon)$ executions of quantum sampling in the first step suffice.

Overall, $\tilde{O}(n^{2.5}/\epsilon)$ quantum queries to O_K suffice to ensure that we obtain an estimate of ξ_K within multiplicative error $\epsilon/2$ with success probability at least 0.99. Since (4.9) ensures that $\overline{\text{Vol}(K')}$ estimates $\text{Vol}(K')$ up to multiplicative error $\epsilon/2$ with probability at least 0.7, $\frac{\overline{\text{Vol}(K')}}{2D\xi_K}$ estimates $\text{Vol}(K)$ up to multiplicative error $\epsilon/2 + \epsilon/2 = \epsilon$ with success probability $0.99 \cdot 0.7 > 2/3$. $\square$

4.3.2 Inner Product between Stationary States of Consecutive Steps. We now show that the inner product between stationary states of consecutive steps is at least a constant. More precisely, we have the following:

LEMMA 4.2. *Let $|\pi_i\rangle$ be the stationary distribution state of the quantum walk W_i for $i \in [m]$ defined in Equation (3.28). For $n \geq 2$, we have $\langle \pi_i | \pi_{i+1} \rangle > 1/3$ for $i \in [m-1]$.*

PROOF. Recall that the stationary distribution π_i of step i has density proportional to $e^{-a_i x_0}$ as discussed in Section 4.1. The corresponding stationary distribution state is $|\pi_i\rangle = \int_{K'} dx \sqrt{\frac{e^{-a_i x_0}}{Z(a_i)}} |x\rangle$. Lovász and Vempala [53, Lemma 3.2] proved that $a^{n+1}Z(a)$ is log-concave (noting that the dimension of K' is $n+1$). This implies that

$$\sqrt{a_i^{n+1}Z(a_i)}\sqrt{a_{i+1}^{n+1}Z(a_{i+1})} \leq \left(\frac{a_i + a_{i+1}}{2}\right)^{n+1} Z\left(\frac{a_i + a_{i+1}}{2}\right). \quad (4.13)$$

Now, we have

$$\langle \pi_i | \pi_{i+1} \rangle = \int_{K'} dx \frac{\exp\left(-\frac{a_i + a_{i+1}}{2} x_0\right)}{\sqrt{Z(a_i)}\sqrt{Z(a_{i+1})}} \quad (4.14)$$

$$\geq \left(\frac{2\sqrt{a_i}\sqrt{a_{i+1}}}{a_i + a_{i+1}}\right)^{n+1} \frac{\int_{K'} dx \exp\left(-\frac{a_i + a_{i+1}}{2} x_0\right)}{Z\left(\frac{a_i + a_{i+1}}{2}\right)} \quad (4.15)$$

$$= \left(\frac{2\sqrt{a_i}\sqrt{a_{i+1}}}{a_i + a_{i+1}}\right)^{n+1} \quad (4.16)$$

$$= \left(\frac{2\sqrt{a_i}\sqrt{a_i(1 - \frac{1}{\sqrt{n}})}}{a_i + a_i(1 - \frac{1}{\sqrt{n}})}\right)^{n+1} = \left(\frac{2\sqrt{1 - \frac{1}{\sqrt{n}}}}{2 - \frac{1}{\sqrt{n}}}\right)^{n+1}, \quad (4.17)$$

where the inequality follows from Equation (4.13). When $n = 2$ or $n = 3$, the above inequality holds. When $n \geq 4$, to lower bound the above quantity, we use the fact that $\sqrt{1 - 1/\sqrt{n}} \geq 1 - \frac{1}{2\sqrt{n}} - \frac{1}{2n}$. Hence, for $n \geq 2$, we have

$$\langle \pi_i | \pi_{i+1} \rangle \geq \left(\frac{2 - \frac{1}{\sqrt{n}} - \frac{1}{n}}{2 - \frac{1}{\sqrt{n}}} \right)^{n+1} = \left(1 - \frac{\frac{1}{n}}{2 - \frac{1}{\sqrt{n}}} \right)^{n+1} \geq \left(1 - \frac{1}{(2 - \frac{1}{\sqrt{2}})n} \right)^{n+1} > \frac{1}{3}$$

as claimed. $\square$

4.3.3 Chebyshev Cooling and Nondestructive Mean Estimation. Now, we briefly review the classical framework for Chebyshev cooling and discuss how to adapt it to quantum algorithms. Suppose we want to compute the expectation of a product

$$V = \prod_{i=1}^m V_i \quad (4.18)$$

of independent random variables. The following theorem of Dyer and Frieze [23] upper bounds the number of samples from V_i that suffices to estimate $\mathbb{E}[V]$ with bounded relative error.

PROPOSITION 4.1 ([23, SECTION 4.1]). *Let $V_1, \dots, V_m$ be independent random variables such that $\frac{\mathbb{E}[V_i^2]}{\mathbb{E}[V_i]^2} \leq B$ for all $i \in [m]$. Let $X_j^{(1)}, \dots, X_j^{(k)}$ be k samples of V_j for $j \in [m]$, and define $\bar{X}_j = \frac{1}{k} \sum_{\ell=1}^k X_j^{(\ell)}$. Let $V = \prod_{j=1}^m V_j$ and $X = \prod_{j=1}^m \bar{X}_j$. Then, taking $k = 16Bm/\epsilon^2$ ensures that*

$$\Pr \left[(1 - \epsilon)\mathbb{E}[V] \leq X \leq (1 + \epsilon)\mathbb{E}[V] \right] \geq \frac{3}{4}. \quad (4.19)$$

With standard techniques, the probability can be boosted to $1 - \delta$ with a $\log(1/\delta)$ overhead.

In applications such as volume estimation [52] and estimating partition functions [69], the samples are produced by a random walk. Let the mixing time for each random walk be at most T . Then the total complexity for estimating $\mathbb{E}[V]$ with success probability $1 - \delta$ is $\tilde{O}(TBm \log(1/\delta)/\epsilon^2)$. Replacing the random walk with a quantum walk can potentially improve the mixing time; see Section 1.3.2 for relevant literature. In particular, Montanaro [57] proposed a quantum algorithm for the simulated annealing framework with complexity $\tilde{O}(TBm \log(1/\delta)/\epsilon)$, which has a quadratic improvement in precision. Note that the dependence on T was not improved, as multiple copies of quantum states were prepared for the mean estimation (which uses measurements). In this article, we use the quantum Chebyshev inequality (see Theorem 2.3) to estimate the expectation of V_i in a nondestructive manner, which, together with Theorem 2.1, achieves complexity $\tilde{O}(\sqrt{TBm} \log(1/\delta)/\epsilon)$.

Recall that the random variables V_i (determined by the cooling schedule) satisfy Equation (4.8). The following lemma uses this property of the simulated annealing procedure to show that the quantum Chebyshev inequality can be used to estimate the mean of V_i on the distribution π_i , which gives an estimate of the ratio $\frac{Z(a_{i+1})}{Z(a_i)}$ in the volume estimation algorithm. We first show that our random variables can be made to satisfy the conditions of Theorem 2.3, and then we outline how the corresponding circuit can be implemented. A detailed error analysis is deferred to Section 4.3.4. To make the mean estimation nondestructive, we use the following theorem, originally posited by Harrow and Wei [34, Theorem 6]. We cite below a version proved by Cornelissen and Hamoudi [19], which fixes some potential issues with the original presentation, arising from assumptions on the output of amplitude estimation.

THEOREM 4.1 ([19, PROPOSITION B.2]). *Given state $|\psi\rangle$ and reflections $R_\psi = 2|\psi\rangle\langle\psi| - I$ and $R = 2P - I$, and any $\eta > 0$, there exists a quantum algorithm that outputs $\tilde{a}$, an approximation to $a = \langle\psi|P|\psi\rangle$, so that*

$$|\tilde{a} - a| \leq \frac{\sqrt{a(1-a)}}{M} + \frac{1}{M^2}, \quad (4.20)$$

with probability at least $1 - \eta$, and $O(\log(1/\eta)M)$ uses of R_ψ and R . Moreover, the algorithm restores the state ψ with probability at least $1 - \eta$. Finally, if $a \leq 1/(4M^2)$, then $\tilde{a} = 0$ with probability at least $1 - \eta$.

LEMMA 4.3. *Given $\tilde{O}(\log(1/\delta)/\epsilon)$ copies of the quantum states $|\pi_{i-1}\rangle$, there exists a quantum algorithm that outputs an estimate of $\mathbb{E}_{\pi_i}[V_i]$ (in Equation (4.7)) with relative error less than ϵ with probability at least $1 - O(\delta)$ using $\tilde{O}(C \log(1/\delta)/\epsilon)$ oracle calls, where C oracle calls are required to implement a sampler for $|\pi_i\rangle$. Moreover, this quantum algorithm is nondestructive, i.e., the initial copies of quantum states $|\pi_{i-1}\rangle$ are restored after the computation with probability at least $1 - O(\delta)$.*

PROOF. We apply the quantum Chebyshev inequality (Theorem 2.3). For the random variables V_i , we let μ_i denote their mean and σ_i^2 their variance. From Equation (4.8), $\sqrt{\sigma_i^2 + \mu_i^2}/\mu_i \leq \sqrt{8} < 3$. For a small constant c , we use $\log(1/\delta)/c^2$ copies of $|\pi_{i-1}\rangle$ to create copies of $|\pi_i\rangle$ using $\pi/3$ -amplitude amplification. We now use a quantum circuit that given $|x\rangle|0\rangle$ computes $|x\rangle|e^{a_i x_0 - a_{i-1} x_0}\rangle$, and then apply a circuit U_{median} that computes the median of all the ancilla registers:

$$U_{\text{median}}|0\rangle|a_1\rangle \cdots |a_s\rangle = |\text{median}\{a_1, \dots, a_s\}\rangle|a_1\rangle \cdots |a_s\rangle. \quad (4.21)$$

By the classical Chebyshev inequality, we measure $\hat{\mu}_i$ such that $|\hat{\mu}_i - \mu_i| \leq c\mu_i$ with probability at least $1 - \delta$. Thus, the probability that $\hat{\mu}_i/(1-c) < \mu$ is less than δ . Taking $H = \hat{\mu}_i/(1-c)$, our variables satisfy the conditions of the quantum Chebyshev inequality. To output an estimate of the mean with relative error at most ϵ , the quantum Chebyshev inequality now requires $\tilde{O}(\log(1/\delta)/\epsilon)$ calls to a sampler for the state $|\pi_i\rangle$, which we construct using $\pi/3$ -amplitude amplification on copies of $|\pi_{i-1}\rangle$. By the union bound, the probability of failure of the whole procedure is $O(\delta)$.

To be more specific, we replace $U|0\rangle$ in BasicEst (Algorithm 1) by $U_{i-1,l}|\pi_{i-1}\rangle$ (where $U_{i-1,l}$ is the circuit transforming the l^{th} copy of $|\pi_{i-1}\rangle$ to $|\pi_i\rangle$), and replace Q by $-U_{i-1,l}(\Pi_{i-1} - \Pi_{i-1}^\perp)U_{i-1,l}^\dagger(\Pi_i - \Pi_i^\perp)$ (where $\Pi_i = |\pi_i\rangle\langle\pi_i|$ and $\Pi_i^\perp = I - \Pi_i$ for all $i \in [m]$). The quantum circuit for nondestructive BasicEst is shown in Figure 4. Here, we run $\Theta(\log(1/\delta))$ executions of amplitude estimation (Figure 2) in parallel. Note that by Equation (2.10), each amplitude estimation returns a state $\frac{e^{i\theta_p}}{\sqrt{2}}|\tilde{\theta}_p\rangle - \frac{e^{-i\theta_p}}{\sqrt{2}}|-\tilde{\theta}_p\rangle$. We use an ancilla register and apply the unitary

$$U_{\sin^2}|\theta\rangle|0\rangle := |\theta\rangle|\sin^2 \theta\rangle, \quad (4.22)$$

because $\sin^2(\tilde{\theta}_p) = \sin^2(-\tilde{\theta}_p) = \tilde{p}$, the ancilla register becomes $|\tilde{p}\rangle$, where $\tilde{p}$ estimates p well as claimed in Theorem 2.2. We then take the median of such $\Theta(\log(1/\delta))$ executions using Equation (4.21), and finally run the inverse of $U_{\sin^2}$ gates and amplitude estimations. The correctness follows from the proof of Theorem 2.3 in Reference [33].

To assure non-destructiveness, we replace every application of Quantum Amplitude Estimation with Nondestructive Mean Estimation as in Theorem 4.1. The resulting guarantees on the error are the same as with the original amplitude estimation algorithm. To ensure an overall success probability of $1 - O(\delta)$, it suffices to perform each instance of Nondestructive Amplitude Estimation with success probability $1 - \tilde{O}(\delta)$. Note that since we estimate an unweighted mean, $2P - I$ with $P = H|0\rangle\langle 0|H$ can be implemented as HR_0H , where R_0 is a reflection around the $|0\rangle$ state. Finally,




Fig. 4. The quantum circuit for nondestructive BasicEst.

we show in Corollary 4.1 that R_{π_i} (the reflection around $|\pi_i\rangle$) can be implemented using the same number of oracle calls and gates as that required to sample π_i (up to polylogarithmic factors). $\square$

A detailed error analysis is given in the next subsection (see Lemma 4.9).

4.3.4 Error Analysis. In this section, we analyze the error incurred by both the quantum Chebyshev inequality (Line 3) and $\pi/3$ -amplitude amplification (Line 5) in Algorithm 3.

LEMMA 4.4. *For $\epsilon_1 < 1$, given $\tilde{O}(\log(1/\delta)/\epsilon_1)$ copies of a state $|\tilde{\pi}_{i-1}\rangle$ such that $\| |\tilde{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| \leq \epsilon_1$, there exists a quantum procedure (using $\pi/3$ -amplitude amplification and the quantum Chebyshev inequality) that outputs a $\tilde{V}_i$ such that $|\tilde{V}_i - \mathbb{E}_{\pi_i}[V_i]| \leq \epsilon_1 \mathbb{E}_{\pi_i}[V_i]$ (where $\mathbb{E}_{\pi_i}[V_i]$ is defined in Equation (4.7)) with success probability $1 - \delta^4$ using $\tilde{O}(n^{3/2} \log(1/\delta)/\epsilon_1 + n^{3/2} \log(1/\epsilon'))$ calls to the membership oracle and returns $\tilde{O}(\log(1/\delta)/\epsilon_1)$ copies of final states $|\tilde{\pi}_i\rangle$ such that $\| |\tilde{\pi}_i\rangle - |\pi_i\rangle \| = O(\epsilon_1 + \delta + \epsilon')$.*

We first show that $\pi/3$ -amplitude amplification can be used to rotate $|\pi_i\rangle$ into $|\pi_{i-1}\rangle$ with error ϵ' using $\tilde{O}(\log(1/\epsilon'))$ oracle calls. This procedure is used as a subroutine in a mean estimation circuit that estimates the mean of the random variable V_i using multiple approximate copies of $|\pi_{i-1}\rangle$. We ensure that the measurement probabilities are highly peaked so that the state is not disturbed very much. Finally, $\pi/3$ -amplitude estimation is used again to rotate the approximate copies of the state $|\pi_{i-1}\rangle$ to approximate copies of the state $|\pi_i\rangle$.

Large effective spectral gap. Consider an ergodic, reversible Markov chain (Ω, p) with transition matrix P and a unique stationary distribution with density π . Let $a(x)$ be a probability measure over Ω such that the Markov chain mixes to its stationary distribution with a corresponding probability

density $\pi(x)$ within a total variation distance of ϵ within t steps. Furthermore, let $a(x)$ be a warm start for $\pi(x)$ (see Definition 4.1). From the definition of the transition matrix $P(x, y) = \langle x|P|y \rangle = p_{y \rightarrow x}$.

The discriminant matrix D defined in Equation (3.11) is related to the transition matrix as $P = D_\pi D D_\pi^{-1}$, as shown in Equation (3.27). For a hit-and-run walk, the transition matrix P represents a convolution with an L_2 normalized function (corresponding to the square root of the density $p_{x \rightarrow y}$). Bounded subsets of $L^2(\Omega)$ are, therefore, mapped by P to other bounded subsets, and hence P is compact. Since D is connected to P by a similarity relation, D is a compact Hermitian operator over $L_2(\Omega)$ and thus has a countable set of real eigenvalues λ_i and corresponding orthonormal eigenvectors (eigenfunctions) $v_i \in L_2(\Omega)$. Orthonormality implies that $\int_\Omega v_i(x) v_j(x) dx = \delta_{i,j}$. Notice that

$$PD_\pi v_i = D_\pi D(v_i) = \lambda_i D_\pi v_i; \quad (4.23)$$

thus, $f_i = D_\pi v_i$ is an eigenvector of P' with eigenvalue λ_i . The eigenvectors f_i may not be orthogonal under the standard inner product on $L_2(\Omega)$. However, we can define an inner product

$$\langle f, g \rangle_\pi := \langle D_\pi^{-1} f, D_\pi^{-1} g \rangle = \int_\Omega \frac{f(x)g(x)}{\pi(x)} dx \quad (4.24)$$

over the space $L_2(\Omega)$. It is easy to see that $\langle f_i, f_j \rangle_\pi = \langle v_i, v_j \rangle = \delta_{i,j}$. A corresponding norm can be defined as $\|f\|_\pi = \sqrt{\langle f, f \rangle_\pi}$.

It can be verified that $\sqrt{\pi}(x)$ is an eigenfunction of D with eigenvalue 1. Thus, the stationary state $\pi(x)$ is an eigenfunction of the transition operator P with eigenvalue 1. Since P is stochastic, this is the leading eigenvalue. The eigenvalues of P are thus $1, \lambda_1, \lambda_2, \dots$ with corresponding eigenfunctions $\pi(x), f_2(x), f_3(x), \dots$. From the orthonormality of the f under $\langle \cdot, \cdot \rangle_\pi$, for any function g in $L_2(\Omega)$, we have

$$g = \sum_{i=1}^{\infty} \langle g, f_i \rangle_\pi f_i = \langle g, \pi \rangle_\pi + \sum_{i=2}^{\infty} \langle g, f_i \rangle_\pi f_i \quad (4.25)$$

$$= \left(\int_\Omega \frac{g(x)\pi(x)}{\pi(x)} dx \right) \pi + \sum_{i=2}^{\infty} \langle g, f_i \rangle_\pi f_i \quad (4.26)$$

$$= \left(\int_\Omega g(x) dx \right) \pi + \sum_{i=2}^{\infty} \langle g, f_i \rangle_\pi f_i. \quad (4.27)$$

Since a is a probability density, $a = \pi + \sum_{i=2}^{\infty} \langle a, f_i \rangle_\pi f_i$. After t steps of the Markov chain M on a , we obtain the state $P^t a = \pi + \sum_{i=2}^{\infty} \lambda_i^t \langle a, f_i \rangle_\pi f_i$. Since the walk mixes to total variation distance ϵ , we have $\|P^t a - \pi\|_1 \leq \epsilon$. Furthermore, since a is a warm start,

$$\|P^t a - \pi\|_\pi = \|P^t(a - \pi)\|_\pi = \int_\Omega P^t(a(x) - \pi(x)) \cdot \frac{P^t(a(x) - \pi(x))}{\pi(x)} dx \quad (4.28)$$

$$\leq \int_\Omega |P^t(a(x) - \pi(x))| \cdot \left| \frac{P^t(a(x) - \pi(x))}{\pi(x)} \right| dx \quad (4.29)$$

$$\leq \int_\Omega |P^t(a(x) - \pi(x))| \cdot O(1) dx \quad (4.30)$$

$$\leq O(1) \cdot \|P^t(a - \pi)\|_1 \quad (4.31)$$

$$= O(1) \cdot \|P^t a - \pi\|_1 = O(\epsilon). \quad (4.32)$$

Consequently, $\|\sum_{i=2}^{\infty} \lambda_i^t \langle a, f_i \rangle_\pi f_i\|_\pi = O(\epsilon)$ and from the orthonormality of f , $\langle a, f_i \rangle_\pi \lambda_i^t = O(\epsilon)$. If $1 > \lambda_i \geq 1 - \frac{1}{\Omega(t)}$, then $\lambda_i^t = \Omega(1)$ and $\langle a, f_i \rangle_\pi = O(\epsilon)$.

The above analysis indicates that if a probability density a (that is a warm start) mixes in t steps under a Markov chain (Ω, p) , then it has small overlap with each of the “bad” eigenfunctions (with spectral gap less than $\frac{1}{\Omega(t)}$). Thus, P effectively has a large spectral gap when it acts on a .

Corresponding to a , consider the quantum states

$$|a\rangle := \int_{\Omega} \sqrt{a(x)} |x\rangle dx, \quad |\phi_a\rangle := \int_{\Omega} \int_{\Omega} \sqrt{a_x p_{x \rightarrow y}} |x\rangle |y\rangle dx dy. \quad (4.33)$$

For an eigenvector v_i of D (with eigenvalue λ_i), define the state $|v_i\rangle := \int_{\Omega} v_i(x) dx = \int_{\Omega} \frac{f_i(x)}{\sqrt{\pi(x)}} dx$.

Then the walk operator W has the corresponding eigenvector $|u_i\rangle = (I - (\lambda_i - i\sqrt{1 - \lambda_i^2})S)T|v_i\rangle$ following the proof of Theorem 3.1. Let $C_i := \lambda_i - i\sqrt{1 - \lambda_i^2}$; then $\langle \phi_a | u_i \rangle = \langle \phi_a | T | v_i \rangle - C_i \langle \phi_a | ST | u_i \rangle$. Furthermore,

$$\langle \phi_a | T | v_i \rangle = \langle a | v_i \rangle = \int_{\Omega} \frac{\sqrt{a(x)} f_i(x)}{\sqrt{\pi(x)}} dx, \quad (4.34)$$

and

$$\langle \phi_a | ST | v_i \rangle = \left(\int_{\Omega} \sqrt{a_x p_{x \rightarrow y}} \langle y | \langle x | \right) \left(\int_{\Omega} \sqrt{v_{x'} p_{x' \rightarrow y'}} |x'\rangle |y'\rangle \right) \quad (4.35)$$

$$= \int_{\Omega} \sqrt{a_x} \left(\int_{\Omega} \sqrt{p_{x \rightarrow y} p_{y \rightarrow x}} v_i(y) dy \right) dx \quad (4.36)$$

$$= \int_{\Omega} \sqrt{a_x} (Dv_i)(x) dx \quad (4.37)$$

$$= \lambda_i \int_{\Omega} \sqrt{a_x} v_i(x) dx \quad (4.38)$$

$$= \lambda_i \langle a | v_i \rangle. \quad (4.39)$$

We have $\langle \phi_a | u_i \rangle = (1 - \lambda_i C_i) \langle a | v_i \rangle$ and, therefore,

$$|\langle \phi_a | u_i \rangle| = (1 - \lambda_i C_i) \langle a | v_i \rangle = \sqrt{(1 - \lambda_i^2)^2 + (1 - \lambda_i)^2} \langle a | v_i \rangle \leq 2 |\langle a | v_i \rangle|. \quad (4.40)$$

In addition,

$$\langle a | v_i \rangle = \int_{\Omega} \frac{\sqrt{a(x)} f_i(x)}{\sqrt{\pi(x)}} dx = \int_{\Omega} \sqrt{\frac{\pi(x)}{a(x)}} \frac{a(x) f_i(x)}{\pi(x)} dx. \quad (4.41)$$

The above discussion establishes the following proposition indicating that if a distribution with density $a(x)$ mixes fast and the stationary distribution with density $\pi(x)$ has a bounded L_2 -norm with respect to $a(x)$, then the quantum walk operator W acting on the subspace spanned by $|\pi\rangle$ and $|a\rangle$ has a large effective spectral gap.

PROPOSITION 4.2. *Let $M = (\Omega, p)$ be an ergodic reversible Markov chain with a transition operator P and unique stationary state with a corresponding density $\pi \in L_2(\Omega)$. Let $\{(\lambda_i, f_i)\}$ be the set of eigenvalues and eigenfunctions of P , and let $|u_i\rangle$ be the eigenvectors of the corresponding quantum walk operator W . Let $a \in L_2(\Omega)$ be a probability density that is a warm start for π and mixes up to total variation distance ϵ in t steps of M . Furthermore, assume that $\int_{\Omega} \frac{\pi(x)}{a(x)} \pi(x) dx \leq c$ for some*

constant c . Define

$$|a\rangle = \int_{\Omega} \sqrt{a(x)} |x\rangle dx, \quad (4.42)$$

$$|\phi_a\rangle = \int_{\Omega} \sqrt{a(x)} \int_{\Omega} \sqrt{p_{x \rightarrow y}} |x\rangle |y\rangle dx dy. \quad (4.43)$$

Then $\langle \phi_a | u_i \rangle = O(\epsilon^{1/2})$ for all i such that $1 > \lambda_i \geq 1 - \frac{1}{\Omega(t)}$.

PROOF. Define $S = \{x | \frac{\pi(x)}{a(x)} \geq \sqrt{\frac{c}{\epsilon}}\}$. Because $\int_{\Omega} \frac{\pi(x)^2}{a(x)^2} a(x) dx = \int_{\Omega} \frac{\pi(x)}{a(x)} \pi(x) dx \leq c$, Markov's inequality implies that $\int_S a(x) dx \leq \epsilon$.

We now define the quantum state $|a'\rangle$ such that $\langle x | a' \rangle = \langle x | a \rangle$ if $x \notin S$ and $\langle a | x' \rangle = 0$ otherwise, and $|\phi_{a'}\rangle = T|a'\rangle$. Then,

$$\| |\phi_a\rangle - |\phi_{a'}\rangle \| = \left\| \int_S \sqrt{a(x)} T|x\rangle dx \right\| = \sqrt{\int_{\Omega} a(x) dx} = \sqrt{\epsilon}. \quad (4.44)$$

From Equations (4.40) and (4.41), if $1 > \lambda_i \geq 1 - \frac{1}{O(t)}$, then

$$|\langle \phi_{a'} | u_i \rangle| \leq \left| 2 \int_{\Omega} \sqrt{\frac{\pi(x)}{a(x)}} \frac{a(x) f_i(x)}{\pi(x)} dx \right| \leq \frac{2c^{1/4} \langle a, f_i \rangle_{\pi}}{\epsilon^{1/4}} \leq 2c^{1/4} \epsilon^{3/4}. \quad (4.45)$$

Finally,

$$\langle \phi_a | u_i \rangle = \langle \phi_{a'} | u_i \rangle + \langle \phi_a - \phi_{a'} | u_i \rangle \leq 2c^{1/4} \epsilon^{3/4} + \sqrt{\epsilon} = O(\sqrt{\epsilon}) \quad (4.46)$$

if $1 > \lambda_i \geq 1 - \frac{1}{\Omega(t)}$. Hence, the result follows. $\square$

Warmness of π_{i+1} with respect to π_i . We show that density π_i mixes to π_{i+1} under the walk W_{i+1} and vice versa. To apply Theorem 2.5, we show that the two distributions are warm with respect to each other.

The L_2 -norm of a distribution with density $\pi_1 \in L_2(\Omega)$ with respect to another with density $\pi_2 \in L_2(\Omega)$ is defined as

$$\|\pi_1/\pi_2\| = \mathbb{E}_{X \sim \pi_1} \left[\frac{\pi_1(X)}{\pi_2(X)} \right] = \int_{\Omega} \frac{\pi_1(x)}{\pi_2(x)} \pi_1(x) dx. \quad (4.47)$$

Definition 4.1 (Warm Start). A density $\pi_1 \in L_2(\Omega)$ is said to be a warm start for $\pi_2 \in L_2(\Omega)$ if the L_2 -norm $\|\pi_1/\pi_2\|$ is bounded by a constant.

We have the following results on the L_2 -norms of π_{i+1} and π_i :

LEMMA 4.5 ([53, LEMMA 4.4]). *The L_2 -norm of the probability distribution with density $\pi_i = \frac{e^{-a_i x_0}}{Z(a_i)}$ with respect to that with density $\pi_{i+1} = \frac{e^{-a_{i+1} x_0}}{Z(a_{i+1})}$ is at most 8.*

LEMMA 4.6. *The L_2 -norm of the probability distribution with density $\pi_{i+1} = \frac{e^{-a_{i+1} x_0}}{Z(a_{i+1})}$ with respect to that with density $\pi_i = \frac{e^{-a_i x_0}}{Z(a_i)}$ is at most e .*

PROOF. Since $a^n Z(a)$ is a log-concave function [53, Lemma 3.2], we have

$$\mathbb{E}_{X \sim \pi_{i+1}} \left[\frac{\pi_{i+1}(X)}{\pi_i(X)} \right] = \frac{\int_{K'} e^{(a_i - a_{i+1})x_0} e^{-a_{i+1}x_0} dx \int_{K'} e^{-a_i x_0} dx}{\int_{K'} e^{-a_{i+1}x_0} dx \int_{K'} e^{-a_{i+1}x_0} dx}$$

$$= \frac{Z(2a_{i+1} - a_i)Z(a_i)}{Z(a_{i+1})^2} \quad (\text{definition of } Z) \quad (4.48)$$

$$\leq \left(\frac{a_{i+1}^2}{a_i(2a_{i+1} - a_i)} \right)^n \quad (\text{logconcavity of } a^n Z(a)) \quad (4.49)$$

$$\leq \left(\frac{\left(1 - \frac{1}{\sqrt{n}}\right)^2}{1 - \frac{2}{\sqrt{n}}} \right)^n \quad (\text{definition of } a_i) \quad (4.50)$$

$$\leq \left(1 + \frac{2}{n}\right)^n < e^2, \quad (4.51)$$

where Equation (4.51) holds, because $1 + \frac{1}{n} - \frac{2}{\sqrt{n}} \leq (1 + \frac{2}{n})(1 - \frac{2}{\sqrt{n}})$ as long as $n \geq 16$. $\square$

Error analysis of $\pi/3$ -amplitude amplification. Consider a simulated annealing procedure that follows a sequence of Markov chains $M_1, M_2, \dots$ with stationary states $\mu_1, \mu_2, \dots$. Consider an alternative walk operator of the form

$$W'_i = U_i^\dagger S U_i R_{\mathcal{A}} U_i^\dagger S U_i R_{\mathcal{A}}, \quad (4.52)$$

where $R_{\mathcal{A}}$ denotes the reflection about the subspace $\mathcal{A} := \text{span}\{|x\rangle|0\rangle : x \in K\}$ and S is the swap operator. We have $U_i|x\rangle|0\rangle = \int_{y \in K} \sqrt{p_{x \rightarrow y}^{(i)}} |x\rangle|y\rangle dy$ where $p^{(i)}$ is the transition probability corresponding to the i th chain.

The W'_i operator is related to the walk operator $W_i = S(2\Pi_i - I)$ via conjugation by U_i , i.e., $W_i = U_i W'_i U_i^\dagger$. Thus, W'_i has the same eigenvalues as W_i , and if $|u_j\rangle$ is an eigenvector of W_i with eigenvalue λ_j , then $|v\rangle = U_i^\dagger |u_j\rangle$ is an eigenvector of W'_i with the same eigenvalue λ_j . For any classical distribution f , we define $|f\rangle = \int_{\Omega} \sqrt{f(x)} |x\rangle dx$ and $|\phi_f^{(i)}\rangle = \int_{\Omega} \sqrt{f(x)} |x\rangle \int_{\Omega} \sqrt{p_{x \rightarrow y}^{(i)}} |y\rangle dy dx$. Since $|\phi_{\pi_i}^{(i)}\rangle$ is a stationary state of W_i with eigenvalue 1, it follows that $|\pi_i\rangle|0\rangle$ is an eigenvalue of W_i with eigenvalue 1.

In each stage of the volume estimation algorithm, we sample from a state with density $\pi_i(x) = \frac{e^{-a_i x_0}}{Z(a_i)}$. Each such distribution is the stationary state of a hit-and-run walk with the corresponding target density. Thus, the corresponding state $|\pi_i\rangle$ is the stationary state of the corresponding walk operators W_i and W'_i . Both W_i and W'_i can be implemented using a constant number of U_i gates.

From Lemma 4.2, we know that the inner product $\langle \pi_i | \pi_{i+1} \rangle$ between the states at any stage of the algorithm is at least $\frac{1}{3}$. This implies that the inner product between $|\pi_i\rangle|0\rangle$ and $|\pi_{i+1}\rangle|0\rangle$ is also at least $\frac{1}{3}$. In the following, we abuse notation by sometimes writing only $|\pi_i\rangle$ to denote $|\pi_i\rangle|0\rangle$, as it is easy to tell from context whether the ancilla register should be present.

Lemma 2.1 in Section 2.2 indicates that $\pi/3$ -amplitude amplification can be used to rotate the state $|\pi_i\rangle$ to $|\pi_{i+1}\rangle$ if we can implement the rotation unitaries

$$R_i = \omega |\pi_i\rangle\langle \pi_i| + (I - |\pi_i\rangle\langle \pi_i|) \quad \text{and} \quad R_{i+1} = \omega |\pi_{i+1}\rangle\langle \pi_{i+1}| + (I - |\pi_{i+1}\rangle\langle \pi_{i+1}|).$$

To implement these rotations, we use the fact that π_i and π_{i+1} are the eigenvectors of the operators W'_i and W'_{i+1} with eigenvalue 1, respectively. We show the following lemmas, which are adapted variants of Lemma 2 and Corollary 2 in Reference [74]:

LEMMA 4.7. *Let W be a unitary operator with a unique leading eigenvector $|\psi_0\rangle$ with eigenvalue 1. Denote the remaining eigenvectors by $|\psi_j\rangle$ with corresponding eigenvalues $e^{2\pi i \xi_j}$. For any $\Delta \in (0, 1]$*




Fig. 5. The quantum phase estimation circuit. Here W is a unitary operator with eigenvector $|\psi_j\rangle$; in $\pi/3$ -amplitude estimation it is the quantum walk operator W'_i in Equation (4.52).

and $\epsilon_2 < 1/2$, define $a := \log(1/\Delta)$ and $c := \log(1/\sqrt{\epsilon_2})$. There exists a quantum circuit V that uses ac ancilla qubits and invokes the controlled- W gate $2^a c$ times such that

$$V|\psi_0\rangle|0\rangle^{\otimes ac} = |\psi_0\rangle|0\rangle^{\otimes ac} \quad (4.53)$$

and

$$V|\psi_j\rangle|0\rangle^{\otimes ac} = \sqrt{1 - \epsilon_2(j)}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2(j)}|\psi_j\rangle|0\rangle^{\otimes ac}, \quad (4.54)$$

where $|\chi_j\rangle$ is orthogonal to $|0\rangle^{\otimes ac}$ for all $|\psi_j\rangle$ such that $\xi_j \geq \Delta$, and $\epsilon_2(j) \leq \epsilon_2$ for all j .

PROOF. Consider a quantum phase estimation circuit U with a ancilla qubits that invokes the controlled- W gate 2^a times (see Figure 5). The phase estimation circuit first creates an equal superposition over a ancilla qubits using Hadamard gates. For $k = 0, \dots, a-1$, we apply a controlled- W^k operator to the input register, controlled by the $(a-k)$ th register. Finally, the inverse quantum Fourier transform is applied on the ancilla registers. Then,

$$U|\psi_j\rangle|0\rangle^{\otimes a} = |\psi_j\rangle \otimes \text{QFT}^\dagger \left(\frac{1}{\sqrt{2^a}} \sum_{m=0}^{2^a-1} e^{2\pi i m \xi_j} |m\rangle \right) \quad (4.55)$$

$$= |\psi_j\rangle \otimes \frac{1}{2^a} \sum_{m, m'=0}^{2^a-1} e^{2\pi i m (\xi_j - m'/2^a)} |m'\rangle. \quad (4.56)$$

The amplitude corresponding to $|0\rangle$ on the ancilla registers is

$$a_{j,0} := \frac{1}{2^a} \sum_{m=0}^{2^a-1} e^{2\pi i m \xi_j} = \frac{1 - e^{2\pi i 2^a \xi_j}}{2^a (1 - e^{2\pi i \xi_j})} \quad (4.57)$$

for $j \neq 0$, and $a_{0,0} = 1$. If $j \neq 0$, then

$$|a_{j,0}| = \left| \frac{1 - e^{2\pi i 2^a \xi_j}}{2^a (1 - e^{2\pi i \xi_j})} \right| \leq \left| \frac{1}{2^{a-1} (1 - e^{2\pi i \xi_j})} \right| \leq \frac{1}{2^{a+1} |\xi_j|}. \quad (4.58)$$

Thus, $|a_{j,0}| \leq \frac{1}{2}$ if $\xi_j \geq \Delta$. Using c copies of the circuit (resulting in ac ancilla registers and $2^a c$ controlled- W gates), the amplitude for 0 in all the ancilla registers if $\xi_j \geq \Delta$ is at most $\frac{1}{2^c} = \sqrt{\epsilon}$. $\square$

COROLLARY 4.1. Let W be a unitary operator with a unique leading eigenvector $|\psi_0\rangle$ with eigenvalue 1. Denote the remaining eigenvectors by $|\psi_j\rangle$ with corresponding eigenvalues $e^{2\pi i \xi_j}$. For any $\Delta \in (0, 1]$ and $\epsilon_2 < 1/2$, define $a := \log(1/\Delta)$ and $c := \log(1/\sqrt{\epsilon_2})$. For any constant $\alpha \in \mathbb{C}$, there

exists a quantum circuit $\tilde{R}$ that uses ac ancilla qubits and invokes the controlled- W gate $2^{a+1}c$ times such that

$$\tilde{R}|\psi_0\rangle|0\rangle^{\otimes ac} = (R|\psi_0\rangle)|0\rangle^{\otimes ac} \quad (4.59)$$

(where $R = \alpha|\psi_0\rangle\langle\psi_0| - (I - |\psi_0\rangle\langle\psi_0|)$) and

$$\|\tilde{R}|\psi_j\rangle|0\rangle^{\otimes ac} - (R|\psi_j\rangle)|0\rangle^{\otimes ac}\| \leq \sqrt{\epsilon_2} \quad (4.60)$$

for $j \neq 0$ such that $\xi_j \geq \Delta$.

PROOF. Let $\tilde{R} := V^\dagger(I \otimes Q)V$ where V is the quantum circuit in Lemma 4.7 and $Q := \alpha|0\rangle\langle 0|^{\otimes ac} + (I - |0\rangle\langle 0|^{\otimes ac})$. Then, we have

$$\tilde{R}|\psi_0\rangle|0\rangle^{\otimes ac} = V^\dagger(I \otimes Q)|\psi_0\rangle|0\rangle^{\otimes ac} = \alpha|\psi_0\rangle|0\rangle^{\otimes ac} = R|\psi_0\rangle|0\rangle^{\otimes ac}. \quad (4.61)$$

For $j \neq 0$ such that $\xi_j \geq \Delta$,

$$\tilde{R}|\psi_j\rangle|0\rangle^{\otimes ac} = V^\dagger(I \otimes Q)(\sqrt{1 - \epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac}) \quad (4.62)$$

$$= V^\dagger(\sqrt{1 - \epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}\alpha|\psi_j\rangle|0\rangle^{\otimes ac}) \quad (4.63)$$

$$= V^\dagger(|\psi_j\rangle \otimes (\sqrt{1 - \epsilon_2}|\chi_j\rangle + \sqrt{\epsilon_2}|0\rangle^{\otimes ac}) + \sqrt{\epsilon_2}(\alpha - 1)|\psi_j\rangle|0\rangle^{\otimes ac}) \quad (4.64)$$

$$= |\psi_j\rangle|0_j\rangle + V^\dagger\sqrt{\epsilon_2}(\alpha - 1)|\psi_j\rangle|0\rangle^{\otimes ac}. \quad (4.65)$$

Thus, $\|\tilde{R}|\psi_j\rangle|0\rangle^{\otimes ac} - (R|\psi_j\rangle)|0\rangle^{\otimes ac}\| \leq \|V^\dagger\sqrt{\epsilon_2}(\alpha - 1)|\psi_j\rangle|0\rangle^{\otimes ac}\| \leq \sqrt{\epsilon_2}$. $\square$

Finally, we prove the following lemma for analyzing the error incurred by $\pi/3$ -amplitude amplification in our quantum volume estimation algorithm:

LEMMA 4.8. *Starting from $|\pi_i\rangle$, we can obtain a state $|\tilde{\pi}_{i+1}\rangle$ such that $\| |\pi_{i+1}\rangle - |\tilde{\pi}_{i+1}\rangle \| \leq \epsilon$ using $\tilde{O}(n^{3/2} \log(1/\epsilon))$ calls to the controlled walk operators W'_i, W'_{i+1} . This results in $\tilde{O}(n^{3/2} \log(1/\epsilon))$ calls to the membership oracle O_K .*

PROOF. From Theorem 2.5, Lemma 4.5, and Lemma 4.6, we find that

- $\pi_i(x)$ mixes up to total variation distance ϵ_1 in $O(n^3 \log^5 \frac{n}{\epsilon_1})$ steps of the Markov chain M_{i+1} , and
- $\pi_{i+1}(x)$ mixes up to total variation distance ϵ_1 in $O(n^3 \log^5 \frac{n}{\epsilon_1})$ steps of the Markov chain M_i .

From Proposition 4.2, we find the following:

- $|\pi_i\rangle = |\pi'_i\rangle + |e_1\rangle$ where $|\pi'_i\rangle$ lies in the space of eigenvectors $|v_j^{(i+1)}\rangle$ of W'_{i+1} such that $\lambda_j^{(i+1)} = 1$ or $\lambda_j^{(i+1)} \leq 1 - \frac{1}{O(n^3 \log^5(n/\epsilon_1))}$, and $\| |e_1\rangle \| \leq \epsilon_1$; and
- $|\pi_{i+1}\rangle = |\pi'_{i+1}\rangle + |e_2\rangle$ where $|\pi'_{i+1}\rangle$ lies in the space of eigenvectors $|v_j^{(i)}\rangle$ of W'_i such that $\lambda_j^{(i)} = 1$ or $\lambda_j^{(i)} \leq 1 - \frac{1}{O(n^3 \log^5(n/\epsilon_1))}$, and $\| |e_2\rangle \| \leq \epsilon_1$.

Note that $|\pi_i\rangle$ and $|\pi_{i+1}\rangle$ are simply the leading eigenvectors of W_i and W_{i+1} , respectively. Thus, both $|\pi_i\rangle$ and $|\pi_{i+1}\rangle$ lie ϵ_1 close to the “good” subspaces corresponding to W'_i (respectively, W'_{i+1}), which are spanned by eigenvectors $|v_j^{(i)}\rangle$ (respectively, $|v_j^{(i+1)}\rangle$) with eigenvalues $e^{2\pi i \xi_j^{(i)}}$ (respectively, $e^{2\pi i \xi_j^{(i+1)}}$) such that $\xi_j^{(i)} = 0$ or $\xi_j^{(i)} \geq \frac{1}{O(n^{3/2} \log^{5/2}(n/\epsilon_1))}$. Each state that occurs during $\pi/3$ -amplitude amplification to rotate $|\pi_i\rangle$ to $|\pi_{i+1}\rangle$ or vice versa is a linear combination of $|\pi_i\rangle$ and $|\pi_{i+1}\rangle$ and is thus also close to the good subspaces of W'_i and W'_{i+1} .

Applying Corollary 4.1 with $\Delta = \frac{1}{n^{3/2} \ln^{5/2}(n/\epsilon_1)}$ and $\epsilon_2 = \epsilon_1^2$, we can implement a quantum operators $\tilde{R}_i, \tilde{R}_{i+1}$ such that $\|R_i - \tilde{R}_i\| \leq 2\epsilon_1$ and $\|R_{i+1} - \tilde{R}_{i+1}\| \leq 2\epsilon_1$, using $O(n^{3/2} \log^{5/2}(n/\epsilon_1) \log(1/\epsilon_1))$ calls to the controlled- W'_i and controlled- W'_{i+1} operators, respectively.

The above shows how to approximately implement R_i and R_{i+1} . If these operators could be implemented perfectly, then Lemma 2.1 and Lemma 4.2 show that we can prepare a state $|\tilde{\pi}_{i+i}\rangle$ such that $\langle \pi_{i+1} | \tilde{\pi}_{i+1} \rangle \leq 1 - (2/3)^{3^m}$ by applying m recursive levels of $\pi/3$ -amplitude amplification to $|\pi_i\rangle$, using 3^m calls to $R_i, R_i^\dagger, R_{i+1}, R_{i+1}^\dagger$. Since $\|\pi_{i+1} - \tilde{\pi}_{i+1}\| = \sqrt{2(1 - \langle \pi_{i+1} | \tilde{\pi}_{i+1} \rangle)}$, after $O(\log(1/\epsilon_2))$ calls to the rotation gates, we obtain a final state with error ϵ_2 . However, each rotation gate can cause an error of ϵ_1 by itself. By making $O(n^{3/2} \log^{5/2}(n/\epsilon_1) \log(1/\epsilon_1) \log(1/\epsilon_2))$ calls to controlled- W'_i and controlled- W'_{i+1} operators, we obtain a final error of $O(\epsilon_1 \log(1/\epsilon_2) + \epsilon_2)$. Choosing $\epsilon_2 = \epsilon/2$ and $\epsilon_1 = \epsilon/(2 \ln(2/\epsilon))$ gives the result. $\square$

Error analysis for the quantum Chebyshev inequality. We also analyze the error from the quantum Chebyshev inequality (Theorem 2.3), giving a robust version of Lemma 4.3.

LEMMA 4.9. *Suppose we have $\tilde{O}(\log(1/\delta)/\epsilon)$ copies of a state $|\tilde{\pi}_{i-1}\rangle$ such that $\| |\tilde{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| \leq \epsilon$. Then the quantum Chebyshev inequality can be used to output $\tilde{V}_i$ such that $|\tilde{V}_i - \mathbb{E}_{\pi_i}[V_i]| \leq O(\epsilon) \mathbb{E}_{\pi_i}[V_i]$ with success probability $1 - \delta^4$ using $\tilde{O}(n^{3/2} \log(1/\delta)/\epsilon)$ calls to the membership oracle. The output state $|\hat{\pi}_{i-1}\rangle$ satisfies $\| |\hat{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| = O(\epsilon + \delta)$.*

PROOF. The error-free version of this lemma was proven in Lemma 4.3. Here, we focus on the error analysis. The quantum Chebyshev inequality uses an implementation of $US_0U^\dagger S_i$ where U is a unitary operator satisfying $U|\pi_{i-1}\rangle = |\pi_i\rangle$. From Lemma 4.8, using $\log(1/\epsilon_2)$ iterations of $\pi/3$ -amplitude amplification ($U_{\log 1/\epsilon_2}$ in Equation (2.5)) instead of U induces an error of ϵ_2 and uses $O(n^{3/2} \log(1/\epsilon_2))$ oracle calls. Using approximate phase estimation as in Corollary 4.1 and Lemma 4.8, Π_{i-1} and Π_i can be implemented up to error ϵ_3 using $O(n^{3/2} \log(1/\epsilon_3))$ oracle calls. Thus, each block corresponding to Theorem 2.2 induces an error of $O(\epsilon_2 + \epsilon_3)$, and the final state before the median is measured has an error of $O(\epsilon + \epsilon_2 + \epsilon_3)$. Therefore, using $O(\log(1/\delta_1)/\epsilon)$ copies of $|\tilde{\pi}_{i-1}\rangle$ returns a sample $\tilde{V}_i$ such that $|\tilde{V}_i - \mathbb{E}_{\pi_i}[V_i]| \leq O(\epsilon_2 + \epsilon_3 + \epsilon) \mathbb{E}_{\pi_i}[V_i]$ with success probability $1 - \delta_1$. Performing a measurement with success probability $1 - \delta_1$ implies that the posterior state has an overlap $\sqrt{1 - \delta_1}$ with the initial state. This induces an error of magnitude at most $\sqrt{2(1 - \sqrt{1 - \delta_1})} = O(\delta_1^{1/4})$.

The measurement on the $\log(1/\delta)/c$ copies of $|\tilde{\pi}_{i-1}\rangle$ used to estimate $\hat{\mu}$ has relative error at most c with probability $1 - \delta$. This causes an error $O(\delta_1^{1/4})$ in addition to the error ϵ_2 from $\pi/3$ -amplitude amplification.

Finally, note that the basic amplitude estimation circuit (analyzed in Theorem 2.2) is a subroutine of the quantum Chebyshev inequality (Theorem 2.3), and uncomputing the block corresponding to Theorem 2.2 induces an error of $O(\epsilon_2 + \epsilon_3)$, giving an overall error of $O(\epsilon_2 + \epsilon_3 + \epsilon + \delta^{1/4})$. The result follows by taking $\epsilon_2 = \epsilon_3 = \epsilon$ and $\delta_1 = \delta^4$. $\square$

We finally prove Lemma 4.4 here.

PROOF. Lemma 4.9 is used to estimate the mean with $\epsilon = \epsilon_1$ and leaves a posterior state $|\hat{\pi}_{i-1}\rangle$ such that $\| |\hat{\pi}_{i-1}\rangle - |\pi_{i-1}\rangle \| = O(\epsilon_1 + \delta)$. We can then use $\pi/3$ -amplitude amplification to rotate this state into $|\tilde{\pi}_i\rangle$, adding error $O(\epsilon')$ at the cost of $O(n^{3/2} \log(1/\epsilon'))$. This completes the proof. $\square$

4.4 Quantum Algorithms for Rounding Logconcave Densities

We first define roundedness of logconcave density functions as follows:

Definition 4.2. A logconcave density function f is said to be c -rounded if

- (1) The level set of f of probability $1/8$ contains a ball of radius r ;
- (2) $\mathbb{E}_f(|x - z_f|) \leq R^2$, where z_f is the centroid of f , i.e., $z_f := \mathbb{E}_f(x)$;

and $R/r \leq c\sqrt{n}$.

In the previous section, we assumed that the distributions π_i sampled during the hit-and-run walk are $O(1)$ -rounded (i.e., well-rounded). From Theorem 2.5, this implies that the hit-and-run walk for the distribution π_i mixes from a warm start in time $\tilde{O}(n^3)$. In this subsection, we show how the distributions π_i can be transformed to satisfy this condition.

Following the classical discussion in Reference [51], we actually show a stronger condition: the distributions are transformed to be in “near-isotropic” position. A density function f is said to be in *isotropic* position if

$$\mathbb{E}_f[x] = 0 \quad \text{and} \quad \mathbb{E}_f[xx^T] = I. \quad (4.66)$$

The latter equation is equivalent to $\int_{\mathbb{R}^n} (u^T x)^2 f(x) dx = |u|^2$ for every vector $u \in \mathbb{R}^n$. We say that K is *near-isotropic* up to a factor of c if

$$\frac{1}{c} \leq \int_{\mathbb{R}^n} (u^T (x - z_f))^2 f(x) dx \leq c \quad (4.67)$$

for every unit vector $u \in \mathbb{R}^n$.

The following lemma shows that logconcave density functions in isotropic position are also $O(1)$ -rounded:

LEMMA 4.10 ([54, LEMMA 5.13]). *Every isotropic logconcave density is $(1/e)$ -rounded.*

The following lemma shows that any logconcave density function can be put into isotropic position by applying an affine transformation, generalizing the same result for uniform distributions by Rudelson [64]:

LEMMA 4.11 ([51, LEMMA 2.2]). *Let f be a logconcave function in $\mathbb{R}^n$ such that there is no linear subspace $S \subseteq \mathbb{R}^n$ such that $\int_S f(x) dx > 1/2$, and let $X^1, \dots, X^k$ be independent random points from the corresponding distribution. There is a constant C_0 such that if $k > C_0 t^3 \ln n$, then the transformation $g(x) = T^{-1/2}x$, where*

$$\bar{X} = \frac{1}{k} \sum_{i=1}^k X^i, \quad T = \frac{1}{k} \sum_{i=1}^k (X^i - \bar{X})(X^i - \bar{X})^T \quad (4.68)$$

puts f in 2-isotropic position with probability at least $1 - 1/2^t$.

From Lemma 4.11, $k = \lceil C_0 n \ln^5 n \rceil = \tilde{\Theta}(n)$ samples from a logconcave density f suffice to put it into near-isotropic position. However, efficiently obtaining samples from a density π_i requires it to be well-rounded to start with. To overcome this difficulty, we interlace the rounding with the stages of the volume estimation algorithm where in each stage, we obtain an affine transformation that puts the density to be sampled in the next stage into isotropic position. The density π_0 is very close to an exponential distribution (since it is concentrated inside the convex body) and can hence be sampled without resorting to a random walk.

To show that samples from π_i can be used to transform π_{i+1} into isotropic position, we use the following lemma:

LEMMA 4.12 ([37, LEMMA 4.3]). *Let f and g be logconcave densities over K with centroids z_f and z_g , respectively. Then for any $u \in \mathbb{R}^n$,*

$$\mathbb{E}_f[(u \cdot (x - z_f))^2] \leq 16 \mathbb{E}_f \left[\frac{f}{g} \right] \mathbb{E}_g[(u \cdot (x - z_g))^2]. \quad (4.69)$$

We now have the following proposition:

PROPOSITION 4.3. *If affine transformation S_i puts π_i in near-isotropic position, then it also puts π_{i+1} in near-isotropic position.*

PROOF. Let S_i put π_i in 2-isotropic position. Applying Lemma 4.12 with $f = \pi_{i+1}, g = \pi_i$, we have that for any unit vector $u \in \mathbb{R}^n$,

$$\mathbb{E}_{\pi_{i+1}}[(u \cdot (x - z_{\pi_{i+1}}))^2] \leq 16 \mathbb{E}_{\pi_{i+1}} \left[\frac{\pi_{i+1}}{\pi_i} \right] \mathbb{E}_{\pi_i}[(u \cdot (x - z_{\pi_i}))^2] \leq 32e^2, \quad (4.70)$$

since $\mathbb{E}_{\pi_{i+1}} \left[\frac{\pi_{i+1}}{\pi_i} \right] \leq e^2$ from Lemma 4.6. Again applying Lemma 4.12

$$\frac{1}{2} \leq \mathbb{E}_{\pi_i}[(u \cdot (x - z_{\pi_i}))^2] \leq \mathbb{E}_{\pi_i} \left[\frac{\pi_i}{\pi_{i+1}} \right] \mathbb{E}_{\pi_{i+1}}[(u \cdot (x - z_{\pi_{i+1}}))^2] \quad (4.71)$$

$\mathbb{E}_{\pi_i} \left[\frac{\pi_i}{\pi_{i+1}} \right] \leq 8$ from Lemma 4.5. Therefore,

$$\frac{1}{2} \leq 128e^2 \mathbb{E}_{\pi_{i+1}}[(u \cdot (x - z_{\pi_{i+1}}))^2]. \quad (4.72)$$

Thus, $E_{\pi_{i+1}}$ is also put in near-isotropic position. $\square$

We finally have the main result of this section:

PROPOSITION 4.4. *At each stage i of Algorithm 4, the affine transformation puts the distribution π_{i+1} in near-isotropic position using an additional $\tilde{O}(n^{2.5})$ quantum queries to O_K .*

PROOF. Since π_0 is nearly an exponential distribution, it can be sampled without using a random walk and thus the proposition is true for $i = 0$. Assume that the proposition is true for $1, 2, \dots, k$. Then an affine transformation can be found to put π_k in near-isotropic position. Thus, a classical hit-and-run walk starting from π_{k-1} converges to π_k in $\tilde{O}(n^3)$ steps. By the analysis in Section 4.3.4, a quantum sample $|\pi_{k-1}\rangle$ can be rotated to $|\pi_k\rangle$ using $\tilde{O}(n^{1.5})$ quantum queries. $\tilde{O}(n)$ such samples suffice to compute the covariance matrix T in Equation (4.68), which puts π_k in 2-isotropic position. By Proposition 4.3, this also puts π_{k+1} in near-isotropic position. This concludes the proof. $\square$

Rounding the convex body as a preprocessing step. Consider applying only the rounding part of Algorithm 4. By Proposition 4.4, the final affine transformation puts the density $\pi_m \propto e^{-a_m x_0}$ in near-isotropic position. Since $a_m \leq \epsilon^2/n$, we have

$$(1 - \epsilon^2) \mathbb{E}_{K'}[|X - \bar{X}|^2] \leq \int_{K'} \frac{e^{-a_m x_0} |x - \bar{x}|^2}{Z(a_m)} dx \leq 2n; \quad (4.73)$$

thus, $\mathbb{E}_{K'}[|X - \bar{X}|^2] \leq 2n/(1 - \epsilon^2)$. From Reference [51, Lemma 3.3], all but an ϵ -fraction of the body is contained inside a ball of radius $O(\sqrt{n})$. Combined with our assumption that $B_2(0, 1) \subseteq K'$, this shows that S_{m+1} puts the convex body K' in well-rounded position.

5 IMPLEMENTATION OF THE QUANTUM HIT-AND-RUN WALK

Due to the precision of representing real numbers, the implementation of volume estimation algorithms in practice requires to walk in a discrete domain that is a subset of $\mathbb{R}^n$. It is known that walks only taking local steps within a short distance (such as the grid walk and the ball walk) can be discretized with good approximation by dividing $\mathbb{R}^n$ into small hypercubes and walking on their centers (see, e.g., Reference [28]), but such error analysis does not automatically apply to hit-and-run walks for which we did not find existing classical discretizations. We emphasize the discretization

ALGORITHM 4: Volume estimation of convex K with interlaced rounding.**Input:** Membership oracle O_K for K .**Output:** ϵ -multiplicative approximation of $\text{Vol}(K)$.

- 1 Set $m = \Theta(\sqrt{n} \log(n/\epsilon))$ to be the number of iterations of simulated annealing and $a_i = 2n(1 - \frac{1}{\sqrt{n}})^i$ for $i \in [m]$. Let π_i be the probability distribution over K' with density proportional to $e^{-a_i x_0}$;
Set error parameters $\delta, \epsilon' = \Theta(\epsilon/m^2)$, $\epsilon_1 = \epsilon/2m$; let $k = \tilde{\Theta}(\sqrt{n}/\epsilon)$ be the number of copies of stationary states for applying the quantum Chebyshev inequality; let $l = \tilde{\Theta}(n)$ be the number of copies of stationary states needed to obtain the affine transformation S_i ; Prepare $k + l$ (approximate) copies of $|\pi_0\rangle$, denoted $|\tilde{\pi}_0^{(1)}\rangle, \dots, |\tilde{\pi}_0^{(k+l)}\rangle$;
- for** $i \in [m]$ **do**
 - 2 Use the quantum Chebyshev inequality on the k copies of the state $|\tilde{\pi}_{i-1}\rangle$ with parameters ϵ_1, δ to estimate the expectation $\mathbb{E}_{\pi_i}[V_i]$ (in Equation (4.7)) as $\tilde{V}_i$ (Lemma 4.9 and Figure 4). The post-measurement states are denoted $|\hat{\pi}_{i-1}^{(1)}\rangle, \dots, |\hat{\pi}_{i-1}^{(k)}\rangle$;
 - 3 Use the l copies of the state $|\pi_{i-1}\rangle$ to nondestructively⁸ obtain the affine transformation $S_i = T = \frac{1}{l} \sum_{q=1}^l (X^q - \bar{X})(X^q - \bar{X})^T$ where the X_q are samples from the density π_{i-1} and $\bar{X} = \frac{1}{l} \sum_{q=1}^l X^q$. The post-measurement states are denoted $|\hat{\pi}_{i-1}^{(k+1)}\rangle, \dots, |\hat{\pi}_{i-1}^{(k+l)}\rangle$;
 - 4 Apply $\pi/3$ -amplitude amplification with error ϵ' (Section 2.2 and Lemma 4.8) and affine transformation S_i to map $|S_i \hat{\pi}_{i-1}^{(1)}\rangle, \dots, |S_i \hat{\pi}_{i-1}^{(k+l)}\rangle$ to $|S_i \tilde{\pi}_i^{(1)}\rangle, \dots, |S_i \tilde{\pi}_i^{(k+l)}\rangle$, using the quantum hit-and-run walk;
 - 5 Invert S_i to get $k + l$ (approximate) copies of the stationary distribution $|\pi_i\rangle$ for use in the next iteration;
- 6 Compute an estimate $\widehat{\text{Vol}}(K') = n! v_n (2n)^{-(n+1)} \tilde{V}_1 \dots \tilde{V}_m$ of the volume of K' , where v_n is the volume of the n -dimensional unit ball;
- 7 Use $\widehat{\text{Vol}}(K')$ to estimate the volume of K as $\widehat{\text{Vol}}(K)$ (Section 4.3.1).

in contrast to most classical treatments for two reasons: (1) Quantum algorithms are typically presented in a circuit model, in contrast to the RAM model used by classical algorithms. Continuous variables in the circuit model correspond to registers of infinite size, preventing a clear analysis of the resources of the algorithm in terms of gate count. Specifically, to obtain the performance of the algorithm in reality, we must show that $\text{poly}(\log(1/\epsilon))$ bit registers suffice. (2) Standard methods of preparing walk operators corresponding to classical Markov Chains (see, for example, Reference [70]) rely on the sparsity of the transition matrix. In the case of geometric random walks sparsity is not well-defined in the continuous case and may not hold even for discretizations (for example, the hit-and-run walk has a non-zero transition density to any point in the convex body.) The efficient preparation of quantum states corresponding to classical distributions is not always a trivial operation, and there has been research [35] about preparing common distributions for quantum Monte Carlo methods. Most existing general procedures come without provable guarantees on the resources required for sufficiently accurate samples; we provide here a simple analysis for the cost of implementing the hit-and-run walk via the Grover-Rudolph method [30].

In this section, we introduce a discretized quantum hit-and-run walk and give an explicit analysis of its implementation. The basic idea of the discretization is to represent the coordinates with

⁸Similar to Lemma 4.3, we do not directly measure the states; instead, we use a quantum circuit to (classically) compute the affine transformation S_i and apply it to the convex body coherently for the next iteration. Note that the quantum register holding the affine transformation will be in some superposition, but by using $O(\log n)$ copies and taking the mean (as in Lemma 4.3), the amplitude of the correct affine transformation will be arbitrarily close to 1.

rational numbers. We approximate K by a set of discretized points in K and define a Markov chain on these points (see Section 5.1). We use a two-level discretization: the hit-and-run process is performed with a coarser discretization and then a point in a finer discretization of the coarse grid is chosen uniformly at random as the actual point to jump to. This ensures that the starting and ending points (in the coarser discretization) of one jump are far from the boundary so that a small perturbation does not change the length of the chord induced by the two points significantly. Then in Section 5.2, the discrete conductance can be bounded by bounding the distance between the discrete and continuous transition probabilities as well as the distance between the discrete and continuous subset measures. In Section 5.3, we prove that the quantum gate complexity of implementing the discretized quantum hit-and-run walk is $\tilde{O}(n)$, the same overhead as for implementing classical hit-and-run walks.

5.1 Discretization of the Hit-and-run Walk

For a convex body $K \subseteq \mathbb{R}^n$, we let K_ϵ denote the set of vectors in K whose coordinates can be represented by some fixed-point representation using $\log(1/\epsilon)$ bits. (Note that this ϵ is different from the multiplicative error in the problem definition. However, this ϵ is not the dominating error and the overhead is only logarithmic.) We call K_ϵ an ϵ -discretization of K . The finite set K_ϵ provides an ϵ -net for K . We also define $(\mathbb{R}^n)_\epsilon$ as a ϵ -discretization of $\mathbb{R}^n$.

We consider a Markov chain whose states are the points in K_ϵ . For any $v \in \mathbb{R}^n$, we define the ϵ -box $b_\epsilon(v) := \{x \in \mathbb{R}^n : x(i) \in [v(i) - \epsilon/2, v(i) + \epsilon/2], \forall i \in [n]\}$. Let $\bar{K}_\epsilon$ be the continuous set formed by the ϵ -boxes of the points in K_ϵ , i.e., $\bar{K}_\epsilon = \bigcup_{x \in K_\epsilon} b_\epsilon(x)$. For two distinct points $u, v \in \mathbb{R}^n$, we denote by ℓ_{uv} the line through them. For a line $\ell \subseteq \mathbb{R}^n$, let $\ell(\bar{K}_\epsilon)$ be the segment of ℓ contained in $\bar{K}_\epsilon$, i.e., $\ell(\bar{K}_\epsilon) = \{x \in \ell : x \in \bar{K}_\epsilon\}$. In addition, for $u \in \ell$, we define $\ell(\bar{K}_\epsilon, u, \epsilon')$ as the ϵ' -discretization of $\ell(\bar{K}_\epsilon)$ starting from u , i.e., $\ell(\bar{K}_\epsilon, u, \epsilon') = \{x \in \ell(\bar{K}_\epsilon) : |x - u| = k\epsilon' \text{ for some } k \in \{0, 1, \dots\}\}$. Analogous to the distribution π_f for the continuous-space case, we define its corresponding discrete distribution $\hat{\pi}_f$ with $\hat{\pi}_f(S) = \sum_{x \in S} f(x) / \sum_{x \in (\mathbb{R}^n)_\epsilon} f(x)$.

To implement the hit-and-run walk (see Section 2.4), we sample a uniformly random direction from a point u . We achieve this by sampling n coordinates according to the standard normal distribution from the corresponding coordinate of u and normalizing the new point to have unit length; the uniformity of such sampling is well known; see, for example, References [56, 58]. (A one-line proof is that this distribution is invariant under orthogonal transformations, but the uniform distribution on the n -dimensional unit sphere B_n is the unique distribution that satisfies this property. Although the Gaussian distributions we sample from are discretized, the invariance under orthogonal transformations holds approximately, so we have approximate uniformity.) Let this normalized point be v , so that the sampled direction is ℓ_{uv} . Note that the coordinate we sample from is discrete. The directions we can sample form a discrete set denoted $L(u, \epsilon')$, where ϵ' is the precision for sampling directions.

Now, we compute the probability that a specific direction is sampled. After the normalization procedure described in the above paragraph, the sampled point u will “snap” to a point in $(\mathbb{R}^n)_\epsilon$, i.e., v . Consider $\bigcup_{v: b_{\epsilon'}(v) \cap B_n \neq \emptyset} b_\epsilon(v)$, where B_n is the n -dimensional unit sphere. We use the $(n-1)$ -dimensional volume (surface area) of this body to approximate that of B_n , with up to a $\sqrt{2}$ enlargement factor due to the fact that ϵ -boxes have sharp corners. Thus, the number of points that v can snap to is in the range $[n \text{Vol}(B_n)/\epsilon^{n-1}, \sqrt{2}n \text{Vol}(B_n)/\epsilon^{n-1}]$, which is also the range of $|L(u, \epsilon')|$. To make the lines in $L(u, \epsilon')$ cover every ϵ -box on the boundary of $\sqrt{n}B_n$ (so that it is possible to sample all the points in K_ϵ), we need $\epsilon' \leq \epsilon/\sqrt{n}$.

Let $L := |L(u, \epsilon')|$. We label the lines in $L(u, \epsilon')$ as $\{\ell_1, \dots, \ell_L\}$ (ordered arbitrarily). For each $i \in [L]$, let v_i be the point after normalization. Intuitively, v_i approximates a point on the “surface”




Fig. 6. Constructing a hyperpyramid. The inner circle represents the unit ball and the outer circle represents the ball of radius $\sqrt{n}$. The grids represents the ϵ' -discretization of $\mathbb{R}^n$; each grid is an ϵ' -box. The shaded boxes are points where a direction “snaps” to after normalization, and the dashed edges of $b_{\epsilon'}(v_i)$ is its “outer face.” The hyperpyramid P_i is represented by a circular sector.

of the unit ball around u (see Figure 6). There are *hyperfaces* of $b_{\epsilon'}(v_i)$ that are out-facing and not adjacent to any ϵ' -box in $\bigcup_{v: b_{\epsilon'}(v) \cap B_n \neq \emptyset} b_{\epsilon}(v)$ (as an illustration, see dashed edges in Figure 6). For all points v'' in these hyperfaces, the line segments from u through v'' of length $\sqrt{n}$ form a set, which we refer to as a *hyperpyramid*, denoted by P_i . The apex of each hyperpyramid is u , and the base of each hyperpyramid is a subset of the hyperspherical surface. Intuitively, the bases of $P_1, \dots, P_L$ form a partition of the “surface” of the ball of radius $\sqrt{n}$ around u , and therefore, $\{P_1, \dots, P_L\}$ forms a partition of the ball of radius $\sqrt{n}$ around u .

5.2 Conductance Lower Bound on the Discretized Hit-and-run Walk

The discretized hit-and-run walk on K_{ϵ} described above can be summarized as Algorithm 5.

Note that we have used a two-level discretization of K , as illustrated in Figure 7. The first level is a coarser discretization $K_{\sqrt{\epsilon}n^{1/4}}$ and the second level is a finer discretization K_{ϵ} . We first choose a temporary point v'' in $K_{\sqrt{\epsilon}n^{1/4}}$. Then, we choose a point v uniformly at random in $b_{\sqrt{\epsilon}n^{1/4}}(v'') \cap (\mathbb{R}^n)_{\epsilon}$ to jump to. The purpose of this two-level discretization is to avoid having a small change of the original point u cause a huge difference in $\ell_{uv}(\bar{K}_{\epsilon})$ (when u is very close to the boundary).

We first compute the transition probability of the discretized hit-and-run walk. The following lemma states a lower bound on the transition probability of the discretized hit-and-run walk, which




Fig. 7. A demonstration of the two-level discretization of K . The thicker grid represents the coarser discretization $K_{\sqrt{\epsilon}n^{1/4}}$ and the thinner grid represents the finer discretization K_ϵ . When v'' is chosen from $K_{\sqrt{\epsilon}n^{1/4}}$, an actual point v to jump is chosen uniformly at random in $b_{\sqrt{\epsilon}n^{1/4}}(v'') \cap (\mathbb{R}^n)_\epsilon$ marked by the points in the shaded region.

ALGORITHM 5: One step of the discretized hit-and-run walk.

Input: Current point $u \in K_\epsilon$.

- 1 Uniformly sample a line $\ell \in L(u, \epsilon)$ by independently sampling n coordinates around u according to the standard normal distribution and then normalizing to unit length;
 - 2 Sample a point v' in $\ell(\bar{K}_\epsilon, u, \epsilon')$ according to f ;
 - 3 Let $v'' \in K_{\sqrt{\epsilon}n^{1/4}}$ that is closest to v' ;
 - 4 Output a uniform sample v in $b_{\sqrt{\epsilon}n^{1/4}}(v'') \cap (\mathbb{R}^n)_\epsilon$.
-

will be used to establish the relationship between the transition probabilities of discretized and continuous hit-and-run walks as in Equation (5.17). This relationship eventually leads to a lower bound on the conductance of the discretized hit-and-run walk as in Theorem 5.1.

LEMMA 5.1. *The transition probabilities defined by Algorithm 5 satisfy*

$$P_{uv} \geq \sum_{\substack{v' \in \ell(\bar{K}_\epsilon, u, \epsilon'), \ell \in L(u, \epsilon): \\ \ell(\bar{K}_\epsilon, u, \epsilon') \cap b_{\sqrt{\epsilon}n^{1/4}}(v) \neq \emptyset}} \frac{\epsilon^{n-1}(\sqrt{\epsilon})^n f(v')}{\sqrt{2}n^{1+n/4} \text{Vol}(B_n)(\sqrt{n})^{n-1} \hat{\mu}_f(\ell(\bar{K}_\epsilon, u, \epsilon'))}, \quad (5.1)$$

where for any $S \subseteq \mathbb{R}^n$, we define

$$\hat{\mu}_f(S) := \sum_{x \in S} f(x). \quad (5.2)$$

PROOF. First note that the probability of a line $\ell \in L(u, \epsilon)$ being sampled is at least $\frac{\epsilon^{n-1}}{\sqrt{2n} \text{Vol}(\text{B}_n)(\sqrt{n})^{n-1}}$. Along ℓ , the probability of sampling v' is $f(v')/\hat{\mu}_f(\ell(\bar{K}_\epsilon, u, \epsilon'))$, and the probability of choosing v in $b_{\sqrt{\epsilon}n^{1/4}}(v'') \cap K_\epsilon$ is $(\sqrt{\epsilon})^n/n^{n/4}$. $\square$

According to the definition in Equation (2.1), the conductance of any subset $S \subseteq K_\epsilon$ is

$$\phi(S) = \frac{\sum_{u \in S} \sum_{v \in K_\epsilon \setminus S} P_{uv} \hat{\pi}_f(u)}{\min\{\hat{\pi}_f(S), \hat{\pi}_f(K_\epsilon \setminus S)\}}, \quad (5.3)$$

where $\hat{\pi}_f$ is defined as $\hat{\pi}_f(A) = \sum_{x \in A} f(x)$. The conductance of the Markov chain is then

$$\phi = \min_{S \subseteq K_\epsilon} \phi(S). \quad (5.4)$$

Now, we prove the main theorem of this section, which shows that the conductance of the discretized hit-and-run walk does not differ significantly from that of the continuous hit-and-run walk.

THEOREM 5.1. *Let K_ϵ be the discretization of convex body K that contains a unit ball and is contained in a ball with radius $R \leq \sqrt{n}$. Let the density function be $f(x) = e^{-a^T x}$ having support K where $a = (1, 0, \dots, 0)$. Let $\epsilon' \leq \sqrt{\epsilon}n^{-3/4}$. For $S \subseteq K_\epsilon$ such that $\hat{\pi}_f(S) \leq 1/2$, we have*

$$\phi(S) \geq \frac{1}{10^{16} n \sqrt{n} \ln\left(\frac{2n\sqrt{n}}{\hat{\pi}_f(S)}\right)} - \epsilon. \quad (5.5)$$

PROOF. This proof closely follows that of Reference [52, Theorem 6.9]. We first consider the transition probability for the continuous hit-and-run walk in K . For $u, v \in K$, recall that

$$P'_u(b_\epsilon(v)) = \frac{2}{n \text{Vol}(\text{B}_n)} \int_{b_\epsilon(v)} \frac{f(x) dx}{\mu_f(u, x)|x - u|^{n-1}}, \quad (5.6)$$

where $\mu_f(u, x)$ is a shorthand for $\mu_f(\ell_{ux}(\bar{K}_\epsilon))$. We compare $P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v))$ with P_{uv} for $u \in K_\epsilon$ and $v \in K_{\sqrt{\epsilon}n^{1/4}}$. To this end, we use $\hat{\mu}_f$ to approximate μ_f : for each ℓ , we have

$$\epsilon' \hat{\mu}_f(\ell(\bar{K}_\epsilon, u, \epsilon')) \leq e^{\epsilon'} \mu_f(\ell(\bar{K}_\epsilon)). \quad (5.7)$$

Consider each hyperpyramid P_i defined in Section 5.1 whose associated line through its apex is ℓ_i and $\ell_i(\bar{K}_\epsilon, u, \epsilon') \cap b_{\sqrt{\epsilon}n^{1/4}}(v) \neq \emptyset$. Note that the distance between each $u \in K_\epsilon$ and the boundary of $\bar{K}_\epsilon$ is at least $\epsilon/2$. Inside each hyperpyramid, the length of the chords through u can differ by a factor at most 2. For each $\ell \subset P_i$, $\hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) \leq 2e^{\epsilon'} \hat{\mu}_f(\ell(\bar{K}_\epsilon, u, \epsilon'))$. Together with Equation (5.7), it follows that

$$\epsilon' \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) \leq 2e^{2\epsilon'} \mu_f(\ell(\bar{K}_\epsilon)) \quad (5.8)$$

for all $\ell \subset P_i$. Define $c_i := |\ell_i(\bar{K}_\epsilon, u, \epsilon') \cap b_{\sqrt{\epsilon}n^{1/4}}(v)|$ (the number of points in this set) and $d_i := |\ell_i(\bar{K}_\epsilon) \cap b_{\sqrt{\epsilon}n^{1/4}}(v)|$ (the length of this line). Note that $c_i \leq d_i/\epsilon'$. We further partition P_i into c_i sets $Q_{i,1}, \dots, Q_{i,c_i}$ along the direction of ℓ_i so that the distance between the hyperplanes that separate adjacent sets is at most ϵ' . For each $j \in [c_i]$, we have

$$\begin{aligned} \frac{\epsilon^{n-1} f(v')}{n \text{Vol}(\text{B}_n)(\sqrt{n})^{n-1} \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon'))} &= \frac{\epsilon^{n-1} f(v') \epsilon' |v' - u|^{n-1}}{\epsilon' n \text{Vol}(\text{B}_n)(\sqrt{n})^{n-1} \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) |v' - u|^{n-1}} \\ &\geq \frac{f(v') \text{Vol}(Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v))}{2\epsilon' n \text{Vol}(\text{B}_n) \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) |v' - u|^{n-1}}, \end{aligned} \quad (5.9)$$

where we have used the fact that the distance between adjacent $Q_{i,j}$ and $Q_{i,j+1}$ can be bounded from below by $|Q_{i,j} \cap \ell_i|/(1 + \epsilon'/2) \geq |Q_{i,j} \cap \ell_i|/2$.

Now, we consider the integration in $Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v)$. We use $f(v)$ to approximate $f(v')$, which causes a relative error at most $e^{\sqrt{\epsilon}n^{1/4}}$, and use $|v' - u|^{n-1}$ to approximate $|x - u|^{n-1}$ for all $x \in Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v)$, which causes a relative error at most e provided $\epsilon' \leq \sqrt{\epsilon}n^{-3/4}$ (noting that the distance between x and u is at most $\sqrt{\epsilon}n^{1/4}$). We have

$$\begin{aligned} & \int_{Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v)} \frac{f(x) dx}{n \text{Vol}(\text{B}_n) \mu_f(u, x) |x - u|^{n-1}} \\ & \leq \frac{2e^{\sqrt{\epsilon}n^{1/4}+2\epsilon'+1} f(v')}{\epsilon' n \text{Vol}(\text{B}_n) \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) |v' - u|^{n-1}} \int_{Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v)} dx \end{aligned} \quad (5.10)$$

$$= \frac{2e^{\sqrt{\epsilon}n^{1/4}+2\epsilon'+1} f(v') \text{Vol}(Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v))}{\epsilon' n \text{Vol}(\text{B}_n) \hat{\mu}_f(\ell_i(\bar{K}_\epsilon, u, \epsilon')) |v' - u|^{n-1}}, \quad (5.11)$$

where the inequality follows from Equation (5.8). Let $i_1, \dots, i_t$ be the indices such that $P_{i_j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v) \neq \emptyset$ for $j \in [t]$. We use $\bigcup_{j \in [t]} P_{i_j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v)$ as a partition to approximate $b_{\sqrt{\epsilon}n^{1/4}}(v)$, which causes a relative error at most $(1 + \epsilon)^n$ for $\text{Vol}(b_{\sqrt{\epsilon}n^{1/4}}(v))$. We have

$$\int_{b_{\sqrt{\epsilon}n^{1/4}}(v)} \frac{f(x) dx}{\mu_f(u, x) |x - u|^{n-1}} \leq (1 + \epsilon)^n \sum_{j \in [t]} \int_{b_{\sqrt{\epsilon}n^{1/4}}(v) \cap P_{i_j}} \frac{f(x) dx}{\mu_f(u, x) |x - u|^{n-1}}.$$

Hence,

$$\begin{aligned} & \frac{(\sqrt{\epsilon})^n}{n^{n/4}} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) \\ & = \frac{2(\sqrt{\epsilon})^n}{n^{1+n/4} \text{Vol}(\text{B}_n)} \int_{b_{\sqrt{\epsilon}n^{1/4}}(v)} \frac{f(x) dx}{\mu_f(u, x) |x - u|^{n-1}} \end{aligned} \quad (5.12)$$

$$\leq \frac{2(\sqrt{\epsilon})^n (1 + \epsilon)^n}{n \text{Vol}(\text{B}_n)} \sum_{j \in [t]} \int_{b_{\sqrt{\epsilon}n^{1/4}}(v) \cap P_{i_j}} \frac{f(x) dx}{\mu_f(u, x) |x - u|^{n-1}} \quad (5.13)$$

$$= \frac{2(\sqrt{\epsilon})^n (1 + \epsilon)^n}{n \text{Vol}(\text{B}_n)} \sum_{j \in [t]} \sum_{k \in [c_{i_j}]} \int_{b_{\sqrt{\epsilon}n^{1/4}}(v) \cap Q_{i_j,k}} \frac{f(x) dx}{\mu_f(u, x) |x - u|^{n-1}} \quad (5.14)$$

$$\leq \sum_{j \in [t]} \sum_{k \in [c_{i_j}]} \frac{4(1 + \epsilon)^n e^{\sqrt{\epsilon}n^{1/4}+2\epsilon'+1} (\sqrt{\epsilon})^n f(v') \text{Vol}(Q_{i,j} \cap b_{\sqrt{\epsilon}n^{1/4}}(v))}{\epsilon' n \text{Vol}(\text{B}_n) \hat{\mu}_f(\ell_{i_j}(\bar{K}_\epsilon, u, \epsilon')) |u - v'|^{n-1}} \quad (5.15)$$

$$\leq 4(1 + \epsilon)^n e^{\sqrt{\epsilon}n^{1/4}+2\epsilon'+1} \sum_{j \in [t]} \sum_{k \in [c_{i_j}]} \frac{2e^{n-1} (\sqrt{\epsilon})^n f(v')}{n \text{Vol}(\text{B}_n) (\sqrt{n})^{n-1} \hat{\mu}(\ell_{i_j}(\bar{K}_\epsilon, u, \epsilon'))} \quad (5.16)$$

$$= 4(1 + \epsilon)^n e^{\sqrt{\epsilon}n^{1/4}+2\epsilon'+1} P_{uv} \leq e^{5+2\epsilon'} P_{uv}, \quad (5.17)$$

where the second inequality follows from Equation (5.11), and the last inequality holds when $\epsilon \leq 1/n$.

For $u \in K_\epsilon$ and $v \in K_{\sqrt{\epsilon}n^{1/4}}$, we approximate $\int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x)$ by $\epsilon^n P_{uv}$. Note that for all $u' \in b_\epsilon(u)$, we have $|u' - v|^n \leq e|u - v|^n$. Also, the lengths of ℓ_{uv} and $\ell_{u'v}$ can differ by at most a factor of 2. As a result, $\hat{\pi}_f(\ell_{u'v}(\bar{K}_\epsilon, u, \epsilon')) \leq 2\hat{\pi}_f(\ell_{uv}(\bar{K}_\epsilon, u', \epsilon'))$. It follows that $P_{uv} \geq P_{u'v}/(2e)$.

Therefore,

$$\int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x) \leq \int_{x \in b_\epsilon(u)} \frac{2e^{5+2\epsilon'} n^{n/4}}{(\sqrt{\epsilon})^n} P_{xv} d\pi_f(x) \quad (5.18)$$

$$\leq \frac{2e^{5+2\epsilon'+\epsilon} n^{n/4}}{(\sqrt{\epsilon})^n} P_{uv} \hat{\pi}_f(u) \epsilon^n. \quad (5.19)$$

Next, for the relationship between $\hat{\pi}_f$ and π_f , we consider the sets $\bar{K}_\epsilon \cap K$, $\bar{K}_\epsilon \setminus K$, and $K \setminus \bar{K}_\epsilon$ separately. Without loss of generality, assume $\hat{\pi}_f(S) \leq \hat{\pi}_f(K_\epsilon \setminus S)$. We partition S as $S_1 \cup S_2$, where $S_1 = \{x \in S : b_\epsilon(x) \subseteq K\}$ and $S_2 = S \setminus S_1$. We also define $\bar{S}_1 := \bigcup_{x \in S_1} b_\epsilon(x)$ and $\bar{S}_2 := \bigcup_{x \in S_2} b_\epsilon(x)$. For $\bar{S}_1$, we use $f(v)$ to approximate $f(x)$ for all $x \in b_\epsilon(v)$; it follows that

$$\hat{\pi}_f(S_1) \leq e^{2\epsilon} \pi_f(\bar{S}_1) \quad \text{and} \quad \pi_f(\bar{S}_1) \leq e^{2\epsilon} \hat{\pi}_f(S_1). \quad (5.20)$$

For S_2 , we have

$$\hat{\pi}_f(S_2) \leq 2e^{2\epsilon} \pi_f(\bar{S}_2 \cap K). \quad (5.21)$$

Combining the above inequalities, we have

$$\hat{\pi}_f(S) = \hat{\pi}_f(S_1) + \hat{\pi}_f(S_2) \leq e^{2\epsilon} \pi_f(\bar{S}_1) + 2e^{2\epsilon} \pi_f(\bar{S}_2 \cap K) \leq 3\pi_f(K \cap \bar{S}). \quad (5.22)$$

Now, we bound the numerator of the conductance: $\sum_{u \in S} \sum_{v \in K_\epsilon \setminus S} P_{uv} \hat{\pi}_f(u)$. For $u \in S$ and $v \in K \setminus S$, we consider four cases. First, when $b_\epsilon(u), b_\epsilon(v) \subseteq K$, we have

$$P_{uv} \hat{\pi}_f(u) \geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x). \quad (5.23)$$

Second, when $b_\epsilon(u) \subseteq K$ and $b_\epsilon(v) \not\subseteq K$, we have

$$P_{uv} \hat{\pi}_f(u) \geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x) \quad (5.24)$$

$$\geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v) \cap K) d\pi_f(x). \quad (5.25)$$

Third, when $b_\epsilon(u) \not\subseteq K$ and $b_\epsilon(v) \subseteq K$, we have

$$P_{uv} \hat{\pi}_f(u) \geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x) \quad (5.26)$$

$$\geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u) \cap K} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x). \quad (5.27)$$

Fourth, when $b_\epsilon(u) \not\subseteq K$ and $b_\epsilon(v) \not\subseteq K$, we have

$$P_{uv} \hat{\pi}_f(u) \geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u)} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v)) d\pi_f(x) \quad (5.28)$$

$$\geq \frac{(\sqrt{\epsilon})^n}{2e^{5+2\epsilon'+\epsilon} n^{n/4}} \int_{x \in b_\epsilon(u) \cap K} P'_u(b_{\sqrt{\epsilon}n^{1/4}}(v) \cap K) d\pi_f(x). \quad (5.29)$$

We also need to consider the set $K \setminus \bar{K}_\epsilon$. There exists a small subset $E \subseteq K \setminus \bar{K}_\epsilon$ such that $\pi_f(E) \leq \epsilon \pi_f(S)$. We need to consider the transition from E to $\subseteq K \setminus \bar{K}_\epsilon \setminus E$: We have $\int_{x \in E} P'_x(K \setminus \bar{K}_\epsilon \setminus$

E) $d\pi_f(x) \leq \pi_f(E) \leq \epsilon\pi_f(S)$. Putting everything together, we have

$$\begin{aligned} \sum_{u \in S} \sum_{v \in K_\epsilon \setminus S} P_{uv} \hat{\pi}_f(u) + \int_{x \in E \cap K} P'_x(K \setminus \bar{K}_\epsilon \setminus E) d\pi_f(x) \\ \geq \frac{1}{2e^{5+2\epsilon'+\epsilon}} \int_{x \in \bar{S} \cap K \cup E} P'_x(K \setminus (\bar{S} \cap K \cup E)) d\pi_f(x), \end{aligned} \quad (5.30)$$

which further implies that

$$\begin{aligned} \sum_{u \in S} \sum_{v \in K_\epsilon \setminus S} P_{uv} \hat{\pi}_f(u) &\geq \frac{1}{2e^{5+2\epsilon'+\epsilon}} \int_{x \in \bar{S} \cap K \cup E} P'_x(K \setminus (\bar{S} \cap K \cup E)) d\pi_f(x) - \epsilon\pi_f(S) \\ &\geq \frac{1}{2e^{5+2\epsilon'+\epsilon}} \int_{x \in \bar{S} \cap K \cup E} P'_x(K \setminus (\bar{S} \cap K \cup E)) d\pi_f(x) - \epsilon e^\epsilon \hat{\pi}_f(S). \end{aligned} \quad (5.31)$$

By Proposition 2.3, we have

$$\phi(S) = \frac{\sum_{u \in S} \sum_{v \in K_\epsilon \setminus S} P_{uv} \hat{\pi}_f(u)}{\hat{\pi}_f(S)} \quad (5.32)$$

$$\geq \frac{1}{2e^{5+2\epsilon'+\epsilon}} \frac{\int_{x \in \bar{S} \cap K \cup E} P'_x(K \setminus (\bar{S} \cap K \cup E)) d\pi_f(x)}{\hat{\pi}_f(S)} - \frac{\epsilon}{2e^{4+2\epsilon'}} \quad (5.33)$$

$$\geq \frac{1}{6e^{5+2\epsilon'+\epsilon}} \frac{\int_{x \in \bar{S} \cap K \cup E} P'_x(K \setminus (\bar{S} \cap K \cup E)) d\pi_f(x)}{\pi_f(\bar{S} \cap K) + \pi_f(E)} - \frac{\epsilon}{2e^{5+2\epsilon'}} \quad (5.34)$$

$$\geq \frac{1}{10^{14} e^{5+2\epsilon'+\epsilon} n \sqrt{n} \ln\left(\frac{n \sqrt{n}}{\pi_f(\bar{S} \cap K)}\right)} - \frac{\epsilon}{2e^{5+2\epsilon'}} \quad (5.35)$$

$$\geq \frac{1}{10^{14} e^{5+2\epsilon'+\epsilon} n \sqrt{n} \ln\left(\frac{n \sqrt{n}}{(1-e^{-\epsilon}/2)e^\epsilon \hat{\pi}_f(S)}\right)} - \frac{\epsilon}{2e^{5+2\epsilon'}}, \quad (5.36)$$

where the second inequality follows from Equation (5.22), and the third inequality follows from Proposition 2.3. The above inequality can then be simplified to

$$\phi(S) \geq \frac{1}{10^{16} n \sqrt{n} \ln\left(\frac{2n \sqrt{n}}{\hat{\pi}_f(S)}\right)} - \epsilon, \quad (5.37)$$

which is exactly the claim in Theorem 5.1. $\square$

The mixing time for the discrete hit-and-run walk can be bounded by the following corollary.

COROLLARY 5.1. *Let K_ϵ be the discretization of convex body K that contains a unit ball and is contained in a ball with radius $R \leq \sqrt{n}$. Let the density function be $f(x) = e^{-a^T x}$ having support K where $a = (1, 0, \dots, 0)$. Let $\epsilon' \leq \sqrt{\epsilon} n^{-3/4}$. Let the initial distribution be σ and the distribution after m steps be σ^m . If $\sum_{x \in K_\epsilon} \frac{\sigma(x)}{\hat{\pi}_f(x)} \sigma(x) \leq M$, then after*

$$m \geq 10^{33} n^3 \ln^2 \frac{M n \sqrt{n}}{\epsilon} \ln \frac{M}{\epsilon} \quad (5.38)$$

steps, we have $d_{TV}(\sigma^m, \hat{\pi}_f) \leq \epsilon$.

PROOF. First note that, since $\sum_{x \in K_\epsilon} \frac{\sigma(x)}{\hat{\pi}_f(x)} \sigma(x) \leq M$, the set $S = \{x : \frac{\sigma(x)}{\hat{\pi}_f(x)} > \frac{2M}{\epsilon}\}$ has measure $\sigma(S) \leq \epsilon/2$. Then a random point in K_ϵ can be thought of as being generated with probability

$1 - \epsilon/2$ from a distribution σ' satisfying $\frac{\sigma'(S')}{\hat{\pi}_f(S')} \leq 2M/\epsilon$ for any subset $S' \subseteq K_\epsilon$ and with probability $\epsilon/2$ from some other distribution. As a consequence of Theorem 5.1, for any such subset S' with $\hat{\pi}_f(S') = p$, the conductance of S' is at least

$$\Phi_p = \frac{1}{10^{16} n \sqrt{n} \ln(2n\sqrt{n}/p)} - \epsilon. \quad (5.39)$$

For the purpose of analysis, we use $p = \frac{\epsilon^2}{8M}$. When ϵ is reasonably small (say, $\epsilon \leq \frac{1}{2 \cdot 10^{16} n \sqrt{n} \ln(Mn\sqrt{n}/\epsilon)}$), the ϵ term in the conductance bound can be ignored with an additional $1/2$ factor. Then, we have $\Phi_p \geq \frac{1}{2 \cdot 10^{16} n \sqrt{n} \ln(2n\sqrt{n}/p)}$. By the condition that $\sigma'(S') \leq (2M/\epsilon)\hat{\pi}_f(S')$, as well as the way a random point in K_ϵ is generated, Proposition 2.2 implies that

$$d_{TV}(\sigma^{(m)}, \hat{\pi}_f) \leq \frac{\epsilon}{2} + \left(1 - \frac{\epsilon}{2}\right) \left(\frac{\epsilon}{2} + \frac{4M}{\epsilon} \left(1 - \frac{\Phi_p^2}{2}\right)^m\right). \quad (5.40)$$

Therefore, after the claimed number of steps, the total variation distance is at most ϵ . $\square$

As the uniform distribution is a special case of a log-concave distribution, the proof of Theorem 5.1 also applies to this case. More specifically, we use Proposition 2.4 in Equation (5.35), which yields the following stronger corollary.

COROLLARY 5.2. *Let K_ϵ be the discretization of a convex body K that contains a unit ball and is contained in a ball with radius $R \leq \sqrt{n}$. Let $\epsilon' \leq \sqrt{\epsilon} n^{-3/4}$. The conductance of the hit-and-run walk in K_ϵ with uniform distribution satisfies*

$$\phi \geq \frac{1}{2^{26} n \sqrt{n}} - \epsilon. \quad (5.41)$$

Note that Corollary 5.2 is stronger than Theorem 5.1, because Equation (5.41) is independent of $S \subseteq K_\epsilon$. This corollary is informative and is not used in this article.

5.3 Implementing the Quantum Walk Operators

We now describe how to implement the discretized quantum walk. Following Equation (1.2), consider a convex body K such that $B_2(0, r) \subseteq K \subseteq B_2(0, R)$. Each stage of the volume estimation algorithm involves a hit-and-run walk over the convex body with target density e^{-ax_0} . To use techniques from Reference [74] to obtain a speedup in mixing time, we implement the quantum walk operator W corresponding to an ϵ -discretized version of this walk Algorithm 5.

Let $|x\rangle$ be the register for the state of the walk, and U be a unitary that satisfies $U|x\rangle|0\rangle = |x\rangle|p_x\rangle$ for all $|x\rangle$ (recall that $|p_x\rangle = \sum_{y \in K_\epsilon} \sqrt{p_{x \rightarrow y}} |y\rangle$ where $p_{x \rightarrow y}$ is the probability of a transition from x to y). Since the state of the hit-and-run walk is given by points on an ϵ -grid that can be restricted to $B_2(0, R)$, there are $(\frac{2R}{\epsilon})^n$ possible values of x and thus $|x\rangle$ can be represented using $n \log(\frac{2R}{\epsilon})$ qubits. In the rest of the section, we abuse notation by letting x refer to both a point on the grid and its corresponding bit representation. Then the quantum walk operator [74] can be realized as

$$W' = U^\dagger SUR_{\mathcal{A}} U^\dagger SUR_{\mathcal{A}}, \quad (5.42)$$

where $R_{\mathcal{A}}$ is the reflection around the subspace $\mathcal{A} = \text{span}\{|x\rangle|0\rangle \mid x \in K_\epsilon\}$ and S is the swap operator. It thus remains to implement the operator U .

Continuous case. We first explain a continuous version of the implementation before explaining how it can be discretized. Given an input $|x\rangle$, consider n real ancilla registers, each in the state $\int_0^1 |z\rangle dz$. Given a pair of uniformly distributed random variables ξ_1, ξ_2 , the Box-Muller transform

$$\phi_1 = \sqrt{-2\delta^2 \ln \xi_1} \cos 2\pi \xi_2, \quad (5.43)$$

$$\phi_2 = \sqrt{-2\delta^2 \ln \xi_1} \sin 2\pi \xi_2 \quad (5.44)$$

yields two variables ϕ_1, ϕ_2 that are distributed according to a univariate normal distribution with mean 0 and variance δ^2 . Thus, applying the unitary mapping

$$|\xi_1\rangle|\xi_2\rangle \mapsto \left| \sqrt{-4 \ln \xi_1} \cos 2\pi \xi_2 \right\rangle \left| \sqrt{-4 \ln \xi_1} \sin 2\pi \xi_2 \right\rangle \quad (5.45)$$

to $\int_0^1 |z\rangle dz \otimes \int_0^1 |z\rangle dz$ yields the state $\int_{\mathbb{R}} \frac{1}{\sqrt{4\pi}} e^{-z^2/4} |z\rangle dz \otimes \int_{\mathbb{R}} \frac{1}{\sqrt{4\pi}} e^{-z^2/4} |z\rangle dz$. With n such registers, we have the state

$$\int_{\mathbb{R}^n} \frac{1}{\sqrt{4\pi}} e^{-(\sum_{i=1}^n z_i^2/4)} |z\rangle dz. \quad (5.46)$$

We now compute the unit vector (direction) corresponding to each x in a different ancilla register, and uncompute the Gaussian registers. Since $\frac{1}{\sqrt{4\pi}} e^{-(\sum_{i=1}^n x_i^2/4)}$ is independent of the direction of the vector z , we obtain a uniform distribution over all the directions on the n -dimensional sphere S^n given by

$$\sqrt{\frac{n\pi^{n/2}}{\Gamma(n + \frac{1}{2})}} \int_{S^n} |u\rangle du. \quad (5.47)$$

Corresponding to each direction u , the line $\{x + tu : t \in \mathbb{R}\}$ intersects the convex body K at two points with parameters t_1, t_2 . These points as well as the length $l(u) = |t_1 - t_2|$ can be determined within error ϵ using $O(\log \frac{1}{\epsilon})$ calls to the membership oracle. We must now map each direction $|u\rangle$ to a superposition proportional to $\int_{t_1}^{t_2} e^{a^T(x+tu)/2} |x + tu\rangle dt = \int_{t_1}^{t_2} e^{a_0(x_0+tu_0)/2} |x + tu\rangle dt$. Since the exponential distribution is efficiently integrable, this can be easily effected by making a variable change starting from the state $\int_0^1 |z\rangle dz$. The normalization factor is

$$A := \sqrt{\frac{a_0 u_0}{e^{-a_0 x_0} (e^{-a_0 t_1} - e^{-a_0 t_2})}}. \quad (5.48)$$

Consider the variable change $f: [0, 1] \rightarrow [t_1, t_2]$ such that $\frac{df^{-1}(t)}{dt} = A e^{a_0(x_0+tu_0)/2}$, $f(0) = t_1$, $f(1) = t_2$. Applying f to $\int_0^1 |z\rangle dz$ produces $\int_{t_1}^{t_2} A e^{a_0(x_0+tu_0)/2} |t\rangle dt$, which can be transformed to $\int_{t_1}^{t_2} e^{a_0(x_0+tu_0)/2} |x + tu\rangle dt$ with an operation controlled on the input register x . This produces the appropriate superposition over points corresponding to each direction.

Discrete case. The operator U can be implemented in a discrete setting using a similar process to the continuous case with two main changes:

- Instead of a continuous uniform variable $\int_0^1 |z\rangle dz$, we use a discrete uniform distribution. We can create a uniform distribution on a grid with spacing ϵ as follows. We take n sets of ancilla registers, each consisting of $\log(1/\epsilon)$ registers initialized to the state 0. We apply Hadamard gates to each of these registers, giving the superposition $\bigotimes_{i=1}^n \sqrt{\epsilon} \sum_{z_i=0}^{1/\epsilon-1} |z_i\rangle$. Each $|z\rangle$ can be mapped to $|z\epsilon\rangle$, producing the required uniform distribution over the grid.

- Applying a bijective mapping to a discrete uniform distribution simply relabels the states, so the change of variable methods used in the continuous setting cannot be used to construct the Gaussian and exponential superpositions. We use instead the Grover-Rudolph method [30] that prepares states with amplitudes corresponding to efficiently integrable probability distributions. Exponential distributions can be analytically integrated, and an n -dimensional Gaussian variable is a product of n univariate standard normal distributions, each of which can be efficiently integrated by Monte Carlo methods.

Given a point $u \in K_\epsilon$ and a line $l(u, \epsilon)$ to be approximately uniformly sampled, we determine the range of points in $l(\bar{K}_\epsilon, u, \epsilon')$ using binary search with the membership oracle and prepare an exponential superposition as described above. We apply a unitary mapping to compute the closest point $v'' \in K_{\sqrt{\epsilon}n^{1/4}}$. Finally, corresponding to each point v'' , we generate a uniform distribution over an ϵ grid in $b_{\sqrt{\epsilon}n^{1/4}} \cap (\mathbb{R}_n)_\epsilon$ by applying the Hadamard transform to $\log(n^{1/4}/\sqrt{\epsilon})$ qubits.

Overall, this implementation of the discretized quantum hit-and-run walk operator gives the following.

THEOREM 5.2. *The gate complexity of implementing an operator $\tilde{U}$ such that $\|\tilde{U} - U\| = O(\epsilon)$ where $U|x\rangle|0\rangle = |x\rangle \sum_{y \in K_\epsilon} \sqrt{p_{x \rightarrow y}}|y\rangle$ is $\tilde{O}(n \log(\frac{1}{\epsilon}))$. The corresponding quantum walk operator W can be implemented using a constant number of calls to U .*

6 QUANTUM LOWER BOUNDS FOR VOLUME ESTIMATION

6.1 A Quantum Lower Bound in n

In this subsection, we prove the following quantum query lower bound in n for volume estimation:

THEOREM 6.1. *Suppose $0 < \epsilon < \sqrt{2} - 1$. Estimating the volume of K with multiplicative precision ϵ requires $\Omega(\sqrt{n})$ quantum queries to the membership oracle O_K defined in Equation (1.3).*

PROOF. We prove Theorem 6.1 by reduction from the Hamming weight problem. In Reference [59] by Nayak and Wu, it is shown that if we are given an oracle $O_s: |i, b\rangle \mapsto |i, b \oplus s_i\rangle$ for an input n -bit string $s = (s_1, \dots, s_n) \in \{0, 1\}^n$, and given the promise that the Hamming weight of s is either 0 or 1, it takes $\Omega(\sqrt{n})$ quantum queries to decide which is the case.

To establish an $\Omega(\sqrt{n})$ lower bound for volume estimation, for an n -bit string $s \in \{0, 1\}^n$ with Hamming weight $|s|_{\text{Ham}} \leq 1$, we consider the convex body $K = \times_{i=1}^n [0, 2^{s_i}]$. The volume of K is $2^{|s|_{\text{Ham}}} \in \{1, 2\}$, and membership in K is determined by the function

$$\text{MEM}_s(x) := \begin{cases} 1 & \text{if for each } i \in [n], 0 \leq x_i \leq 2^{s_i}, \\ 0 & \text{otherwise.} \end{cases} \quad (6.1)$$

The corresponding membership oracle O_K (defined in Equation (1.3)) can be simulated by querying O_s using Algorithm 6.

We now prove that for any positive integer k and $s \in \{0, 1\}^n$ with $|s|_{\text{Ham}} \leq 1$, if there is a k -query algorithm that computes the volume with access to MEM_s , then there is a k -query algorithm for deciding whether $|s|_{\text{Ham}} > 0$ with access to O_s . We first show that Algorithm 6 simulates the oracle MEM_s . In the for loop of Line 1, we know that $y_i = 1$ if and only if $1 < x_i \leq 2$, which is inside the convex body if $s_i = 1$. The case $|y|_{\text{Ham}} > 1$ implies that there exist two distinct coordinates i, j such that $x_i, x_j > 1$, which implies that x lies outside the convex body. Now, we are left with the cases $|y|_{\text{Ham}} = 1$ or 0. In Line 9, $y_i = 1$ implies $1 < x_i \leq 2$, which lies in the convex body if and only if $s_i = O_s(i) = 1$. Also, $|y| = 0$ implies that for every coordinate i , $0 \leq x_i \leq 1$, which lies in the body for all s .

ALGORITHM 6: Simulating MEM_S with one query to O_S .

Input: A vector $x = (x_1, \dots, x_n) \in \mathbb{R}^n$.
Output: $\text{MEM}_S(x)$.

```

1 for  $i = 1, \dots, n$  do
2   if  $x_i > 2$  or  $x_i < 0$  then
3      $\perp$  Return 0;
4   Set  $y_i = 1$  if  $x_i > 1$  and 0 otherwise;
5 if  $|y|_{\text{Ham}} > 1$  then
6    $\perp$  Return 0;
7 else
8   if  $|y|_{\text{Ham}} = 1$  then
9     Find  $i$  such that  $y_i = 1$ . Return  $O_S(i)$ ;
10  else
11     $\perp$  Return 1.

```

Finally, if there is a k -query algorithm that computes an estimate $\widetilde{\text{Vol}}(K)$ of the volume of K up to multiplicative precision $0 < \epsilon < \sqrt{2} - 1$, then $s = \lceil \log_2 \widetilde{\text{Vol}}(K) \rceil$, where $\lceil \cdot \rceil$ returns the nearest integer. This immediately gives a k -query algorithm that decides whether $|s|_{\text{Ham}} = 0$ or 1. Since there is an $\Omega(\sqrt{n})$ quantum query lower bound for this task, the $\Omega(\sqrt{n})$ lower bound on volume estimation follows. $\square$

6.2 An Optimal Quantum Lower Bound in $1/\epsilon$

In this subsection, we prove:

THEOREM 6.2. *Suppose $1/n \leq \epsilon \leq 1/3$. Estimating the volume of K with multiplicative precision ϵ requires $\Omega(1/\epsilon)$ quantum queries to the membership oracle O_K defined in Equation (1.3).*

Comparing with Theorem 1.1, this shows that our quantum algorithm for volume estimation is optimal in $1/\epsilon$ up to poly-logarithmic factors.

The proof constructs a convex body whose volume encodes the Hamming weight of a string. A membership oracle for this convex body can be implemented by querying the bits of the string. Then the tight lower bound of Nayak and Wu on the quantum query complexity of approximating the Hamming weight [59] implies a lower bound on the query complexity of volume estimation.

We construct the convex body by attaching hyperpyramids to the faces of the n -dimensional unit hypercube. The axis of each hyperpyramid is aligned with the axis of the face of the hypercube it corresponds to, and the height of the hyperpyramid is $1/2$. More concretely, if the unit hypercube is $H_n := [-1/2, 1/2]^n$, then the two hyperpyramids on the face perpendicular to the i th axis are

$$P_{i,+} := \{x : x_i \geq 1/2, |x_k| + |x_i| \leq 1 \ \forall k \in [n]/\{i\}\}, \quad (6.2)$$

$$P_{i,-} := \{x : x_i \leq -1/2, |x_k| + |x_i| \leq 1 \ \forall k \in [n]/\{i\}\}. \quad (6.3)$$


We denote the convex body with all hyperpyramids attached by

$$C_n := H_n \cup \left(\bigcup_{i=1}^n (P_{i,+} \cup P_{i,-}) \right). \quad (6.4)$$

For illustration, the three-dimensional convex body C_3 is shown in Figure 8.

We first prove:

LEMMA 6.1. *C_n is convex for all $n \in \mathbb{N}$.*

Fig. 8. The convex body C_3 .

PROOF. It suffices to show that if $x, y \in C_n$ and $\alpha \in [0, 1]$, then $\alpha x + (1 - \alpha)y \in C_n$. We consider three cases:

Case 1: $x, y \in H_n$. This case is straightforward as H_n is convex, hence $\alpha x + (1 - \alpha)y \in H_n \subset C_n$.

Case 2: $x \in \bigcup_{i=1}^n (P_{i,+} \cup P_{i,-})$, $y \in H_n$. Let $i^* \in [n]$ such that $x \in P_{i^*,+} \cup P_{i^*, -}$. Then by Equations (6.2) and (6.3), $|x_{i^*}| \geq 1/2$ and $|x_i| + |x_{i^*}| \leq 1 \ \forall i \in [n] \setminus \{i^*\}$, which implies $|x_i| \leq 1/2 \ \forall i \in [n] \setminus \{i^*\}$. Also note that $y \in H_n$ implies $|y_i| \leq 1/2 \ \forall i \in [n]$. Therefore,

$$|\alpha x_i + (1 - \alpha)y_i| \leq \alpha|x_i| + (1 - \alpha)|y_i| \leq \frac{\alpha}{2} + \frac{1 - \alpha}{2} = \frac{1}{2} \quad \forall i \in [n] \setminus \{i^*\}. \quad (6.5)$$

If $|\alpha x_{i^*} + (1 - \alpha)y_{i^*}| \leq 1/2$, then $\alpha x + (1 - \alpha)y \in H_n \subseteq C_n$. If $|\alpha x_{i^*} + (1 - \alpha)y_{i^*}| > 1/2$, then

$$|\alpha x_{i^*} + (1 - \alpha)y_{i^*}| + |\alpha x_i + (1 - \alpha)y_i| \leq \alpha(|x_{i^*}| + |x_i|) + (1 - \alpha)(|y_{i^*}| + |y_i|) \quad (6.6)$$

$$\leq \alpha + (1 - \alpha) = 1 \quad \forall i \in [n] \setminus \{i^*\}. \quad (6.7)$$

Therefore, by Equations (6.2) and (6.3), we have $\alpha x + (1 - \alpha)y \in P_{i^*,+} \cup P_{i^*, -} \subset C_n$. In any case, we always have $\alpha x + (1 - \alpha)y \in C_n$.

Case 3: $x, y \in \bigcup_{i=1}^n (P_{i,+} \cup P_{i,-})$. Let $i^*, j^* \in [n]$ such that $x \in P_{i^*,+} \cup P_{i^*, -}$ and $y \in P_{j^*,+} \cup P_{j^*, -}$. If $i^* = j^*$, then the proof is identical to that of Case 2, and we omit the details here. It remains to consider the case $i^* \neq j^*$. Then, we have $|x_i|, |y_i| \leq 1/2 \ \forall i \in [n] \setminus \{i^*, j^*\}$. In addition,

$$\begin{aligned} |\alpha x_{i^*} + (1 - \alpha)y_{i^*}| + |\alpha x_{j^*} + (1 - \alpha)y_{j^*}| &\leq \alpha(|x_{i^*}| + |x_{j^*}|) + (1 - \alpha)(|y_{j^*}| + |y_{i^*}|) \\ &\leq \alpha + (1 - \alpha) = 1 \end{aligned} \quad (6.8)$$

by Equations (6.2) and (6.3). This means that at most one of $|\alpha x_{i^*} + (1 - \alpha)y_{i^*}|$ and $|\alpha x_{j^*} + (1 - \alpha)y_{j^*}|$ can be more than $1/2$. If neither of them is more than $1/2$, then $\alpha x + (1 - \alpha)y \in H_n \subset C_n$. If exactly one of them is more than $1/2$, say $|\alpha x_{i^*} + (1 - \alpha)y_{i^*}| > 1/2$ and $|\alpha x_{j^*} + (1 - \alpha)y_{j^*}| \leq 1/2$, then $\alpha x + (1 - \alpha)y \in P_{i^*,+} \cup P_{i^*, -} \subset C_n$. In any case, we always have $\alpha x + (1 - \alpha)y \in C_n$. $\square$

We use the following lower bound on the quantum query complexity of approximating the Hamming weight:

PROPOSITION 6.1 ([59]). Suppose we are given the quantum oracle $O_s|i\rangle|0\rangle = |i\rangle|s_i\rangle \ \forall i \in [n]$ for some $s \in \{0, 1\}^n$. Let $0 \leq l < l' \leq n$ be two integers, $\Delta = |l - l'|$, and $m \in \{l, l'\}$ such that $|\frac{n}{2} - m|$ is

maximized. Then the quantum query complexity of determining whether s has Hamming weight at most l or at least l' is $\Theta(\sqrt{n/\Delta} + \sqrt{m(n-m)/\Delta})$.

Now, we can prove Theorem 6.2.

PROOF. Given a binary string $s \in \{0, 1\}^n$, we consider the convex body

$$C_s := H_n \cup \left(\bigcup_{i: s_i=1} (P_{i,+} \cup P_{i,-}) \right). \quad (6.9)$$

By Lemma 6.1 and the fact that each hyperpyramid is the intersection of C_n and the convex spaces $\{x : x_i \geq 1/2\}$ or $\{x : x_i \leq 1/2\}$, C_s is also convex. Furthermore, a query to the membership oracle in Equation (1.3) for C_s can be implemented using one query to the binary string oracle O_s : queries to points outside C_n or inside H_n are trivially answered with 0 and 1, respectively, whereas queries to points in $P_{i,+} \cup P_{i,-}$ should return s_i . Also note that for each $i \in [n]$, the volume of the hyperpyramid $P_{i,+}$ is

$$\text{Vol}(P_{i,+}) = \int_0^{1/2} (1-2t)^{n-1} dt = \frac{1}{2n}, \quad (6.10)$$

since the intersection of $P_{i,+}$ and $\{x : x_i = 1/2 + t\}$ is an $(n-1)$ -dimensional hypercube with side-length $1-2t$ and hence volume $(1-2t)^{n-1}$. By symmetry, we also have $\text{Vol}(P_{i,-}) = \frac{1}{2n}$. Therefore,

$$\text{Vol}(C_s) = \text{Vol}(H_n) + \sum_{i: s_i=1} (\text{Vol}(P_{i,+}) + \text{Vol}(P_{i,-})) \quad (6.11)$$

$$= 1 + |s|_{\text{Ham}} \cdot \frac{2}{2n} = 1 + \frac{|s|_{\text{Ham}}}{n}. \quad (6.12)$$

In other words, estimating the volume of C_s with multiplicative error ϵ is equivalent to the Hamming distance problem with $\Delta = 4\epsilon n$. Taking $m = \frac{n}{2} + \Delta$ in Proposition 6.1, we find that the quantum query complexity of estimating the volume of C_s is at least

$$\Omega\left(\sqrt{\frac{n}{\epsilon n}} + \frac{\sqrt{n^2/4 - \epsilon^2 n^2}}{\epsilon n}\right) = \Omega\left(\frac{1}{\epsilon}\right) \quad (6.13)$$

for any $1/n \leq \epsilon \leq 1/3$. □

Remark 6.3. The same proof strategy implies a classical lower bound of $\Omega(1/\epsilon^2)$ for volume estimation if we replace Proposition 6.1 by its folklore classical counterpart. In particular, this shows that our quantum algorithm in Theorem 1.1 achieves a provable quadratic quantum speedup in $1/\epsilon$.

Remark 6.4. Although the proofs of both theorems consider well-rounded convex bodies, this assumption can be simply waived by assuming known multiplicative rescaling factors $c_1, \dots, c_n$ along all the n directions. The proofs follow from the same arguments.

APPENDIX

A PROOF DETAILS FOR THE THEORY OF CONTINUOUS-SPACE QUANTUM WALKS

In this Appendix, we prove Theorem 3.1, restated below:

THEOREM 3.1. *Let*

$$D := \int_{\Omega} \int_{\Omega} dx dy \sqrt{p_{x \rightarrow y} p_{y \rightarrow x}} |x\rangle\langle y| \quad (3.11)$$

denote the discriminant operator of p . Let Λ be the set of eigenvalues of D , so that $D = \int_{\Lambda} d\lambda \lambda |\lambda\rangle\langle\lambda|$. Then the eigenvalues of the quantum walk operator W in Equation (3.10) are ± 1 and $\lambda \pm i\sqrt{1 - \lambda^2}$ for all $\lambda \in \Lambda$.

We first prove the following lemma:

LEMMA A.1. For any $\lambda \in \Lambda$, we have $|\lambda| \leq 1$.

PROOF. Since λ is an eigenvalue of D , we have $D|\lambda\rangle = \lambda|\lambda\rangle$. As a result, we have

$$|\lambda|\delta(0) = |\lambda|\langle\lambda|\lambda\rangle \quad (\text{A.1})$$

$$= |\langle\lambda|D|\lambda\rangle| \quad (\text{A.2})$$

$$= \left| \int_{\Omega} \int_{\Omega} dx dy \sqrt{p_{y \rightarrow x} p_{x \rightarrow y}} \langle\lambda|x\rangle\langle y|\lambda\rangle \right| \quad (\text{A.3})$$

$$\leq \sqrt{\left(\int_{\Omega} \int_{\Omega} dx dy p_{y \rightarrow x} |\langle y|\lambda\rangle|^2 \right) \left(\int_{\Omega} \int_{\Omega} dx dy p_{x \rightarrow y} |\langle\lambda|x\rangle|^2 \right)} \quad (\text{by Cauchy-Schwarz}) \quad (\text{A.4})$$

$$= \sqrt{\left(\int_{\Omega} dy |\langle y|\lambda\rangle|^2 \right) \left(\int_{\Omega} dx |\langle\lambda|x\rangle|^2 \right)} \quad (\text{by } \int_{\Omega} dy p_{x \rightarrow y} = 1) \quad (\text{A.5})$$

$$= \int_{\Omega} dx \langle\lambda|x\rangle\langle x|\lambda\rangle \quad (\text{A.6})$$

$$= \langle\lambda| \left(\int_{\Omega} dx |x\rangle\langle x| \right) |\lambda\rangle \quad (\text{by Equation (3.5)}) \quad (\text{A.7})$$

$$= \delta(0). \quad (\text{A.8})$$

Hence, the result follows. $\square$

PROOF OF THEOREM 3.1. Define an isometry

$$T := \int_{\Omega} dx |\phi_x\rangle\langle x| = \int_{\Omega} \int_{\Omega} dx dy \sqrt{p_{x \rightarrow y}} |x, y\rangle\langle x|. \quad (\text{A.9})$$

Then,

$$TT^{\dagger} = \int_{\Omega} \int_{\Omega} dx dy |\phi_x\rangle\langle x|y\rangle\langle\phi_y| = \int_{\Omega} dx |\phi_x\rangle\langle\phi_x| = \Pi, \quad (\text{A.10})$$

and

$$T^{\dagger}T = \int_{\Omega} \int_{\Omega} dx dy |x\rangle\langle\phi_x| \langle\phi_y| \langle y| \quad (\text{A.11})$$

$$= \int_{\Omega} \int_{\Omega} \int_{\Omega} \int_{\Omega} dx dy da db \langle x|y\rangle\langle a|b\rangle \sqrt{p_{x \rightarrow a} p_{y \rightarrow b}} |x\rangle\langle y| \quad (\text{A.12})$$

$$= \int_{\Omega} \int_{\Omega} dx da p_{x \rightarrow a} |x\rangle\langle x| \quad (\text{A.13})$$

$$= \int_{\Omega} dx |x\rangle\langle x| \quad (\text{A.14})$$

$$= I. \quad (\text{A.15})$$

Furthermore,

$$T^\dagger ST = \int_{\Omega} \int_{\Omega} dx dy |x\rangle \langle \phi_x| S |\phi_y\rangle \langle y| \quad (\text{A.16})$$

$$= \int_{\Omega} \int_{\Omega} \int_{\Omega} \int_{\Omega} dx dy da db \langle x, a | S | y, b \rangle \sqrt{p_{x \rightarrow a} p_{y \rightarrow b}} |x\rangle \langle y| \quad (\text{A.17})$$

$$= \int_{\Omega} \int_{\Omega} dx da \sqrt{p_{x \rightarrow a} p_{a \rightarrow x}} |x\rangle \langle a| \quad (\text{A.18})$$

$$= D. \quad (\text{A.19})$$

As a result, for any $\lambda \in \Lambda$, we have

$$WT|\lambda\rangle = S(2\Pi - I)T|\lambda\rangle = (2STT^\dagger T - ST)|\lambda\rangle = (2ST - ST)|\lambda\rangle = ST|\lambda\rangle. \quad (\text{A.20})$$

Similarly, we have

$$\begin{aligned} WST|\lambda\rangle &= S(2\Pi - I)ST|\lambda\rangle \\ &= (2STT^\dagger ST - S^2T)|\lambda\rangle = (2STD - T)|\lambda\rangle = (2\lambda S - I)T|\lambda\rangle. \end{aligned} \quad (\text{A.21})$$

By Lemma A.1, $|\lambda| \leq 1$. As a result, we have

$$W(I - (\lambda + i\sqrt{1 - \lambda^2})S)T|\lambda\rangle = WT|\lambda\rangle - (\lambda + i\sqrt{1 - \lambda^2})WST|\lambda\rangle \quad (\text{A.22})$$

$$= ST|\lambda\rangle - (\lambda + i\sqrt{1 - \lambda^2})(2\lambda S - I)T|\lambda\rangle \quad (\text{A.23})$$

$$= (S - (\lambda + i\sqrt{1 - \lambda^2})(2\lambda S - I))T|\lambda\rangle \quad (\text{A.24})$$

$$= (\lambda + i\sqrt{1 - \lambda^2})(I - (\lambda + i\sqrt{1 - \lambda^2})S)T|\lambda\rangle; \quad (\text{A.25})$$

in other words, $\lambda + i\sqrt{1 - \lambda^2}$ is an eigenvalue of W with eigenvector $(I - (\lambda + i\sqrt{1 - \lambda^2})S)T|\lambda\rangle$. Similarly, we have

$$W(I - (\lambda - i\sqrt{1 - \lambda^2})S)T|\lambda\rangle = (\lambda - i\sqrt{1 - \lambda^2})(I - (\lambda - i\sqrt{1 - \lambda^2})S)T|\lambda\rangle, \quad (\text{A.26})$$

i.e., $\lambda - i\sqrt{1 - \lambda^2}$ is an eigenvalue of W with eigenvector $(I - (\lambda - i\sqrt{1 - \lambda^2})S)T|\lambda\rangle$.

Finally, note that for any vector $|u\rangle$ in the orthogonal complement of the space $\text{span}_{\lambda \in \Lambda}\{T|\lambda\rangle, ST|\lambda\rangle\}$, W simply acts as $-S$, since

$$\Pi = TT^\dagger = \int_{\Lambda} d\lambda T|\lambda\rangle \langle \lambda| T^\dagger, \quad (\text{A.27})$$

which projects onto $\text{span}_{\lambda \in \Lambda}\{T|\lambda\rangle\}$. In this orthogonal complement subspace, the eigenvalues are ± 1 , because $S^2 = I$. $\square$

ACKNOWLEDGMENTS

We thank anonymous reviewers for helpful suggestions on earlier versions of this article, including one who pointed out a minor technical issue in the nondestructive amplitude estimation circuit (Figure 4). T. L. and C. W. thank Nai-Hui Chia for helpful discussions. T. L. thanks Ruizhe Zhang for helpful discussions about large effective spectral gaps of Markov chains.

REFERENCES

- [1] Dorit Aharonov, Andris Ambainis, Julia Kempe, and Umesh Vazirani. 2001. Quantum walks on graphs. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*. 50–59. arXiv:[quant-ph/0012090](https://arxiv.org/abs/quant-ph/0012090).
- [2] Dorit Aharonov and Amnon Ta-Shma. 2003. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*. 20–29. arXiv:[quant-ph/0301023](https://arxiv.org/abs/quant-ph/0301023).

- [3] Gorjan Alagic and Alexander Russell. 2005. Decoherence in quantum walks on the hypercube. *Phys. Rev. A* 72 (2005), 062304. arXiv:[quant-ph/0501169](https://arxiv.org/abs/quant-ph/0501169).
- [4] Andris Ambainis, Eric Bach, Ashwin Nayak, Ashvin Vishwanath, and John Watrous. 2001. One-dimensional quantum walks. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*. 37–49. <https://doi.org/10.1145/380752.380757>
- [5] Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. 2020. Convex optimization using quantum oracles. *Quantum* 4 (2020), 220. arXiv:[1809.00643](https://arxiv.org/abs/1809.00643).
- [6] Simon Apers and Alain Sarlette. 2019. Quantum fast-forwarding: Markov chains and graph property testing. *Quant. Info. Comput.* (2019). arXiv:[1804.02321](https://arxiv.org/abs/1804.02321).
- [7] David Applegate and Ravi Kannan. 1991. Sampling and integration of near log-concave functions. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*. 156–163.
- [8] Srinivasan Arunachalam, Vojtech Havlicek, Giacomo Nannicini, Kristan Temme, and Pawel Wocjan. 2021. Simpler (classical) and faster (quantum) algorithms for gibbs partition functions. In *Proceedings of the IEEE International Conference on Quantum Computing and Engineering*. IEEE, 112–122. arXiv:[2009.11270](https://arxiv.org/abs/2009.11270).
- [9] Imre Bárány and Zoltán Füredi. 1987. Computing the volume is difficult. *Discrete Comput. Geom.* 2, 4 (1987), 319–326.
- [10] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. 1997. Strengths and weaknesses of quantum computing. *SIAM J. Comput.* 26, 5 (1997), 1510–1523. arXiv:[quant-ph/9701001](https://arxiv.org/abs/quant-ph/9701001).
- [11] Sergio Boixo and Rolando D. Somma. 2015. Quantum algorithms for simulated annealing. Retrieved from <https://arxiv.org/abs/1512.03806>.
- [12] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. 2002. Quantum amplitude amplification and estimation. *Contemp. Math.* 305 (2002), 53–74. arXiv:[quant-ph/0005055](https://arxiv.org/abs/quant-ph/0005055).
- [13] Gilles Brassard, Peter Høyer, and Alain Tapp. 1998. Quantum counting. In *Proceedings of the 25th International Colloquium on Automata, Languages, and Programming*. 820–831. arXiv:[quant-ph/9805082](https://arxiv.org/abs/quant-ph/9805082).
- [14] Shouvanik Chakrabarti, Andrew M. Childs, Tongyang Li, and Xiaodi Wu. 2020. Quantum algorithms and lower bounds for convex optimization. *Quantum* 4 (2020), 221. arXiv:[1809.01731](https://arxiv.org/abs/1809.01731).
- [15] Shantanav Chakraborty, Kyle Luh, and Jérémie Roland. 2020. Analog quantum algorithms for the mixing of Markov chains. *Phys. Rev. A* 102, 2 (2020), 022423. arXiv:[1904.11895](https://arxiv.org/abs/1904.11895).
- [16] Yuansi Chen. 2021. An almost constant lower bound of the isoperimetric coefficient in the KLS conjecture. *Geom. Funct. Anal.* 31, 1 (2021), 34–61. arXiv:[2011.13661](https://arxiv.org/abs/2011.13661).
- [17] Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. 2003. Exponential algorithmic speedup by quantum walk. In *Proceedings of the 35th ACM Symposium on Theory of Computing*. 59–68. <https://doi.org/10.1145/780542.780552> arXiv:[quant-ph/0209131](https://arxiv.org/abs/quant-ph/0209131).
- [18] Andrew M. Childs, Tongyang Li, Jin-Peng Liu, Chunhao Wang, and Ruizhe Zhang. 2022. Quantum Algorithms for Sampling Log-Concave Distributions and Estimating Normalizing Constants. Retrieved from <https://arxiv.org/abs/2210.06539>.
- [19] Arjan Cornelissen and Yassine Hamoudi. 2022. A Sublinear-Time Quantum Algorithm for Approximating Partition Functions. Retrieved from <https://arxiv.org/abs/2207.08643>.
- [20] Ben Cousins and Santosh Vempala. 2014. A cubic algorithm for computing gaussian volume. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1215–1228. arXiv:[1306.5829](https://arxiv.org/abs/1306.5829).
- [21] Ben Cousins and Santosh Vempala. 2015. Bypassing KLS: Gaussian cooling and an $O^*(n^3)$ volume algorithm. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*. 539–548. arXiv:[1409.6011](https://arxiv.org/abs/1409.6011).
- [22] Christopher M. Dawson and Michael A. Nielsen. 2006. The Solovay-Kitaev algorithm. *Quant. Info. Comput.* 6, 1 (2006), 81–95. arXiv:[quant-ph/0505030](https://arxiv.org/abs/quant-ph/0505030).
- [23] Martin Dyer and Alan Frieze. 1991. Computing the volume of convex bodies: A case where randomness provably helps. *Prob. Combinat. Appl.* 44 (1991), 123–170.
- [24] Martin Dyer, Alan Frieze, and Ravi Kannan. 1991. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM* 38, 1 (1991), 1–17.
- [25] Martin E. Dyer and Alan M. Frieze. 1988. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.* 17, 5 (1988), 967–974.
- [26] György Elekes. 1986. A geometric inequality and the complexity of computing volume. *Discrete Comput. Geom.* 1, 4 (1986), 289–292.
- [27] Edward Farhi and Sam Gutmann. 1998. Quantum computation and decision trees. *Phys. Rev. A* 58, 2 (1998), 915–928. arXiv:[quant-ph/9706062](https://arxiv.org/abs/quant-ph/9706062).
- [28] Alan Frieze and Ravi Kannan. 1999. Log-sobolev inequalities and sampling from log-concave distributions. *Ann. Appl. Probabil.* 9, 1 (1999), 14–26.
- [29] Martin Grötschel, László Lovász, and Alexander Schrijver. 2012. *Geometric Algorithms and Combinatorial Optimization*. Vol. 2. Springer Science & Business Media.

- [30] Lov Grover and Terry Rudolph. 2002. Creating superpositions that correspond to efficiently integrable probability distributions. Retrieved from <https://arxiv.org/abs/quant-ph/0208112>.
- [31] Lov K. Grover. 1997. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* 79, 2 (1997), 325. arXiv:quant-ph/9706033.
- [32] Lov K. Grover. 2005. A different kind of quantum search. Retrieved from <https://arxiv.org/abs/quant-ph/0503205>.
- [33] Yassine Hamoudi and Frédéric Magniez. 2019. Quantum chebyshev's inequality and applications. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming*. 69:1–69:16. arXiv:1807.06456.
- [34] Aram W. Harrow and Annie Y. Wei. 2020. Adaptive quantum simulated annealing for bayesian inference and estimating partition functions. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*. 193–212. arXiv:1907.09965.
- [35] Adam Holmes and A. Y. Matsuura. 2020. Efficient quantum circuits for accurate state preparation of smooth, differentiable functions. In *Proceedings of the IEEE International Conference on Quantum Computing and Engineering (QCE'20)*. 169–179. <https://doi.org/10.1109/QCE49297.2020.00030>
- [36] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh S. Vempala. 2021. Reducing isotropy and volume to KLS: An $O(n^3\psi^2)$ volume algorithm. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 961–974. arXiv:2008.02146.
- [37] Adam T. Kalai and Santosh Vempala. 2006. Simulated annealing for convex optimization. *Math. Oper. Res.* 31, 2 (2006), 253–266.
- [38] Ravi Kannan, László Lovász, and Miklós Simonovits. 1997. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algor.* 11, 1 (1997), 1–50.
- [39] Leonid G. Khachiyan. 1988. On the complexity of computing the volume of a polytope. *Izvestia Akad. Nauk SSSR, Eng. Cybernet.* 3 (1988), 216–217.
- [40] Leonid G. Khachiyan. 1989. The problem of computing the volume of polytopes is NP-hard. *Uspekhi Mat. Nauk* 44, 3 (1989), 199–200.
- [41] Yin Tat Lee, Aaron Sidford, and Santosh S. Vempala. 2018. Efficient convex optimization with membership oracles. In *Proceedings of the 31st Conference on Learning Theory*. 1292–1294. arXiv:1706.07357.
- [42] Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. 2015. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*. 1049–1065. arXiv:1508.04874.
- [43] Yin Tat Lee and Santosh S. Vempala. 2017. Eldan's stochastic localization and the KLS hyperplane conjecture: An improved lower bound for expansion. In *Proceedings of the 58th Annual Symposium on Foundations of Computer Science*. 998–1007. arXiv:1612.01507.
- [44] Yin Tat Lee and Santosh S. Vempala. 2017. The Kannan-Lovász-Simonovits conjecture. *Curr. Dev. Math.* 2017, 1 (2017), 1–36. arXiv:1807.03465.
- [45] Yin Tat Lee and Santosh S. Vempala. 2018. Convergence rate of riemannian hamiltonian Monte Carlo and faster polytope volume computation. In *Proceedings of the 50th Annual Symposium on Theory of Computing*. 1115–1121. arXiv:1710.06261.
- [46] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. 2017. *Markov Chains and Mixing Times*. American Mathematical Society.
- [47] Tongyang Li and Xiaodi Wu. 2019. Quantum query complexity of entropy estimation. *IEEE Trans. Info. Theory* 65, 5 (2019), 2899–2921. arXiv:1710.06025.
- [48] László Lovász. 1999. Hit-and-run mixes fast. *Math. Program.* 86, 3 (1999), 443–461.
- [49] László Lovász and Miklós Simonovits. 1990. The mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*. 346–354.
- [50] László Lovász and Miklós Simonovits. 1993. Random walks in a convex body and an improved volume algorithm. *Random Struct. Algor.* 4, 4 (1993), 359–412.
- [51] László Lovász and Santosh Vempala. 2006. Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*. 57–68.
- [52] László Lovász and Santosh Vempala. 2006. Hit-and-run from a corner. *SIAM J. Comput.* 35, 4 (2006), 985–1005.
- [53] László Lovász and Santosh Vempala. 2006. Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *J. Comput. Syst. Sci.* 72, 2 (2006), 392–417.
- [54] László Lovász and Santosh Vempala. 2007. The geometry of logconcave functions and sampling algorithms. *Random Struct. Algor.* 30, 3 (2007), 307–358.
- [55] Frédéric Magniez, Ashwin Nayak, Jérémie Roland, and Miklos Santha. 2011. Search via quantum walk. *SIAM J. Comput.* 40, 1 (2011), 142–164. <https://doi.org/10.1137/090745854>. arXiv:quant-ph/0608026.
- [56] George Marsaglia. 1972. Choosing a point from the surface of a sphere. *Ann. Math. Stat.* 43, 2 (1972), 645–646.

- [57] Ashley Montanaro. 2015. Quantum speedup of Monte Carlo methods. *Proc. Roy. Soc. A* 471, 2181 (2015), 20150301. arXiv:[1504.06987](#).
- [58] Mervin E. Muller. 1959. A note on a method for generating points uniformly on n -dimensional spheres. *Commun. ACM* 2, 4 (1959), 19–20.
- [59] Ashwin Nayak and Felix Wu. 1999. The quantum query complexity of approximating the median and related statistics. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*. 384–393. arXiv:[quant-ph/9804066](#).
- [60] Davide Orsucci, Hans J. Briegel, and Vedran Dunjko. 2018. Faster quantum mixing for slowly evolving sequences of Markov chains. *Quantum* 2 (2018), 105. arXiv:[1503.01334](#).
- [61] Luis Rademacher and Santosh Vempala. 2008. Dispersion of mass and the complexity of randomized geometric algorithms. *Adv. Math.* 219, 3 (2008), 1037–1069. arXiv:[cs/0608054](#).
- [62] Peter C. Richter. 2007. Almost uniform sampling via quantum walks. *New J. Phys.* 9, 3 (2007), 72. arXiv:[quant-ph/0606202](#).
- [63] Peter C. Richter. 2007. Quantum speedup of classical mixing processes. *Phys. Rev. A* 76, 4 (2007), 042306. arXiv:[quant-ph/0609204](#).
- [64] Mark Rudelson. 1999. Random vectors in the isotropic position. *J. Funct. Anal.* 164, 1 (1999), 60–72. arXiv:[math/9608208](#).
- [65] Jun John Sakurai and Jim Napolitano. 2014. *Modern Quantum Mechanics*. Pearson Harlow.
- [66] Robert L. Smith. 1984. Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Oper. Res.* 32, 6 (1984), 1296–1308.
- [67] Rolando D. Somma, Sergio Boixo, and Howard Barnum. 2007. Quantum simulated annealing. Retrieved from <https://arxiv.org/abs/0712.1008>.
- [68] Rolando D. Somma, Sergio Boixo, Howard Barnum, and Emanuel Knill. 2008. Quantum simulations of classical annealing processes. *Phys. Rev. Lett.* 101, 13 (2008), 130504. arXiv:[0804.1571](#).
- [69] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. 2009. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *J. ACM* 56, 3 (2009), 1–36. <https://doi.org/10.1145/1516512.1516520>. arXiv:[cs/0612058](#).
- [70] Mario Szegedy. 2004. Quantum speed-up of Markov chain-based algorithms. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*. 32–41.
- [71] Kristan Temme, Tobias J. Osborne, Karl G. Vollbrecht, David Poulin, and Frank Verstraete. 2011. Quantum metropolis sampling. *Nature* 471, 7336 (2011), 87–90. arXiv:[0911.3635](#).
- [72] Santosh Vempala. 2005. Geometric random walks: A survey. In *Combinatorial and Computational Geometry*. Mathematical Sciences Research Institute Publications, Vol. 52. MSRI, 573–612.
- [73] Chunhao Wang. 2018. *Computational Problems Related to Open Quantum Systems*. Ph.D. Dissertation. University of Waterloo.
- [74] Pawel Wocjan and Anura Abeyesinghe. 2008. Speedup via quantum sampling. *Phys. Rev. A* 78, 4 (2008), 042336. arXiv:[0804.4259](#).
- [75] Pawel Wocjan, Chen-Fu Chiang, Daniel Nagaj, and Anura Abeyesinghe. 2009. Quantum algorithm for approximating partition functions. *Phys. Rev. A* 80, 2 (2009), 022340. arXiv:[0811.0596](#).
- [76] Theodore J. Yoder, Guang Hao Low, and Isaac L. Chuang. 2014. Fixed-point quantum search with an optimal number of queries. *Phys. Rev. Lett.* 113, 21 (2014), 210501. arXiv:[1409.3305](#).
- [77] Man-Hong Yung and Alán Aspuru-Guzik. 2012. A quantum–quantum metropolis algorithm. *Proc. Natl. Acad. Sci. U.S.A.* 109, 3 (2012), 754–759. arXiv:[1011.1468](#).

Received 28 April 2022; revised 31 December 2022; accepted 8 March 2023