QUANTUM GLOBAL SMOOTH NONCONVEX OPTIMIZATION UNDER WEAK LOCALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We develop a fully discrete quantum simulated-annealing framework for global optimization of smooth non-convex functions under a weak localization assumption: a known bounded region contains an $\varepsilon_{\mathrm{loc}}$-optimal solution of the original problem. This assumption is strictly weaker than the dissipativity conditions commonly imposed in Langevin-based analyses. Unlike continuous-space formulations that defer discretization to implementation, our algorithm is defined end-to-end on a finite discrete space and can be implemented with finite registers using only function and gradient access. We provide an explicit spectral gap lower bound for the underlying discrete walk, which in turn gives a computable phase-gap guarantee for the quantum walk and makes the resulting complexity fully quantifiable under weak localization. While a polynomial gap cannot be guaranteed uniformly over this broad problem class without additional structural assumptions, we derive an unconditional and computable lower bound that can be supplied directly to the algorithm. This yields a complexity guarantee without requiring stronger structural assumptions. Overall, our framework makes the spectral gap limitations of weakly localized non-convex optimization explicit while reducing the annealing costs through adaptive temperature selection.

1 INTRODUCTION

Global optimization of a smooth non-convex function is difficult because local information need not reveal the location of a global minimizer. Langevin and simulated-annealing methods approach this problem through Gibbs distributions, $\pi_\beta(dx) \propto e^{-\beta f(x)}\,dx$, which, when well defined, increasingly favor low-objective regions as the inverse temperature β grows. This sampling-to-optimization paradigm underlies classical analyses of Langevin optimization (Raginsky et al., 2017; Xu et al., 2018). Quantum sampling and simulated-annealing algorithms seek to reduce the cost of preparing these distributions through quantum walks and coherent state preparation (Somma et al., 2008; Harrow and Wei, 2020; Childs et al., 2022; Ozgul et al., 2024). For objectives on $\mathbb{R}^d$, however, smoothness alone does not guarantee integrability of $e^{-\beta f}$, localization of a global minimizer within a known bounded region, or confinement of the associated Langevin dynamics. Langevin-based analyses therefore commonly impose dissipativity, for example,

$$\langle \nabla f(x), x\rangle \geq m\|x\|^2 - b, \qquad m > 0, \tag{1}$$

which provides radial confinement and quantitative tail control (Raginsky et al., 2017; Xu et al., 2018; Ozgul et al., 2024). Such assumptions facilitate sampling on an unbounded domain, but exclude objectives whose low-value regions can be localized without requiring growth at infinity. Our goal is to develop a quantum global-optimization framework for this broader setting.

Weak localization. For a target accuracy $\varepsilon > 0$, the algorithm is given a bounded domain Ω with

$$\min_{x\in\Omega} f(x) \leq f_\star + \varepsilon_{\mathrm{loc}}, \qquad f_\star := \inf_{x\in\mathbb{R}^d} f(x), \qquad 0 \leq \varepsilon_{\mathrm{loc}} < \varepsilon.$$

Thus, Ω contains an approximately global solution, with no growth condition imposed outside it. The domain and its localization guarantee are inputs, not computed by the algorithm. We define the Gibbs distribution directly on Ω,

$$\pi_{\beta,R}(dx) \propto e^{-\beta f(x)}\mathbf{1}_\Omega(x)\,dx,$$

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

where R is the domain-size parameter from Section 3. Dissipativity with known constants yields a bounded region containing a global minimizer, but the converse fails. This formulation separates localization, Gibbs concentration, and sampling complexity: localization connects restricted and global optima, concentration determines a sufficient inverse temperature, and chain geometry controls the sampling cost. Weak localization does not ensure efficient mixing, as smooth bounded-domain objectives may still have bottlenecks with arbitrarily small spectral gaps.

From localized Gibbs distributions to finite quantum walks. Continuous-space sampling does not define a finite-register algorithm: discretization and numerical approximation introduce distribution, state-preparation, and quantum-walk errors that must also be controlled. Thus, we work directly on a finite grid: discrete Gibbs laws are stationary for ideal finite MALA chains, finite-precision quantum errors are budgeted explicitly, and lifting separates discretization from numerical implementation errors. Instead of assuming an unknown gap, we derive a computable spectral-gap lower bound for each discrete MALA chain, providing the quantum-walk phase-gap certificate without claiming polynomial-time performance for arbitrary weakly localized objectives.

Contributions. Our main contributions are:

1. *Global nonconvex optimization beyond dissipativity.* We extend Langevin-based quantum optimization beyond dissipativity using weak localization, a provably strictly weaker assumption requiring only a known bounded region containing an approximately global solution, without growth conditions at infinity.
2. *A fully discrete and rigorous construction.* Rather than formulating the Langevin-based quantum walk in continuous space and treating discretization as a routine implementation detail, we define the annealing states, MALA chains, and quantum walks directly on a finite grid. We prove exact discrete detailed balance, finite-register implementability, and an explicit discrete-to-continuous approximation guarantee, so that the analyzed Markov chain is the one actually implemented by the quantum algorithm.
3. *No hidden mixing assumptions.* We show that smoothness and weak localization alone cannot guarantee an inverse-polynomial phase gap. Rather than imposing additional mixing or functional-inequality assumptions on the target or intermediate Gibbs distributions, we derive a computable lower bound on the spectral gap of every queried discrete MALA chain, and hence on the quantum-walk phase gap. Thus, our final complexity contains no unknown quantities like Cheeger constant or inaccessible spectral-gap parameter. Combined with an explicit inverse-temperature choice, this yields a fully quantified guarantee for finding an ε-global minimizer.

Roadmap. In Section 3, we introduce the localized Gibbs distribution $\pi_{\beta,R}$, and discretize it to $\pi^h_{\beta,R}$ on Ω_h with grid size h. The lifted distribution $\widetilde{\pi}^h_{\beta,R}$ from Eq. (6) connects the discrete and continuous settings, with discretization error $\|\widetilde{\pi}^h_{\beta,R} - \pi_{\beta,R}\|_{\mathrm{TV}} \le \xi$. Section 4 constructs explicit and adaptive annealing paths through the discrete Gibbs family $\{\pi^h_{\theta,R}\}_{0\le\theta\le\beta}$, with the adaptive path requiring fewer stages. In Section 5 each intermediate distribution is realized as the stationary distribution of a finite MALA chain and quantized using Szegedy walks. The full adaptive localized Q-MALA procedure is given in Algorithm 1. This yields Theorem 10, which produces an output distribution ν_h satisfying $\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \le \zeta$, with complexity controlled by a spectral-gap lower bound $\underline{\delta}$ analyzed in Section 6. For the adaptive path, χ denotes the total failure probability of the overlap-estimation procedure used to construct the annealing schedule. Finally, Section 7 combines the sampling, discretization, finite-temperature, and localization errors. In Theorem 13, we choose parameters so that the final output is an ε-global minimizer with the query complexity in Eq. (31).

2 Related Work

Non-convex optimization. Non-convex optimization has been extensively studied in both optimization theory and machine learning (Nesterov, 2013). Classical complexity results highlight the intrinsic difficulty of the problem, motivating a focus on weaker solution concepts such as approximate stationary points (Hollender and Zampetakis, 2025). First-order methods, including stochastic gradient descent and its variants (Gower et al., 2019; Johnson and Zhang, 2013), are widely used due to their scalability, with theoretical guarantees on convergence to first- and second-order stationary

points under smoothness assumptions (Ghadimi and Lan, 2013; Jin et al., 2017). Techniques for escaping saddle points—such as perturbation methods and negative curvature exploitation—have further improved the theoretical understanding of these methods (Lu et al., 2019; 2020).

Langevin sampling and annealing. The unadjusted Langevin algorithm (ULA) discretizes Langevin diffusion, whereas the Metropolis-adjusted Langevin algorithm (MALA) adds an acceptance correction to preserve the target distribution. Langevin and stochastic-gradient Langevin methods admit global-optimization guarantees under smoothness, dissipativity, and appropriate oracle assumptions (Raginsky et al., 2017; Xu et al., 2018). Other analyses use log-Sobolev inequalities without requiring convexity (Vempala and Wibisono, 2019), or establish non-convex MALA bounds through conductance and higher-order regularity (Mangoubi and Vishnoi, 2019). Alternatively, Chen et al. (2024) analyze global optimization through Lyapunov potentials and hitting times, rather than relying exclusively on approximation of the full Gibbs distribution. More recently, Guo et al. (2025) establish non-asymptotic annealed-Langevin guarantees without imposing log-concavity, log-Sobolev, or Poincaré inequalities. Their complexity depends on the Wasserstein action of an interpolating distributional path, under smoothness, moment, finite-action, and minimizer-localization assumptions.

Quantum algorithms for optimization and sampling. Quantum-walk algorithms accelerate Markov-chain sampling and simulated annealing (Somma et al., 2008; Wocjan and Abeyesinghe, 2008). Harrow and Wei (2020) develop adaptive cooling schedules for coherent Gibbs-state preparation, while quantum fast-forwarding accelerates reversible Markov-chain discriminant dynamics (Apers and Sarlette, 2019). For continuous optimization, Li and Zhang (2022) combine quantum annealing with quantum hit-and-run for approximately convex objectives, a structural assumption distinct from weak localization. Childs et al. (2022) develop quantum MALA and underdamped-Langevin algorithms for strongly log-concave sampling and normalizing-constant estimation. Ozgul et al. (2024) study quantum MALA, ULA, and stochastic ULA for smooth dissipative non-logconcave targets, using a reversible reference chain to analyze unadjusted and stochastic implementations. Complementarily, Ozgul et al. (2026) accelerate stochastic Langevin and Hamiltonian Monte Carlo through quantum mean and gradient estimation without requiring reversibility or exact gradients. Their non-logconcave sampling guarantees use a log-Sobolev inequality, with optimization applications to structured non-convex objectives. Other approaches use different mechanisms or performance criteria. Herman et al. (2025) analyze real-space adiabatic quantum optimization for structured non-convex families through Schrödinger spectra. Quantum Bayesian optimization improves cumulative regret through quantum reward estimation under coherent oracle access (Dai et al., 2023). These online guarantees are not directly comparable with offline global-optimization bounds without aligning the performance criteria and oracle models.

3 Problem Formulation and Assumptions

Let $f : \mathbb{R}^d \to \mathbb{R}$ and $f^\star := \inf_{x\in\mathbb{R}^d} f(x)$. For $\varepsilon > 0$, we call $\widehat{x}$ an ε-global minimizer if $f(\widehat{x}) \le f^\star + \varepsilon$. Neither convexity nor uniqueness of the global minimizer is assumed. Throughout, $d \ge 2$.

Assumption 1 (Smoothness)**.** *The function f is differentiable and L-smooth: $\|\nabla f(x) - \nabla f(y)\| \le L\|x - y\|$ for all $x, y \in \mathbb{R}^d$. A known bound $G_0 \ge \|\nabla f(0)\|$ is available.*

Assumption 2 (Weak localization)**.** *For a known radius $R > 0$ and localization error $\varepsilon_{\mathrm{loc}} \ge 0$, the closed ball $\Omega := B(0, R)$ satisfies $f^\star_R := \min_{x\in\Omega} f(x) \le f^\star + \varepsilon_{\mathrm{loc}}$. For the requested optimization accuracy, $0 \le \varepsilon_{\mathrm{loc}} < \varepsilon$.*

Assumption 2 covers both constrained optimization ($\varepsilon_{\mathrm{loc}} = 0$) and unconstrained optimization when the supplied search region contains an exact or approximate global solution. It makes no assertion about f outside Ω and does not require a full-space Gibbs distribution.

Proposition 1 (Weak localization is strictly weaker than dissipativity)**.** *Let f be continuous and satisfy the dissipativity condition Eq. (1) with parameters $m > 0$ and $b \ge 0$. Then f admits a global minimizer $x^\star$ with $\|x^\star\| \le \sqrt{b/m}$. So Assumption 2 holds with $\varepsilon_{\mathrm{loc}} = 0$ for every $R \ge \sqrt{b/m}$.*

The converse fails even for bounded globally smooth nonconvex objectives. More details can be found in Appendix B.1. Additionally, weak localization takes an admissible search radius R as part of the problem specification, whereas dissipativity derives such a radius from a global growth

condition through (m, b). Our complexity keeps the cost of this supplied radius explicit through $G_R = G_0 + LR$ and $\widehat{D}_R = 2RG_R$.

3.1 Localized Gibbs law and basic bounds

For $\beta \geq 0$, define the Gibbs distribution directly on Ω:

$$\pi_{\beta,R}(dx) := \frac{e^{-\beta f(x)}\mathbf{1}_\Omega(x)}{Z_{\beta,R}}dx, \qquad Z_{\beta,R} := \int_\Omega e^{-\beta f(y)}dy. \tag{2}$$

Compactness makes $Z_{\beta,R}$ finite without dissipativity or any full-space tail condition.

Lemma 2 (Localized gradient and oscillation bounds)**.** *Under Assumptions 1 and 2, define*

$$G_R := G_0 + LR. \tag{3}$$

Then, for all $x, y \in \Omega$, we have $\|\nabla f(x)\| \leq G_R$ and $|f(x) - f(y)| \leq G_R\|x - y\|$. Consequently,

$$D_R := \max_{x\in\Omega} f(x) - \min_{x\in\Omega} f(x) \leq 2RG_R =: \widehat{D}_R. \tag{4}$$

Thus, G_R controls the local variation of f, while $\widehat{D}_R$ gives an explicit upper bound on its oscillation over Ω. We use $\widehat{D}_R$ to construct the annealing schedule and bound its length.

3.2 Oracle and discretization model

The algorithm queries quantum oracles that have coherent finite-precision access to $\mathcal{O}_f : |x, z\rangle \mapsto |x, z \oplus f(x)\rangle$ and $\mathcal{O}_{\nabla f} : |x, z\rangle \mapsto |x, z \oplus \nabla f(x)\rangle$, with working precisions fixed before the run. All annealing states, Markov chains, and quantum walks are defined on the finite grid. The continuous Gibbs law Eq. (2) is used only for the final optimization guarantee, with the discrete-to-continuous error controlled separately in total variation.

Computational grid. Let $0 < h \leq R/\sqrt{d}$ and define

$$\Lambda_h := h\mathbb{Z}^d, \qquad C_x := x + [0, h)^d, \qquad \Omega_h := \{x \in \Lambda_h : C_x \subseteq \Omega\}.$$

We use the finite set Ω_h as the discrete state space of the quantum algorithm, with each $x \in \Omega_h$ representing the cell C_x. Each cell has volume h^d and diameter $\sqrt{d}\,h$. Further geometric and finite-register properties of this discretization are collected in Appendix B.3.

Discrete Gibbs. The target distribution of the finite-state Markov chains is the discrete Gibbs law

$$\pi^h_{\beta,R}(x) := \Pr[X_h = x] := e^{-\beta f(x)}/\Big(\sum_{z\in\Omega_h} e^{-\beta f(z)}\Big), \qquad x \in \Omega_h. \tag{5}$$

Comparison with the continuous law. Define the region covered by admissible cells as $\widehat{\Omega}_h := \bigcup_{x\in\Omega_h} C_x \subseteq \Omega$. For $y \in \mathbb{R}^d$, let $x(y) \in \Lambda_h$ denote the corner of the unique cell containing y. For comparison with the continuous Gibbs distribution, we lift any probability mass function p on Ω_h to the piecewise-constant density $\widetilde{p}_h(y) := p(x)/h^d$, for $y \in C_x, x \in \Omega_h$, with $\widetilde{p}_h(y) = 0$ outside $\widehat{\Omega}_h$. Hence, for Eq. (5), with $\widetilde{\pi}^h_{\beta,R}(y) = 0$ outside $\widehat{\Omega}_h$,

$$\widetilde{\pi}^h_{\beta,R}(y) = \pi^h_{\beta,R}(x(y))/h^d, \qquad y \in \widehat{\Omega}_h. \tag{6}$$

The following two lemmas isolate the discretization effects that will be used in Section 7. We use

$$\omega_h := G_R\sqrt{d}\,h \tag{7}$$

to denote the maximum objective variation at the scale of one grid cell. Lemma 3 controls the variation of the objective within a single cell.

Lemma 3. *Let $0 < h \leq R/\sqrt{d}$, then for every $x \in \Omega_h$ and all $y_1, y_2 \in C_x$,*

$$|f(y_1) - f(y_2)| \;\leq\; G_R\|y_1 - y_2\| \;\leq\; G_R\,\mathrm{diam}(C_x) \;=\; \omega_h.$$

Lemma 4 controls the distributional discrepancy between the discrete Gibbs target and the continuous localized Gibbs law, and the proof is in Appendix B.4.1.

Lemma 4. *Let Assumptions 1 and 2 hold, let $\beta > 0$, and let $\xi \in (0,1)$ be a target discretization budget. We have $\|\widetilde{\pi}^h_{\beta,R} - \pi_{\beta,R}\|_{\mathrm{TV}} \le \xi$, if the grid spacing satisfies*

$$h \le \min\{\xi/(6\sqrt{d}\,\beta G_R),\ \xi R/(2d^{3/2}),\ R/\sqrt{d}\}. \tag{8}$$

Gradient rounding. We use coordinatewise b-bit rounding

$$\widehat{\nabla} f(x) := \mathrm{rd}_b(\nabla f(x)), \qquad \|\widehat{\nabla} f(x) - \nabla f(x)\|_\infty \le 2^{-b-1},$$

where rd_b denotes round-to-nearest onto $2^{-b}\mathbb{Z}$. For $G_R > 0$, we choose b so that $\sqrt{d}\,2^{-b-1} \le G_R$. An explicit sufficient choice is given in Proposition 14 of Appendix B.4.2.

Initial state preparation. At $\theta = 0$, the discrete Gibbs law is uniform on the admissible grid, thus is independent of the objective function:

$$|\Pi_0\rangle := |\pi^h_{0,R}\rangle = \frac{1}{\sqrt{|\Omega_h|}} \sum_{x\in\Omega_h} |x\rangle. \tag{9}$$

For any $\varepsilon_{\mathrm{st}} \in (0,1)$ and h satisfying Eq. (8), the uniform state $|\Pi_0\rangle$ can be coherently prepared to state error $\varepsilon_{\mathrm{st}}$, and its reflection can be implemented to operator-norm error at most $2\varepsilon_{\mathrm{st}}$, using no calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$; see Lemma 15 in Appendix B.5.

4 Inverse-temperature paths: explicit and adaptive

Let $\beta > 0$ be the target inverse temperature. For $\theta \in [0,\beta]$, let $Z^h_{\theta,R} := \sum_{z\in\Omega_h} e^{-\theta f(z)}$ and define the discrete Gibbs distribution and its coherent encoding by

$$\pi^h_{\theta,R}(x) := e^{-\theta f(x)}/Z^h_{\theta,R}, \qquad |\pi^h_{\theta,R}\rangle := \sum_{x\in\Omega_h} \sqrt{\pi^h_{\theta,R}(x)}\,|x\rangle. \tag{10}$$

Then $|\pi^h_{0,R}\rangle = |\Pi_0\rangle$ and $|\pi^h_{\beta,R}\rangle$ is the target discrete Gibbs state. Following Harrow and Wei (2020), we consider both explicit and adaptive inverse-temperature schedules connecting these endpoints while maintaining constant successive-state overlap.

Explicit baseline. Recall that $\widehat{D}_R := 2RG_R$ is the explicit upper bound on the objective oscillation D_R. If $\widehat{D}_R = 0$, f is constant on Ω and the restricted optimization problem is trivial. Otherwise set

$$M := \left\lceil \beta \widehat{D}_R \right\rceil, \qquad \beta_j := j\beta/M, \qquad j = 0,1,\dots,M, \tag{11}$$

and write $\pi^h_j := \pi^h_{\beta_j,R}$. Then $(\beta_{j+1} - \beta_j)D_R \le 1$.

Lemma 5 (Explicit constant-overlap)**.** *For every $j = 0,\dots,M-1$, we have $|\langle \pi^h_j | \pi^h_{j+1}\rangle| \ge e^{-1/2}$. Thus the explicit discrete path is slowly varying and has $M = O(1 + \beta R G_R)$ stages.*

Adaptive schedule. Define the shifted nonnegative Hamiltonian and its initial mean energy by

$$H_\beta(x) := \beta\big(f(x) - f^\star_R\big), \qquad A_{\beta,R} := \mathbb{E}_{X\sim\pi^h_{0,R}}[H_\beta(X)]. \tag{12}$$

Then

$$0 \le A_{\beta,R} \le \beta D_R \le \beta \widehat{D}_R. \tag{13}$$

Parameterize the discrete family by $t \in [0,1]$ as $\Pi_t := \pi^h_{t\beta,R}$, with coherent encoding $|\Pi_t\rangle$.

Lemma 6 (Adaptive constant-overlap)**.** *With H_β, $A_{\beta,R}$, and $\{\Pi_t\}_{t\in[0,1]}$ defined above, there exists a schedule*

$$0 = t_0 < t_1 < \cdots < t_\ell = 1 \quad \textit{such that} \quad \big|\langle \Pi_{t_i} \mid \Pi_{t_{i+1}}\rangle\big|^2 \ge e^{-2} \quad \textit{for every } i,$$

and $\ell = O\Big(\sqrt{(1 + A_{\beta,R})\log(2 + A_{\beta,R})}\Big) = \widetilde{O}(1 + \sqrt{A_{\beta,R}})$.

See Appendix C.2, the lemma follows from the adaptive cooling-schedule argument of Štefankovič et al. (2009). Following Harrow and Wei (2020), the schedule is found online by nondestructive overlap estimation. Let $\mathcal{E}$ denote the event that all probes succeed with $\Pr(\mathcal{E}) \geq 1 - \chi$, and this will be used in Theorem 10.

Queried inverse temperatures. Let $\mathcal{B} \subset (0, \beta]$ denote the finite set of inverse temperatures at which a Markov chain is queried by the explicit or adaptive route, and set $\theta_{\min} := \min \mathcal{B}$. For the explicit and adaptive routes, respectively, the precise queried sets and the deterministic bound $\theta_{\min} \geq \frac{\beta}{2+\beta \widehat{D}_R}$ are given in Appendix F.1. The endpoint $\theta = 0$ is not included in $\mathcal{B}$: no Markov chain is run there, and the reflection about $|\Pi_0\rangle$ is implemented directly using Lemma 15.

5 Quantum MALA sampling on the finite discrete

For each $\theta \in \mathcal{B}$, we construct a finite-state Metropolis–Hastings chain on Ω_h with target $\pi^h_{\theta,R}$ from Eq. (10), and quantize it using the corresponding Szegedy walk. Our discrete proposal preserves the drift and noise scales of continuous MALA while being defined entirely on finite registers.

For reference, continuous MALA at inverse temperature θ uses the Langevin dynamics

$$y = x - \eta_\theta \nabla f(x) + \sigma_\theta g, \qquad g \sim \mathcal{N}(0, I_d), \qquad \sigma_\theta^2 = 2\eta_\theta/\theta. \tag{14}$$

Accordingly, our finite discrete proposal is

$$y = x + m_\theta(x) + \xi_\theta, \tag{15}$$

where $m_\theta(x) \in h\mathbb{Z}^d$ and ξ_θ are the rounded drift and discrete Gaussian increment defined below. a detailed continuous-to-discrete correspondence is given in Appendix D.1.

5.1 The discrete Metropolis chain

Conditions on the step size. The step sizes $\{\eta_\theta\}_{\theta \in \mathcal{B}}$ are prescribed inputs, for every $\theta \in \mathcal{B}$,

$$(\eta 1)\ \ \eta_\theta \leq \theta R^2/(2d), \qquad (\eta 2)\ \ \eta_\theta G_R \geq h, \qquad (\eta 3)\ \ \eta_\theta \leq 1/(\theta G_R^2. \tag{16}$$

Here G_R is defined in Eq. (3). The window is nonempty for the parameter choices used in Theorem 13, and its feasibility and the roles of these conditions are detailed in Appendix D.2.

5.1.1 Discrete drift and Gaussian noise

Rounded drift. For each $\theta \in \mathcal{B}$, let η_θ satisfy Eq. (16). For $x \in \Omega_h$, define

$$m_\theta(x) := h \cdot \mathrm{round}\Big(-\eta_\theta \widehat{\nabla f}(x)/h\Big) \in h\mathbb{Z}^d. \tag{17}$$

Coordinatewise rounding and Eq. (37) imply

$$\|m_\theta(x)/h\|_\infty \leq r_\theta := \left\lfloor \eta_\theta (G_R + 2^{-b-1})/h + 1/2 \right\rfloor. \tag{18}$$

The corresponding rounding bounds are proved in Appendix D.3.

Discrete Gaussian noise. Let $\sigma_\theta^2 = 2\eta_\theta/\theta$ as in Eq. (14). To obtain a finite-register proposal, let K_θ denote the coordinatewise cutoff radius of the discrete Gaussian increment measured in grid units, and let $\mathcal{D}_{h,\theta}$ denote the corresponding finite set of admissible noise displacements:

$$K_\theta := \max\left\{ r_\theta + 1,\ \left\lceil c\,\frac{\sigma_\theta}{h}\sqrt{\log\big(1/\xi_{\mathrm{impl}}\big)} \right\rceil \right\}, \qquad \mathcal{D}_{h,\theta} := h\left\{ k \in \mathbb{Z}^d : \|k\|_\infty \leq K_\theta \right\}, \tag{19}$$

where $c > 0$ is a sufficiently large absolute constant and $\xi_{\mathrm{impl}} \in (0,1)$ is the implementation-accuracy parameter. The two terms in K_θ play different roles: $r_\theta + 1$ guarantees that the displacements for local grid moves remain in the proposal support for irreducibility (see Appendix D.5), while the second term truncates the noise for finite registers implementation (see Appendix D.3).

We draw $\xi_\theta \in \mathcal{D}_{h,\theta}$ according to

$$\nu_\theta(z) := \exp\big(-\theta\|z\|^2/(4\eta_\theta)\big) \left[\sum_{u \in \mathcal{D}_{h,\theta}} \exp\big(-\theta\|u\|^2/(4\eta_\theta)\big) \right]^{-1}, \qquad z \in \mathcal{D}_{h,\theta}. \tag{20}$$

and set $\nu_\theta(z) = 0$ otherwise. Since $\mathcal{D}_{h,\theta}$ is state-independent, its normalizing constant cancels in the Metropolis–Hastings ratio. Finite-register details are given in Appendix D.3.

5.1.2 Proposal and acceptance

By construction, the proposal Eq. (15) lies in $h\mathbb{Z}^d$. Proposals with $y \notin \Omega_h$ are rejected. For $y \in \Omega_h$, accept with probability

$$\alpha_\theta(x,y) := \min\left\{1, \frac{\pi^h_{\theta,R}(y)\,\nu_\theta\big(x-y-m_\theta(y)\big)}{\pi^h_{\theta,R}(x)\,\nu_\theta\big(y-x-m_\theta(x)\big)}\right\}. \tag{21}$$

The resulting transition matrix on Ω_h is

$$P_\theta(x,y) = \nu_\theta\big(y-x-m_\theta(x)\big)\alpha_\theta(x,y), \qquad y \in \Omega_h \setminus \{x\}, \tag{22}$$

with $P_\theta(x,x) = 1 - \sum_{y\neq x} P_\theta(x,y)$.

The self-loop collects rejected proposals, including proposals leaving Ω_h or having zero reverse-proposal probability. We use the lazy chain $\overline{P}_\theta = (I + P_\theta)/2$. If the reverse displacement lies outside $\mathcal{D}_{h,\theta}$, its proposal probability is zero and the move is rejected by the Metropolis–Hastings filter; see Remark 17. Because $\pi^h_{\theta,R}$ in Eq. (5) enters the Metropolis–Hastings ratio only through $e^{-\theta(f(y)-f(x))}$, the acceptance probability Eq. (21) requires no partition function.

We next establish the properties of the discrete MALA chain: exact reversibility and stationarity from Lemma 7, ergodicity and a positive phase gap from Lemma 8, and coherent implementation with $O(1)$ oracle queries per transition from Proposition 9. Proofs are at Appendix D.4 - D.6.

Lemma 7 (Exact discrete detailed balance)**.** *For every $\theta \in (0,\beta]$ and every $x, y \in \Omega_h$,*

$$\pi^h_{\theta,R}(x)P_\theta(x,y) = \pi^h_{\theta,R}(y)P_\theta(y,x).$$

Hence $\overline{P}_\theta$ is reversible and has $\pi^h_{\theta,R}$ as its exact discrete stationary distribution, and the coherent state $|\pi^h_{\theta,R}\rangle$ is exactly a zero-phase eigenstate of the associated Szegedy walk.

Lemma 8 (Irreducibility and aperiodicity)**.** *Let $0 < h \le R/\sqrt{d}$ and let $\mathcal{D}_{h,\theta}$ satisfy Eq. (19). Then, for every $\theta \in (0,\beta]$, the lazy chain $\overline{P}_\theta$ is irreducible and aperiodic on Ω_h. In particular, its spectral gap, and hence the corresponding Szegedy phase gap, is strictly positive.*

Proposition 9 (Coherent transition with $O(1)$ oracle calls)**.** *Let $T \ge 1$ bound the total number of controlled MALA-transition-state calls and let $\xi_{\mathrm{impl}} \in (0,1)$. For every required inverse temperature θ, U_θ and $U_\theta^\dagger$ admit controlled implementations with operator-norm error at most ξ_{impl}/T per call. Hence any sequence of at most T such calls has total implementation error at most ξ_{impl}.*

5.2 Phase gap and annealing schedules

Let W_θ be the Szegedy walk of the finite reversible chain $\overline{P}_\theta$, restricted to the invariant subspace generated by its transition states Wocjan and Abeyesinghe (2008). By Lemmas 7 and 8, $|\pi^h_{\theta,R}\rangle$ is an isolated zero-phase eigenstate. Define

$$\Delta(\theta;\eta_\theta) := \min\big\{|\varphi| :\ e^{i\varphi} \in \mathrm{spec}(W_\theta),\ \varphi \neq 0\big\}, \tag{23}$$

where the spectrum is restricted to this subspace.

Phase gap versus classical spectral gap. Let $\delta_\theta := 1 - \lambda_2(\overline{P}_\theta)$ be the spectral gap of the lazy reversible chain. Since $\mathrm{spec}(\overline{P}_\theta) \subseteq [0,1]$, the standard Szegedy spectral mapping gives

$$\Delta(\theta;\eta_\theta) = 2\arccos(1-\delta_\theta), \qquad 2\sqrt{\delta_\theta} \le \Delta(\theta;\eta_\theta) \le \pi\sqrt{\delta_\theta}. \tag{24}$$

Thus phase detection has $\widetilde{O}(1/\sqrt{\delta_\theta})$ gap dependence; the lower bound in Eq. (24) controls its required resolution, while the upper bound transfers spectral-gap obstructions to the quantum walk. Section 6 analyzes these gaps explicitly.

For the explicit schedule from Eq. (11) and the adaptive schedule with search grid $\mathcal{G}_\beta$ defined in Appendix F.1, define Δ_{path} and Δ_{all}, respectively:

$$\Delta_{\mathrm{path}} := \min_{1\le j\le M} \Delta(\beta_j;\eta_{\beta_j}), \qquad \Delta_{\mathrm{all}} := \min_{t\in\mathcal{G}_\beta} \Delta(t\beta;\eta_{t\beta}). \tag{25}$$

The endpoint $t = 0$ is not queried by either route, and its reflection is implemented directly using Lemma 15. The monotonicity and adjacent-grid overlap properties required by the adaptive search are given in Appendix F.1.

Known gap for phase detection. Because the phase-detection resolution must be fixed before execution, the algorithm uses a known lower bound

$$0 < \underline{\delta} \le \min_{\theta \in \mathcal{B}} \delta_\theta. \tag{26}$$

By Eq. (24), we have $\Delta(\theta; \eta_\theta) \ge 2\sqrt{\delta_\theta} \ge 2\sqrt{\underline{\delta}}$, so the common resolution $2\sqrt{\underline{\delta}}$ suffices for every queried temperature. $\mathcal{B}$ is fixed in advance: it is the explicit schedule for the nonadaptive route and the full search grid for the adaptive route, regardless of which grid points are eventually visited. Proposition 12 gives an admissible choice of $\underline{\delta}$, while Δ_{path} and Δ_{all} denote the actual gaps.

5.3 Sampling Guarantee on the Discrete Domain

Let ν_h be the measurement distribution produced by the quantum sampler and $\pi^h_{\beta,R}$ the discrete Gibbs target. We require

$$\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \le \zeta, \tag{27}$$

where all sampling-implementation errors are absorbed into ζ. For the adaptive route, the schedule-construction failure probability χ is accounted for separately.

Setup for Q-MALA. Fix $\zeta, \chi \in (0, 1)$. In addition to the stated oracles, suppose we are given $0 < h \le R/\sqrt{d}$, step sizes $\{\eta_\theta\}_{\theta \in \mathcal{B}}$ satisfying Eq. (16), an explicit gap bound $\underline{\delta}$ satisfying Eq. (26), and displacement sets $\{\mathcal{D}_{h,\theta}\}_{\theta \in \mathcal{B}}$ from Eq. (19). For the explicit schedule and adaptive search grid in Eqs. (11) and (58), write

$$M_{\mathrm{exp}} := K_\beta = \left\lceil 1 + \beta \widehat{D}_R \right\rceil, \qquad M_{\mathrm{ad}} := \sqrt{(1 + A_{\beta,R}) \log(2 + A_{\beta,R})}.$$

Let $\mathcal{E}$ be the adaptive success event defined after Lemma 6, and let $\bar{\ell}$ be the computable schedule-length cap from Eq. (65), with $\bar{\ell} = O\left(\sqrt{(1 + \beta \widehat{D}_R) \log(2 + \beta \widehat{D}_R)}\right)$. Let ν_h denote the output law and $\nu^{\mathcal{E}}_h$ its conditional law given $\mathcal{E}$. Choose b by Proposition 14, initialize via Lemma 15 with $\varepsilon_{\mathrm{st}} \le \zeta/4$, and implement the controlled transition-state unitaries of Proposition 9 with total error $\xi_{\mathrm{impl}} \le \zeta/8$. Set implementation precisions from a deterministic a priori call bound. Use M_{exp} for explicit reflection accuracies and $\bar{\ell}$ for adaptive probe and reflection accuracies, aborting beyond $\bar{\ell}$ adaptive steps. This substitution adds only $\log \bar{\ell} = O(\log K_\beta)$ factors.

Theorem 10 (Localized Q-MALA sampling)**.** *Under the preceding setup and precision conventions, the following sampling guarantees hold.*

1. ***Explicit schedule.*** *Traversing Eq. (11) and measuring the position register yields* $\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \le \zeta$ *unconditionally, using* $\widetilde{O}\left(\frac{1+\beta \widehat{D}_R}{\sqrt{\underline{\delta}}}\right)$ *controlled MALA-transition-state calls.*
2. ***Adaptive schedule.*** *On* $\mathcal{E}$*, with* $\Pr(\mathcal{E}) \ge 1 - \chi$*, Algorithm 1 traverses a constant-overlap schedule of length* $O(M_{\mathrm{ad}})$ *and achieves* $\|\nu^{\mathcal{E}}_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \le \zeta$*, using* $\widetilde{O}\left(\frac{1+\sqrt{\beta D_R}}{\sqrt{\underline{\delta}}}\right)$ *controlled MALA-transition-state calls. Unconditionally,* $\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \le \zeta + \chi$.

The same bounds apply to queries to each of $\mathcal{O}_f$ *and* $\mathcal{O}_{\nabla f}$*, since each controlled transition-state call uses* $O(1)$ *queries to each oracle (Proposition 9).*

Explicit versus adaptive. Since $A_{\beta,R} \le \beta \widehat{D}_R$, the adaptive route replaces the explicit schedule factor $O(1 + \beta \widehat{D}_R)$ by $\widetilde{O}(1 + \sqrt{A_{\beta,R}})$. Both bounds use the gap parameter $\underline{\delta}$, so the difference comes from the number of annealing stages rather than the per-stage gap cost. Thus, when $\beta \widehat{D}_R$ is large, the adaptive construction yields a square-root-type reduction in the annealing cost, up to logarithmic factors. A detailed comparison is deferred to Appendix F.4.

6 PHASE GAP LOWER BOUND

No guaranteed polynomial lower bound. Spectral-gap dependence is inherent to quantum simulated annealing: slowly varying Markov-chain constructions require lower bounds on the gaps along the annealing path, which remain explicit runtime parameters in Wocjan and Abeyesinghe (2008) and Harrow and Wei (2020). In the non-logconcave Langevin setting, Ozgul et al. (2024, Remark 10) likewise note that the relevant Cheeger and log-Sobolev constants may scale exponentially with dimension. Proposition 11 exhibits an exponentially small spectral gap within our weak-localization class, ruling out any uniform polynomial gap assumption.

Proposition 11. *Weak localization alone does not imply uniform inverse-polynomial spectral or phase gaps. For every $d \geq 2$, $R > 0$, and $H > 0$, there exists a globally smooth objective satisfying Assumptions 1–2 with $\varepsilon_{\mathrm{loc}} = 0$ and $D_R = H$. For every $h \leq R/(32\sqrt{d})$, every queried inverse temperature $\theta > 0$, and every admissible step size satisfying $2hK_\theta \leq R/4$, the corresponding lazy MALA chain and Szegedy walk satisfy, for an absolute constant $C > 0$,*

$$\delta_\theta \leq \exp(Cd - \theta D_R/2)\,, \qquad \Delta(\theta;\eta_\theta) \leq \pi \exp(Cd/2 - \theta D_R/4)\,. \tag{28}$$

Moreover, if $D_R \geq 3\bar{\varepsilon}$, β is chosen according to Eq. (30), and $2hK_\beta \leq R/4$, then

$$\Delta_{\mathrm{path}}, \Delta_{\mathrm{all}} \leq \pi \exp(-dD_R/(4\bar{\varepsilon})) = \exp(-\Omega(dD_R/\bar{\varepsilon}))\,.$$

Spectral gap lower bound. Despite Proposition 11, every queried finite-state chain admits a computed spectral-gap lower bound.

Proposition 12 (Explicit spectral gap bound)**.** *Under the conditions of Section 5.1, with b chosen as in Proposition 14 and h satisfying Eq. (8) and $h \leq R/(4\sqrt{d})$, every $\theta \in \mathcal{B}$ satisfies*

$$\delta_\theta \;\geq\; e^{-9} c_0^d \frac{h^2}{d^2R^2} \left(\frac{h}{4\sigma_\theta}\right)^d e^{-2\theta \widehat{D}_R} \;\geq\; e^{-9} \frac{h^2}{d^2R^2} \left(\frac{c_0 h\sqrt{d}}{4R}\right)^d e^{-2\beta \widehat{D}_R} =: \underline{\delta}, \tag{29}$$

where $c_0 \in (0,1)$ is absolute and $c_0 = 10^{-3}$ is valid. Thus $\underline{\delta}$ is a priori computable and satisfies $0 < \underline{\delta} \leq \min_{\theta \in \mathcal{B}} \delta_\theta$.

The factor $(h/(4\sigma_\theta))^d$ is a nearest-neighbour comparison loss, not an additional approximation error, and may make $\underline{\delta}$ conservative in high dimension, see proof in Appendix E.2.

7 FROM LOCALIZED GIBBS SAMPLING TO GLOBAL OPTIMIZATION

Fix $\varepsilon > \varepsilon_{\mathrm{loc}}$ and $p_f \in (0,1)$, and set $\bar{\varepsilon} := \varepsilon - \varepsilon_{\mathrm{loc}}$. The reduction combines weak localization with Gibbs concentration: grid rounding consumes objective-value accuracy, while sampling, discretization, and adaptive schedule construction contribute to the failure probability. Appendix G gives the error accounting and supporting Gibbs-tail and sublevel-volume bounds.

Parameter choices. Let $p_f \in (0,1)$ denote the target failure probability, and $\widehat{D}_R > \bar{\varepsilon}$. Define $\log_+(u) := \max\{0, \log u\}$ and choose

$$\beta := \frac{4}{\bar{\varepsilon}} \left[\log \frac{2}{p_f} + d \log_+\left(\frac{8RG_R}{\bar{\varepsilon}}\right)\right]. \tag{30}$$

Appendix G.4 shows the necessary of the β scaling. Set $\xi = \zeta = \chi = p_f/8$, choose h satisfying Eqs. (8) and (71), and let $\underline{\delta}$ satisfy Eq. (26). Use the setup and precision conventions of Theorem 10.

Theorem 13 (Quantum global optimization under weak localization)**.** *Let f satisfy Assumptions 1–2. If $\widehat{D}_R \leq \bar{\varepsilon}$, any $\widehat{x} \in \Omega_h$ is ε-optimal, requiring no quantum-walk steps. Otherwise, under the preceding parameter choices, either annealing schedule returns $\widehat{x} \in \Omega_h$ satisfying $\Pr[f(\widehat{x}) \leq f^\star + \varepsilon] \geq 1 - p_f$, with respective query bounds to each of $\mathcal{O}_f$ and $\mathcal{O}_{\nabla f}$:*

$$C_{\mathrm{ad}} = \widetilde{O}\left(\left(1 + \sqrt{\beta \widehat{D}_R}\right)\Big/\sqrt{\underline{\delta}}\right), \qquad C_{\mathrm{exp}} = \widetilde{O}\left(\left(1 + \beta \widehat{D}_R\right)\Big/\sqrt{\underline{\delta}}\right). \tag{31}$$

On $\mathcal{E}$, the sharper bound $\widetilde{O}((1 + \sqrt{A_{\beta,R}})/\sqrt{\underline{\delta}})$ holds, where $A_{\beta,R} \leq \beta D_R \leq \beta \widehat{D}_R$.

Both routes use the same explicit gap lower bound $\underline{\delta}$ from Proposition 12 and Eq. (29), eliminating unknown Cheeger or spectral-gap terms. The prescribed step sizes satisfy Eq. (16) and are absorbed into $\underline{\delta}$; no step-size optimization is claimed.

AI USE STATEMENT

The idea of applying quantum simulated annealing to non-convex global optimization under assumptions weaker than dissipativity originated with the authors. We used ChatGPT 5.6 Sol to refine proofs, with particular attention to the localization assumption and the discretization arguments for continuous quantum states. These tools also assisted with language editing and improving the clarity of the exposition. The authors have reviewed all AI-assisted work and are responsible for verifying the correctness of all presented results and for the final revision and organization of the manuscript. We take full responsibility for the final content of this work.

ETHICS STATEMENT

This work is theoretical and does not involve human participants, personal data, or experiments on deployed systems. We are not aware of any direct ethical concerns arising from the results presented in this paper.

REPRODUCIBILITY STATEMENT

This work is theoretical and does not rely on empirical experiments or proprietary data. We provide complete proofs of all main results and technical claims in the paper and appendices, including the assumptions and intermediate lemmas needed to verify the analysis.

REFERENCES

Simon Apers and Alain Sarlette. Quantum fast-forwarding: Markov chains and graph property testing. *Quantum Information & Computation*, 19(3–4):181–213, 2019.

August Y. Chen, Ayush Sekhari, and Karthik Sridharan. Langevin dynamics: A unified perspective on optimization via Lyapunov potentials. arXiv preprint arXiv:2407.04264, 2024.

Andrew M. Childs, Tongyang Li, Jin-Peng Liu, Chunhao Wang, and Ruizhe Zhang. Quantum algorithms for sampling log-concave distributions and estimating normalizing constants. In *Advances in Neural Information Processing Systems*, volume 35, pages 23205–23217, 2022.

Zhongxiang Dai, Gregory Kang Ruey Lau, Arun Verma, Yao Shu, Bryan Kian Hsiang Low, and Patrick Jaillet. Quantum bayesian optimization. *Advances in neural information processing systems*, 36:20179–20207, 2023.

Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of Markov chains. *The Annals of Applied Probability*, 1(1):36–61, 1991.

Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.

Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. In *International conference on machine learning*, pages 5200–5209. PMLR, 2019.

Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions, 2002. arXiv:quant-ph/0208112.

Wei Guo, Molei Tao, and Yongxin Chen. Provable benefit of annealed Langevin Monte Carlo for non-log-concave sampling. In *International Conference on Learning Representations*, 2025.

Thomas Häner, Martin Roetteler, and Krysta M. Svore. Optimizing quantum circuits for arithmetic, 2018. arXiv:1805.12445.

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

Dylan Herman, Guneykan Ozgul, Anuj Apte, Junhyung Lyle Kim, Anupam Prakash, Jiayu Shen, and Shouvanik Chakrabarti. Mechanisms for quantum advantage in global optimization of non-convex functions. arXiv preprint arXiv:2510.03385, 2025.

Alexandros Hollender and Manolis Zampetakis. The computational complexity of finding stationary points in non-convex optimization: A. hollender, m. zampetakis. *Mathematical Programming*, 213(1):281–341, 2025.

Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *International conference on machine learning*, pages 1724–1732. PMLR, 2017.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.

Tongyang Li and Ruizhe Zhang. Quantum speedups of optimizing approximately convex functions with applications to logarithmic regret stochastic convex bandits. In *Advances in Neural Information Processing Systems*, volume 35, pages 3152–3164. Curran Associates, Inc., 2022.

Songtao Lu, Mingyi Hong, and Zhengdao Wang. Pa-gd: On the convergence of perturbed alternating gradient descent to second-order stationary points for structured nonconvex optimization. In *International Conference on Machine Learning*, pages 4134–4143. PMLR, 2019.

Songtao Lu, Meisam Razaviyayn, Bo Yang, Kejun Huang, and Mingyi Hong. Finding second-order stationary points efficiently in smooth nonconvex linearly constrained optimization problems. *Advances in Neural Information Processing Systems*, 33:2811–2822, 2020.

Oren Mangoubi and Nisheeth K. Vishnoi. Nonconvex sampling with the Metropolis-adjusted Langevin algorithm. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2259–2293. PMLR, 2019.

Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

Guneykan Ozgul, Xiantao Li, Mehrdad Mahdavi, and Chunhao Wang. Stochastic quantum sampling for non-logconcave distributions and estimating partition functions. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 38953–38982. PMLR, 2024.

Guneykan Ozgul, Xiantao Li, Mehrdad Mahdavi, and Chunhao Wang. Quantum speedups for sampling and non-convex optimization with stochastic oracles. In *21st Conference on the Theory of Quantum Computation, Communication and Cryptography*, volume 389 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:25, 2026.

Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: A nonasymptotic analysis. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1674–1703. PMLR, 2017.

Rolando D. Somma, Sergio Boixo, Howard Barnum, and Emanuel Knill. Quantum simulations of classical annealing processes. *Physical Review Letters*, 101(13):130504, 2008.

Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *Journal of the ACM*, 56(3):18:1–18:36, 2009. doi: 10.1145/1516512.1516520.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices. *Advances in neural information processing systems*, 32, 2019.

Pawel Wocjan and Anura Abeyesinghe. Speed-up via quantum sampling. *Physical Review A*, 78(4): 042336, 2008. doi: 10.1103/PhysRevA.78.042336.

Pan Xu, Jinghui Chen, Difan Zou, and Quanquan Gu. Global convergence of langevin dynamics based algorithms for nonconvex optimization. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

A Detailed annealing-path comparison

For localized Q-MALA, Π_0 is uniform on Ω and $H_\beta = \beta(f - f_R^\star)$. The adaptive construction estimates $|\langle \Pi_{t_i} | \Pi_{t'} \rangle|^2$ nondestructively and binary-searches for a next constant-overlap temperature. The explicit alternative uses $(\beta_{j+1} - \beta_j) D_R \le 1$, enforced through the known upper bound $\widehat{D}_R$.

For soft Q-ULA, restoring the inverse temperature absorbed into the potential by Ozgul et al. (2024), the states are

$$\mu_0^\lambda(x) \propto e^{-\|x\|^2/(2\sigma_1^2)}, \qquad \mu_i^\lambda(x) \propto e^{-\beta F_\lambda(x) - \|x\|^2/(2\sigma_i^2)}, \quad \sigma_{i+1}^2 = (1+\alpha)\sigma_i^2. \tag{32}$$

The Q-ULA path therefore contains a permanent soft penalty in F_λ and a temporary Gaussian warm-start term. Its terminal regularized state has constant overlap with $\pi_{\beta,\lambda}$ rather than equaling it, requiring a final transition.

Table 1: Annealing constructions used in the paper.

Feature	Localized Q-MALA	Soft Q-ULA via Ozgul et al. (2024)
State family	$e^{-t\beta f}\mathbf{1}_\Omega$	$e^{-\beta F_\lambda - \|x\|^2/(2\sigma_i^2)}$
Annealed parameter	$t \uparrow 1$	$\sigma_i^2 \uparrow \infty$
Schedule	Adaptive overlap search; explicit constant spacing is a baseline	Geometric Gaussian deconfinement
Start and support	Uniform on fixed compact Ω	Gaussian; full-space target with truncation bookkeeping
Endpoint	Exact target $\pi_{\beta,R}$	Constant-overlap warm state, then target
Overlap control	Free-energy schedule or bounded oscillation D_R	Dissipativity and log-Sobolev moment bounds
Stage count	$\widetilde{O}(1+\sqrt{A_{\beta,R}})$ adaptive; $O(1+\beta\widehat{D}_R)$ explicit	$\widetilde{O}(c_{\mathrm{LSI}}^{-1}\sqrt{dL/m})$ as reported by Ozgul et al. (2024); not claimed here

The constructions address different domain models. The compact family needs no full-space normalizability, while the Gaussian-deconfinement analysis obtains its nominal schedule using dissipativity and functional-inequality control. We do not use the stochastic-gradient Q-ULA variant.

B Proofs for Problem Setup and Discretization

B.1 Weak Localization versus Dissipativity

Proof of Proposition 1. We first show that dissipativity implies exact weak localization. For $x = r\hat{x}$ with $r \ge 1$ and $\|\hat{x}\| = 1$,

$$\frac{d}{dr} f(r\hat{x}) = \frac{1}{r}\langle \nabla f(r\hat{x}), r\hat{x} \rangle \ge mr - \frac{b}{r}.$$

Integrating from 1 to r gives

$$f(r\hat{x}) \ge f(\hat{x}) + \frac{m}{2}(r^2 - 1) - b \log r.$$

Since f is continuous on the unit sphere, the right-hand side tends uniformly to $+\infty$ as $r \to \infty$. Thus f is coercive and, by continuity, attains its global minimum at some $x^\star$.

Differentiability gives $\nabla f(x^\star) = 0$. Substituting $x^\star$ into Eq. (1) yields

$$0 = \langle \nabla f(x^\star), x^\star \rangle \ge m\|x^\star\|^2 - b,$$

so

$$\|x^\star\| \le \sqrt{b/m}.$$

Hence every $R \ge \sqrt{b/m}$ contains a global minimizer, and therefore Assumption 2 holds with $\varepsilon_{\mathrm{loc}} = 0$.

To show that the converse fails, and hence that weak localization is strictly weaker than dissipativity, consider

$$f(x) = 1 - e^{-\|x\|^2}.$$

Its Hessian is

$$\nabla^2 f(x) = 2e^{-\|x\|^2} I - 4e^{-\|x\|^2} xx^\top,$$

whose operator norm is uniformly bounded, so f is globally smooth. Its radial Hessian eigenvalue is

$$2e^{-r^2}(1-2r^2),$$

which is negative for $r > 1/\sqrt{2}$, so f is nonconvex. Moreover, $x = 0$ is its unique global minimizer, so exact weak localization holds for every $R > 0$.

On the other hand,

$$\langle \nabla f(x), x\rangle = 2\|x\|^2 e^{-\|x\|^2} \longrightarrow 0 \qquad \text{as } \|x\| \to \infty.$$

For every $m > 0$ and finite b,

$$m\|x\|^2 - b \longrightarrow +\infty,$$

so the dissipativity inequality cannot hold. Thus weak localization does not imply dissipativity. □

B.2 Localized Gradient and Oscillation Bounds

Proof of Lemma 2. For every $x \in \Omega$, Assumption 1 and $G_0 \geq \|\nabla f(0)\|$ imply

$$\|\nabla f(x)\| \leq \|\nabla f(x) - \nabla f(0)\| + \|\nabla f(0)\| \leq L\|x\| + G_0 \leq LR + G_0 = G_R. \tag{33}$$

Moreover, for any $x, y \in \Omega$, convexity of the ball implies

$$y + t(x-y) \in \Omega \qquad \text{for every } t \in [0,1].$$

Hence, by the fundamental theorem of calculus and Eq. (33),

$$\begin{aligned} |f(x) - f(y)| &= \left| \int_0^1 \big\langle \nabla f\big(y + t(x-y)\big), x - y \big\rangle \, dt \right| \\ &\leq \int_0^1 \big\| \nabla f\big(y + t(x-y)\big) \big\| \, \|x - y\| \, dt \\ &\leq G_R \|x - y\|. \end{aligned} \tag{34}$$

Thus f is G_R-Lipschitz on Ω.

Finally, since $\operatorname{diam}(\Omega) = 2R$,

$$D_R = \max_{x\in\Omega} f(x) - \min_{x\in\Omega} f(x) \leq G_R \operatorname{diam}(\Omega) = 2RG_R.$$

□

B.3 Computational Grid Geometry and Encoding

Recall from Section 3.2 that

$$\Lambda_h = h\mathbb{Z}^d, \qquad C_x = x + [0,h)^d, \qquad \Omega_h = \{x \in \Lambda_h : C_x \subseteq \Omega\}.$$

The cells $\{C_x\}_{x\in\Lambda_h}$ tile $\mathbb{R}^d$, with

$$\operatorname{Vol}(C_x) = h^d, \qquad \operatorname{diam}(C_x) = \sqrt{d}\, h.$$

Recall also

$$\widehat{\Omega}_h := \bigcup_{x\in\Omega_h} C_x,$$

and define the uncovered boundary shell

$$S_h := \Omega \setminus \widehat{\Omega}_h.$$

For $y \in \mathbb{R}^d$, let $x(y) \in \Lambda_h$ be the corner of the unique cell containing y. Then every $z \in C_{x(y)}$ satisfies

$$\|z - y\| \leq \sqrt{d}\, h.$$

Consequently, if $\|y\| \le R - \sqrt{d}\,h$, then $C_{x(y)} \subseteq B(0,R) = \Omega$. Hence

$$\{x \in \Lambda_h : \|x\| \le R - \sqrt{d}\,h\} \subseteq \Omega_h, \qquad B(0, R - \sqrt{d}\,h) \subseteq \widehat{\Omega}_h. \tag{35}$$

Every $x \in \Omega_h$ satisfies $\|x\| \le R$. Thus each coordinate x_j/h takes at most $2\lfloor R/h \rfloor + 1$ integer values, and a position can be encoded using

$$d\left\lceil \log_2\left(2\left\lfloor \frac{R}{h} \right\rfloor + 1\right)\right\rceil = O\left(d \log\left(1 + \frac{R}{h}\right)\right) \tag{36}$$

qubits.

Equivalently, positions may be represented in the regular bounding box

$$\{-\lfloor R/h \rfloor, \dots, \lfloor R/h \rfloor\}^d.$$

Membership in Ω_h is tested by checking whether the entire cell C_x lies in $B(0,R)$; for the Euclidean ball this reduces to comparing the squared norm of the farthest corner of C_x with R^2.

B.4 Discrete Approximation of the Localized Gibbs Distribution

B.4.1 Total-Variation Approximation of the Gibbs Distribution

Proof of Lemma 4. Write p for the density of $\pi_{\beta,R}$ and $q := \widetilde{\pi}^h_{\beta,R}$, and let $p_{\widehat{\Omega}_h}$ denote the density of $\pi_{\beta,R}$ conditioned on $\widehat{\Omega}_h$. Let

$$S_h := \Omega \setminus \widehat{\Omega}_h$$

denote the uncovered boundary shell.

Since q is supported on $\widehat{\Omega}_h$, the triangle inequality gives

$$\|q - p\|_{\mathrm{TV}} \le \|q - p_{\widehat{\Omega}_h}\|_{\mathrm{TV}} + \|p_{\widehat{\Omega}_h} - p\|_{\mathrm{TV}}.$$

Because conditioning removes exactly the mass of the boundary shell S_h,

$$\|p_{\widehat{\Omega}_h} - p\|_{\mathrm{TV}} = \pi_{\beta,R}(S_h).$$

We first control the error inside the covered region $\widehat{\Omega}_h$. For $y \in \widehat{\Omega}_h$, both densities are positive and

$$\frac{q(y)}{p_{\widehat{\Omega}_h}(y)} = e^{-\beta(f(x(y)) - f(y))} \frac{\displaystyle\int_{\widehat{\Omega}_h} e^{-\beta f(z)}\,dz}{\displaystyle\int_{\widehat{\Omega}_h} e^{-\beta f(x(z))}\,dz}.$$

By Lemma 3,

$$|f(x(y)) - f(y)| \le \omega_h,$$

and hence

$$e^{-\beta\omega_h} \le e^{-\beta(f(x(y)) - f(y))} \le e^{\beta\omega_h}.$$

The same cell-oscillation bound gives, pointwise on $\widehat{\Omega}_h$,

$$e^{-\beta\omega_h} e^{-\beta f(z)} \le e^{-\beta f(x(z))} \le e^{\beta\omega_h} e^{-\beta f(z)}.$$

Integrating over $\widehat{\Omega}_h$ therefore yields

$$e^{-\beta\omega_h} \int_{\widehat{\Omega}_h} e^{-\beta f(z)}\,dz \le \int_{\widehat{\Omega}_h} e^{-\beta f(x(z))}\,dz \le e^{\beta\omega_h} \int_{\widehat{\Omega}_h} e^{-\beta f(z)}\,dz,$$

so the normalization-factor ratio also belongs to $[e^{-\beta\omega_h}, e^{\beta\omega_h}]$. Consequently,

$$e^{-2\beta\omega_h} \le \frac{q(y)}{p_{\widehat{\Omega}_h}(y)} \le e^{2\beta\omega_h},$$

and in particular

$$q(y) \geq e^{-2\beta\omega_h} p_{\widehat{\Omega}_h}(y) \qquad \text{for all } y \in \widehat{\Omega}_h.$$

Since q and $p_{\widehat{\Omega}_h}$ are probability densities on the same support,

$$\|q - p_{\widehat{\Omega}_h}\|_{\mathrm{TV}} = \int_{\widehat{\Omega}_h} \big(p_{\widehat{\Omega}_h}(y) - q(y)\big)_+ \, dy.$$

Using the pointwise lower bound on q,

$$\begin{aligned}\|q - p_{\widehat{\Omega}_h}\|_{\mathrm{TV}} &\leq \int_{\widehat{\Omega}_h} \big(1 - e^{-2\beta\omega_h}\big)\, p_{\widehat{\Omega}_h}(y)\, dy \\ &= 1 - e^{-2\beta\omega_h}.\end{aligned}$$

Thus

$$\|q - p\|_{\mathrm{TV}} \leq 1 - e^{-2\beta\omega_h} + \pi_{\beta,R}(S_h).$$

Notice that this argument only uses ratios of Gibbs weights, so no sign assumption on f is required.

It remains to control the Gibbs mass of the uncovered boundary shell. Set

$$s := 1 - \frac{\sqrt{d}\, h}{R}.$$

Under the choice Eq. (8), we have $0 < s < 1$. Define the scaling map

$$T_s(y) := sy.$$

For every $y \in \Omega = B(0, R)$,

$$\|T_s(y)\| = s\|y\| \leq sR = R - \sqrt{d}\, h.$$

Hence, by Eq. (35),

$$T_s(\Omega) = B(0, R - \sqrt{d}\, h) \subseteq \widehat{\Omega}_h.$$

Moreover,

$$\|T_s(y) - y\| = (1 - s)\|y\| \leq (1 - s)R = \sqrt{d}\, h.$$

Since f is G_R-Lipschitz on Ω,

$$|f(T_s(y)) - f(y)| \leq G_R \sqrt{d}\, h = \omega_h.$$

Therefore

$$e^{-\beta f(T_s(y))} \geq e^{-\beta\omega_h} e^{-\beta f(y)}.$$

Using $T_s(\Omega) \subseteq \widehat{\Omega}_h$ and the change of variables $z = T_s(y)$, we obtain

$$\begin{aligned}\int_{\widehat{\Omega}_h} e^{-\beta f(z)}\, dz &\geq \int_{T_s(\Omega)} e^{-\beta f(z)}\, dz \\ &= s^d \int_{\Omega} e^{-\beta f(T_s(y))}\, dy \\ &\geq s^d e^{-\beta\omega_h} \int_{\Omega} e^{-\beta f(y)}\, dy.\end{aligned}$$

Dividing by

$$Z_{\beta,R} = \int_{\Omega} e^{-\beta f(y)}\, dy$$

gives

$$\pi_{\beta,R}(\widehat{\Omega}_h) \geq s^d e^{-\beta\omega_h}.$$

Hence

$$\pi_{\beta,R}(S_h) = 1 - \pi_{\beta,R}(\widehat{\Omega}_h) \leq 1 - s^d e^{-\beta\omega_h}.$$

Using

$$1 - ab = (1 - a) + a(1 - b) \leq (1 - a) + (1 - b) \qquad (a, b \in [0, 1]),$$

we obtain

$$\pi_{\beta,R}(S_h) \le (1 - s^d) + (1 - e^{-\beta\omega_h}).$$

Since $s \in [0, 1]$,

$$1 - s^d \le d(1 - s) = \frac{d^{3/2}h}{R},$$

and since $1 - e^{-u} \le u$ for $u \ge 0$,

$$1 - e^{-\beta\omega_h} \le \beta\omega_h.$$

Therefore

$$\pi_{\beta,R}(S_h) \le \beta\omega_h + \frac{d^{3/2}h}{R}.$$

Combining the cell-quantization and boundary-shell bounds gives

$$\begin{aligned}\left\|\widetilde{\pi}^h_{\beta,R} - \pi_{\beta,R}\right\|_{\mathrm{TV}} &\le 1 - e^{-2\beta\omega_h} + \beta\omega_h + \frac{d^{3/2}h}{R} \\ &\le 3\beta\omega_h + \frac{d^{3/2}h}{R},\end{aligned}$$

where the last inequality again uses $1 - e^{-u} \le u$.

Finally, by Eq. (7),

$$\omega_h = G_R\sqrt{d}\, h.$$

The first condition in Eq. (8),

$$h \le \frac{\xi}{6\sqrt{d}\,\beta G_R},$$

implies

$$3\beta\omega_h = 3\beta G_R\sqrt{d}\, h \le \frac{\xi}{2},$$

while the second condition,

$$h \le \frac{\xi R}{2d^{3/2}},$$

implies

$$\frac{d^{3/2}h}{R} \le \frac{\xi}{2}.$$

Consequently,

$$\left\|\widetilde{\pi}^h_{\beta,R} - \pi_{\beta,R}\right\|_{\mathrm{TV}} \le \frac{\xi}{2} + \frac{\xi}{2} = \xi,$$

as claimed. □

B.4.2 Finite-Precision Gradient Representation

Proposition 14. *Let $b \in \mathbb{N}$ be the gradient-rounding precision, and let rd_b denote round-to-nearest onto $2^{-b}\mathbb{Z}$. Define the rounded gradient coordinatewise by $\widehat{\nabla} f(x) := \mathrm{rd}_b(\nabla f(x))$. Then, for every $x \in \Omega_h$,*

$$\|\widehat{\nabla} f(x) - \nabla f(x)\|_\infty \le 2^{-b-1}. \tag{37}$$

Moreover, in the nontrivial regime $G_R > 0$, it suffices to choose

$$b \ge \left\lceil \max\left\{0,\, \log_2\left(\frac{\sqrt{d}}{2G_R}\right)\right\}\right\rceil, \tag{38}$$

which guarantees

$$\sqrt{d}\, 2^{-b-1} \le G_R. \tag{39}$$

Proof of Proposition 14. The pointwise bound in Eq. (37) follows immediately from round-to-nearest with grid spacing 2^{-b}.

The lower bound Eq. (38) gives

$$2^{-b} \le \frac{2G_R}{\sqrt{d}},$$

and hence

$$\sqrt{d}\,2^{-b-1} \le G_R,$$

which proves Eq. (39). □

B.5 Preparation of the Uniform Domain State

Lemma 15 (Coherent uniform state on the admissible grid)**.** *Let $\varepsilon_{\mathrm{st}} \in (0,1)$ and let h satisfy Eq. (8). Then there is a quantum circuit U_0 such that*

$$U_0|0\rangle|0\rangle_{\mathrm{anc}} = |\psi_0\rangle|0\rangle_{\mathrm{anc}}, \qquad \mathrm{supp}|\psi_0\rangle \subseteq \Omega_h, \qquad \big\| \,|\psi_0\rangle - |\Pi_0\rangle\, \big\| \le \varepsilon_{\mathrm{st}}. \tag{40}$$

The circuit uses no calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$ and has gate and workspace complexity polynomial in d, $\log(R/h)$, and $\log(1/\varepsilon_{\mathrm{st}})$. Consequently,

$$U_0(2|0\rangle\langle 0| - I)U_0^\dagger$$

implements the reflection about $|\Pi_0\rangle$ up to operator-norm error at most $2\varepsilon_{\mathrm{st}}$, with no objective-oracle or gradient-oracle query.

Proof of Lemma 15. The circuit prepares the coherent uniform state on the ball and then projects onto full cells. No ball–cell intersection volume is computed at any point.

Step 1: uniform state on the ball. Prepare, to state error ϵ_0,

$$|B_R\rangle := \mathrm{Vol}\big(B(0,R)\big)^{-1/2} \int_{B(0,R)} |y\rangle\, dy$$

coordinate by coordinate from its exact conditional marginals, in the controlled-rotation form of Grover and Rudolph (2002). After $y_1, \dots, y_{j-1}$ have been prepared, put $R_j^2 := R^2 - \sum_{i<j} y_i^2$; then

$$p(y_j \mid y_{<j}) \;\propto\; \big(R_j^2 - y_j^2\big)^{(d-j)/2}, \qquad -R_j \le y_j \le R_j,$$

whose cumulative distribution function is a regularized incomplete-beta function after the affine change of variable $z = (y_j/R_j + 1)/2$. The rotation angle for coordinate j is obtained by evaluating that function in a fixed-point register.

The evaluation is elementary because the exponent is a half-integer. With $t = y_j/R_j$, the standard reduction of $\int_0^t (1-u^2)^k\, du$ and $\int_0^t (1-u^2)^{k+1/2}\, du$ shows that the required integral is a polynomial in t of degree $d-j+1$ when $d-j$ is even, and is $P(t)\sqrt{1-t^2} + c\arcsin t$ with P a polynomial of degree $d-j$ and c a rational constant when $d-j$ is odd. The coefficients depend only on (d,j) and are computed classically offline, so evaluating either branch to p bits costs $O(d)$ reversible fixed-point multiply–adds together with at most one square root and one arcsine, and reversible fixed-point circuits evaluate those to p bits with $\mathrm{poly}(p)$ gates (Häner et al., 2018). Over the d coordinates the cost is $\mathrm{poly}(d,p)$.

This is not the construction the lemma is avoiding, although both are preparations from integrated marginals: the integrals required here are one-dimensional and have the closed form just displayed, whereas the alternative requires coherent access to the d-dimensional intersection volumes $\mathrm{Vol}(C_x \cap \Omega)$, for which no such routine is available. Nor can Step 1 be replaced by Hadamards on the bounding box followed by postselection onto $B(0,R)$: that succeeds with probability $\mathrm{Vol}(B_R)/(2R)^d = (c/\sqrt{d})^{\,d}$, which is exponentially small.

Step 2: cell label and admissibility. Decrease h if necessary to a power-of-two multiple of the position-register unit, preserving Eq. (8). Then reversibly compute the cell label $x = x(y) \in \Lambda_h$, offset $u = y - x \in [0,h)^d$, and

$$a(x) := \mathbf{1}\{C_x \subseteq \Omega\}.$$

Because h is a power of two times the fixed-point unit, as fixed after Eq. (8), the map $y \leftrightarrow (x, u)$ is an exact split of each coordinate register into high and low bits; $a(x)$ is then a comparison against R^2 of the squared norm of the farthest corner of C_x. All of this is finite arithmetic on the position registers and requires no query to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$.

Step 3: projection onto admissible cells. Project on $a(x) = 1$ by fixed-point amplitude amplification. By Eq. (35),

$$p_{\mathrm{acc}} = \frac{\mathrm{Vol}(\widehat{\Omega}_h)}{\mathrm{Vol}(\Omega)} \geq \Big(1 - \frac{\sqrt{d}\, h}{R}\Big)^d,$$

and Eq. (8) with $\xi \in (0, 1)$ gives $\sqrt{d}\, h/R \leq 1/(2d)$, so $p_{\mathrm{acc}} \geq (1 - 1/(2d))^d \geq 1/2$ and $O(\log(1/\varepsilon_{\mathrm{st}}))$ amplification rounds suffice.

Step 4: the accepted state is uniform on Ω_h. On the accepted subspace every cell is a full translate of the same cube $[0, h)^d$, so the change of variables of Step 2 factorizes exactly,

$$\frac{1}{\sqrt{|\Omega_h|h^d}} \sum_{x \in \Omega_h} |x\rangle \int_{[0,h)^d} |u\rangle\, du = |\Pi_0\rangle \otimes |U_h\rangle.$$

The offset register therefore carries the same state $|U_h\rangle$ for every x; it is uncomputed by the inverse of Step 2 and discarded. This is the point at which the amplitudes on Ω_h come out uniform without any boundary-cell volume having been evaluated.

Precision and cost. Take the coordinate precision

$$p = O\bigg(\log \frac{R}{h} + \log \frac{d}{\varepsilon_{\mathrm{st}}}\bigg)$$

and ϵ_0 small enough that the ball-state preparation, the reversible arithmetic and the amplification contribute $O(\varepsilon_{\mathrm{st}})$ in total. By Step 1 the gate and workspace complexity is $\mathrm{poly}(d, \log(R/h), \log(1/\varepsilon_{\mathrm{st}}))$, and no call to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$ is made.

Reflection. Let

$$R_0 := 2|\Pi_0\rangle\langle\Pi_0| - I, \qquad \widetilde{R}_0 := U_0(2|0\rangle\langle 0| - I)U_0^\dagger.$$

Since U_0 prepares $|\psi_0\rangle$ with $\big\||\psi_0\rangle - |\Pi_0\rangle\big\| \leq \varepsilon_{\mathrm{st}}$,

$$\|\widetilde{R}_0 - R_0\| = 2\big\||\psi_0\rangle\langle\psi_0| - |\Pi_0\rangle\langle\Pi_0|\big\| \leq 2\big\||\psi_0\rangle - |\Pi_0\rangle\big\| \leq 2\varepsilon_{\mathrm{st}}.$$

Thus the reflection also requires no calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$. □

C Proofs for the Inverse-Temperature Annealing Paths

C.1 Explicit Constant-Overlap Schedule

Discrete partition-function. For the overlap calculation below, define

$$Z^h_{\theta,R} := h^d \sum_{z \in \Omega_h} e^{-\theta f(z)}, \qquad \theta \geq 0. \tag{41}$$

The factor h^d makes $Z^h_{\theta,R}$ the Riemann-sum analogue of the continuous partition function $Z_{\theta,R}$ in Eq. (2); it cancels from the normalized discrete Gibbs law in Eq. (5). Equivalently,

$$\pi^h_{\theta,R}(x) = \frac{h^d e^{-\theta f(x)}}{Z^h_{\theta,R}}.$$

Proof of Lemma 5. Fix j and let $\gamma = \beta_{j+1} - \beta_j$. Shifting f by the constant $f^\star_R$ does not change any discrete Gibbs distribution, so on Ω_h we have

$$0 \leq f(x) - f^\star_R \leq D_R.$$

Using the discrete partition function defined in Eq. (41) together with Eq. (5),

$$\langle \pi_j^h | \pi_{j+1}^h \rangle = \frac{Z^h_{\beta_j+\gamma/2,R}}{\sqrt{Z^h_{\beta_j,R} Z^h_{\beta_j+\gamma,R}}} = \frac{\mathbb{E}_{X\sim\pi_j^h}\left[e^{-\gamma(f(X)-f_R^\star)/2}\right]}{\mathbb{E}_{X\sim\pi_j^h}\left[e^{-\gamma(f(X)-f_R^\star)}\right]^{1/2}}.$$

The denominator is at most one, while the numerator is at least $e^{-\gamma D_R/2}$. Since the schedule satisfies $\gamma D_R \le 1$,

$$\langle \pi_j^h | \pi_{j+1}^h \rangle \ge e^{-1/2}.$$

Finally, $M = \lceil \beta \widehat{D}_R \rceil = O(1 + \beta R G_R)$ by Eq. (4). □

C.2 Adaptive Constant-Overlap Schedule

Proof of Lemma 6. Let

$$Z(t) = \mathbb{E}_{X\sim\Pi_0}\left[e^{-tH_\beta(X)}\right], \qquad \phi(t) = \log Z(t).$$

Then

$$\phi'(t) = -\mathbb{E}_{X\sim\Pi_t}[H_\beta(X)], \qquad \phi''(t) = \mathrm{Var}_{X\sim\Pi_t}(H_\beta(X)) \ge 0.$$

Hence ϕ is convex.

Define

$$g(t) := \phi(t) - t.$$

Since the subtraction of a linear function does not change the midpoint convexity gap,

$$g\left(\frac{s+t}{2}\right) - \frac{g(s)+g(t)}{2} = \phi\left(\frac{s+t}{2}\right) - \frac{\phi(s)+\phi(t)}{2}.$$

Moreover, the squared overlap between two Gibbs states satisfies

$$|\langle \Pi_s \mid \Pi_t \rangle|^2 = \frac{Z((s+t)/2)^2}{Z(s)Z(t)} = \exp\left(2\phi\left(\frac{s+t}{2}\right) - \phi(s) - \phi(t)\right).$$

Therefore, the SVV midpoint condition

$$g\left(\frac{s+t}{2}\right) \ge \frac{g(s)+g(t)}{2} - 1$$

implies

$$|\langle \Pi_s \mid \Pi_t \rangle|^2 \ge e^{-2}.$$

Thus a cooling schedule satisfying the SVV midpoint condition for g also gives the desired constant-overlap schedule for the Gibbs family Π_t.

Moreover,

$$g'(t) = -\mathbb{E}_{\Pi_t}[H_\beta] - 1 < 0,$$

so g is decreasing. Its convexity is separate and already established above: $g'' = \phi'' = \mathrm{Var}_{\Pi_t}(H_\beta) \ge 0$. Hence g is a decreasing convex function on $[0,1]$.

Applying Lemma 4.3 of Štefankovič, Vempala, and Vigoda to the decreasing convex function g on $[0,1]$ gives

$$\ell \le 1 + \sqrt{[g(0) - g(1)] \log \frac{g'(0)}{g'(1)}}.$$

Let

$$F_{\beta,R} := -\log Z(1).$$

Then

$$g(0) - g(1) = 1 + F_{\beta,R}.$$

By Jensen's inequality,

$$F_{\beta,R} = -\log \mathbb{E}_{\Pi_0}[e^{-H_\beta}] \le \mathbb{E}_{\Pi_0}[H_\beta] = A_{\beta,R}.$$

Furthermore,

$$\frac{g'(0)}{g'(1)} = \frac{1 + A_{\beta,R}}{1 + \mathbb{E}_{\Pi_1}[H_\beta]} \le 1 + A_{\beta,R}.$$

Since $\phi''(t) \ge 0$, we have

$$\mathbb{E}_{\Pi_0}[H_\beta] \ge \mathbb{E}_{\Pi_1}[H_\beta].$$

Therefore,

$$1 \le \frac{g'(0)}{g'(1)} \le 1 + A_{\beta,R},$$

and hence

$$0 \le \log \frac{g'(0)}{g'(1)} \le \log(1 + A_{\beta,R}).$$

Therefore

$$\ell = O\left(\sqrt{(1 + A_{\beta,R}) \log(2 + A_{\beta,R})}\right).$$

□

D Proofs for the Discrete Q-MALA Construction

D.1 Continuous-to-Discrete MALA Correspondence

The table below summarizes how the components of continuous MALA are represented in our finite-state construction.

continuous	**discrete**	**defined in**
$-\eta_\theta \nabla f(x)$	$m_\theta(x)$	Eq. (17)
$\mathcal{N}(0, \sigma_\theta^2 I_d)$	ν_θ on $\mathcal{D}_{h,\theta}$	Eq. (20)
target $\pi_{\theta,R}$	$\pi^h_{\theta,R}$	Eq. (10)
proposal density	$\nu_\theta(y - x - m_\theta(x))$	Eq. (22)
boundary rejection on Ω	boundary rejection on Ω_h	Eq. (22)

D.2 Step-Size Feasibility and High-Temperature Control

We first verify that the window in Eq. (16) is nonempty at every queried temperature. Compatibility of $(\eta 2)$ and $(\eta 3)$ follows from Eq. (8), since

$$h\theta G_R \le h\beta G_R \le \frac{\xi}{6\sqrt{d}} < 1.$$

For $(\eta 2)$ and $(\eta 1)$, using $\theta \ge \theta_{\min} \ge \beta/(2 + \beta \widehat{D}_R)$ from Eq. (61) and $h \le \xi R/(2d^{3/2})$ from Eq. (8), it suffices that

$$\frac{\beta \widehat{D}_R}{2 + \beta \widehat{D}_R} \ge \frac{2\xi}{\sqrt{d}}. \tag{42}$$

For the choices $\xi = p_f/8$ and $\beta\bar{\varepsilon} \ge 4 \log 2$ in Theorem 13, Eq. (42) holds whenever $\widehat{D}_R > \bar{\varepsilon}$; when $\widehat{D}_R \le \bar{\varepsilon}$, the restricted optimization problem is handled directly without a quantum-walk step. Thus an admissible η_θ exists for every $\theta \in \mathcal{B}$.

Role of the lower bound $(\eta 2)$. Condition $(\eta 2)$ couples the grid scale to the Langevin step size. Together with $(\eta 3)$,

$$h \le \eta_\theta G_R \le \sqrt{\frac{\eta_\theta}{\theta}} = \frac{\sigma_\theta}{\sqrt{2}}.$$

This scale relation is used in the local transition comparison underlying the explicit spectral-gap bound in Proposition 12.

D.3 Rounded Drift and Finite-Register Proposal Implementation

Drift-rounding bounds. By the definition Eq. (17) and coordinatewise nearest-integer rounding,

$$\left\|m_\theta(x) + \eta_\theta \widehat{\nabla f}(x)\right\|_\infty \le \frac{h}{2}. \tag{43}$$

Moreover, Eq. (37) and Eq. (3) give

$$\|\widehat{\nabla f}(x)\|_\infty \le G_R + 2^{-b-1},$$

and hence

$$\frac{\|m_\theta(x)\|_\infty}{h} \le \frac{\eta_\theta(G_R + 2^{-b-1})}{h} + \frac{1}{2}.$$

Since $m_\theta(x)/h \in \mathbb{Z}^d$, its ℓ_∞ norm is an integer, which yields Eq. (18).

Lemma 16 (Finite-register cost)**.** *For each inverse temperature θ, define*

$$K_\theta := \max\left\{r_\theta + 1,\ \left\lceil c\,\frac{\sigma_\theta}{h}\sqrt{\log\big(1/\xi_{\mathrm{impl}}\big)}\right\rceil\right\}.$$

From Appendix B.3, we know a position $x \in \Omega_h$ can be stored using $O\big(d \log\left(1 + \frac{R}{h}\right)\big)$ qubits. A displacement $z \in \mathcal{D}_{h,\theta}$ can be stored using

$$d\,\lceil \log_2(2K_\theta + 1)\rceil = O\left(d \log\left(2 + \frac{\eta_\theta(G_R + 2^{-b}) + \sigma_\theta\sqrt{\log(1/\xi_{\mathrm{impl}})}}{h}\right)\right)$$

qubits. Moreover, the discrete-noise state factorizes over coordinates, and its preparation is independent of the current position. Hence its preparation requires no calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$.

Proof of Lemma 16. The position-register bound follows from Eq. (36). It remains to bound the displacement register. By Eq. (19),

$$\mathcal{D}_{h,\theta} = h\{-K_\theta, \dots, K_\theta\}^d,$$

so each coordinate of a displacement has $2K_\theta + 1$ possible values. Therefore the displacement register requires

$$d\,\lceil \log_2(2K_\theta + 1)\rceil$$

qubits. From Eq. (18),

$$r_\theta + 1 \le \frac{\eta_\theta(G_R + 2^{-b-1})}{h} + \frac{3}{2},$$

while

$$\left\lceil c\frac{\sigma_\theta}{h}\sqrt{\log(1/\xi_{\mathrm{impl}})}\right\rceil \le c\frac{\sigma_\theta}{h}\sqrt{\log(1/\xi_{\mathrm{impl}})} + 1.$$

Since c is a fixed absolute constant and $2^{-b-1} \le 2^{-b}$, it follows that

$$K_\theta = O\left(1 + \frac{\eta_\theta(G_R + 2^{-b}) + \sigma_\theta\sqrt{\log(1/\xi_{\mathrm{impl}})}}{h}\right),$$

and consequently

$$d\,\lceil \log_2(2K_\theta + 1)\rceil = O\left(d \log\left(2 + \frac{\eta_\theta(G_R + 2^{-b}) + \sigma_\theta\sqrt{\log(1/\xi_{\mathrm{impl}})}}{h}\right)\right).$$

Relative to a box of a constant number of standard deviations, the enlargement therefore costs an additive $O(d \log\log(1/\xi_{\mathrm{impl}}))$ qubits.

Finally, writing $z = hk$ with $k \in \{-K_\theta, \dots, K_\theta\}^d$, Eq. (20) gives

$$\nu_\theta(hk) = \prod_{j=1}^{d} \frac{\exp\left(-\frac{\theta h^2 k_j^2}{4\eta_\theta}\right)}{\sum_{\ell=-K_\theta}^{K_\theta} \exp\left(-\frac{\theta h^2 \ell^2}{4\eta_\theta}\right)}.$$

Thus the coherent noise state factorizes as

$$|\nu_\theta\rangle = \bigotimes_{j=1}^{d} \left(\frac{1}{\sqrt{Z_\theta^{(1)}}} \sum_{k=-K_\theta}^{K_\theta} \exp\left(-\frac{\theta h^2 k^2}{8\eta_\theta}\right) |hk\rangle \right),$$

where

$$Z_\theta^{(1)} = \sum_{k=-K_\theta}^{K_\theta} \exp\left(-\frac{\theta h^2 k^2}{4\eta_\theta}\right).$$

This state depends only on the fixed parameters θ, η_θ, h, and K_θ, and not on the current position x. Its preparation therefore requires no calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$. □

D.4 Exact Detailed Balance and Stationarity

Remark 17 (Finite proposal support and the reverse move)**.** *Although $\mathcal{D}_{h,\theta}$ is symmetric, the state-dependent drift means that a supported forward proposal need not have a supported reverse displacement: if $y = x + m_\theta(x) + \xi_\theta$, then the reverse displacement is*

$$x - y - m_\theta(y) = -\xi_\theta - m_\theta(x) - m_\theta(y),$$

which may lie outside $\mathcal{D}_{h,\theta}$. In that case $\nu_\theta(x - y - m_\theta(y)) = 0$, so the Metropolis–Hastings acceptance probability is zero and the proposal is rejected. Hence finite proposal support is handled by the Metropolis–Hastings filter through automatic rejection of moves whose reverse proposal has zero probability. Lemma 7 shows that this preserves exact reversibility with respect to $\pi^h_{\theta,R}$, although the additional rejections may reduce the spectral gap.

Proof of Lemma 7. This is the standard Metropolis–Hastings detailed-balance argument. Define

$$q_\theta(x, y) := \nu_\theta\big(y - x - m_\theta(x)\big).$$

For $x \neq y$ in Ω_h, by Eq. (21) and Eq. (22),

$$\begin{aligned}
\pi^h_{\theta,R}(x) P_\theta(x, y) &= \pi^h_{\theta,R}(x) q_\theta(x, y) \min\left\{1, \frac{\pi^h_{\theta,R}(y) q_\theta(y, x)}{\pi^h_{\theta,R}(x) q_\theta(x, y)}\right\} \\
&= \min\Big\{\pi^h_{\theta,R}(x) q_\theta(x, y),\ \pi^h_{\theta,R}(y) q_\theta(y, x)\Big\} \\
&= \pi^h_{\theta,R}(y) P_\theta(y, x).
\end{aligned}$$

The case $x = y$ is immediate. Thus P_θ satisfies detailed balance with respect to $\pi^h_{\theta,R}$, and adding laziness preserves detailed balance. Therefore $\overline{P}_\theta$ is reversible and has stationary distribution exactly $\pi^h_{\theta,R}$.

For a reversible Markov chain, the coherent encoding of its stationary distribution is the eigenvalue 1, equivalently zero-phase, state of the associated quantum walk. Hence $|\pi^h_{\theta,R}\rangle$ is exactly a zero-phase eigenstate of the Szegedy walk associated with $\overline{P}_\theta$. □

D.5 Irreducibility and Aperiodicity

Proof of Lemma 8. Step 1: single-coordinate moves are available. The argument uses only that $\mathcal{D}_{h,\theta}$ contains the box of radius $h(r_\theta + 1)$, so it is unaffected by the enlargement of the second term in Eq. (19).

Let $y = x \pm he_i$ with $x, y \in \Omega_h$. The forward displacement is

$$y - x - m_\theta(x) = h(\pm e_i - m_\theta(x)/h)$$

and the reverse one is

$$x - y - m_\theta(y) = h(\mp e_i - m_\theta(y)/h).$$

By Eq. (18) both have ℓ_∞ norm at most $h(r_\theta + 1) \leq hK_\theta$, so both lie in $\mathcal{D}_{h,\theta}$ and both receive positive ν_θ-weight. In particular such a pair never falls in the outer shell of Remark 17, so the

acceptance ratio Eq. (21) has a nonzero numerator. Since the target mass function is positive on Ω_h, the acceptance probabilities are then positive in both directions. Therefore, whenever $x, y \in \Omega_h$ differ by one grid step in a single coordinate,

$$P_\theta(x, y) > 0, \qquad P_\theta(y, x) > 0.$$

Step 2: Connectivity of Ω_h under single-coordinate moves. We now show that every $x \in \Omega_h$ can be connected to the origin by repeatedly moving one nonzero coordinate one grid step toward zero.

Recall in Section 3.2 that for $x \in \Lambda_h$, its associated cell is

$$C_x = x + [0, h)^d.$$

For $t \in h\mathbb{Z}$, define

$$\varphi(t) := \max\{t^2, (t + h)^2\},$$

and let

$$g(x) := \sum_{i=1}^{d} \varphi(x_i).$$

Since the supremum of z_i^2 over $z_i \in [x_i, x_i + h)$ is $\varphi(x_i)$, we have

$$g(x) = \sup_{z \in C_x} \|z\|^2.$$

Therefore,

$$x \in \Omega_h \quad \Longleftrightarrow \quad g(x) \le R^2.$$

Since $h \le R/\sqrt{d}$, the origin belongs to Ω_h,

$$g(0) = dh^2 \le R^2.$$

Now take any $x \in \Omega_h$ with $x \neq 0$. Choose a coordinate i with $x_i \neq 0$. Since $x_i \in h\mathbb{Z}$, either $x_i \ge h$ or $x_i \le -h$. Move this coordinate one step toward zero, i.e.,

$$x' = \begin{cases} x - he_i, & x_i > 0, \\ x + he_i, & x_i < 0. \end{cases}$$

This move cannot increase g. If $x_i \ge h$, then

$$\varphi(x_i - h) = x_i^2 \le (x_i + h)^2 = \varphi(x_i),$$

while if $x_i \le -h$, then

$$\varphi(x_i) = x_i^2,$$

and, since both $x_i + h$ and $x_i + 2h$ have absolute value at most $|x_i|$,

$$\varphi(x_i + h) = \max\{(x_i + h)^2, (x_i + 2h)^2\} \le x_i^2 = \varphi(x_i).$$

All other coordinates are unchanged, so

$$g(x') \le g(x) \le R^2.$$

Hence moving one coordinate toward zero never leaves the discrete domain, thus $x' \in \Omega_h$.

Moreover, each such move decreases the ℓ_1-distance to the origin by exactly h:

$$\|x'\|_1 = \|x\|_1 - h.$$

Thus, after finitely many single-coordinate moves, we reach 0, while remaining inside Ω_h throughout. Reversing the path connects 0 to any $x \in \Omega_h$. Therefore, Ω_h is connected under single-coordinate moves.

Combining the two steps, every pair of states in Ω_h is connected by a path of transitions with positive probability. Hence $\overline{P}_\theta$ is irreducible. Moreover, by definition

$$\overline{P}_\theta(x, x) \ge \frac{1}{2} \qquad \text{for every } x \in \Omega_h$$

, thus $\overline{P}_\theta$ is aperiodic.

By Lemma 7, $\overline{P}_\theta$ is also reversible with stationary distribution $\pi^h_{\theta,R}$. Since Ω_h is finite and $\overline{P}_\theta$ is irreducible and aperiodic, the eigenvalue 1 is simple and every other eigenvalue has modulus strictly smaller than 1. Therefore its spectral gap δ_θ is strictly positive.

Finally, by the spectral mapping relation Eq. (24) for the Szegedy walk of a reversible chain,

$$\Delta(\theta;\eta_\theta) \geq 2\sqrt{\delta_\theta} > 0$$

Wocjan and Abeyesinghe (2008). Hence the corresponding Szegedy phase gap is strictly positive. □

D.6 Coherent Implementation of the Transition

Proof of Proposition 9. We follow the coherent MALA construction of Childs et al. (2022, Algorithm 10 and Lemma C.6): prepare the proposal, compute the Metropolis acceptance amplitude coherently, and then uncompute the work registers. The only difference here is that the proposal uses the discrete displacement state and is restricted to Ω_h.

The displacement state

$$|\nu_\theta\rangle = \sum_{z\in\mathcal{D}_{h,\theta}} \sqrt{\nu_\theta(z)}\,|z/h\rangle$$

is independent of the position register and therefore requires no oracle query. Since $\mathcal{D}_{h,\theta}$ is the product set in Eq. (19) and the weight $\exp(-\theta\|z\|^2/(4\eta_\theta))$ in Eq. (20) factorizes over coordinates, $|\nu_\theta\rangle$ is a tensor product of d one-dimensional states supported on $\{-K_\theta,\dots,K_\theta\}$, with K_θ as in Lemma 16. Each such state has an explicit finite cumulative distribution and can be prepared sequentially with gate cost polynomial in $\log K_\theta$ and the requested accuracy. Thus this preparation is fixed and state-independent; in particular, it does not require preparing a position-dependent or d-dimensional Gibbs state.

Starting from $|x\rangle|k\rangle$, one query to $\mathcal{O}_{\nabla f}$ gives $\nabla f(x)$. Reversible fixed-point arithmetic then computes $m_\theta(x)$ from Eq. (17) and

$$y = x + m_\theta(x) + hk$$

in a second position register. Membership of y in Ω_h is checked without an oracle query by comparing $\sum_i \max\{y_i^2,(y_i+h)^2\}$ with R^2; by the corner characterization in the proof of Lemma 8, this is equivalent to $C_y \subseteq \Omega$.

A second query to $\mathcal{O}_{\nabla f}$ at y gives $m_\theta(y)$, which determines the reverse displacement $x-y-m_\theta(y)$ and whether it belongs to $\mathcal{D}_{h,\theta}$. Two queries to $\mathcal{O}_f$ give $f(x)$ and $f(y)$. When the reverse displacement lies in $\mathcal{D}_{h,\theta}$, the logarithm of the Metropolis ratio in Eq. (21) is

$$-\theta\big(f(y)-f(x)\big) + \frac{\theta}{4\eta_\theta}\left(\|y-x-m_\theta(x)\|^2 - \|x-y-m_\theta(y)\|^2\right). \tag{44}$$

If the reverse displacement is outside $\mathcal{D}_{h,\theta}$, its proposal probability is zero and the move is rejected; this is the outer-shell case of Remark 17, and the comparison test costs no oracle query. Otherwise, a controlled rotation implements the acceptance amplitude $\sqrt{\alpha_\theta(x,y)}$. Rejected moves, together with branches excluded by the admissibility tests, are routed to the self-loop amplitude.

Finally, all function, gradient, and arithmetic work registers are uncomputed, using the same oracle calls in reverse. Hence the forward computation uses two queries to $\mathcal{O}_{\nabla f}$ and two queries to $\mathcal{O}_f$, and uncomputation uses the same number. The laziness step requires only an additional ancilla and no further oracle query. Thus one application of the transition-state unitary, or its inverse, uses $O(1)$ objective and gradient oracle queries.

Here U_θ denotes the exact transition-state unitary of the discrete kernel already defined using the rounded gradient $\widehat{\nabla f}$ of Proposition 14. The remaining finite-precision errors arise only from the coherent implementation of this fixed discrete kernel.

For a deterministic call budget T, choose the oracle precision and the reversible-arithmetic, discrete-Gaussian-preparation, and gate-synthesis precisions so that every controlled transition-state application satisfies

$$\|\widehat{U}_t - U_t\| \leq \frac{\xi_{\mathrm{impl}}}{T}.$$

Increasing these working precisions changes only gate and workspace costs and does not increase the number of calls to $\mathcal{O}_f$ or $\mathcal{O}_{\nabla f}$.

Therefore, for any sequence of $s \le T$ controlled transition-state applications, the standard hybrid argument gives

$$\begin{aligned}\left\|\widehat{U}_s \cdots \widehat{U}_1 - U_s \cdots U_1\right\| &\le \sum_{t=1}^{s} \|\widehat{U}_t - U_t\| \\ &\le s\,\frac{\xi_{\mathrm{impl}}}{T} \le \xi_{\mathrm{impl}}.\end{aligned}$$

This proves both the claimed implementation accuracy and the $O(1)$ oracle-query cost per transition-state application. □

E Phase-Gap Analysis

E.1 Double-Well Obstruction and the Local-Proposal Regime

Construction. Choose once and for all an even function $\phi \in C_c^\infty((-1/2, 1/2))$ such that

$$0 \le \phi \le 1, \qquad \phi(0) = 1, \qquad \phi(t) \ge \frac{1}{2} \quad \text{for } |t| \le \frac{1}{4}.$$

Let $c_\psi := \|\phi''\|_\infty$ and define the rescaled bump

$$\psi_R(t) := \phi(t/R).$$

Then ψ_R is even, vanishes for $|t| \ge R/2$, satisfies $\psi_R \ge 1/2$ on $[-R/4, R/4]$, and obeys $\|\psi_R''\|_\infty \le c_\psi/R^2$. For $H > 0$, set

$$f(x) := H\psi_R(x_1). \tag{45}$$

Thus the two slabs $\{x_1 \le -R/2\}$ and $\{x_1 \ge R/2\}$ have zero energy, while every point with $|x_1| \le R/4$ has energy at least $H/2$.

We use the following local-proposal condition:

$$2hK_\theta \le \frac{R}{4}. \tag{46}$$

Proof of Proposition 11. Since

$$\nabla^2 f(x) = \mathrm{diag}\big(H\psi_R''(x_1), 0, \ldots, 0\big),$$

f is globally L-smooth with

$$L = H\|\psi_R''\|_\infty \le \frac{c_\psi H}{R^2}.$$

Evenness gives $\psi_R'(0) = 0$, hence $G_0 = 0$. Moreover $0 \le f \le H$, $f(0) = H$, and $f(x) = 0$ whenever $|x_1| \ge R/2$. Therefore

$$f^\star = f_R^\star = 0, \qquad \varepsilon_{\mathrm{loc}} = 0, \qquad D_R = H.$$

Thus the construction satisfies Assumptions 1–2. With the admissible smoothness constant above, $G_R = LR$ and

$$D_R \le \widehat{D}_R = 2LR^2 \le 2c_\psi D_R.$$

Let

$$S := \{x \in \Omega_h : x_1 < 0\}, \qquad g := \mathbf{1}_S,$$

and write

$$Q_\theta(A, B) := \sum_{x \in A} \sum_{y \in B} \pi_{\theta,R}^h(x)\overline{P}_\theta(x, y).$$

For a finite reversible chain, the Rayleigh–Ritz characterization of the spectral gap gives

$$\delta_\theta \le \frac{\mathcal{E}_\theta(g, g)}{\mathrm{Var}_{\pi_{\theta,R}^h}(g)} = \frac{Q_\theta(S, S^c)}{\pi_{\theta,R}^h(S)\pi_{\theta,R}^h(S^c)};$$

see, for example, Diaconis and Stroock (1991). In particular,

$$\delta_\theta \le \frac{Q_\theta(S, S^c)}{\pi^h_{\theta,R}(S)\pi^h_{\theta,R}(S^c)}. \tag{47}$$

Step 1: a local proposal cannot cross the barrier without starting in its central slab. Any accepted off-diagonal transition has

$$y - x = m_\theta(x) + \xi_\theta, \qquad \xi_\theta \in \mathcal{D}_{h,\theta}.$$

By Eq. (19) and Eq. (18),

$$\|y - x\|_\infty \le h(K_\theta + r_\theta) \le 2hK_\theta \le \frac{R}{4},$$

because $r_\theta \le K_\theta - 1$. Hence, if $x \in S$ and $y \in S^c$, then $x_1 < 0 \le y_1$ implies $-R/4 \le x_1 < 0$. Define

$$U := \{x \in \Omega_h : -R/4 \le x_1 < 0\}.$$

Since $f \ge H/2$ on U and the total transition probability out of any x is at most one,

$$Q_\theta(S, S^c) \le \pi^h_{\theta,R}(U). \tag{48}$$

Step 2: both sides of the barrier carry non-negligible stationary mass. Let

$$W_\pm := \{x \in \Omega_h : \pm x_1 \ge R/2\}.$$

On $W_- \cup W_+$ we have $f = 0$. Since $0 \le f \le H$,

$$\pi^h_{\theta,R}(U) \le \frac{|\Omega_h|}{|W_-|} e^{-\theta H/2},$$

whereas

$$\pi^h_{\theta,R}(S) \ge \frac{|W_-|}{|\Omega_h|}, \qquad \pi^h_{\theta,R}(S^c) \ge \frac{|W_+|}{|\Omega_h|}.$$

Combining these estimates with Eq. (47) and Eq. (48) gives

$$\delta_\theta \le \left(\frac{|\Omega_h|}{\min\{|W_-|, |W_+|\}}\right)^3 e^{-\theta H/2}. \tag{49}$$

It remains to lower-bound the two lattice counts. Assume $h \le R/(32\sqrt{d})$ and set $c_\pm := \pm(3R/4)e_1$. If $y \in B(c_+, R/16)$ and $x(y)$ is the corner of the grid cell containing y, then

$$\|x(y) - c_+\| \le \|x(y) - y\| + \|y - c_+\| \le \sqrt{d}\, h + R/16 \le 3R/32.$$

Consequently $x(y)_1 > R/2$. Moreover every $z \in C_{x(y)}$ satisfies

$$\|z - c_+\| \le \|z - x(y)\| + \|x(y) - c_+\| \le \sqrt{d}\, h + 3R/32 \le R/8,$$

so $\|z\| \le 3R/4 + R/8 = 7R/8 < R$. Hence $x(y) \in W_+$ and $C_{x(y)} \subseteq \Omega$. Because the cells are disjoint and have volume h^d,

$$|W_+|h^d \ge \mathrm{Vol}(B(0, R/16)).$$

The same argument with c_- gives the identical bound for W_-. On the other hand $|\Omega_h|h^d = \mathrm{Vol}(\widehat{\Omega}_h) \le \mathrm{Vol}(B(0, R))$. Therefore

$$\frac{|W_\pm|}{|\Omega_h|} \ge 16^{-d}. \tag{50}$$

Substituting Eq. (50) into Eq. (49) yields

$$\delta_\theta \le 16^{3d} e^{-\theta H/2} = e^{C_2 d} e^{-\theta D_R/2}, \qquad C_2 := 3\log 16.$$

Using the Szegedy spectral mapping Eq. (24), which is the standard relation used in quantum sampling from reversible chains (Wocjan and Abeyesinghe, 2008),

$$\Delta(\theta; \eta_\theta) \le \pi\sqrt{\delta_\theta} \le \pi e^{C_2 d/2} e^{-\theta D_R/4}.$$

This proves Eq. (28) with $C = C_2$.

Step 3: specialization to the optimization temperature. The endpoint $\theta = \beta$ belongs to both the explicit path and the adaptive search grid, so

$$\Delta_{\mathrm{path}}, \Delta_{\mathrm{all}} \le \Delta(\beta; \eta_\beta).$$

Let $r := D_R/\bar{\varepsilon}$. Since $\widehat{D}_R \ge D_R$, the temperature choice Eq. (30) implies

$$\beta\bar{\varepsilon} \ge 4d \log_+(4r),$$

and hence

$$\frac{\beta D_R}{4} \ge dr \log_+(4r).$$

For $r \ge 3$,

$$r\left(\log_+(4r) - \frac{1}{4}\right) \ge \frac{C_2}{2},$$

so

$$\frac{\beta D_R}{4} - \frac{C_2 d}{2} \ge \frac{dD_R}{4\bar{\varepsilon}}.$$

Therefore, whenever Eq. (46) holds at $\theta = \beta$,

$$\Delta_{\mathrm{path}}, \Delta_{\mathrm{all}} \le \pi \exp\left(-\frac{dD_R}{4\bar{\varepsilon}}\right),$$

which proves the endpoint bound in Proposition 11. □

Why the local-proposal condition is needed, and why it is non-vacuous. Condition Eq. (46) is a condition on the prescribed MALA proposal, not an additional structural assumption on f. It is essential for this particular bottleneck argument: if the proposal support is wide enough to jump directly from W_- to W_+, then both endpoint energies and gradients are zero, so such a supported jump is not penalized by the central barrier and the estimate Eq. (48) need not hold.

The condition is nevertheless compatible with the standing step-size window. For any fixed $\theta > 0$ in the construction, take the smallest admissible drift scale

$$\eta_\theta := \frac{h}{G_R}.$$

For sufficiently small h, this choice satisfies $(\eta 1)$ and $(\eta 3)$, while $(\eta 2)$ holds with equality. Moreover

$$\sigma_\theta = \sqrt{\frac{2h}{\theta G_R}},$$

and, using $\sqrt{d}\, 2^{-b-1} \le G_R$ and $d \ge 2$,

$$r_\theta \le \left\lfloor 1 + \frac{1}{\sqrt{d}} + \frac{1}{2} \right\rfloor \le 2.$$

Hence

$$2hK_\theta \le 2\max\left\{3h, c\sqrt{\frac{2h}{\theta G_R}}\sqrt{\log(1/\xi_{\mathrm{impl}})} + h\right\} \xrightarrow[h\downarrow 0]{} 0.$$

Thus Eq. (46) holds for all sufficiently fine discretizations. In particular, the exponentially small-gap family in the proposition is not created by an inconsistent choice of step size.

E.2 Explicit Spectral-Gap Lower Bound

Let $\overline{P}_h^{\mathrm{nn}}$ denote the lazy nearest-neighbour walk on Ω_h: with probability $1/2$ it stays at the current state, and otherwise it chooses a coordinate $i \in \{1, \dots, d\}$ and a sign uniformly and attempts the move $x \mapsto x \pm he_i$, staying put if the proposed point does not lie in Ω_h. The chain is symmetric and therefore reversible with respect to the uniform distribution on Ω_h. By Step 2 of the proof of Lemma 8, Ω_h is connected under these single-coordinate moves, so the chain is irreducible. Denote its spectral gap by δ_h^{nn}.

Lemma 18. *Let $\theta \in \mathcal{B}$, let $\{\eta_\theta\}$ satisfy Eq. (16), let h satisfy Eq. (8), and choose b as in Proposition 14. Then*

$$\delta_\theta \;\geq\; e^{-9}\left(\frac{h}{4\sigma_\theta}\right)^d e^{-2\theta \widehat{D}_R}\delta_h^{\mathrm{nn}}. \tag{51}$$

Proof of Lemma 18. Write $\pi := \pi_{\theta,R}^h$ and let $\widetilde{\pi}$ be uniform on Ω_h. Since every nearest-neighbour edge of $\overline{P}_h^{\mathrm{nn}}$ is also an edge of $\overline{P}_\theta$, the Dirichlet forms satisfy

$$\mathcal{E}_{\overline{P}_h^{\mathrm{nn}}}(g,g) \leq A\,\mathcal{E}_{\overline{P}_\theta}(g,g), \qquad A := \max_{y=x\pm he_i} \frac{\widetilde{\pi}(x)\overline{P}_h^{\mathrm{nn}}(x,y)}{\pi(x)\overline{P}_\theta(x,y)}.$$

Moreover, since the oscillation of f on Ω is at most D_R,

$$e^{-\theta D_R} \leq \frac{\pi(x)}{\widetilde{\pi}(x)} \leq e^{\theta D_R},$$

and hence $\mathrm{Var}_\pi(g) \leq e^{\theta D_R}\mathrm{Var}_{\widetilde{\pi}}(g)$. By the variational characterization of the spectral gap,

$$\delta_\theta \geq A^{-1}e^{-\theta D_R}\delta_h^{\mathrm{nn}}. \tag{52}$$

It remains to lower-bound the transition probability across a nearest-neighbour edge $y = x \pm he_i$. Set

$$z_f := y - x - m_\theta(x), \qquad z_r := x - y - m_\theta(y).$$

By Step 1 of Lemma 8, both displacements belong to $\mathcal{D}_{h,\theta}$. From Eq. (43), Eq. (39), and $(\eta 3)$,

$$\|m_\theta(x)\| \leq 2\eta_\theta G_R + \frac{\sqrt{d}\,h}{2} \leq \sqrt{2}\,\sigma_\theta + \frac{\sqrt{d}\,h}{2}.$$

Also, $(\eta 2)$–$(\eta 3)$ give $h \leq \sigma_\theta/\sqrt{2}$. Therefore

$$\frac{\|z_f\|^2}{2\sigma_\theta^2}, \frac{\|z_r\|^2}{2\sigma_\theta^2} \leq E_d, \qquad E_d := \frac{1}{4}\left(3 + \frac{\sqrt{d}}{2}\right)^2.$$

Let $t := h/\sigma_\theta \leq 1/\sqrt{2}$. Bounding the normalizing sum in Eq. (20) by the full lattice sum gives

$$\begin{aligned}\sum_{u\in\mathcal{D}_{h,\theta}} e^{-\|u\|^2/(2\sigma_\theta^2)} &\leq \left(1 + 2\sum_{k\geq 1} e^{-t^2k^2/2}\right)^d \\ &\leq \left(1 + \frac{\sqrt{2\pi}}{t}\right)^d \leq \left(\frac{13\sigma_\theta}{4h}\right)^d.\end{aligned}$$

Consequently,

$$\nu_\theta(z_f) \geq e^{-E_d}\left(\frac{4h}{13\sigma_\theta}\right)^d, \qquad \frac{\nu_\theta(z_r)}{\nu_\theta(z_f)} \geq e^{-E_d}.$$

By Lemma 2, $|f(y)-f(x)| \leq G_R h$, while $(\eta 2)$–$(\eta 3)$ imply $h \leq 1/(\theta G_R)$. Hence $\theta|f(y)-f(x)| \leq 1$, and

$$\alpha_\theta(x,y) \geq e^{-1-E_d}.$$

Since $\overline{P}_h^{\mathrm{nn}}(x,y) = 1/(4d)$ and $\overline{P}_\theta(x,y) = \frac{1}{2}\nu_\theta(z_f)\alpha_\theta(x,y)$,

$$A^{-1} \geq 2d\,e^{-\theta D_R}e^{-1-2E_d}\left(\frac{4h}{13\sigma_\theta}\right)^d \geq e^{-9}e^{-\theta D_R}\left(\frac{h}{4\sigma_\theta}\right)^d.$$

The last inequality holds for every $d \geq 1$. Substituting this into Eq. (52) and using $D_R \leq \widehat{D}_R$ yields

$$\delta_\theta \geq e^{-9}\left(\frac{h}{4\sigma_\theta}\right)^d e^{-2\theta\widehat{D}_R}\delta_h^{\mathrm{nn}},$$

as claimed. □

Lemma 19 (Free discrete walk on the ball)**.** *If $h \le R/(4\sqrt{d})$, then*

$$\delta_h^{\mathrm{nn}} \ge c_0^d \frac{h^2}{d^2 R^2}, \tag{53}$$

where one may take $c_0 = 10^{-3}$.

Proof of Lemma 19. Set $\rho := R/h \ge 4\sqrt{d}$ and identify $x \in \Lambda_h$ with

$$u_j := x_j/h + \tfrac{1}{2} \in \mathbb{Z} + \tfrac{1}{2}.$$

Then $x \in \Omega_h$ exactly when

$$\Phi(u) := \left(\sum_{j=1}^{d} (|u_j| + \tfrac{1}{2})^2 \right)^{1/2} \le \rho.$$

Write

$$\Omega := \{u : \Phi(u) \le \rho\}, \qquad N := |\Omega|.$$

In these coordinates, $\overline{P}_h^{\mathrm{nn}}$ moves by $\pm e_j$ with probability $1/(4d)$ and has the uniform stationary distribution. Hence the canonical-path bound of Diaconis and Stroock (1991) gives

$$\delta_h^{\mathrm{nn}} \ge \kappa^{-1}, \qquad \kappa \le \frac{4d}{N} \Lambda \Gamma, \tag{54}$$

where Λ is the maximal path length and Γ the maximal number of ordered pairs routed through one edge.

For $u \in \Omega$, define the contraction

$$\operatorname{sgn}(Tu)_j = \operatorname{sgn}(u_j), \qquad |(Tu)_j| := \left\lfloor \frac{|u_j| - \frac{1}{2}}{4} \right\rfloor + \tfrac{1}{2}.$$

Coordinatewise,

$$|(Tu)_j| + \tfrac{1}{2} \le \frac{1}{4}(|u_j| + \tfrac{1}{2}) + 1,$$

so

$$\Phi(Tu) \le \frac{1}{4}\Phi(u) + \sqrt{d}. \tag{55}$$

Thus, for $u, v \in \Omega$,

$$\Phi(Tu) + \Phi(Tv) \le \frac{\rho}{2} + 2\sqrt{d} \le \rho.$$

Route (u, v) by three coordinatewise monotone legs:

$$u \longrightarrow Tu \longrightarrow Tv \longrightarrow v.$$

The first and third legs remain in Ω because no coordinate magnitude increases. Along the middle leg,

$$|w_j| + \tfrac{1}{2} \le (|(Tu)_j| + \tfrac{1}{2}) + (|(Tv)_j| + \tfrac{1}{2}),$$

and therefore $\Phi(w) \le \Phi(Tu) + \Phi(Tv) \le \rho$. Hence every assigned path stays inside Ω.

Furthermore,

$$\sum_j (|u_j| + \tfrac{1}{2}) \le \sqrt{d}\, \Phi(u) \le \sqrt{d}\, \rho,$$

and

$$|u_j| - |(Tu)_j| \le \frac{3}{4}(|u_j| + \tfrac{1}{2}) + \frac{3}{4}.$$

Since $d \le \sqrt{d}\, \rho/4$, each contraction leg has length at most $\sqrt{d}\, \rho$, while

$$\sum_j \big(|(Tu)_j| + |(Tv)_j|\big) \le \sqrt{d}\, \big(\Phi(Tu) + \Phi(Tv)\big) \le \sqrt{d}\, \rho.$$

Thus

$$\Lambda \le 3\sqrt{d}\,\rho. \tag{56}$$

It remains to bound the congestion. Fix an edge $e = (z, z + e_i)$. The map T is at most 4-to-1 in each coordinate. If e is used on the first or third leg, the edge fixes all but at most $i-1$ preimages of T, one coordinate of u, and the unrestricted endpoint; hence these two legs contribute at most

$$2 \cdot 4^d (2\rho + 1)N$$

ordered pairs.

For the middle leg, let

$$P_k := \#\left\{ w \in (\mathbb{Z} + \tfrac{1}{2})^k : \sum_{j=1}^{k} (|w_j| + \tfrac{1}{2})^2 \le \rho^2 \right\}.$$

The edge determines $(Tv)_{<i}$ and $(Tu)_{>i}$, so the number of pairs using e on this leg is at most $4^{d-1} P_i P_{d-i+1}$. Therefore

$$\Gamma \le 4^{d-1} \max_{1 \le i \le d} P_i P_{d-i+1} + 2 \cdot 4^d (2\rho + 1)N. \tag{57}$$

A volume comparison gives

$$P_k \le \omega_k \left(\frac{9}{8}\rho\right)^k, \qquad N \ge \omega_d \left(\frac{5}{8}\rho\right)^d.$$

Moreover,

$$\omega_i \omega_{d-i+1} \le 2^{(d+3)/2} \omega_d,$$

and hence

$$P_i P_{d-i+1} \le 4(2.6)^d \rho N.$$

Substitution into Eq. (57) yields

$$\Gamma \le 7 \cdot 11^d \rho N.$$

Together with Eq. (54) and Eq. (56),

$$\kappa \le 84\, d^{3/2} 11^d \rho^2 \le d^2\, 10^{3d} \rho^2.$$

Since $\rho = R/h$,

$$\delta_h^{\mathrm{nn}} \ge 10^{-3d} \frac{h^2}{d^2 R^2},$$

which proves Eq. (53) with $c_0 = 10^{-3}$. □

Proof of Proposition 12. Lemma 18 lower-bounds the spectral gap of the discrete MALA chain in terms of that of the lazy nearest-neighbour walk, while Lemma 19 gives an explicit lower bound for the latter. Substituting the second bound into the first yields Eq. (29).

By $(\eta 1)$ in Eq. (16),

$$\sigma_\theta^2 = \frac{2\eta_\theta}{\theta} \le \frac{R^2}{d},$$

and hence

$$\left(\frac{h}{4\sigma_\theta}\right)^d \ge \left(\frac{h\sqrt{d}}{4R}\right)^d.$$

Moreover, every queried temperature satisfies $\theta \le \beta$, so

$$e^{-2\theta \widehat{D}_R} \ge e^{-2\beta \widehat{D}_R}.$$

Substituting these inequalities into Eq. (29) proves the claim. □

The factor $(h/(4\sigma_\theta))^d$ in Eq. (29) arises from comparison with the lazy nearest-neighbour walk rather than from an additional approximation error. Indeed, $(\eta 2)$ gives

$$\sigma_\theta^2 = \frac{2\eta_\theta}{\theta} \geq \frac{2h}{\theta G_R},$$

and therefore

$$\left(\frac{h}{4\sigma_\theta}\right)^d \leq \left(\frac{\theta G_R h}{32}\right)^{d/2} \leq \left(\frac{\xi}{192\sqrt{d}}\right)^{d/2},$$

where the last inequality uses $\theta \leq \beta$ and Eq. (8). Thus the certified lower bound may be conservative in high dimension; this comparison loss need not be intrinsic to the true spectral gap.

F Quantum Annealing and Localized Q-MALA Sampling

F.1 Queried inverse temperatures and adaptive search grid

Each annealing route queries only finitely many positive inverse temperatures.

For the explicit route of Eq. (11), define

$$\mathcal{B} := \{\beta_j : 1 \leq j \leq M\}.$$

For the adaptive route, the schedule search is restricted to the normalized grid

$$\mathcal{G}_\beta := \left\{\frac{k}{K_\beta} : k = 1, \ldots, K_\beta\right\}, \qquad K_\beta := \left\lceil 1 + \beta \widehat{D}_R \right\rceil, \qquad \delta_g := K_\beta^{-1}. \tag{58}$$

A grid point $t \in \mathcal{G}_\beta$ corresponds to inverse temperature $\theta = t\beta$, and for the adaptive route we therefore define

$$\mathcal{B} := \{t\beta : t \in \mathcal{G}_\beta\}.$$

The choice of K_β ensures $\beta \widehat{D}_R \delta_g \leq 1$, which is sufficient for the adjacent-grid constant-overlap guarantee used in the adaptive search.

Overlap structure of the adaptive grid. Recall from Eq. (12) that

$$\Pi_t(x) \propto e^{-tH_\beta(x)}, \qquad 0 \leq H_\beta(x) \leq \beta D_R \leq \beta \widehat{D}_R.$$

For $0 \leq t \leq s \leq 1$, let

$$p(t, s) := |\langle \Pi_t \mid \Pi_s \rangle|^2.$$

Differentiating gives

$$\begin{aligned} \partial_s \log p(t, s) &= \mathbb{E}_{\Pi_s}[H_\beta] - \mathbb{E}_{\Pi_{(t+s)/2}}[H_\beta] \leq 0, \\ \partial_t \log p(t, s) &= \mathbb{E}_{\Pi_t}[H_\beta] - \mathbb{E}_{\Pi_{(t+s)/2}}[H_\beta] \geq 0, \end{aligned} \tag{59}$$

because

$$\frac{d}{du}\mathbb{E}_{\Pi_u}[H_\beta] = -\operatorname{Var}_{\Pi_u}(H_\beta) \leq 0.$$

Hence $p(t, s)$ decreases as the endpoints separate. Moreover,

$$|\partial_t \log p(t, s)|,\ |\partial_s \log p(t, s)| \leq \beta \widehat{D}_R.$$

Therefore, whenever $t + \delta_g \leq 1$, $\beta \widehat{D}_R \delta_g \leq 1$ implies

$$p(t, t + \delta_g) \geq e^{-1}. \tag{60}$$

Thus every nonterminal grid point has a constant-overlap successor, while Eq. (59) provides the monotonicity used in the adaptive search analysis.

In either route, define

$$\theta_{\min} := \min \mathcal{B} = \begin{cases} \beta/M, & \text{for the explicit route,} \\ \beta/K_\beta, & \text{for the adaptive route.} \end{cases} \tag{61}$$

Since

$$M = \left\lceil \beta \widehat{D}_R \right\rceil \le 1 + \beta \widehat{D}_R$$

for the explicit route and

$$K_\beta = \left\lceil 1 + \beta \widehat{D}_R \right\rceil \le 2 + \beta \widehat{D}_R$$

for the adaptive route, both routes satisfy the deterministic bound

$$\theta_{\min} \ge \frac{\beta}{2 + \beta \widehat{D}_R}. \tag{62}$$

The endpoint $\theta = 0$ belongs to neither queried set. The corresponding uniform state is prepared directly and its reflection is implemented from the preparation unitary of Lemma 15, so no Markov chain is required at zero inverse temperature.

F.2 Adaptive Localized Q-MALA Procedure

Following the adaptive quantum annealing construction of Harrow and Wei (2020), we locate the next inverse temperature by binary search using nondestructive overlap estimation.

Algorithm 1 Adaptive localized Q-MALA

Require: $\mathcal{O}_f, \mathcal{O}_{\nabla f}, L, G_0, R$, target β, sampling accuracy ζ, failure probability χ, grid spacing $h \le R/\sqrt{d}$, step-size schedule $\{\eta_\theta\}_{\theta \in \mathcal{B}}$ satisfying Eq. (16), explicit spectral gap bound $\underline{\delta}$ satisfying Eq. (26), displacement sets $\mathcal{D}_{h,\theta}$ as in Eq. (19)

1: Set the grid $\mathcal{G}_\beta$ and spacing $\delta_g = K_\beta^{-1}$ as in Eq. (58); compute the a priori schedule-length bound $\bar{\ell}$ of Eq. (65) and the precisions $\bar{N}_{\text{test}}, \bar{N}_{\text{rot}}, \eta_{\text{test}}, \epsilon_\star$ of Eq. (66)–Eq. (67).
2: Prepare the uniform discrete state $|\Pi_0\rangle$ on Ω_h using Lemma 15 with $\varepsilon_{\text{st}} \le \zeta/4$; set $t \leftarrow 0$.
3: **while** $t < 1$ **do**
4: **Probe.** For a candidate grid point $t' > t$, estimate $p(t,t') = |\langle \Pi_t | \Pi_{t'} \rangle|^2$ by nondestructive amplitude estimation to additive accuracy ε_0, median-amplified to failure probability η_{test}, and *accept* iff the estimate is at least τ, with ε_0 and τ the absolute constants of Eq. (64). Each probe uses approximate reflections about $|\Pi_t\rangle|0\rangle$ and $|\Pi_{t'}\rangle|0\rangle$, implemented to operator-norm error $\epsilon_\star$ by Eq. (63).
5: **Binary search.** Starting from the accepted probe $t' = t + \delta_g$ and the rejected upper end (or $t' = 1$ if accepted), halve the index interval between an accepted lower probe and a rejected upper probe until they are adjacent; let t' be the final accepted grid point. No monotonicity of the estimated predicate is assumed.
6: **Advance.** Use the slowly-varying state-conversion procedure of Section 5, whose selective phase rotations are built from the same approximate reflections Eq. (63), to map the current state to an approximation of $|\Pi_{t'}\rangle$, and set $t \leftarrow t'$.
7: **Abort** and declare failure if the total number of probes ever exceeds $\bar{N}_{\text{test}}$.
8: **end while**
9: Measure the position register and return the resulting point.

F.3 Proof of the Localized Q-MALA Sampling Theorem

Proof of Theorem 10. Write

$$|\Pi_t\rangle := |\pi^h_{t\beta,R}\rangle, \qquad 0 \le t \le 1.$$

For every $t > 0$, Lemmas 7 and 8 give a finite reversible MALA chain whose Szegedy walk has $|\Pi_t\rangle$ as its stationary qsample. By the standard phase-detection construction of Wocjan and Abeyesinghe (2008), a selective phase rotation about $|\Pi_t\rangle$ can therefore be implemented to operator-norm error $\epsilon_\star$ using

$$O\left(\frac{1}{2\sqrt{\underline{\delta}}} \log \frac{1}{\epsilon_\star}\right) \tag{63}$$

controlled walk applications. The circuit must fix its phase-estimation resolution in advance, and $2\sqrt{\underline{\delta}}$ is an admissible choice because $\Delta(t\beta; \eta_{t\beta}) \ge 2\sqrt{\delta_{t\beta}} \ge 2\sqrt{\underline{\delta}}$ by Eq. (24) and Eq. (26).

The rotation at $t = 0$ is implemented directly from the preparation unitary of Lemma 15. Each controlled Szegedy-walk application uses only $O(1)$ controlled applications of the corresponding transition-state unitary and its inverse. Thus the controlled-walk bounds below are also controlled transition-state call bounds up to an absolute constant factor.

We repeatedly use the slowly-varying-chain state-conversion theorem of Wocjan and Abeyesinghe (2008): a path of m states with constant consecutive squared overlap can be traversed to accuracy $O(\zeta)$ using $O(m \log(m/\zeta))$ selective phase rotations. If each rotation is implemented to error $\epsilon_\star$, a standard hybrid argument adds at most the number of rotations times $\epsilon_\star$ to the state error.

Explicit schedule. For the schedule Eq. (11), Lemma 5 gives

$$|\langle \pi_{j-1}^h | \pi_j^h \rangle|^2 \geq e^{-1}.$$

Hence M stages require $O(M \log(M/\zeta))$ selective rotations. Taking $\epsilon_\star = \Theta(\zeta/[M \log(M/\zeta)])$ and charging each rotation at Eq. (63) gives

$$T_{\exp} = O\left(\frac{M}{\sqrt{\underline{\delta}}} \log^2 \frac{M}{\zeta}\right) = O\left(\frac{M_{\exp}}{\sqrt{\underline{\delta}}} \log^2 \frac{M_{\exp}}{\zeta}\right),$$

because $M \leq M_{\exp}$.

Since $M_{\exp}$, $\underline{\delta}$, and ζ are fixed before the run, the preceding bound is a deterministic a priori bound on the number of controlled walk applications. By the constant-overhead relation above, it also gives, up to an absolute constant factor, a deterministic a priori bound on the total number of controlled transition-state applications. We use this bound as T in Proposition 9.

Together with the initial-state and implementation budgets in the theorem statement, choosing the absolute constants in the error allocation sufficiently small makes the final state ζ-close to $|\Pi_1\rangle$.

Adaptive schedule. Let

$$p(t, s) := |\langle \Pi_t | \Pi_s \rangle|^2.$$

Lemma 6 gives an existential schedule $0 = t_0^\star < \cdots < t_\ell^\star = 1$ satisfying

$$p(t_i^\star, t_{i+1}^\star) \geq e^{-2}, \qquad \ell = O(M_{\mathrm{ad}}).$$

We show that Algorithm 1 constructs a grid schedule of the same order of length.

Let $\delta_g = K_\beta^{-1}$. From Eq. (59) and $0 \leq H_\beta \leq \beta \widehat{D}_R$,

$$|\log p(t', s') - \log p(t, s)| \leq \beta \widehat{D}_R \big(|t - t'| + |s - s'|\big).$$

Since $\beta \widehat{D}_R \, \delta_g \leq 1$, changing an endpoint by at most k grid spacings changes p by at most a factor e^k. In particular, Eq. (60) gives $p(t, t + \delta_g) \geq e^{-1}$.

For each overlap probe, use additive accuracy

$$|\widehat{p} - p(t, t')| \leq \varepsilon_0 := \tfrac{1}{4} e^{-4}, \qquad \text{accept iff} \qquad \widehat{p} \geq \tau := \tfrac{3}{4} e^{-4}. \tag{64}$$

On the event that all probes succeed, every pair with $p(t, t') \geq e^{-4}$ is accepted, while every accepted pair has

$$p(t, t') \geq \tfrac{1}{2} e^{-4} =: c_0.$$

Thus the search never stalls because the next grid point always has overlap at least e^{-1}.

It remains only to compare the constructed schedule with the existential one. For each $t_i^\star$, let u_i be the largest grid point not exceeding $t_i^\star$ and put $v_i = \max\{u_i - \delta_g, 0\}$. Then $t_i^\star - v_i \leq 2\delta_g$, and hence, using the preceding Lipschitz estimate and the monotonicity in Eq. (59),

$$p(v_i, u_{i+1}) \geq e^{-2} p(t_i^\star, t_{i+1}^\star) \geq e^{-4}.$$

Suppose inductively that the current constructed point satisfies $t^{(i)} \geq v_i$. If $t^{(i)} < u_{i+1}$, then monotonicity gives $p(t^{(i)}, u_{i+1}) \geq e^{-4}$, so u_{i+1} must be accepted. The binary search returns an accepted grid point whose successor is rejected; it therefore cannot lie strictly below u_{i+1}, since otherwise that successor would still have overlap at least e^{-4} and would also be accepted. Thus $t^{(i+1)} \geq u_{i+1} \geq v_{i+1}$. If $t^{(i)} \geq u_{i+1}$ the conclusion is immediate. Consequently the algorithm

reaches 1 after at most $\ell + 1 = O(M_{\mathrm{ad}})$ outer iterations, and every accepted transition has squared overlap at least c_0.

To choose all precisions before the run, use the computable upper bound

$$\bar{\ell} := C\sqrt{(1+\beta\widehat{D}_R)\log(2+\beta\widehat{D}_R)} \geq \ell, \tag{65}$$

and set

$$\bar{N}_{\mathrm{test}} := C'(\bar{\ell}+1)\log K_\beta, \qquad \bar{N}_{\mathrm{rot}} := C'(\bar{\ell}+1)\log\frac{\bar{\ell}+1}{\zeta}, \qquad \eta_{\mathrm{test}} := \frac{\chi}{\bar{N}_{\mathrm{test}}}, \tag{66}$$

together with

$$\epsilon_\star := \Theta\left(\frac{\zeta}{\bar{N}_{\mathrm{test}}\log(1/\eta_{\mathrm{test}}) + \bar{N}_{\mathrm{rot}}}\right). \tag{67}$$

Because the algorithm aborts once the number of probes exceeds $\bar{N}_{\mathrm{test}}$, and because $\bar{N}_{\mathrm{rot}}$ is an a priori bound on the number of selective phase rotations used for state conversion, these quantities, together with the gap bound $\underline{\delta}$, give a deterministic a priori upper bound on the total number of controlled walk applications. By the constant-overhead relation above, they therefore also give a deterministic a priori upper bound on the total number of controlled transition-state applications. We use this upper bound as T in Proposition 9.

A union bound gives simultaneous success of all nondestructive probes with probability at least $1-\chi$. Since each outer iteration uses $O(\log K_\beta)$ probes, schedule construction costs

$$O\left(\frac{\ell}{\sqrt{\underline{\delta}}}\log\frac{\ell}{\zeta}\log\frac{K_\beta\bar{\ell}}{\zeta\chi}\right).$$

The choice Eq. (67) makes the total error from approximate reflections $O(\zeta)$ by the hybrid argument. Combining this with the initial-state, state-conversion, and implementation errors gives final state error at most ζ after adjusting absolute constants.

Using $\ell = O(M_{\mathrm{ad}})$, $\log\bar{\ell} = O(\log K_\beta)$, and absorbing only logarithmic factors, the total adaptive cost is

$$O\left(\frac{M_{\mathrm{ad}}}{\sqrt{\underline{\delta}}}\log K_\beta\,\log^2\frac{K_\beta M_{\mathrm{ad}}}{\zeta}\,\log\frac{K_\beta M_{\mathrm{ad}}}{\chi}\right).$$

Finally, Eq. (13) implies

$$M_{\mathrm{ad}} = \widetilde{O}(1+\sqrt{A_{\beta,R}}) \leq \widetilde{O}(1+\sqrt{\beta D_R}),$$

giving the desired bounds.

In either case, measuring a state that is ζ-close in Euclidean norm to $|\Pi_1\rangle$ produces a classical distribution ν_h satisfying

$$\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \leq \zeta,$$

since trace distance cannot increase under measurement. In the adaptive case the construction succeeds with probability at least $1-\chi$; allowing an arbitrary output on the failure event gives unconditionally

$$\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \leq \zeta + \chi.$$

Proposition 9 converts controlled-walk complexity to the stated $\mathcal{O}_f$- and $\mathcal{O}_{\nabla f}$-query complexity. □

F.4 Comparison of the Explicit and Adaptive Sampling Bounds

At a fixed certified gap parameter $\underline{\delta}$, the difference between the two bounds in Theorem 10 is the annealing-schedule length. The adaptive bound is smaller whenever

$$\sqrt{(1+A_{\beta,R})\log(2+A_{\beta,R})} \;\lesssim\; 1+\beta\widehat{D}_R. \tag{68}$$

Since $A_{\beta,R} \le \beta \widehat{D}_R$, this condition holds up to absolute constants. For $A_{\beta,R} = \Theta(\beta \widehat{D}_R) \gg 1$, the ratio between the explicit and adaptive schedule factors is

$$\Theta\left(\sqrt{\frac{\beta \widehat{D}_R}{\log(\beta \widehat{D}_R)}}\right) = \widetilde{\Theta}\left(\sqrt{\beta \widehat{D}_R}\right),$$

and the gain is larger when $A_{\beta,R} \ll \beta \widehat{D}_R$.

For a direct comparison under the same queried-temperature set, refine the explicit schedule by taking

$$M = K_\beta = \left\lceil 1 + \beta \widehat{D}_R \right\rceil.$$

This preserves the hypothesis of Lemma 5, since

$$(\beta/M) D_R \le \frac{\beta \widehat{D}_R}{1 + \beta \widehat{D}_R} \le 1.$$

The resulting explicit temperatures coincide with the adaptive search grid $\mathcal{G}_\beta$. Hence the minimum phase gap over this refined explicit route is the same minimum used over the adaptive search grid. Moreover, K_β differs from the original explicit schedule length $\lceil \beta \widehat{D}_R \rceil$ by at most one stage, so this refinement does not change its asymptotic complexity. Thus, under a common gap certification, the adaptive advantage comes from reducing the number of annealing stages, rather than from assigning the adaptive route a more favorable per-stage gap.

G Details of the sampling-to-optimization reduction

Reduction and error budgets. Set $\bar{\varepsilon} := \varepsilon - \varepsilon_{\mathrm{loc}} > 0$, so weak localization reduces the ε-optimization target to $f_R^\star + \bar{\varepsilon}$. If $X \sim \nu_h$ and Y conditional on X is uniform on its cell C_X, then $Y \sim \widetilde{\nu}_h$ and Lemma 3 gives $|f(X) - f(Y)| \le \omega_h$. Hence it suffices to control the lifted sample at accuracy $a := \bar{\varepsilon} - \omega_h$. Moreover, $\|\widetilde{\nu}_h - \pi_{\beta,R}\|_{\mathrm{TV}} \le \|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} + \|\widetilde{\pi}^h_{\beta,R} - \pi_{\beta,R}\|_{\mathrm{TV}}$. Thus sampling, discretization, and, for the adaptive route, schedule construction enter the total-variation budget in Theorem 10 and Lemma 4, while finite temperature contributes the remaining Gibbs tail.

Finite-temperature Gibbs tail. For $\tau > 0$, define

$$S_\tau := \{x \in \Omega : f(x) \le f_R^\star + \tau\}, \qquad v_\tau := \frac{\mathrm{Vol}(S_\tau)}{\mathrm{Vol}(\Omega)}. \tag{69}$$

For $0 < \tau < a$, comparing the Gibbs mass of S_τ with that of $\{f > f_R^\star + a\}$ yields the factor $e^{-\beta(a-\tau)}/v_\tau$. The next theorem adds total-variation error, yielding a high-probability localized version of the standard Langevin sampling-to-optimization decomposition into sampling and Gibbs-model errors (Raginsky et al., 2017; Xu et al., 2018).

Theorem 20 (Finite-temperature Gibbs tail)**.** *Let ν satisfy $\|\nu - \pi_{\beta,R}\|_{\mathrm{TV}} \le \alpha_{\mathrm{TV}}$. For every $a := \bar{\varepsilon} - \omega_h > 0$ and every $0 < \tau < a \le \bar{\varepsilon}$ with $v_\tau > 0$,*

$$\Pr_{X \sim \nu}\left[f(X) > f_R^\star + a\right] \le \alpha_{\mathrm{TV}} + e^{-\beta(a-\tau)}/v_\tau. \tag{70}$$

To transfer Theorem 20 to grid point, apply it to $Y \sim \widetilde{\nu}_h$ with $a = \bar{\varepsilon} - \omega_h$. We ensure $a \ge \bar{\varepsilon}/2$ by

$$\omega_h \le \bar{\varepsilon}/2 \qquad \Longleftrightarrow \qquad h \le \bar{\varepsilon}/(2G_R\sqrt{d}). \tag{71}$$

The total-variation error is bounded above, while Eq. (71) and Eq. (8) are compatible upper bounds on h. Smoothness lower-bounds the sublevel volume v_τ in Eq. (69), even for a boundary minimizer.

Lemma 21 (Sublevel volume under smoothness)**.** *Under Assumption 1 and 2, for every $\tau > 0$,*

$$v_\tau \ge \min\left\{1, (\tau/2RG_R)^d\right\}, \quad \textit{and consequently,} \quad v_{\bar{\varepsilon}/2} \ge \min\left\{1, (\bar{\varepsilon}/4RG_R)^d\right\}. \tag{72}$$

G.1 Finite-Temperature Gibbs Concentration

Proof of Theorem 20. Let

$$B_a := \{x \in \Omega : f(x) > f_R^\star + a\}.$$

Since

$$\|\nu - \pi_{\beta,R}\|_{\mathrm{TV}} \le \alpha_{\mathrm{TV}},$$

the definition of total variation gives

$$\nu(B_a) \le \pi_{\beta,R}(B_a) + \alpha_{\mathrm{TV}}.$$

It remains to bound the Gibbs probability of B_a. By definition,

$$\pi_{\beta,R}(B_a) = \frac{\int_{B_a} e^{-\beta f(x)}\, dx}{Z_{\beta,R}}.$$

For every $x \in B_a$,

$$f(x) > f_R^\star + a,$$

and hence

$$\int_{B_a} e^{-\beta f(x)}\, dx \le e^{-\beta(f_R^\star + a)} \operatorname{Vol}(\Omega).$$

On the other hand, for

$$S_\tau = \{x \in \Omega : f(x) \le f_R^\star + \tau\},$$

we have, for every $x \in S_\tau$,

$$e^{-\beta f(x)} \ge e^{-\beta(f_R^\star + \tau)}.$$

Therefore,

$$\begin{aligned} Z_{\beta,R} &= \int_\Omega e^{-\beta f(x)}\, dx \\ &\ge \int_{S_\tau} e^{-\beta f(x)}\, dx \\ &\ge e^{-\beta(f_R^\star + \tau)} \operatorname{Vol}(S_\tau) \\ &= e^{-\beta(f_R^\star + \tau)} v_\tau \operatorname{Vol}(\Omega). \end{aligned}$$

Combining the numerator and denominator bounds yields

$$\pi_{\beta,R}(B_a) \le \frac{e^{-\beta(a-\tau)}}{v_\tau}.$$

Consequently,

$$\Pr_{X \sim \nu}[f(X) > f_R^\star + a] = \nu(B_a) \le \alpha_{\mathrm{TV}} + \frac{e^{-\beta(a-\tau)}}{v_\tau},$$

which proves Eq. (70). □

G.2 Sublevel-Set Volume Bound

Proof of Lemma 21. Let $x_R^\star \in \arg\min_{x \in \Omega} f(x)$, which exists by continuity and compactness. By Eq. (3), f is G_R-Lipschitz on Ω. If $G_R = 0$, then f is constant on Ω and $v_\tau = 1$. Assume $G_R > 0$ and set

$$r := \frac{\tau}{G_R}.$$

Every $x \in \Omega \cap B(x_R^\star, r)$ satisfies

$$f(x) - f_R^\star \le G_R \|x - x_R^\star\| \le \tau,$$

so this intersection is contained in S_τ.

If $r \geq 2R$, then the oscillation bound $D_R \leq 2RG_R \leq \tau$ implies $S_\tau = \Omega$ and $v_\tau = 1$. Suppose $r < 2R$. We claim that $B(0,R) \cap B(x_R^\star, r)$ contains a Euclidean ball of radius $r/2$. If $\|x_R^\star\| < r/2$, use the ball $B(0, r/2)$. Otherwise, let $u = x_R^\star/\|x_R^\star\|$ and use

$$B\Big(x_R^\star - \frac{r}{2}u, \frac{r}{2}\Big).$$

The latter lies in $B(x_R^\star, r)$, and for every point y in it,

$$\|y\| \leq \Big\|x_R^\star - \frac{r}{2}u\Big\| + \frac{r}{2} = \|x_R^\star\| \leq R.$$

Thus

$$\mathrm{Vol}(S_\tau) \geq \mathrm{Vol}(B(0, r/2)),$$

and dividing by $\mathrm{Vol}(B(0,R))$ gives

$$v_\tau \geq \Big(\frac{r}{2R}\Big)^d = \Big(\frac{\tau}{2RG_R}\Big)^d.$$

Combining the two cases proves Eq. (72); setting $\tau = \bar{\varepsilon}/2$ gives the second term. □

G.3 Proof of the Global Optimization Theorem

Proof of Theorem 13. If $\widehat{D}_R \leq \bar{\varepsilon}$, then for any $\widehat{x} \in \Omega_h$,

$$f(\widehat{x}) \leq f_R^\star + D_R \leq f^\star + \varepsilon_{\mathrm{loc}} + \widehat{D}_R \leq f^\star + \varepsilon,$$

so the claim is immediate. Hence assume $\widehat{D}_R > \bar{\varepsilon}$.

For the adaptive sampler, Theorem 10 with $\zeta = \chi = p_f/8$ gives

$$\|\nu_h - \pi^h_{\beta,R}\|_{\mathrm{TV}} \leq \zeta + \chi = \frac{p_f}{4}.$$

Since h satisfies Eq. (8) with $\xi = p_f/8$, Lemma 4 and preservation of total variation under the lift give

$$\|\widetilde{\nu}_h - \pi_{\beta,R}\|_{\mathrm{TV}} \leq \frac{3p_f}{8}.$$

Let $\widehat{x} \sim \nu_h$, and conditionally on $\widehat{x}$, let Y be uniform on $C_{\widehat{x}}$. Then $Y \sim \widetilde{\nu}_h$, and by Lemma 3,

$$|f(\widehat{x}) - f(Y)| \leq \omega_h.$$

Using weak localization, $f_R^\star \leq f^\star + \varepsilon_{\mathrm{loc}}$, so

$$\{f(\widehat{x}) > f^\star + \varepsilon\} \subseteq \{f(Y) > f_R^\star + \bar{\varepsilon} - \omega_h\}.$$

Set

$$a := \bar{\varepsilon} - \omega_h.$$

By Eq. (71), $a \geq \bar{\varepsilon}/2$. Applying Theorem 20 with $\tau = \bar{\varepsilon}/4$ gives

$$\Pr[f(\widehat{x}) > f^\star + \varepsilon] \leq \frac{3p_f}{8} + \frac{e^{-\beta(a-\tau)}}{v_\tau}.$$

Since $a - \tau \geq \bar{\varepsilon}/4$, Lemma 21 implies

$$\frac{e^{-\beta(a-\tau)}}{v_\tau} \leq \exp\Big(-\frac{\beta\bar{\varepsilon}}{4} + d\log_+\frac{8RG_R}{\bar{\varepsilon}}\Big) \leq \frac{p_f}{2},$$

where the last inequality follows from Eq. (30). Therefore

$$\Pr[f(\widehat{x}) > f^\star + \varepsilon] \leq \frac{7p_f}{8} < p_f,$$

which proves the optimization guarantee.

For the complexity, Theorem 10 gives

$$\widetilde{O}\left(\frac{1 + \sqrt{A_{\beta,R}}}{\sqrt{\underline{\delta}}}\right).$$

Since

$$A_{\beta,R} \leq \beta D_R \leq \beta \widehat{D}_R,$$

taking β equal to the sufficient value in Eq. (30) yields Eq. (31). The explicit-schedule bound in Eq. (31) follows similarly from the first part of Theorem 10. □

G.4 QUADRATIC BENCHMARK FOR THE INVERSE-TEMPERATURE SCALING

Consider

$$f(x) = \frac{1}{2}\|x\|^2, \qquad \Omega = B(0, R),$$

and suppose $0 < 2\bar{\varepsilon} \le R^2$. If $Z \sim \mathcal{N}(0, \beta^{-1} I_d)$, then $\pi_{\beta,R}$ is the law of Z conditioned on $\|Z\| \le R$, while $2\beta f(Z) \sim \chi^2_d$. Hence

$$\Pr_{X \sim \pi_{\beta,R}}[f(X) \le \bar{\varepsilon}] = \frac{\Pr\big(\chi^2_d \le 2\beta\bar{\varepsilon}\big)}{\Pr(\chi^2_d \le \beta R^2)}.$$

Whenever the localization radius retains a constant fraction of the Gaussian mass, i.e.,

$$\Pr\big(\chi^2_d \le \beta R^2\big) \ge c$$

for an absolute constant $c > 0$, constant-probability $\bar{\varepsilon}$-optimality requires a constant quantile of χ^2_d to be at most $2\beta\bar{\varepsilon}$. Since every fixed nonzero quantile of χ^2_d is $\Theta(d)$,

$$\beta = \Omega\left(\frac{d}{\bar{\varepsilon}}\right).$$

Thus the linear $d/\bar{\varepsilon}$ scaling cannot in general be removed by sharpening the sampling-to-optimization reduction alone; the additional logarithmic factors in Eq. (30) may still be improvable.