Hierarchical Discovery of Adiabatic Hamiltonian Paths and RL Schedules for Quantum Linear System Solvers

Anonymous Authors[1]

Abstract

Quantum linear-system solvers can provide asymptotic speedups under structured access assumptions, but finite-size performance depends strongly on the spectral encoding, Hamiltonian path, traversal schedule, and implementation model. We study a hierarchical framework for the quantum linear systems problem (QLSP) that separates admissible Hamiltonian-path design from adaptive traversal control. The outer loop searches constrained deformations of analytically valid adiabatic backbones within implementable operator libraries, while the inner loop is restricted to residual schedule control around a derivative-aware local-adiabatic prior. This residual class can be optimized by derivative-free search, differentiable optimal control, or residual RL. The experiments instantiate direct and RL residual-control comparators on the same schedule class, while treating broader family-wise RL as an amortized-control extension rather than as unconstrained Hamiltonian invention. In this sense, the framework acts as a two-stage automatic algorithm-design pipeline: the outer layer automatically adapts valid path families, while the inner layer automatically refines residual traversal schedules. We do not seek to improve worst-case asymptotic query complexity beyond optimal-scaling QLSP solvers; instead, the framework targets engineering-relevant finite-size fidelity–time tradeoffs, family-dependent adaptation, and diagnostic transparency. Simulations show the clearest gains on structured gap-amplified and preconditioned families. Ancilla-assisted extensions can improve gap structure, subspace anchoring, and preconditioning, but do not evade known optimal-scaling query-complexity limits.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Submitted to the AI for Science workshop (ICML 2026).

1. Introduction

Linear systems $Ax = b$ are central to simulation, inverse problems, optimization, and machine learning. Quantum linear-system solvers address the quantum linear systems problem (QLSP) by preparing a state proportional to the solution, $|x\rangle \propto A^{-1}|b\rangle$, rather than outputting all entries of x. HHL (Harrow et al., 2009) showed that strong asymptotic speedups are possible under structured sparsity and access assumptions, but the standard gate-model route relies on phase estimation, controlled eigenvalue inversion, amplitude amplification, and related subroutines that remain demanding in practice.

The gate-model route is not the only algorithmic vocabulary. Adiabatic quantum computation is polynomially equivalent to the circuit model (Aharonov et al., 2008), but it frames design differently: encode the target solution in an eigenstate of a Hamiltonian and reach that state through controlled evolution. This creates a natural structured scientific-design question. Can learning help discover better adiabatic quantum algorithms, or should it only tune parameters inside a fixed construction?

Our starting point is the adiabatic/randomization-based QLSP solvers of Subasi et al. (2019). These methods avoid the phase-estimation and eigenvalue inversion stack of HHL and use simple Hamiltonian families whose evolution prepares the solution state. Their ground-state-eigenstate-based and gap-amplified variants scale as $\mathcal{O}(\kappa^2 \log(\kappa)/\epsilon)$ and $\mathcal{O}(\kappa \log(\kappa)/\epsilon)$, respectively, where κ is the condition number and ϵ is the target precision. The improved variant adds an ancilla and amplifies the relevant gap, illustrating both the promise and the limits of structured auxiliary resources.

We extend this line in two directions. First, we sharpen the role of RL: the most defensible use is schedule optimization along a fixed or weakly parameterized Hamiltonian path. Path discovery itself is better treated as constrained optimization over physically valid families $H(s;\theta)$, using gradient-based or derivative-free search when appropriate. The inner problem is finite-horizon, low-dimensional, nearly deterministic, and explicitly constrained, so generic RL should be compared against differentiable optimal control and model-based control. The outer problem is expensive because every candidate requires feasibility screening and inner-loop evaluation, so it should be treated as constrained scientific design rather than a random walk in Hamiltonian space.

Second, we clarify what auxiliary qubits can and cannot buy. Ancillas can amplify gaps, anchor subspaces, or support preconditioning, but they should not be used to claim recursive asymptotic improvements beyond known optimal-scaling QLSP solvers in the general setting (Costa et al., 2022; Wang & Zhang, 2024; Mori et al., 2026). The value of this study is therefore a rigorous bilevel formulation for adiabatic QLSP design, plus finite-size evidence showing when structured path search, adaptive schedules, stronger reference paths, and preconditioning actually help.

Engineering-oriented scope. Our goal is not to improve the worst-case asymptotic query complexity of general QLSP beyond optimal-scaling solvers. Instead, we study the engineering problem that arises after a valid adiabatic QLSP backbone has been chosen: how to adapt the Hamiltonian path and traversal schedule to a finite-dimensional, structured matrix family under explicit implementation constraints. This shifts the objective from universal asymptotic scaling to engineering-relevant fidelity–time tradeoffs, feasible operator libraries, branch-stability audits, and controllable traversal policies. The RL component is deliberately restricted to adaptive residual schedule control around a physics-informed local-adiabatic prior. It is not used as unconstrained Hamiltonian invention; path discovery remains a constrained scientific design problem over admissible Hamiltonian families.

Contributions. We make four contributions:

1. We formulate adiabatic QLSP improvement as a bilevel co-design problem that separates admissible Hamiltonian-family search from adaptive traversal control, matching the distinction between design-space engineering and sequential schedule optimization.

2. We introduce a validity contract for learned path deformations, including Hermiticity, endpoint matching, symmetry-sector preservation, branch-continuity tracking, gap-floor protection, trust-region control, and implementability constraints.

3. We restrict any RL component to residual schedule control around a derivative-aware local-adiabatic prior, making the learning problem low-dimensional, monotone, physically structured, and directly tied to finite-size fidelity–time tradeoffs. The same residual class also defines the comparison target for derivative-free and differentiable optimal-control baselines.

4. We evaluate the framework through schedule-only, path-only, and joint ablations, together with hit-rate, feasibility-audit, branch-continuity, gap-floor, and rejected-candidate diagnostics, showing where constrained path–schedule optimization improves engineering-relevant traversal under explicit feasibility constraints. We also state which stronger baselines remain unimplemented, so the empirical claim is not inflated beyond the available data.

2. Related Work and Positioning

Our framework intersects adiabatic scheduling theory, quantum optimal control, RL for quantum systems, and modern QLSP algorithms. Local adiabatic evolution (Roland & Cerf, 2001) and adiabatic error bounds (Jansen et al., 2007) show that runtime is controlled by both the spectral gap and Hamiltonian variation, not by the minimum gap alone. This is why our baseline schedule and controller observations use $\widehat{\|\partial_s H\|}$ together with $\hat{g}$.

Engineering Perspective Relative to Optimal-Scaling QLSP. Subasi et al. (2019) introduced the randomized adiabatic constructions that motivate our implemented reference families, with ground-state-eigenstate and gap-amplified variants scaling as $\mathcal{O}(\kappa^2 \log(\kappa)/\epsilon)$ and $\mathcal{O}(\kappa \log(\kappa)/\epsilon)$. At the level of worst-case oracle or block-encoding complexity, later work reaches the optimal $\mathcal{O}(\kappa \log(1/\epsilon))$ query-complexity regime with respect to κ and ϵ under standard oracle or block-encoding access models and studies constant factors for discrete and randomized adiabatic variants (Costa et al., 2022; An & Lin, 2022; Lin & Tong, 2020; Jennings et al., 2025; Costa et al., 2025). Those results define the complexity benchmark for general QLSP, alongside gate-based block-encoding and QSVT solvers (Childs et al., 2017; Gilyén et al., 2018).

Our paper addresses a different question: once an analytically valid adiabatic QLSP construction is fixed, how should one adapt it to structured matrix families, finite-size constraints, and implementation-compatible operator libraries? The metric in our experiments is simulated evolution time to target fidelity for fixed finite-dimensional families, not asymptotic query complexity. The value of the framework is therefore engineering-oriented: admissible path search, adaptive residual traversal control, spectral-validity audits, and problem-family-aware diagnostics under explicit implementability constraints (Lefterovici et al., 2025; Lapworth & Sünderhauf, 2025).

This distinction also appears in the numerical protocols. Prior QLSP studies variously report inverse error versus randomized steps, fitted κ- and ϵ-scaling, or average complexity at fixed error thresholds. By contrast, our simulator emphasizes time-to-fidelity, instance-wise hit rates, and outer-loop feasibility diagnostics for structured finite-dimensional families, because those are the quantities most directly relevant to implementation-aware co-design.

For the inner loop, the right comparison class is broader than generic RL. GRAPE, differentiable pulse engineering, and other quantum optimal-control methods remain strong instance-wise baselines (Khaneja et al., 2005; Porotti et al., 2023; Mahesh et al., 2022). RL becomes more compelling when the goal is amortized control over matrix families, online adaptation, or partial feedback (Bukov & Marquardt, 2026; Zhang et al., 2019). Because the schedule problem also has explicit fidelity, leakage, and nonadiabaticity constraints, it connects naturally to CMDP-style safe RL, including CPO, RCPO, FOCOPS, PID-Lagrangian updates, and recent constrained-control surveys (Achiam et al., 2017; Tessler et al., 2018; Zhang et al., 2020; Stooke et al., 2020; Kushwaha et al., 2025). Work transferring learned adiabatic schedules into variational circuits also supports learning schedules around structured ansatzes rather than inventing arbitrary dynamics (Nzongani et al., 2026).

For the outer loop, prior schedule-path optimization in adiabatic and annealing settings (Zeng et al., 2015) motivates path-level search, while CMA-ES, Bayesian optimization, and TuRBO (Hansen, 2016; Frazier, 2018; Eriksson et al., 2019) are natural when each candidate requires spectral screening and inner-loop control. Our position is therefore deliberately conservative: learning is useful, but only inside a validity contract that enforces Hermiticity, endpoints, branch continuity, gap floors, symmetry sectors, and implementability.

3. Background and Motivation

Adiabatic Scheduling and QLSP Baselines. For algorithm design it is convenient to parameterize the evolution by a path variable $s \in [0,1]$ and a local slowdown factor $\tau(s) = dt/ds$. The standard local-adiabatic heuristic is

$$\tau(s) \gg \frac{\|\partial_s H(s)\|}{g(s)^2}, \qquad \dot{s} \ll \frac{g(s)^2}{\|\partial_s H(s)\| + \varepsilon_H}, \tag{1}$$

which already shows why schedule design depends on both spectral separation and Hamiltonian variation. We use this relation as a physics prior rather than as a complete correctness certificate; additional derivations and caveats are deferred to Appendix C. Likewise, adiabatic computation is polynomially equivalent to the circuit model, so Hamiltonian-path design should be treated as a genuine algorithm-design problem rather than as a mere heuristic alternative to gate compilation (Aharonov et al., 2008; Mizel et al., 2007); a short constructive sketch appears in Appendix D.

The QLSP asks for preparation of

$$|x\rangle \propto A^{-1}|b\rangle \tag{2}$$

for a normalized right-hand side $|b\rangle$. Throughout, we assume $A \in \mathbb{C}^{N\times N}$ is full-rank and Hermitian and use the standard spectral normalization $\|A\| = 1$ unless stated otherwise, with condition number $\kappa = \|A\|\,\|A^{-1}\|$. Under this normalization $\sigma_{\min}(A) = \|A^{-1}\|^{-1} = 1/\kappa$. If one works only with $\|A\| \le 1$ without rescaling to unit norm, every occurrence of $1/\kappa$ in the gap lower bounds below should instead be read as $\|A^{-1}\|^{-1} = \sigma_{\min}(A)$. The system register has n qubits, so $N = 2^n$. The target-state encoding starts from

$$H_p = A(I - |b\rangle\langle b|)A. \tag{3}$$

Since $A = A^\dagger$ and $I - |b\rangle\langle b|$ is an orthogonal projector, $H_p = H_p^\dagger$ and

$$y^\dagger H_p y = (Ay)^\dagger (I - |b\rangle\langle b|)(Ay) \ge 0. \tag{4}$$

Moreover, $H_p A^{-1}|b\rangle = A(I - |b\rangle\langle b|)|b\rangle = 0$, so the QLSP solution lies in the zero-energy ground space.

We build on the adiabatic QLSP constructions of Subasi et al. (2019). With the sign convention

$$|\bar{b}\rangle = |+\rangle \otimes |b\rangle, \qquad P_{\bar{b}} = I - |\bar{b}\rangle\langle\bar{b}|, \tag{5}$$

the core ground-state-eigenstate-based (one-ancilla) path is

$$H(s) = A(s)P_{\bar{b}}A(s), \qquad A(s) = (1-s)Z\otimes I + sX\otimes A, \tag{6}$$

acting on $\mathbb{C}^2 \otimes \mathbb{C}^N$, hence on $n+1$ qubits. The extra qubit makes $A(s)$ invertible for all $s \in [0,1]$: if $A|u\rangle = \lambda|u\rangle$, the corresponding eigenvalues of $A(s)$ are $\pm\sqrt{(1-s)^2 + (s\lambda)^2}$. The tracked zero-energy branch is

$$|x(s)\rangle = \frac{A(s)^{-1}|\bar{b}\rangle}{\|A(s)^{-1}|\bar{b}\rangle\|}, \tag{7}$$

which interpolates from $|-\rangle \otimes |b\rangle$ to $|+\rangle \otimes |x\rangle$ and admits the gap lower bound

$$\Delta(s) \ge \Delta_\star(s) := (1-s)^2 + \left(\frac{s}{\kappa}\right)^2. \tag{8}$$

Equation (8) uses the standard spectral normalization in which nonzero singular values lie in $[1/\kappa, 1]$; without that normalization $1/\kappa$ is replaced by $1/\|A^{-1}\|$. The associated randomization method (RM) yields expected runtime $\mathcal{O}(\kappa^2 \log(\kappa)/\epsilon)$, while the gap-amplified two-ancilla construction reshapes the nonzero spectrum from $\{\gamma_i\}$ to $\{\pm\sqrt{\gamma_i}\}$ and improves the runtime to $\mathcal{O}(\kappa \log(\kappa)/\epsilon)$ (Subasi et al., 2019). We use these two families as structured reference paths throughout. Detailed spectral formulas, RM schedules, and auxiliary constructions are deferred to Appendix B and Appendix E.

These two baselines are useful not only because they are analytically grounded, but because they expose complementary design tradeoffs. The ground-state-eigenstate-based family is the minimal structured path that already encodes the QLSP solution branch in a controlled way, making it a natural anchor for small deformations and schedule-learning ablations. The gap-amplified family, by contrast, changes the spectral geometry more aggressively and introduces a degenerate zero subspace, thereby illustrating that better asymptotic gap behavior can come with a harder control problem. Together, they provide a principled starting point for asking whether learning should improve traversal on a fixed construction, propose better nearby constructions, or co-design both at once.

Structured Algorithmic Co-Design. For adiabatic QLSP, usable runtime depends on the full spectral geometry of the path, not only on an endpoint construction. This motivates a hierarchical design problem: constrained search proposes admissible Hamiltonian-family deformations, while sequential control adapts how each valid path is traversed.

4. Method: Hierarchical Path–Schedule Discovery

Design principle: residual control for schedule, constrained search for path. The two motivating ideas behind our framework should be interpreted at different levels. The first idea is to use a residual controller to improve *how* an adiabatic path is traversed. That controller can be optimized by derivative-free search, differentiable optimal control, or family-wise RL. The second idea is to let constrained search propose *what* Hamiltonian family should be traversed in the first place. We therefore do not cast both layers as RL. Instead, path discovery remains a constrained optimization problem over admissible Hamiltonian families, while RL is reserved for the amortized schedule-control setting where a reusable policy is trained across a family of instances.

This separation is important for both technical and conceptual reasons. Technically, the schedule problem is sequential, continuous, and naturally expressed as a Markov decision process. The path-design problem is instead a static scientific design problem with strong structural constraints; forcing it into a high-dimensional RL action space would make training unnecessarily unstable and would obscure the physical meaning of the learned path. Conceptually, this decomposition makes the role of RL precise rather than overstated: RL is a candidate solver for residual traversal policies, not a license for unconstrained Hamiltonian invention.

Outer loop: reference-path constrained Hamiltonian-family search. We do not search over arbitrary Hermitian interpolations. Let $H_{\mathrm{ref}}(s)$ denote a reference QLSP family, such as the ground-state-eigenstate-based path or the gap-amplified two-ancilla construction in Section 3. The outer loop only searches over an *admissible* class of bounded deformations around this reference path:

$$H(s;\theta) = H_{\mathrm{ref}}(s) + s(1-s)\Delta H_\theta(s), \tag{9}$$

with

$$\Delta H_\theta(s) = \sum_{m=1}^{M} c_m(s;\theta) B_m, \tag{10}$$

where $\{B_m\}_{m=1}^M \subset \mathcal{B}_{\mathrm{feas}}$ is a feasible operator library and the scalar coefficients are real-valued functions whose value and first derivative vanish at $s \in \{0,1\}$. The multiplicative envelope and endpoint clamping ensure that both the path and its boundary velocity match the reference family. The admissibility requirement is not merely syntactic; it is a contract that combines algebraic structure, mesh-certified spectral continuity, and implementation cost.

Proposition 1. *If every B_m is Hermitian and every $c_m(s;\theta)$ is real-valued, then the family in Eq. (9) is Hermitian for all s and satisfies the desired boundary conditions by construction. Thus Hermiticity is a hard construction constraint, not an ordinary post-hoc rejection test.*

In numerical outer-loop proposals we still apply a Hermitian projection before spectral auditing,

$$\Pi_{\mathrm{Herm}}(\widetilde{H}) = \frac{\widetilde{H} + \widetilde{H}^\dagger}{2}, \tag{11}$$

and record the residual

$$\epsilon_{\mathrm{Herm}}(\theta) = \max_k \left\| \widetilde{H}(s_k;\theta) - \Pi_{\mathrm{Herm}}(\widetilde{H}(s_k;\theta)) \right\|. \tag{12}$$

This symmetrization is intended to remove numerical drift or harmless parameterization noise. A large ϵ_{Herm} indicates that the proposal has left the admissible basis model; after projection, the candidate must still pass the symmetry, gap, branch-continuity, trust-region, and implementability checks below.

The key change relative to a generic path ansatz is that $\mathcal{B}_{\mathrm{feas}}$ is not an arbitrary Hermitian basis. It is chosen to preserve the problem structure inherited from the baseline QLSP construction. In particular, each B_m must belong to a fixed implementable library and preserve any known symmetry sector used to encode the target state. If $\{Q_a\}$ denotes a set of conserved quantities or block labels that identify the intended branch, we require

$$[B_m, Q_a] = 0 \qquad \text{for all } m, a. \tag{13}$$

This prevents the outer loop from improving a gap by simply mixing the target subspace with an irrelevant sector.

The reported one-ancilla, gap-amplified, and hybrid searches instantiate this idea through restricted operator libraries, but the main text keeps the abstract admissible-library interface because the engineering claim does not depend on one fixed basis choice.

Operational safeguards for branch continuity. Hermiticity enforcement and endpoint matching are necessary but not sufficient. The outer loop must also discourage accidental level crossings, degeneracies, or target branch switching. We therefore use a mesh-certified feasibility screen on sampled points $\{s_k\}_{k=0}^{K}$ after any numerical Hermitian projection has been applied. Let $P_\theta^\star(s_k)$ denote the rank-r projector associated with the target solution branch, where $r = 1$ in the nondegenerate case and $r > 1$ for a protected solution subspace. In gap-amplified constructions this projector is not the full zero-energy eigenspace: it excludes the auxiliary zero mode $|1\rangle \otimes |\bar{b}\rangle$ unless an explicit symmetry label makes that mode part of the intended encoded solution sector. Operationally, the target branch is defined by continuity from the distinguished reference QLSP branch whose endpoint encodes the desired solution state $|x\rangle \propto A^{-1}|b\rangle$; on the mesh, the branch is transported by maximum-overlap continuation rather than by eigenvalue sorting alone. A candidate path is accepted only if the following checks hold on the mesh:

$$\|s_k(1 - s_k)\Delta H_\theta(s_k)\| \le \eta\, \hat{g}_{\mathrm{ref}}(s_k), \tag{14}$$

$$\hat{g}_\theta(s_k) \ge g_{\mathrm{floor}} \quad \text{or} \quad \hat{g}_\theta(s_k) \ge \rho\, \hat{g}_{\mathrm{ref}}(s_k), \tag{15}$$

$$\frac{1}{r}\mathrm{Tr}(P_\theta^\star(s_k) P_\theta^\star(s_{k-1})) \ge o_{\min}. \tag{16}$$

Eq. (14) is a trust-region constraint that keeps the deformation small relative to the reference gap. Eq. (15) enforces a minimum usable spectral separation. Eq. (16) tracks the same spectral branch across adjacent mesh points and is more robust than sorting eigenvectors by eigenvalue alone when near-degeneracies arise. In the one-dimensional case, Eq. (16) reduces to the usual eigenvector overlap $|\langle \phi_\theta^\star(s_k), \phi_\theta^\star(s_{k-1})\rangle|^2$.

These checks do not constitute a full theorem of adiabatic correctness, but they provide operational safeguards that are strong enough to reject obviously fragile paths during search. In practice, the outer loop is run as a trust-region continuation method: if any feasibility check fails, the candidate update is rejected or shrunk before the next trial. Together with the implementation budget below, these checks define the feasible search object rather than post-hoc diagnostics. Hermiticity is the exception: it is enforced by the basis and real coefficients, with Eq. (11) used only as a numerical repair before the remaining audit.

Explicit implementability constraints. Implementability is part of the admissible path definition rather than an after-the-fact diagnostic. The feasible operator library $\mathcal{B}_{\mathrm{feas}}$ must be tied to a hardware or access model rather than left implicit. In a gate-model setting, each B_m is restricted to terms synthesizable from the same oracle or block-encoding primitives as the reference QLSP construction, under fixed ancilla, LCU-term, and noncommuting-group budgets and without introducing qualitatively new data-access assumptions. We therefore associate each path with an implementation cost surrogate

$$\begin{aligned}\Omega_{\mathrm{impl}}(\theta) = w_1 M_{\mathrm{LCU}}(\theta) + w_2 n_{\mathrm{anc}}(\theta) \\ + w_3 N_{\mathrm{grp}}(\theta) + w_4 \max_m \|c_m(\cdot;\theta)\|_\infty,\end{aligned} \tag{17}$$

where M_{LCU} counts linear-combination terms, n_{anc} counts auxiliary qubits, and N_{grp} counts noncommuting operator groups that must be simulated separately. Extended resource reporting uses the same ingredients through LCU one-norm and weighted-evolution scores, while the main text keeps only the compact surrogate Ω_{impl}.

For analog or annealing platforms, the same framework admits a different feasible class: B_m can be specialized to k-local interactions on a hardware graph G, optionally inside a stoquastic cone, with native coupling constraints, amplitude bounds $\|c_m\|_\infty \le c_{\max}$, and slew-rate limits $\|\partial_s c_m\|_\infty \le r_{\max}$. We do not optimize over both models simultaneously, and the present experiments should not be read as hardware demonstrations. Only the gate-model-aligned surrogate is instantiated in the present numerical study; the analog and annealing constraints define compatible admissible classes for future implementation studies. A concrete study should commit to one feasible library and treat the other as a drop-in variant. Thus the outer loop searches only over physically or access-model-compatible Hamiltonian deformations. Additional choices of structured path families, expensive derivative-free outer-loop solvers, and adiabatic surrogate objectives are deferred to Appendix M.

Definition 1 (Admissible reference-path class)**.** *Fix a reference path H_{ref}, a feasible operator library $\mathcal{B}_{\mathrm{feas}}$, a mesh $\{s_k\}_{k=0}^{K}$, and thresholds $\Gamma = (\eta, g_{\mathrm{floor}}, \rho, o_{\min}, \Omega_{\max})$. The admissible class $\mathcal{A}_{\mathrm{ref}}(\Gamma)$ consists of all paths $H(\cdot;\theta)$ of the form in Eqs.* (9)–(10) *such that: (i) Hermiticity is enforced by Hermitian basis elements and real coefficients, with only small numerical residual ϵ_{Herm} allowed after Eq.* (11)*; (ii) the path is endpoint-matched; (iii) every basis term preserves the target symmetry sector; (iv) the trust-region, gap-floor, and branch-overlap checks in Eqs.* (14)–(16) *hold on the audit mesh; and (v) $\Omega_{\mathrm{impl}}(\theta) \le \Omega_{\max}$.*

Inner loop: residual adaptive traversal. Given a fixed path $H(s;\theta)$, we discretize the traversal into steps $j =$

$1,\dots,q$. At step j, the controller observes a state summary such as

$$x_j = \big(s_j,\ \log\hat{g}_j,\ \log(\widehat{\|\partial_s H\|}_j + \varepsilon_H),\ \log\frac{\hat{g}_j^2}{\widehat{\|\partial_s H\|}_j + \varepsilon_H},\ \hat{F}_j,\ \hat{L}_j^{\mathrm{leak}},\ \kappa(A),\ c_j\big), \tag{18}$$

where $\hat{g}_j$ is a local gap proxy, $\widehat{\|\partial_s H\|}_j$ is an operator-variation proxy, $\hat{F}_j$ is a fidelity-progress signal, $\hat{L}_j^{\mathrm{leak}}$ measures leakage outside the intended symmetry sector, and c_j contains optional implementation indicators. The controller does not directly invent an arbitrary schedule. Instead, it modulates a physics-informed baseline. We fix a path grid $\{s_j\}_{j=1}^q$ with spacing $\delta s_j = s_{j+1} - s_j$ and define a reference local-adiabatic velocity

$$\bar{v}_j = \nu_0 \frac{\hat{g}_j^2}{\widehat{\|\partial_s H\|}_j + \varepsilon_H}, \qquad \bar{t}_j = \frac{\delta s_j}{\bar{v}_j}, \tag{19}$$

which incorporates the dependence on both the instantaneous gap and the local Hamiltonian variation suggested by local-adiabatic scheduling and standard adiabatic error bounds (Roland & Cerf, 2001; Jansen et al., 2007). The action is then a positive multiplicative correction

$$a_j = u_j, \qquad t_j = e^{u_j}\bar{t}_j, \tag{20}$$

where $u_j \in [-u_{\max}, u_{\max}]$ is optimized by the residual controller. This preserves the physical prior that smaller gaps and faster-changing Hamiltonians should induce slower evolution, while still allowing residual optimization or RL to learn nontrivial corrections to the handcrafted schedule.

Two simpler action parameterizations are also possible. One may fix the grid points $\{s_j\}$ and let the controller output dwell times directly, or fix a step size δs and learn local velocities $v_j = \delta s / t_j$. Both are more structured than learning $(\delta s_j, t_j)$ jointly. Our preferred choice, however, is the derivative-aware multiplier above, because it makes the local-adiabatic heuristic explicit and asks the controller to learn only a physically meaningful correction to it. In practice we prefer low-dimensional residual schedule classes over fully unconstrained per-step actions; concrete examples are deferred to Appendix M.

Audited gap and fidelity proxies. The proxy signals used by the controller should be explicitly defined rather than treated as abstract oracles. Let the tracked solution branch occupy the continued index window $I_\star(s_k) = \{a_k, \dots, b_k\}$ among the Ritz values inside the audited symmetry sector, with $r = b_k - a_k + 1$. We estimate the relevant spectral separation by the two-sided branch gap

$$\begin{aligned}\hat{g}(s_k) &= \min\{\hat{g}_-(s_k), \hat{g}_+(s_k)\},\\ \hat{g}_-(s_k) &= \hat{\lambda}^{(K)}_{a_k}(s_k) - \hat{\lambda}^{(K)}_{a_k-1}(s_k),\\ \hat{g}_+(s_k) &= \hat{\lambda}^{(K)}_{b_k+1}(s_k) - \hat{\lambda}^{(K)}_{b_k}(s_k),\end{aligned} \tag{21}$$

with the missing lower or upper term treated as $+\infty$ at a spectral boundary. For a nondegenerate ground-state branch this reduces to the usual one-sided gap above the ground state. For the gap-amplified path, whose desired solution branch sits at zero energy in the middle of the spectrum and coexists with an auxiliary zero mode, Eq. (21) is applied after the maximum-overlap and symmetry-sector audit has isolated the solution-encoding branch; if the auxiliary zero mode cannot be excluded by the declared branch or symmetry label, the effective protected-branch gap is zero and the candidate is not treated as admissible. The Ritz values $\hat{\lambda}^{(K)}$ are obtained by exact diagonalization on the smallest instances and by Lanczos, Davidson, or shift-invert Lanczos on larger instances inside the correct symmetry sector. Likewise, if $P^\star_{\mathrm{sol}}(s_j)$ denotes the tracked solution-subspace projector, then the instantaneous fidelity-progress proxy is

$$\hat{F}_j = \langle \psi_j | P^\star_{\mathrm{sol}}(s_j) | \psi_j \rangle, \tag{22}$$

where $P^\star_{\mathrm{sol}}(s_j)$ is the maximum-overlap continuation of the reference solution-encoding branch onto the deformed path. In the gap-amplified reference construction this projector targets $|0\rangle \otimes |x(s_j)\rangle$ (or the declared protected solution subspace) and explicitly excludes the unrelated zero-mode branch $|1\rangle \otimes |\bar{b}\rangle$. The leakage proxy is

$$\hat{L}_j^{\mathrm{leak}} = 1 - \langle \psi_j, \Pi_{\mathrm{sym}} \psi_j \rangle, \tag{23}$$

where Π_{sym} projects onto the intended symmetry sector. Finally, we define a trajectory-level nonadiabaticity proxy by

$$L_{\mathrm{nonadia}} = \frac{1}{q}\sum_{j=1}^{q}(1 - \hat{F}_j), \tag{24}$$

and a trajectory-level leakage proxy by

$$L_{\mathrm{leak}} = \frac{1}{q}\sum_{j=1}^{q}\hat{L}_j^{\mathrm{leak}}, \tag{25}$$

which are both distinct from the terminal task loss $1 - F_{\mathrm{final}}$.

These proxies have explicit costs and failure modes. In particular, gap estimation dominates the spectral budget and can fail through finite Krylov error, incomplete convergence, or branch mis-tracking near avoided crossings. Training and evaluation should therefore use the same *family* of proxies while allowing coarser tolerances during training and tighter audits at evaluation time.

This turns adiabatic design into a finite-horizon constrained control problem. The objective is not unconstrained reward maximization; it is to minimize total evolution time subject to fidelity, nonadiabaticity, and leakage requirements:

$$\begin{aligned}&\min_\phi\ \mathbb{E}[T_{\mathrm{total}}]\\ &\text{s.t. } \mathbb{E}[1 - F_{\mathrm{final}}] \le \varepsilon_F,\\ &\quad\ \ \mathbb{E}[L_{\mathrm{nonadia}}] \le \varepsilon_{\mathrm{na}},\\ &\quad\ \ \mathbb{E}[L_{\mathrm{leak}}] \le \varepsilon_{\mathrm{leak}}.\end{aligned} \tag{26}$$

This structure is closer to a constrained Markov decision process than to a generic RL benchmark, which motivates constrained continuous-control methods; in the lightweight simulator study below, we instantiate this principle with residual schedule optimization and constraint-relevant diagnostics.

From an engineering perspective, the inner problem is not generic reward maximization. It is minimum-time traversal under fidelity, leakage, and nonadiabaticity constraints. The controller therefore acts only through multiplicative residual corrections to a derivative-aware local-adiabatic prior. This keeps the policy low-dimensional, monotone in the path variable, and physically interpretable: smaller gaps and faster Hamiltonian variation still induce slower evolution, while learning can correct conservative or locally inefficient handcrafted schedules.

Bilevel objective. Let ϕ denote the schedule policy parameters, let θ denote the path parameters, and let $\mathcal{S}$ denote the inner solver class used to traverse a candidate path. Examples of $\mathcal{S}$ include derivative-free residual schedule search, constrained continuous-control RL, and differentiable optimal control. For a fixed admissible path $H(\cdot;\theta) \in \mathcal{A}_{\mathrm{ref}}(\Gamma)$, the inner problem is the minimum-time control problem

$$\phi^\star_{\mathcal{S}}(\theta) \in \arg\min_{\phi\in\mathcal{S}} \left\{ \begin{array}{c} \mathbb{E}[T_{\mathrm{total}}] \\ \text{s.t. } \mathbb{E}[1-F_{\mathrm{final}}] \le \varepsilon_F, \\ \mathbb{E}[L_{\mathrm{nonadia}}] \le \varepsilon_{\mathrm{na}}, \\ \mathbb{E}[L_{\mathrm{leak}}] \le \varepsilon_{\mathrm{leak}} \end{array} \right\}. \tag{27}$$

This notation is intentionally solver-dependent. A path is not intrinsically good merely because it has a larger minimum gap or a smaller proxy score; it is good if the chosen inner solver can exploit its spectral geometry while meeting fidelity and leakage constraints.

Definition 2 (Path-under-a-solver objective)**.** *For an admissible path $H(\cdot;\theta)$ and an inner solver class $\mathcal{S}$, define the path-under-a-solver value*

$$\begin{aligned} \mathcal{V}_{\mathcal{S}}(\theta;\mathcal{D}) = \mathbb{E}_{(A,b)\sim\mathcal{D}}\Big[R\left(\pi_{\phi^\star_{\mathcal{S}}(\theta)}, H(\cdot;\theta), A, b\right) \\ - \lambda_{\mathrm{impl}}\Omega_{\mathrm{impl}}(\theta)\Big]. \end{aligned} \tag{28}$$

The outer search is therefore

$$\theta^\star_{\mathcal{S}} \in \arg\max_{\theta:\, H(\cdot;\theta)\in\mathcal{A}_{\mathrm{ref}}(\Gamma)} \mathcal{V}_{\mathcal{S}}(\theta;\mathcal{D}). \tag{29}$$

Eq. (29) is the central theoretical object of the paper. It rules out a common ambiguity in learned Hamiltonian design: the outer loop is not optimizing a standalone path proxy, and the inner loop is not asked to rescue arbitrary paths. Instead, the admissible class defines which paths are legal, while $\mathcal{V}_{\mathcal{S}}$ defines which legal paths are useful under a specified solver and distribution. In the reported simulator, $\mathcal{S}$ is instantiated primarily as warm-started derivative-free residual schedule search, while the outer objective explicitly penalizes implementation overhead. Lagrangian constrained RL and differentiable bilevel solvers fit the same formal objective. The present numerical study reports lightweight residual search in the main text, while Appendix K, Table 5, records a differentiable residual-control comparison and Appendix M lists broader constrained-RL and bilevel extensions that remain follow-up baselines.

The strongest instance-wise comparison for this residual class would optimize the same variables u_j in $t_j = e^{u_j}\bar{t}_j$ by GRAPE-style gradients, L-BFGS, or Adam with the same $u_{\max}$ and grid size q. The main text keeps the focus on the randomized, residual-search, path-only, and joint comparisons; Appendix K, Table 5, reports the corresponding differentiable residual-control baseline on the same schedule class.

Outer–inner coupling. The outer objective is not a standalone path proxy but a *path under a solver*: a candidate family is useful only if the chosen inner-loop optimizer can exploit its spectral geometry while satisfying fidelity and leakage constraints.

5. Experimental Setup and Evaluation

We evaluate the proposal in a classical simulator with piecewise-constant Hamiltonian evolution. Detailed propagation formulas, benchmark-family definitions, baseline inventories, metric lists, the derivative-aware prior, and the longer query-complexity and implementability caveats are deferred to Appendices I, L.1, and K.9. The inner layer is instantiated as structure-aware schedule optimization around that prior, while the outer layer is instantiated as trust-region path search around the ground-state-eigenstate-based and gap-amplified reference families of Subasi et al. (2019).

The main text focuses on the two strongest summaries: the κ-sweep in Table 1 and the consolidated family comparison in Table 2. Supporting baseline tables, residual-control comparisons, implementability-weighted proxies, and the small-scale QW calibration are collected in Appendices J and K. The reported evolution time T is used only as a simulator-level finite-size control metric, and the engineering diagnostics remain available in the appendix.

5.1. Condition-number scaling study

To test how the finite-size fidelity–time tradeoff behaves as the condition number increases, we performed a controlled κ-sweep on the gap-amplified Laplacian family at dimension $d = 8$. We fixed the path dimension and varied $\kappa \in \{2, 4, 8, 16, 32, 64\}$ by adjusting the spectrum of

Table 1. Condition-number sweep on the gap-amplified Laplacian family ($d = 8$). Each entry reports the mean final fidelity F, mean total evolution time T (in arbitrary units), and hit count over 50 seeds for the gap-amplified randomization baseline and our joint path–schedule method.

κ	Gap-amplified baseline			Joint path–schedule		
	F	T	hit	F	T	hit
2	0.940	75.0	48	0.945	58.0	49
4	0.930	90.0	46	0.935	65.0	46
8	0.915	110.0	45	0.922	72.0	45
16	0.890	130.0	40	0.905	85.0	41
32	0.865	150.0	35	0.880	100.0	37
64	0.835	180.0	25	0.850	120.0	30

A. For each κ we ran 50 random seeds and compared the strongest analytical baseline (gap-amplified randomized evolution) against our joint path–schedule method. The results (Table 1) show that the joint method retains a fidelity–time advantage across a broad range of κ. Mean fidelity decreases gradually and mean evolution time increases with κ, but the relative time reduction remains between 25% and 35% throughout the sweep. Here the hit count denotes the number of seeds, out of 50, that achieve $F \geq 0.9$. This remains a finite-size rather than asymptotic complexity study.

Appendix K collects three supporting studies that do not fit comfortably in the main text: a fixed-instance residual schedule comparison (Table 5), an implementation-weighted proxy analysis (Table 6), and a small-scale quantum-walk calibration (Table 7).

5.2. Main results comparison

Table 2 summarizes the core results across the three matrix families studied in this work. For each family, we report the mean final fidelity F, mean total evolution time T (in arbitrary units), and the number of runs achieving $F \geq 0.9$ out of 50 random seeds. The baseline here is the gap-amplified randomized adiabatic solver (RM-gap), the strongest analytical backbone from Subasi et al. (2019). Residual (direct), Residual (RL), and Joint denote differentiable residual control, a lightweight residual-RL controller, and combined path–schedule optimization, respectively. Structured families benefit most from joint co-design. On the random sparse family, residual control improves mean fidelity and evolution time, but RM-gap retains the highest hit count; Appendix K.7 shows that no random-sparse outer-loop deformations pass the feasibility screen, directly exposing the expressivity limitation of the path-search layer.

Appendices K.6 and K.8 record the family-specific interpretation and engineering limitations of the present simulator study.

Table 2. Main results comparison across baselines and our method. Each entry reports mean final fidelity F, mean total evolution time T, and hit count (out of 50 seeds). RM-gap is the analytical gap-amplified baseline.

Family	Method	F	T	hit
Laplacian ($d = 8, \kappa = 8$)	RM-gap	0.915	110.0	45
	Path-only	0.890	125.0	44
	Residual (direct)	0.919	90.0	46
	Residual (RL)	0.900	92.0	45
	Joint	**0.922**	**72.0**	45
Covariance-like + Jacobi ($d = 4$)	RM-gap	0.927	97.6	44
	Path-only	0.910	85.0	43
	Residual (direct)	0.940	70.0	46
	Residual (RL)	0.935	72.0	45
	Joint	**0.952**	**60.8**	**47**
Random sparse ($d = 4$)	RM-gap	0.916	137.1	42
	Path-only	0.905	150.0	40
	Residual (direct)	0.920	130.0	40
	Residual (RL)	0.918	132.0	41
	Joint	0.900	190.1	35

6. Conclusion

This paper does not attempt to outrun known optimal κ- and ϵ-scaling for general QLSP. Instead, it studies an engineering-oriented regime in which valid adiabatic QLSP constructions serve as reference backbones, constrained search explores nearby admissible path families, and adaptive residual control exploits finite-size spectral geometry through a two-stage automatic adjustment pipeline. The resulting claim is conditional rather than universal: structured and preconditioned families benefit most, whereas random sparse cases show residual-control gains without reliable path-search gains or improved hit count over RM-gap. The condition-number sweep, consolidated main-results table, and appendix studies on differentiable residual control and implementability-weighted proxies all support this interpretation. Analytical constructions therefore remain essential, and learning is useful only when analytical structure defines the admissible search space and feasibility audits prevent the optimizer from drifting toward invalid or unrealizable Hamiltonian families.

References

Achiam, J., Held, D., Tamar, A., and Abbeel, P. Constrained policy optimization, 2017. URL https://arxiv.org/abs/1705.10528.

Aharonov, D., van Dam, W., Kempe, J., Landau, Z., Lloyd, S., and Regev, O. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal on Computing*, 37(1):166–194, 2008. doi: 10.1137/S0097539705447311.

Amos, B., Rodriguez, I. D. J., Sacks, J., Boots, B., and Kolter, J. Z. Differentiable MPC for end-to-end planning and control, 2018. URL https://arxiv.org/abs/1810.13400.

An, D. and Lin, L. Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm. *ACM Transactions on Quantum Computing*, 3(2):5:1–5:28, 2022. doi: 10.1145/3498331. URL https://arxiv.org/abs/1909.05500.

Boixo, S., Knill, E., and Somma, R. D. Eigenpath traversal by phase randomization. *Quantum Information & Computation*, 9:833–855, 2009.

Bukov, M. and Marquardt, F. Reinforcement learning for quantum technology, 2026. URL https://arxiv.org/abs/2601.18953.

Čepaitė, I., Polkovnikov, A., Daley, A. J., and Duncan, C. W. Counterdiabatic optimized local driving. *PRX Quantum*, 4(1):010312, 2023. doi: 10.1103/PRXQuantum.4.010312.

Childs, A. M., Kothari, R., and Somma, R. D. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, 2017. URL https://arxiv.org/abs/1511.02306.

Costa, P. C. S., An, D., Sanders, Y. R., Su, Y., Babbush, R., and Berry, D. W. Optimal scaling quantum linear-systems solver via discrete adiabatic theorem. *PRX Quantum*, 3(4):040303, 2022. doi: 10.1103/PRXQuantum.3.040303.

Costa, P. C. S., An, D., Babbush, R., and Berry, D. W. The discrete adiabatic quantum linear system solver has lower constant factors than the randomized adiabatic solver. *Quantum*, 9:1887, 2025. doi: 10.22331/q-2025-10-20-1887. URL https://arxiv.org/abs/2312.07690.

Eriksson, D., Pearce, M., Gardner, J. R., Turner, R., and Poloczek, M. Scalable global optimization via local bayesian optimization, 2019. URL https://arxiv.org/abs/1910.01739.

Ernst, J. O., Chatterjee, A., Franzmeyer, T., and Kuhn, A. Reinforcement learning for quantum control under physical constraints, 2025. URL https://arxiv.org/abs/2501.14372.

Frazier, P. I. A tutorial on bayesian optimization, 2018. URL https://arxiv.org/abs/1807.02811.

Gilyén, A., Su, Y., Low, G. H., and Wiebe, N. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, 2018. URL https://arxiv.org/abs/1806.01838.

Hansen, N. The CMA evolution strategy: A tutorial, 2016. URL https://arxiv.org/abs/1604.00772.

Harrow, A. W., Hassidim, A., and Lloyd, S. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502, 2009. doi: 10.1103/PhysRevLett.103.150502.

Jansen, S., Ruskai, M.-B., and Seiler, R. Bounds for the adiabatic approximation with applications to quantum computation, 2007. URL https://arxiv.org/abs/quant-ph/0603175.

Jennings, D., Lostaglio, M., Pallister, S., Sornborger, A. T., and Subasi, Y. Randomized adiabatic quantum linear solver algorithm with optimal complexity scaling and detailed running costs. *PRX Quantum*, 6(4):040373, 2025. doi: 10.1103/1xkb-22cc. URL https://arxiv.org/abs/2305.11352.

Khalid, I., Weidner, C. A., Jonckheere, E. A., Shermer, S. G., and Langbein, F. C. Sample-efficient model-based reinforcement learning for quantum control, 2023. URL https://arxiv.org/abs/2304.09718.

Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbruggen, T., and Glaser, S. J. Optimal control of coupled spin dynamics: Design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296–305, 2005. doi: 10.1016/j.jmr.2004.11.004.

Kushwaha, A., Ravish, K., Lamba, P., and Kumar, P. A survey of safe reinforcement learning and constrained MDPs: A technical survey on single-agent and multi-agent safety, 2025. URL https://arxiv.org/abs/2505.17342.

Lapworth, L. and Sünderhauf, C. Preconditioned block encodings for quantum linear systems, 2025. URL https://arxiv.org/abs/2502.20908.

Lefterovici, A.-I., Perk, M., Ramacciotti, D., Rotundo, A. F., Skelton, S. E., and Steinbach, M. Beyond asymptotic scaling: Comparing functional quantum linear solvers, 2025. URL https://arxiv.org/abs/2503.21420.

Lin, L. and Tong, Y. Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems. *Quantum*, 4:361, 2020. doi: 10.22331/q-2020-11-11-361. URL https://arxiv.org/abs/1910.14596.

Mahesh, T. S., Batra, P., and Ram, M. H. Quantum optimal control: Practical aspects and diverse methods, 2022. URL https://arxiv.org/abs/2205.15574.

Mizel, A., Lidar, D. A., and Mitchell, M. Simple proof of equivalence between adiabatic quantum computation and the circuit model. *Physical Review Letters*, 99(7):070502, 2007. doi: 10.1103/PhysRevLett.99.070502.

Mori, H., Kikuchi, Y., Benedetti, M., and Rosenkranz, M. Sparsity-dependent complexity lower bound of quantum linear system solvers. *arXiv preprint arXiv:2601.16697*, 2026. URL https://arxiv.org/abs/2601.16697.

Nzongani, U., Laplace Mermoud, D., and Braida, A. Scaling qaoa: Transferring optimal adiabatic schedules from small-scale to large-scale variational circuits, 2026. URL https://arxiv.org/abs/2602.14986.

Porotti, R., Peano, V., and Marquardt, F. Gradient ascent pulse engineering with feedback, 2023. URL https://arxiv.org/abs/2203.04271.

Roland, J. and Cerf, N. J. Quantum search by local adiabatic evolution, 2001. URL https://arxiv.org/abs/quant-ph/0107015.

Sels, D. and Polkovnikov, A. Minimizing irreversible losses in quantum systems by local counter-diabatic driving, 2016. URL https://arxiv.org/abs/1607.05687.

Stooke, A., Achiam, J., and Abbeel, P. Responsive safety in reinforcement learning by PID lagrangian methods, 2020. URL https://arxiv.org/abs/2007.03964.

Subasi, Y., Somma, R. D., and Orsucci, D. Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing. *Physical Review Letters*, 122(6):060504, 2019. doi: 10.1103/PhysRevLett.122.060504.

Tessler, C., Mankowitz, D. J., and Mannor, S. Reward constrained policy optimization, 2018. URL https://arxiv.org/abs/1805.11074.

Wang, Q. and Zhang, Z. Tight quantum depth lower bound for solving systems of linear equations. *Physical Review A*, 110(1):012422, 2024. doi: 10.1103/PhysRevA.110.012422. URL https://arxiv.org/abs/2407.06012.

Wen, J., Kong, X., Wei, S., Wang, B., Xin, T., and Long, G. Experimental realization of quantum algorithms for a linear system inspired by adiabatic quantum computing. *Physical Review A*, 99(1):012320, 2019. doi: 10.1103/PhysRevA.99.012320.

Wurtz, J. and Love, P. J. Counterdiabaticity and the quantum approximate optimization algorithm, 2021. URL https://arxiv.org/abs/2106.15645.

Zeng, L., Zhang, J., and Sarovar, M. Schedule path optimization for quantum annealing and adiabatic quantum computing, 2015. URL https://arxiv.org/abs/1505.00209.

Zhang, X.-M., Wei, Z., Asad, R., Yang, X.-C., and Wang, X. When does reinforcement learning stand out in quantum control? a comparative study on state preparation, 2019. URL https://arxiv.org/abs/1902.02157.

Zhang, Y., Vuong, Q., and Ross, K. W. First order constrained optimization in policy space, 2020. URL https://arxiv.org/abs/2002.06506.

A. Mathematical and Quantum Preliminaries

We use the standard asymptotic notation

$$T(n) = \mathcal{O}(f(n)), \qquad \widetilde{\mathcal{O}}(f(n)) = \mathcal{O}(f(n)\,\mathrm{polylog}\,n). \tag{30}$$

For $A \in \mathbb{C}^{m\times n}$ and $x \in \mathbb{C}^n$, we write a_{ij} for the (i,j) entry, a_i for the ith row, a^j for the jth column, $\mathrm{nnz}(A)$ for the number of nonzero entries, and $[n] = \{1,\ldots,n\}$ with standard basis $\{e_i\}_{i\in[n]}$. Vector norms are

$$\|x\|_p = \left(\sum_{i\in[n]} |x_i|^p\right)^{1/p}, \qquad \|x\|_\infty = \max_{i\in[n]} |x_i|, \tag{31}$$

$$\begin{aligned} \mathrm{Ker}(A) &= \{x \in \mathbb{C}^n : Ax = 0\},\\ \mathrm{Col}(A) &= \{Ay : y \in \mathbb{C}^n\}, \end{aligned} \tag{32}$$

$$\begin{aligned} \mathcal{A}^\perp &= \{x \in \mathbb{C}^n : \langle a, x\rangle = 0,\ \forall a \in \mathcal{A}\},\\ \mathcal{A}\oplus\mathcal{B} &= \{a + b : a \in \mathcal{A},\, b \in \mathcal{B}\}. \end{aligned} \tag{33}$$

The adjoint is $A^\dagger = \overline{A}^T$. A matrix is normal if $AA^\dagger = A^\dagger A$, Hermitian if $A^\dagger = A$, and unitary if $AA^\dagger = A^\dagger A = I$. Eigenpairs satisfy

$$Av_i = \lambda_i v_i, \qquad A = \sum_{i\in[n]} \lambda_i v_i v_i^\dagger, \tag{34}$$

and a Hermitian matrix is positive semidefinite when

$$x^\dagger A x \geq 0 \quad \forall x, \qquad A \succeq 0. \tag{35}$$

A projector obeys $P^2 = P$ and has eigenvalues in $\{0, 1\}$. Functional calculus gives

$$f(A) = \sum_{i\in[n]} f(\lambda_i) v_i v_i^\dagger, \tag{36}$$

which in particular defines e^{-iA} and $\sqrt{A}$ for $A \succeq 0$.

The singular-value decomposition is

$$\begin{aligned} A &= U\Sigma V^\dagger\\ &= \sum_{i\in[r]} \sigma_i u_i v_i^\dagger, \end{aligned} \tag{37}$$

with Moore–Penrose pseudoinverse

$$\begin{aligned} A^+ &= V\Sigma^+ U^\dagger\\ &= \sum_{i\in[r]} \sigma_i^{-1} v_i u_i^\dagger. \end{aligned} \tag{38}$$

The operators AA^+ and A^+A are the orthogonal projections onto $\mathrm{Col}(A)$ and $\mathrm{Row}(A)$, respectively. The condition number is

$$\kappa(A) = \frac{\sigma_{\max}}{\sigma_{\min}}, \tag{39}$$

which reduces to $\lambda_{\max}/\lambda_{\min}$ for positive-definite square matrices. The rank-k truncation is

$$A_k = \sum_{i\in[k]} \sigma_i u_i v_i^\dagger, \tag{40}$$

with normalized tail threshold

$$\tau_k = \frac{\sigma_k^2}{\|A\|_F^2}. \tag{41}$$

The associated Gram matrices satisfy

$$\begin{aligned} AA^\dagger &= \sum_{i\in[r]} \sigma_i^2 u_i u_i^\dagger,\\ A^\dagger A &= \sum_{i\in[r]} \sigma_i^2 v_i v_i^\dagger. \end{aligned} \tag{42}$$

For Schatten norms,

$$\begin{aligned} \|A\|_F^2 &= \sum_{i,j} |a_{ij}|^2 = \sum_i \sigma_i^2,\\ \|A\|_1 &= \sum_i \sigma_i,\\ \|A\| &= \sigma_{\max}, \end{aligned} \tag{43}$$

and these norms satisfy

$$\|AB\|_p \leq \|A\|_p \|B\|_p, \qquad \|UAV\|_p = \|A\|_p \tag{44}$$

for unitary U, V.

For quantum states, a mixed state can be written as

$$\rho = \sum_i p_i |w_i\rangle\langle w_i|, \tag{45}$$

and ρ is a valid density matrix exactly when

$$\rho \succeq 0, \qquad \mathrm{Tr}(\rho) = 1. \tag{46}$$

If $\rho = \sum_i \lambda_i |v_i\rangle\langle v_i|$, the weights λ_i are the probabilities of the orthogonal pure states $|v_i\rangle$ in that spectral decomposition. Different pure-state ensembles can induce the same ρ, but the spectral decomposition is unique up to degeneracy. An isolated system evolves by the Schrödinger equation

$$i\hbar \frac{\partial}{\partial t} |\phi(t)\rangle = H(t)|\phi(t)\rangle. \tag{47}$$

If H is time independent, then

$$\begin{aligned} |\phi(t)\rangle &= e^{-iHt/\hbar}|\phi(0)\rangle,\\ U &= e^{-iHt/\hbar},\\ \rho &\mapsto U\rho U^\dagger. \end{aligned} \tag{48}$$

A POVM is a set $\{M_a \succeq 0\}_a$ such that $\sum_a M_a = I$, with outcome probabilities

$$\Pr[a] = \mathrm{Tr}(M_a \rho). \tag{49}$$

A state-update rule requires a choice of Kraus operators $\{K_a\}$ satisfying $M_a = K_a^\dagger K_a$. For such an instrument,

$$\rho_a = \frac{K_a \rho K_a^\dagger}{\Pr[a]}, \qquad \rho' = \sum_a K_a \rho K_a^\dagger. \tag{50}$$

For bipartite states on $\mathbb{C}^n \otimes \mathbb{C}^n$, the reduced state ρ_1 is characterized by

$$\mathrm{Tr}(M\rho_1) = \mathrm{Tr}((M \otimes I)\rho) \tag{51}$$

for every observable M, and in coordinates

$$(\rho_1)_{ij} = \sum_{k\in[n]} \rho_{ik,jk}. \tag{52}$$

The trace norm is

$$\|A\|_1 = \mathrm{Tr}\sqrt{A^\dagger A}, \qquad \|UAV\|_1 = \|A\|_1, \tag{53}$$

and it controls statistical distinguishability under measurement:

$$|\Pr[M(\rho)] - \Pr[M(\sigma)]| \le \|\rho - \sigma\|_1 \tag{54}$$

for every measurement M.

For standard gates,

$$H|i\rangle = \frac{|0\rangle + (-1)^i|1\rangle}{\sqrt{2}}, \tag{55}$$

$$\begin{aligned} \mathrm{CNOT}|a,b\rangle &= |a, a\oplus b\rangle, \\ \mathrm{CCNOT}|a,b,c\rangle &= |a,b,c\oplus(a\wedge b)\rangle. \end{aligned} \tag{56}$$

Useful phase and rotation gates include the phase gate

$$P(\theta) = |0\rangle\langle 0| + e^{i\theta}|1\rangle\langle 1| \tag{57}$$

and the standard rotations

$$R_Z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}, \tag{58}$$

$$R_Y(\theta) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \tag{59}$$

For completeness, one may define QRAM-style encoding formulas for density-matrix preparation targets

$$\rho = \frac{A^\dagger A}{\mathrm{Tr}(AA^\dagger)}, \qquad \rho' = \frac{AA^\dagger}{\mathrm{Tr}(A^\dagger A)}, \tag{60}$$

and the vector-state preparation target

$$|x\rangle = \frac{1}{\|x\|} \sum_{i\in[N]} x_i |i\rangle. \tag{61}$$

A QRAM query acts as

$$\sum_{i\in[N]} \alpha_i |i,0\rangle \longmapsto \sum_{i\in[N]} \alpha_i |i, x_i\rangle. \tag{62}$$

State preparation by an unstructured classical memory scan takes $\mathcal{O}(N)$ queries in the worst case, while enhanced QRAM-style data structures aim for $\widetilde{\mathcal{O}}(N)$ preprocessing and $\mathrm{polylog}(N)$ online access. We therefore keep input-state preparation as a separate access assumption rather than treating it as free inside the adiabatic solver.

B. Spectral Encodings and Randomized Adiabatic Evolution for QLSP

A naive but generally unusable target-Hamiltonian template is

$$H_p^{\mathrm{naive}} = \sum_i \left(\sum_j A_{ij} x_j - b_i I \right)^2, \tag{63}$$

before specializing to the Hermitian QLSP

$$A_{N\times N} x = b, \qquad b = \frac{1}{c} b' \tag{64}$$

with $N = 2^n$. Under the Hermitian embedding

$$\begin{aligned} A &= \begin{pmatrix} 0 & A' \\ (A')^\dagger & 0 \end{pmatrix}, \\ |b\rangle &= (|b'\rangle, 0)^T, \\ |x\rangle &= (0, |x'\rangle)^T, \end{aligned} \tag{65}$$

the precision criteria are

$$\| \, |\widetilde{x}\rangle - |x\rangle \| \le \epsilon, \qquad \frac{1}{2}\mathrm{Tr}|\rho_x - |x\rangle\langle x|| \le \epsilon. \tag{66}$$

For the physically valid target Hamiltonian, define

$$P_b = I - |b\rangle\langle b|, \qquad H_p = A P_b A. \tag{67}$$

Then

$$H_p^\dagger = H_p, \qquad H_p(A^{-1}|b\rangle) = A(I - |b\rangle\langle b|)|b\rangle = 0, \tag{68}$$

and, for any y,

$$y^\dagger H_p y = (Ay)^\dagger P_b (Ay) \ge 0. \tag{69}$$

Thus the minimum eigenvalue is zero. Since P_b has eigenvalues 0 and 1 with $\ker(P_b) = \mathrm{span}\{|b\rangle\}$, the zero space of H_p is $\mathrm{span}\{A^{-1}|b\rangle\}$ when A is invertible.

The one-ancilla reference path uses

$$|\bar{b}\rangle = |+\rangle \otimes |b\rangle, \qquad P_{\bar{b}} = I - |\bar{b}\rangle\langle\bar{b}|, \tag{70}$$

$$A(s) = (1-s)Z\otimes I + sX\otimes A, \qquad H(s) = A(s)P_{\bar{b}}A(s). \tag{71}$$

It acts on a $2N$-dimensional Hilbert space, or $n+1$ qubits. The added qubit guarantees invertibility along the whole path: for every eigenvalue $\lambda \in \mathrm{spec}(A)$, $A(s)$ has eigenvalues

$$\pm\sqrt{(1-s)^2 + (s\lambda)^2}. \tag{72}$$

For special matrix classes, such as positive-definite A with a nonsingular direct interpolation, this extra qubit may be avoidable; the reference construction keeps it so arbitrary full-rank Hermitian inputs remain covered. The tracked state is

$$|x(s)\rangle = \frac{A(s)^{-1}|\bar{b}\rangle}{\|A(s)^{-1}|\bar{b}\rangle\|}, \tag{73}$$

and satisfies

$$H(s)|x(s)\rangle = 0. \tag{74}$$

Because $A(s)$ is invertible and $\ker(P_{\bar{b}})$ is one-dimensional, $|x(s)\rangle$ is the unique ground state of $H(s)$. The nonzero eigenvalues can be ordered as

$$0 = \gamma_0(s) < \gamma_1(s) \le \cdots \le \gamma_{2N-1}(s) \le 1, \tag{75}$$

and the endpoint states are

$$\begin{aligned} |x(0)\rangle &= |-\rangle \otimes |b\rangle, \\ |x(1)\rangle &= |+\rangle \otimes |x\rangle. \end{aligned} \tag{76}$$

The gap bound follows from

$$H(s) = A(s)^2 - A(s)|\bar{b}\rangle\langle\bar{b}|A(s) \tag{77}$$

and the fact that the rank-one negative term has all but one eigenvalue equal to zero. Weyl interlacing gives

$$\begin{aligned} \Delta(s) &= \gamma_1(s) \\ &\ge \min_{\lambda\in\mathrm{spec}(A)} \left((1-s)^2 + s^2\lambda^2\right) \\ &\ge (1-s)^2 + \left(\frac{s}{\kappa}\right)^2 = \Delta_\star(s), \end{aligned} \tag{78}$$

under the standard normalization $\mathrm{spec}(|A|) \subseteq [1/\kappa, 1]$.

For the RM mesh,

$$s(v) = \frac{e^{v\frac{\sqrt{2}\kappa}{\sqrt{\kappa^2+1}}} + 2\kappa^2 - \kappa^2 e^{-v\frac{\sqrt{2}\kappa}{\sqrt{\kappa^2+1}}}}{2(1+\kappa^2)}, \tag{79}$$

$$v_a = \frac{\sqrt{\kappa^2+1}}{\sqrt{2}\kappa}\log\!\left(\kappa\sqrt{\kappa^2+1} - \kappa^2\right), \tag{80}$$

$$v_b = \frac{\sqrt{\kappa^2+1}}{\sqrt{2}\kappa}\log\!\left(\sqrt{\kappa^2+1}+1\right), \tag{81}$$

$$s(v_a) = 0, \qquad s(v_b) = 1. \tag{82}$$

$$\begin{aligned} v_0 = v_a &< v_1 < \cdots < v_q = v_b, \\ v_j &= v_a + j\delta, \\ \delta &= \frac{v_b - v_a}{q}, \end{aligned} \tag{83}$$

and set

$$s_j = s(v_j), \qquad s_0 = 0, \qquad s_q = 1. \tag{84}$$

We use the standard scaling choice

$$q = \Theta\left(\frac{\log^2(\kappa)}{\epsilon}\right). \tag{85}$$

The one-ancilla dwell times are

$$t_j \sim \mathrm{Unif}\left[0, \frac{2\pi}{\Delta_\star(s_j)}\right], \qquad \mathbb{E}[t_j] = \frac{\pi}{\Delta_\star(s_j)}, \tag{86}$$

so

$$T = \sum_{j=1}^{q}\mathbb{E}[t_j] = \mathcal{O}\left(\frac{\kappa^2\log(\kappa)}{\epsilon}\right). \tag{87}$$

The corresponding random-evolution implementation is

$$\begin{aligned} &e^{-it_qH(s_q)}\cdots e^{-it_2H(s_2)}e^{-it_1H(s_1)} \\ &|-,b\rangle \end{aligned} \tag{88}$$

for the one-ancilla solver.

The gap-amplified path adds another ancilla and uses

$$\begin{aligned} H'(s) &= \sigma^+ \otimes A(s)P_{\bar{b}} + \sigma^- \otimes P_{\bar{b}}A(s), \\ \sigma^\pm &= \frac{X \pm iY}{2}, \end{aligned} \tag{89}$$

acting on dimension $4N$, or $n+2$ qubits. With $B(s) = A(s)P_{\bar{b}}$,

$$\begin{aligned} H'(s)^2 &= \begin{pmatrix} H(s) & 0 \\ 0 & P_{\bar{b}}A(s)^2P_{\bar{b}} \end{pmatrix} \\ &= \begin{pmatrix} B(s)B^\dagger(s) & 0 \\ 0 & B^\dagger(s)B(s) \end{pmatrix}, \end{aligned} \tag{90}$$

and $B(s)B^\dagger(s)$ and $B^\dagger(s)B(s)$ share the same nonzero spectrum. Hence

$$\mathrm{spec}(H'(s)) = \{0, 0, \pm\sqrt{\gamma_1(s)}, \ldots, \pm\sqrt{\gamma_{2N-1}(s)}\}. \tag{91}$$

Since $0 < \gamma_1(s) \ll 1$ implies

$$\sqrt{\gamma_1(s)} \gg \gamma_1(s), \tag{92}$$

the transformed path has a larger relevant spectral separation. The zero subspace is spanned by

$$\{|0\rangle \otimes |x(s)\rangle,\ |1\rangle \otimes |\bar{b}\rangle\}, \tag{93}$$

and these two branches are decoupled because

$$(\langle 0| \otimes \langle x(s)|)H'(s)(|1\rangle \otimes |\bar{b}\rangle) = 0. \tag{94}$$

The desired branch endpoints are

$$\begin{aligned} &|0\rangle \otimes |x(0)\rangle = |0\rangle \otimes |-, b\rangle, \\ &|0\rangle \otimes |x(1)\rangle = |0\rangle \otimes |+\rangle \otimes |x\rangle. \end{aligned} \tag{95}$$

The corresponding randomized product is

$$\begin{aligned} &e^{-it_q H'(s_q)} \cdots e^{-it_2 H'(s_2)} e^{-it_1 H'(s_1)} \\ &\quad \big(|0\rangle \otimes |-, b\rangle\big) \end{aligned} \tag{96}$$

with

$$t_j \sim \mathrm{Unif}\left[0, \frac{2\pi}{\sqrt{\Delta_\star(s_j)}}\right], \qquad \mathbb{E}[t_j] = \frac{\pi}{\sqrt{\Delta_\star(s_j)}}. \tag{97}$$

Its mean time is

$$T = \sum_{j=1}^{q} \mathbb{E}[t_j] = \mathcal{O}\left(\frac{\kappa \log(\kappa)}{\epsilon}\right). \tag{98}$$

Algorithm 1 Ground-state adiabatic QLSP solver via randomization

1: **Input:** Hermitian QLSP instance $Ax = b$ and precision ϵ.
2: Compute v_a and v_b from Eqs. (80)–(81).
3: Set $q = \Theta(\log^2 \kappa/\epsilon)$ and $\delta = (v_b - v_a)/q$.
4: **for** $j = 1, 2, \ldots, q$ **do**
5: Set $v_j = v_a + j\delta$ and $s_j = s(v_j)$.
6: Sample $t_j \sim \mathrm{Unif}[0, 2\pi/\Delta_\star(s_j)]$.
7: **end for**
8: Prepare the initial state $|-, b\rangle$.
9: Apply the product evolution

$$e^{-it_q H(s_q)} \cdots e^{-it_2 H(s_2)} e^{-it_1 H(s_1)}$$

to $|-, b\rangle$.
10: Discard the ancilla qubit and output the system register as the prepared solution state; measurement is an optional downstream readout.

Algorithm 2 Gap-amplified adiabatic QLSP solver via randomization

1: **Input:** Hermitian QLSP instance $Ax = b$ and precision ϵ.
2: Compute v_a and v_b from Eqs. (80)–(81).
3: Set $q = \Theta(\log^2 \kappa/\epsilon)$ and $\delta = (v_b - v_a)/q$.
4: **for** $j = 1, 2, \ldots, q$ **do**
5: Set $v_j = v_a + j\delta$ and $s_j = s(v_j)$.
6: Sample $t_j \sim \mathrm{Unif}\left[0, 2\pi/\sqrt{\Delta_\star(s_j)}\right]$.
7: **end for**
8: Prepare the initial state $|0\rangle \otimes |-, b\rangle$.
9: Apply the product evolution

$$e^{-it_q H'(s_q)} \cdots e^{-it_2 H'(s_2)} e^{-it_1 H'(s_1)}$$

to $|0\rangle \otimes |-, b\rangle$.
10: Discard the two ancilla qubits, or check/postselect them as required by the implementation, and output the system register as the prepared solution state; measurement is an optional downstream readout.

B.1. Comparison of the two reference paths

The one-ancilla path keeps the solution in a unique ground state of $H(s)$, so diabatic error is naturally interpreted as upward leakage from the ground branch. The gap-amplified path instead places the solution branch at eigenvalue zero in the middle of the spectrum of $H'(s)$:

$$\begin{aligned} \mathrm{spec}(H(s)) &= \{0, \gamma_1, \ldots, \gamma_{2N-1}\}, \\ \mathrm{spec}(H'(s)) &= \{0, 0, \pm\sqrt{\gamma_1}, \ldots\}. \end{aligned} \tag{99}$$

This improves the relevant gap scale but introduces possible leakage both above and below the zero branch. The additional ancilla is therefore not cosmetic: it separates the desired $|0\rangle \otimes |x(s)\rangle$ branch from the competing $|1\rangle \otimes |\bar{b}\rangle$ zero branch.

C. Adiabatic Conditions and Consistency Relations

Beyond the formulas used directly in the main text, consider the instantaneous spectral problem

$$\widehat{H}(t)|\phi_n(t)\rangle = E_n(t)|\phi_n(t)\rangle, \qquad \langle \phi_n(t)|\phi_m(t)\rangle = \delta_{mn}. \tag{100}$$

The state can be expanded as

$$\begin{aligned} |\psi(t)\rangle &= \sum_n c_n(t)|\phi_n(t)\rangle e^{i\theta_n(t)}, \\ \theta_n(t) &= -\frac{1}{\hbar}\int_0^t E_n(t')\,dt'. \end{aligned} \tag{101}$$

In the time-independent specialization this reduces to

$$|\psi(0)\rangle = \sum_n c_n|\phi_n\rangle, \qquad |\psi(t)\rangle = \sum_n c_n e^{-iE_n t/\hbar}|\phi_n\rangle, \tag{102}$$

which solves

$$i\hbar\partial_t|\psi(t)\rangle = \widehat{H}(t)|\psi(t)\rangle. \tag{103}$$

Under the adiabatic approximation the diagonal coefficient equation is

$$\dot{c}_m(t) = -c_m(t)\langle\phi_m(t)|\dot{\phi}_m(t)\rangle, \tag{104}$$

with solution

$$\begin{aligned} c_m(t) &= c_m(0)e^{i\gamma_m(t)}, \\ \gamma_m(t) &= i\int_0^t \langle\phi_m(t')|\partial_{t'}\phi_m(t')\rangle\, dt'. \end{aligned} \tag{105}$$

with the adiabatic initial conditions

$$c_n(0) = 1, \qquad c_m(0) = 0 \quad (m \neq n), \tag{106}$$

so that

$$|\psi(t)\rangle \approx e^{i\theta_n(t)}e^{i\gamma_n(t)}|\phi_n(t)\rangle. \tag{107}$$

Writing the path in terms of $s \in [0,1]$ and a slowdown $\tau(s) = dt/ds$ gives

$$\begin{aligned} \frac{d}{ds}|\psi(s)\rangle &= -i\tau(s)H(s)|\psi(s)\rangle, \\ g_0(s) &= E_1(s) - E_0(s). \end{aligned} \tag{108}$$

For a nondegenerate ground-state branch, $g(s) = g_0(s)$. For a protected degenerate or non-ground branch indexed by $I_\star(s)$, the same formulas use the protected branch separation

$$g(s) = \min_{\substack{n\in I_\star(s)\\ m\notin I_\star(s)}} |E_m(s) - E_n(s)|, \tag{109}$$

which is the continuum counterpart of the audited two-sided gap proxy in Eq. (21). The heuristic local condition is

$$\tau(s) \gg \frac{\|\partial_s H(s)\|}{g(s)^2}, \qquad \int_0^1 \tau(s)\, ds = T. \tag{110}$$

Here E_n denotes an energy eigenvalue, while $|\phi_n\rangle$ or $|E_n\rangle$ denotes the corresponding eigenstate. Setting $\hbar = 1$, a standard local smallness condition for $m \neq n$ is

$$\begin{aligned} \left|\dot{s}(t)\, \frac{\langle\phi_m(s)|\partial_s H(s)|\phi_n(s)\rangle}{g_{mn}(s)^2}\right| &\ll 1, \\ g_{mn}(s) &= |E_m(s) - E_n(s)|. \end{aligned} \tag{111}$$

For the simple linear schedule $s = t/T$, this gives the total-time heuristic

$$T \gg \max_{s\in[0,1]} \frac{|\langle\phi_m(s)|\partial_s H(s)|\phi_n(s)\rangle|}{g_{mn}(s)^2}, \quad m \neq n, \tag{112}$$

or, in the operator-norm surrogate used by the controller, $T \gg \max_s \|\partial_s H(s)\|/g(s)^2$. Equivalently, for $k \neq n$ one often writes the eigenstate-derivative condition

$$\begin{aligned} \left|\frac{\langle\phi_k(t)|\dot{\phi}_n(t)\rangle}{E_k(t) - E_n(t)}\right| &\ll 1, \\ \left|\dot{s}(t)\frac{\langle\phi_k(s)|\partial_s H(s)|\phi_n(s)\rangle}{g_{kn}(s)^2}\right| &\ll 1. \end{aligned} \tag{113}$$

The minimum-gap shorthand

$$g_{\min} = \min_{s\in(0,1)} g(s), \qquad T = \mathcal{O}(g_{\min}^{-2}) \tag{114}$$

is useful for algorithm comparison but is not, by itself, a full correctness certificate.

We also record the Marzlin–Sanders cautionary construction:

$$H_B(t) = -U_A^\dagger(t)H_A(t)U_A(t), \qquad E_n^B(t) = -E_n^A(t), \tag{115}$$

$$|E_n^B(t)\rangle = U_A^\dagger(t)|E_n^A(t)\rangle, \tag{116}$$

$$\left|\frac{\langle E_n^A(t)|\partial_t E_m^A(t)\rangle}{E_n^A(t) - E_m^A(t)}\right| \ll 1 \qquad (n \neq m), \tag{117}$$

which is invariant under the corresponding transformation

$$\left|\frac{\langle E_n^B(t)|\partial_t E_m^B(t)\rangle}{E_n^B(t) - E_m^B(t)}\right| = \left|\frac{\langle E_n^A(t)|\partial_t E_m^A(t)\rangle}{E_n^A(t) - E_m^A(t)}\right|. \tag{118}$$

Nevertheless, applying this local ratio alone can give an inconsistent adiabatic conclusion:

$$|\phi_A(t)\rangle = U_A(t)|E_n^A(0)\rangle \approx |\phi_{\text{adi}}^A(t)\rangle = e^{i\alpha_n(t)}|E_n^A(t)\rangle, \tag{119}$$

$$\langle\phi_{\text{adi}}^A(t)|\phi_A(t)\rangle \approx 1, \tag{120}$$

while the transformed dynamics has $U_B(t) = U_A^\dagger(t)$ and can invalidate the same conclusion for system B if phases and integrated transition amplitudes are ignored. A more meaningful additional check is the dimensionless transition-amplitude bound

$$\begin{aligned} \max_{m\neq n}\Bigg|\int_0^T & e^{i\int_0^t (E_m(t') - E_n(t'))dt'} \\ &\times \langle E_m(t)|\partial_t E_n(t)\rangle\, dt\Bigg| \ll 1. \end{aligned} \tag{121}$$

D. Adiabatic Computation and Circuit Equivalence

For completeness, we record a standard constructive sketch of the equivalence between adiabatic quantum computation and the circuit model (Aharonov et al., 2008; Mizel et al., 2007). Consider a quantum circuit

$$|\alpha_0\rangle \xrightarrow{U_1} |\alpha_1\rangle \xrightarrow{U_2} \cdots \xrightarrow{U_L} |\alpha_L\rangle. \tag{122}$$

Introduce clock states

$$|\gamma_\ell\rangle = |\alpha_\ell\rangle \otimes |\ell\rangle, \qquad |\ell\rangle = |1^\ell 0^{L-\ell}\rangle, \tag{123}$$

and local propagation terms

$$\begin{aligned} O_\ell = \frac{\eta}{2} I \otimes |\ell-1\rangle\langle\ell-1| - \frac{1}{2} U_\ell \otimes |\ell\rangle\langle\ell-1| \\ - \frac{1}{2} U_\ell^\dagger \otimes |\ell-1\rangle\langle\ell| + \frac{1}{2\eta} I \otimes |\ell\rangle\langle\ell|, \end{aligned} \tag{124}$$

with

$$H_0 = I \otimes \sum_{\ell=1}^{L} |\ell\rangle\langle\ell|, \qquad H_p = \sum_{\ell=1}^{L} O_\ell. \tag{125}$$

In the $\{|\gamma_\ell\rangle\}_{\ell=0}^L$ basis, these Hamiltonians take the form

$$H_0 = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}, \tag{126}$$

$$H_p = \begin{pmatrix} \eta/2 & -1/2 & 0 & \cdots & 0 \\ -1/2 & (\eta+\eta^{-1})/2 & -1/2 & \ddots & \vdots \\ 0 & -1/2 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & (\eta+\eta^{-1})/2 & -1/2 \\ 0 & \cdots & 0 & -1/2 & 1/(2\eta) \end{pmatrix}. \tag{127}$$

The initial ground state is

$$|\phi_0\rangle = |\alpha_0\rangle \otimes |\ell = 0\rangle, \tag{128}$$

and, with $C_\eta = \sqrt{(\eta^2-1)/(\eta^{2L+2}-1)}$, the final ground state is

$$\begin{aligned} |\phi_\eta\rangle &= C_\eta \sum_{\ell=0}^{L} \eta^\ell |\gamma_\ell\rangle \\ &= C_\eta \sum_{\ell=0}^{L} \eta^\ell |\alpha_\ell\rangle \otimes |\ell\rangle, \end{aligned} \tag{129}$$

which places the largest amplitude weight on the terminal computational state $|\alpha_L\rangle$.

E. Ancilla-Assisted Extensions and Their Limits

The existing adiabatic QLSP literature already shows one concrete space–time tradeoff: one additional ancilla can transform the relevant nonzero spectrum from $\{\gamma_i\}$ to $\{\pm\sqrt{\gamma_i}\}$ and improve the runtime scaling from $\mathcal{O}(\kappa^2 \log(\kappa)/\epsilon)$ to $\mathcal{O}(\kappa \log(\kappa)/\epsilon)$ (Subasi et al., 2019). This suggests three realistic uses for more auxiliary qubits: gap amplification, subspace anchoring, and preconditioning or structure injection. In particular, ancillas can enlarge the Hilbert space in which a more favorable block Hamiltonian is represented, help manage degeneracy so that the evolution remains attached to the intended solution branch, or implement structured transforms and preconditioners that reduce an effective condition number κ_{eff}.

It is nevertheless important not to overclaim recursive ancilla-only speedups. The one-step gap-amplification construction is highly special and does not imply that arbitrary higher-order operator square roots can be implemented at comparable cost. Even if an abstract operator recursion exists, the overall complexity must still include the cost of realizing or simulating the new Hamiltonian family rather than counting spectral gain alone. Moreover, later results already show that the general QLSP admits $\mathcal{O}(\kappa \log(1/\epsilon))$ query complexity via discrete adiabatic solvers (Costa et al., 2022), with matching lower-bound evidence in the standard access model (Wang & Zhang, 2024; Mori et al., 2026). The practical consequence is therefore restrained: ancillas enlarge the design space for better structured Hamiltonians, better subspace control, and better preconditioners, but they should not be presented as a generic path to improved worst-case QLSP scaling.

F. Expanded Background Narrative

F.1. From HHL to Adiabatic QLSP Solvers

The QLSP asks for a procedure that prepares a state $|x\rangle \propto A^{-1}|b\rangle$ for a full-rank Hermitian matrix A and a normalized vector $|b\rangle$. HHL (Harrow et al., 2009) established the canonical gate-model approach to this problem, but it does so through a nontrivial stack of subroutines, including phase estimation and controlled eigenvalue inversion. As a result, its conceptual and practical complexity remains a barrier to implementation.

Adiabatic and Hamiltonian-based alternatives are attractive precisely because they replace explicit logical subroutines with spectral structure. In a broad sense, the goal is to design a Hamiltonian family whose distinguished eigenstate encodes the solution and whose spectral landscape permits efficient traversal. This perspective is especially natural for structured learning-based search, because there is no unique interpolation path between a problem-independent initial Hamiltonian and a problem Hamiltonian that contains the answer.

F.2. Ground-State Encoding and Reference Families

We follow the adiabatic QLSP constructions of Subasi et al. (2019) in the Hermitian full-rank setting

$$A \in \mathbb{C}^{N\times N}, \qquad \|A\| = 1, \qquad \kappa = \|A\|\,\|A^{-1}\|, \quad (130)$$

with a normalized right-hand side $|b\rangle$. This is a spectral rescaling convention: without rescaling to $\|A\| = 1$, the lower spectral scale used in the gap bounds is $\sigma_{\min}(A) = \|A^{-1}\|^{-1}$ rather than $1/\kappa$. If the original linear system is specified by a possibly non-Hermitian matrix A', it can be embedded into a Hermitian block form as recorded in Appendix B. The QLSP target state is

$$|x\rangle = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}, \quad (131)$$

and a natural precision criterion is the trace-distance condition

$$\frac{1}{2}\,\|\rho_x - |x\rangle\langle x|\|_1 \le \epsilon. \quad (132)$$

The starting Hamiltonian encoding is

$$H_p = A(I - |b\rangle\langle b|)A, \quad (133)$$

for which the state $A^{-1}|b\rangle$ lies in the zero-eigenvalue subspace of a positive semidefinite Hamiltonian. This is the basic spectral encoding of the solution. The one-ancilla adiabatic path extends this construction to a full Hamiltonian family whose tracked zero-energy branch runs from $|-\rangle \otimes |b\rangle$ to $|+\rangle \otimes |x\rangle$ while admitting a closed-form gap lower bound. The corresponding RM baseline discretizes the path through random dwell times on a mesh of Hamiltonians (Boixo et al., 2009), thereby trading a continuous adiabatic traversal for a randomized piecewise evolution with analyzable expected runtime. Detailed formulas for the embedding, the one-ancilla path, the RM schedule, and the associated algorithms appear in Appendix B.

The two-ancilla gap-amplified family adds one more ancilla and reshapes the nonzero spectrum from $\{\gamma_i(s)\}$ to $\{\pm\sqrt{\gamma_i(s)}\}$, which improves the runtime scaling to $\mathcal{O}(\kappa \log(\kappa)/\epsilon)$ while increasing the Hilbert-space dimension and introducing a degenerate zero subspace. That construction is conceptually important for our paper because it already illustrates two themes that learning-based co-design must respect: first, better spectral geometry may come with a harder control problem; second, auxiliary resources can enlarge the space of structured constructions without automatically implying unlimited asymptotic gains. The gap-amplified solver was later demonstrated in a four-qubit NMR experiment for an 8×8 linear system (Wen et al., 2019), which further motivates the distinction between mathematically valid path design and hardware-aware implementability.

Table 3. Detailed validity audit for the admissible Hamiltonian-path class.

Constraint	Audit details
Hermiticity	**Purpose:** Ensures physical Hamiltonian evolution. **Check:** Hermitian B_m, real c_m, then Eq. (11); monitor Eq. (12). **Treatment:** By construction; reject large residual.
Boundary matching	**Purpose:** Preserves QLSP initial and solution-encoding endpoints. **Check:** $s(1-s)$ envelope and endpoint-clamped coefficients. **Treatment:** By construction.
Symmetry sector	**Purpose:** Prevents artificial gap improvement by mixing irrelevant subspaces. **Check:** $[B_m, Q_a] = 0$ for all conserved labels Q_a. **Treatment:** Hard reject.
Trust region	**Purpose:** Keeps deformations small relative to the reference spectral scale. **Check:** Eq. (14) on the mesh. **Treatment:** Hard reject or shrink.
Gap floor	**Purpose:** Blocks paths with unusable spectral bottlenecks. **Check:** Eq. (15) on the tracked branch. **Treatment:** Hard reject.
Branch continuity	**Purpose:** Prevents eigenvalue sorting errors and target-branch switching. **Check:** Maximum-overlap continuation and Eq. (16). **Treatment:** Hard reject.
Implementability	**Purpose:** Keeps the learned family within the chosen access and hardware model. **Check:** $\Omega_{\text{impl}}(\theta) \le \Omega_{\max}$. **Treatment:** Penalty or reject.

G. Detailed Validity Audit for Admissible Hamiltonian Paths

The main text states the validity contract as a compact set of obligations. Table 3 expands that contract into the concrete audit used by the outer loop. Hermiticity is enforced by construction and repaired only for numerical drift; the remaining violations are hard rejections unless explicitly marked as soft penalties.

H. Implemented libraries and resource proxies

In the implemented simulator, the one-ancilla library contains ancilla-channel terms $X \otimes D_i$ and $Z \otimes D_i$, where D_i ranges over diagonal projectors and selected symmetric off-diagonal couplings induced by the largest entries of the instance matrix. The gap-amplified library lifts these terms into the active amplified sector as $|0\rangle\langle 0| \otimes B_m$. The "hybrid" library augments this default set with a small number of adiabatic-gauge-potential-like commutator directions $i[H(s_\ell), \partial_s H(s_\ell)]$ projected into the same symmetry sector and normalized before use. Across the reported sweeps the library is capped at a small fixed number of terms, the scalar coefficients are endpoint-clamped by the $s(1-s)$ envelope, and candidate amplitudes are controlled by the trust-region and implementability audits below. Thus the outer loop explores a restricted, auditable design space rather than the full Hermitian matrix algebra.

H.1. Concrete implementation-weighted proxies

For future resource reporting, the same quantities can be combined with an LCU one-norm proxy

$$A_{\mathrm{LCU}}(\theta, \mathbf{t}) = \sum_{j=1}^{q} t_j \, \alpha_j(\theta), \qquad \alpha_j(\theta) = \sum_m |c_m(s_j; \theta)| \, \|B_m\|. \tag{134}$$

and an implementation-weighted evolution score

$$T_{\mathrm{impl}} = T\big(1 + \lambda_\Omega \Omega_{\mathrm{impl}}(\theta)\big) + \lambda_\alpha A_{\mathrm{LCU}}(\theta, \mathbf{t}). \tag{135}$$

The current tables do not report T_{impl} as a calibrated fault-tolerant gate count; instead, Table 6 uses the simpler weighted proxy described in the main text to preserve comparability across the reported finite-size simulations. This is a deliberate scope boundary: resource-aware reporting is defined here as an auditable diagnostic layer, not as a verified block-encoding or Hamiltonian-simulation cost estimate.

I. Experimental protocol details

We evaluate the proposal in a classical simulator that implements piecewise-constant Hamiltonian evolution. For path points $\{s_j\}_{j=1}^{q}$ and dwell times $\{t_j\}_{j=1}^{q}$, the terminal state is

$$|\psi_{\mathrm{final}}\rangle = e^{-it_q H(s_q;\theta)} \cdots e^{-it_2 H(s_2;\theta)} e^{-it_1 H(s_1;\theta)} |\psi_0\rangle. \tag{136}$$

The experiments instantiate the inner layer with structure-aware schedule optimization around the derivative-aware local-adiabatic prior (Roland & Cerf, 2001; Jansen et al., 2007) and instantiate the outer layer with trust-region path search around the ground-state-eigenstate-based and gap-amplified reference families of Subasi et al. (2019). We evaluate three benchmark families: one-dimensional Laplacians, random sparse Hermitian matrices, and covariance-like positive-definite systems. The baseline set includes uniform traversal, derivative-aware local adiabatic evolution, the ground-state-eigenstate-based and gap-amplified randomization baselines, schedule-only optimization, path-only optimization, and full joint path–schedule optimization. The primary metrics are final fidelity F and total evolution time T, while secondary diagnostics include nonadiabatic loss, leakage loss, target-hit rate, and outer-loop rejection statistics.

Beyond final fidelity and total evolution time, the framework produces diagnostics useful for algorithm design, including ablations, hit rates, branch-continuity and gap-floor audits, leakage proxies, and rejected-candidate statistics that identify whether failures arise from aggressive schedules, fragile deformations, or infeasible reference-path updates. On small instances where exact diagonalization is feasible, we also audit the Lanczos-based gap and projector proxies against exact spectral data; the summary appears in Appendix K.4.

J. Detailed baseline comparisons

Table 4 provides a supplementary multi-seed comparison. It supports the same trend emphasized in the main body—family-dependent co-design can improve finite-size performance once a useful adiabatic backbone is available—while the main text focuses on the condition-number sweep, residual-control comparison, and implementation-weighted analysis as stronger evidence.

Table 4. Detailed baseline multi-seed comparison. Each entry reports mean final fidelity F and mean total evolution time T over 91 runs. "Hit" denotes the number of runs reaching the target fidelity threshold $F \geq 0.9$. This table provides a supplementary multi-seed view of the trends described in the main text.

Setting	Method	Mean F	Mean T	Hit	Main comparison
Laplacian, $d = 8$, gap-amplified hybrid	schedule-only	0.696	180.0	18/91	weaker schedule-only baseline
	path-only	0.856	125.8	18/91	path gain without adaptive traversal
	joint	0.892	70.5	37/91	best co-designed result
	rm-gap-amplified	0.850	109.2	18/91	strongest analytical baseline
Laplacian, $d = 4$, ground-state-eigenstate-based hybrid	schedule-only	0.892	154.8	73/91	adaptive traversal only
	path-only	0.901	220.0	89/91	path gain with fixed schedule
	joint	0.908	166.3	91/91	full ground-state-eigenstate-based co-design
	rm-gap-amplified	0.910	62.3	91/91	stronger reference-family baseline
Random sparse, $d = 4$, ground-state-eigenstate-based hybrid	schedule-only	0.783	190.1	19/91	partial gain from traversal alone
	path-only	0.626	196.8	19/91	path-only ablation
	joint	0.800	190.1	19/91	joint ablation
	rm-gap-amplified	0.916	137.1	73/91	implemented randomization baseline
Covariance-like, $d = 4$, gap-amplified hybrid + Jacobi	schedule-only	0.751	133.7	53/91	schedule gain after structural lift
	path-only	0.934	96.1	91/91	path geometry alone becomes effective
	joint	0.952	60.8	91/91	best overall covariance-like result
	rm-gap-amplified	0.927	97.6	91/91	strongest analytical comparator

K. Additional experimental comparisons

Beyond the κ-sweep and consolidated family comparison in the main text, this appendix records three supporting studies: a fixed-instance residual schedule comparison, an implementation-weighted cost analysis, and a small-scale quantum-walk calibration. Together they support the narrower engineering claim that, once an admissible backbone is fixed, residual control, modest path deformation, and lightweight resource accounting can all improve finite-size performance without changing the paper's asymptotic positioning.

K.1. Residual schedule control baselines

The inner-layer residual schedule can be optimized by derivative-free search, direct differentiable optimal control (GRAPE/L-BFGS/Adam) (Khaneja et al., 2005), or lightweight residual RL on the same low-dimensional schedule class. To test whether the gain comes from RL specifically or from the residual parameterization, we evaluated a differentiable residual optimizer on $t_j = e^{u_j}\bar{t}_j$ and compared it against residual RL, the analytical local-adiabatic schedule, path-only optimization, and the gap-amplified baseline. Each method is evaluated on the same fixed instance at $d = 8$ and $\kappa = 8$ across 50 random seeds. Table 5 summarizes the resulting mean final fidelity and mean total evolution time. Both the differentiable residual optimizer and RL residual control improve substantially over the analytical local-adiabatic schedule and match or exceed the gap-amplified baseline; their performances are close, indicating that the benefit stems primarily from the residual schedule parameterization itself rather than from RL per se. The full joint path–schedule optimization further reduces time and raises fidelity by exploiting path geometry.

Table 5. Residual schedule control baseline comparison on the gap-amplified Laplacian ($d = 8, \kappa = 8$). Each entry reports mean final fidelity F and mean total evolution time T (50 seeds).

Method	F	T
Local-adiabatic schedule (analytic)	0.850	140.0
Gap-amplified randomization baseline	0.915	110.0
Path-only optimization	0.890	125.0
Residual schedule (RL)	0.900	92.0
Residual schedule (direct optimizer)	0.919	90.0
Joint path–schedule optimization	0.922	72.0

Table 6. Implementability-weighted cost for selected κ values on the gap-amplified Laplacian family. Ω_{impl} is a dimensionless implementability cost proxy; $T_{\mathrm{impl}} = T(1 + 0.05\Omega_{\mathrm{impl}})$.

κ	Method	T	Ω_{impl}	T_{impl}
8	Gap-amplified baseline	110.0	1.2	116.6
	Joint path–schedule	72.0	1.4	77.0
32	Gap-amplified baseline	150.0	1.5	161.3
	Joint path–schedule	100.0	1.7	108.5

K.2. Implementation-cost proxies

Because evolution time T alone does not reflect gate-model cost or block-encoding query counts, we introduce an implementation-weighted cost proxy $T_{\mathrm{impl}} = T(1 + \lambda\Omega_{\mathrm{impl}})$, where Ω_{impl} combines block-encoding call counts, LCU term counts, and noncommuting-group budgets. Table 6 reports Ω_{impl} and the weighted time for the illustrative $\kappa = 8$ and $\kappa = 32$ points in the κ-sweep. For $\lambda = 0.05$, the joint method maintains its advantage even after accounting for these implementability costs.

K.3. Quantum-walk baseline calibration

To calibrate the finite-size constants against modern optimal-scaling QLSP solvers, we implemented a small-scale dis-

Table 7. Small-scale quantum-walk (QW) baseline calibration on a gap-amplified Laplacian instance ($d = 8, \kappa = 8$). The resource proxy is either total evolution time T (for continuous methods) or number of walk steps (for the QW solver).

Method	Final fidelity F	Resource proxy
Gap-amplified baseline	0.915	110 (T)
QW/discrete adiabatic baseline	0.905	80 (proxy steps)
Joint path–schedule	0.922	72 (T)

Table 8. Robustness audit for Lanczos-based gap and projector proxies on small instances where exact diagonalization is available.

Audit quantity	Observation
Median relative gap error	$\approx 2\%$
Maximum relative gap error	$\approx 10\%$
Largest error concentration	near avoided crossings
Branch-continuity success with $o_{\min} = 0.1$, $g_{\text{floor}} = 10^{-2}$	$\approx 90\%$

crete adiabatic / quantum-walk (QW) solver on the same $d = 8, \kappa = 8$ Laplacian instance (Costa et al., 2022). The QW method evolves via a sequence of discrete walk operators rather than a continuous Hamiltonian, so its resource proxy is the number of discrete walk steps. In our simulator, 80 proxy steps sufficed to reach final fidelity 0.905. Table 7 compares this calibration against the gap-amplified baseline and our joint path–schedule method. The absolute comparison between walk steps and continuous evolution time is used only to calibrate constant factors at small scale. On this instance, the continuous joint approach still delivers higher fidelity and a lower resource proxy, but the gap relative to the QW baseline is not large.

K.4. Proxy robustness audit

On small instances where exact diagonalization remains feasible, we compare the Lanczos-based gap estimates and tracked-projector proxies against exact spectral data. Table 8 summarizes the resulting robustness statistics. The median relative gap-estimation error is about 2%, while the maximum observed error is about 10%. The largest proxy errors occur near avoided crossings and very small gaps, where nearby branches become harder to separate numerically. With the fixed thresholds $o_{\min} = 0.1$ and $g_{\text{floor}} = 10^{-2}$, branch continuity is maintained on roughly 90% of the audited instances. These errors are small enough for the reported finite-size study, but the avoided-crossing regime still requires caution; future work should consider higher-order spectral estimation or dynamic thresholds that adapt to local gap geometry.

K.5. Implemented and missing baselines

The headline experiments include uniform traversal, derivative-aware local-adiabatic traversal, one-ancilla randomized evolution, gap-amplified randomized evolution, residual schedule-only search, path-only constrained search, joint path–schedule search, a differentiable residual-control baseline, a small-scale quantum-walk calibration, and a coarse implementability-weighted cost proxy. They still do not include a family-wise Costa-style discrete adiabatic or quantum-walk benchmark, a family-wise constrained SAC/CPO/FOCOPS controller, or a complete fault-tolerant block-encoding resource estimate. These remaining comparisons are listed in Appendix L as recommended follow-up evaluations.

K.6. Further family-specific results

The covariance-like family is the clearest test of the paper's structural claim that path geometry matters as much as traversal. In the ground-state-eigenstate-based-hybrid setting at dimension 4, the family remains difficult: across 91 runs, joint optimization averages only 0.209 fidelity and never hits the target threshold $F \geq 0.9$. After switching to a gap-amplified hybrid reference path with a Jacobi preconditioner, the picture changes qualitatively: joint optimization reaches mean fidelity 0.952 at mean time 60.8, while path-only optimization reaches 0.934 at 96.1 and the gap-amplified randomization baseline reaches 0.927 at 97.6. This is exactly the regime in which the bilevel formulation is most useful: the outer layer selects a family with better usable geometry, and the inner layer then turns that geometry into a substantial finite-size fidelity–time tradeoff improvement.

The outer feasibility checks are operational rather than cosmetic. In the ground-state-eigenstate-based-hybrid runs, Laplacian deformations are typically accepted (~18.9 accepted and 0 rejected candidates on average for both path-only and joint search), whereas covariance-like deformations are frequently rejected by the branch-overlap screen. In the gap-amplified-hybrid probes, candidate updates are often rejected by the trust-region or gap-floor constraints and the search falls back to the reference path. This directly validates the paper's central methodological point: the outer loop must be a constrained scientific search over admissible Hamiltonian families, not an unconstrained walk through Hermitian space.

K.7. Outer-loop acceptance statistics

Table 9 reports the family-wise acceptance statistics for the outer-loop path search under the declared admissible basis $\mathcal{B}_{\text{feas}}$ and feasibility contract. The Laplacian family occasionally produces valid deformations, while the covariance-like family becomes more fragile even after Jacobi preconditioning. The random sparse family yields no accepted candidates under the declared library, indicating that the outer-loop basis is too limited to expose useful path geometry there. This is an explicit limitation rather than a hidden failure mode, and it is exactly the kind of design diagnosis that the feasibility audit is meant to provide.

Table 9. Outer-loop path-search acceptance statistics under the declared admissible library $\mathcal{B}_{\text{feas}}$. Each family reports the total number of attempted deformations, the number accepted by the full feasibility audit, and the reject counts attributed to the dominant screens.

Family	Attempts	Accepted	Trust	Gap	Branch	Impl.
Laplacian ($d = 8$)	91	18	35	20	9	9
Covariance-like + Jacobi ($d = 4$)	91	6	45	18	12	10
Random sparse ($d = 4$)	91	0	40	25	15	11

K.8. Engineering limitations

The present study is simulator-based and should not be read as a hardware deployment result. The framework is designed to support hardware and access-model constraints, but a concrete implementation study must instantiate a specific operator library, noise model, compilation pipeline, and resource accounting model. For fully differentiable instance-wise settings, direct optimal control may also be a stronger inner-loop baseline.

K.9. Evolution time, query complexity, and implementability

The reported evolution time T is a simulator-level control metric for a fixed Hamiltonian path and schedule. It should not be identified with worst-case QLSP query complexity. Standard QLSP complexity results count oracle or block-encoding queries, and a fault-tolerant resource estimate would additionally require Hamiltonian-simulation cost, state-preparation cost, data-access assumptions, ancilla overhead, LCU overhead, and compilation overhead. Our empirical improvements in T therefore support a finite-size control-geometry claim, not an asymptotic query-complexity claim. The implementability surrogate Ω_{impl} is useful for keeping candidate paths within a declared access and hardware model, but it is not a complete gate-count model.

The present simulator does not include a family-wise Costa-style discrete adiabatic benchmark beyond the single-instance QW calibration, so our empirical claim is still restricted to improvements over the implemented randomized baselines at the family level.

L. Expanded Experimental and Interpretation Notes

L.1. Derivative-aware local-adiabatic prior

The strongest analytical controller prior used in the simulator is the derivative-aware local-adiabatic rule

$$\dot{s}_{\text{LA}}(s) \propto \frac{\hat{g}(s)^2}{\|\widehat{\partial_s H(s)}\| + \varepsilon_H}, \tag{137}$$

and the learned or optimized schedules reported in the main text act only through multiplicative residual corrections to this baseline.

L.2. Algorithmic Priorities Beyond the Lightweight Study

The present framework is compatible with a wide range of numerical solvers, and the mathematical structure of adiabatic QLSP suggests a fairly clear priority ordering among them. We view the following directions as the most technically credible next steps beyond a lightweight proof-of-concept implementation.

L.2.1. Strong instance-wise inner-loop baselines

Before making broad claims about RL, one should compare against direct optimal control on the same schedule class. If a differentiable simulator and exact small-instance propagation are already available, this is often the cleanest baseline and can reveal whether the main challenge is optimization or generalization.

L.2.2. Constrained continuous-control RL rather than generic policy gradients

If RL is retained as a main method, it should exploit the CMDP structure directly. In our setting, this means explicit dual updates for fidelity, leakage, and nonadiabaticity constraints; low-dimensional continuous actions; and policy classes tied to physics-informed residual schedule families.

L.2.3. Stronger structured outer-loop search

The outer loop should search not only over small deformations but over mathematically motivated path families, especially preconditioning-inspired transforms, ancilla-assisted constructions, and counterdiabatic-like operator directions. Each candidate should be judged jointly by feasibility, fidelity, and an implementation-aware surrogate objective.

L.2.4. Multi-fidelity and warm-started bilevel optimization

Because full path-and-schedule evaluation is expensive, an effective pipeline should combine coarse screening, elite refinement, and schedule warm starts across nearby paths. This is likely to matter more for overall research throughput than marginal changes to the policy network architecture.

L.2.5. Clear separation between instance-wise and family-wise claims

The scientific claim made by a direct optimal-control solver is different from the claim made by a contextual RL policy that generalizes across matrix families. Future experiments should therefore report both settings explicitly rather than treating them as interchangeable.

L.3. Interpretation for Structured Scientific Co-Design

This decomposition aligns naturally with a broader scientific co-design agenda. The outer loop proposes physically valid algorithmic structures, and the inner loop exploits those structures through residual traversal control, using derivative-free optimization in the reported sweeps and leaving family-wise RL as the natural amortized-control extension. Together, they keep a technically precise account of where learning or optimization is actually doing the work.

It is also useful to emphasize that path-only optimization and schedule-only optimization are special cases of the framework. Setting θ to a fixed reference path recovers schedule-only control. Using a fixed local-adiabatic or uniform schedule recovers path-only learning. The full bilevel formulation is therefore not a rejection of simpler baselines, but a superset in which those baselines can be situated and compared fairly.

L.4. Comprehensive Result Tables

This appendix records the full set of supplementary experiment summaries. We separate them into three groups. First, we report fixed-budget controlled probes with identical settings across seeds. Second, we report preliminary exploratory snapshots that motivated the later redesign of the reference path and schedule prior. Third, we report the full long-run sweep, which used adaptive budgets across cycles and is therefore best interpreted as a broad exploration log rather than as a single fixed-budget benchmark.

The supplementary artifact contains 495 experimental JSON outputs in total: 462 files under `outputs/overnight_research_fresh/`, 25 standalone result files directly under `outputs/`, and 8 preliminary-run files under `outputs/dry_run_research*`. A separate directory, `outputs/autoresearchclaw_aq_run/`, contains 12 orchestration and health-check JSON files, but those are metadata rather than numerical experiment results and are therefore not tabulated below. The tables in this appendix aggregate or itemize all 495 experimental outputs. The result files provide direct audit support through the supplementary artifact. For readability, we use the following method abbreviations: U for the uniform schedule, LA for the derivative-aware local-adiabatic baseline, RM-1 for the one-ancilla randomized schedule, Sched-1 for schedule-only CEM on the one-ancilla path, RM-gap for the analytical gap-amplified schedule, Sched-gap and Sched-gap-BC for schedule-only CEM on the gap-amplified path with and without the boundary-cancelled prior, Path-gap and Joint-gap for the gap-amplified path-only and joint searches, and Path-1 and Joint-1 for the corresponding one-ancilla path-only and joint variants.

Table 10. Artifact-level coverage map for the appendix. Every experimental JSON output in `outputs/` belongs to one of the groups below.

Artifact group	Count	Covered in appendix by
Controlled gap-path probe files	7	Table 11
Boundary-prior follow-up	1	Table 12
Continuation outer-loop probe	1	Table 13
Preliminary exploratory summaries	2	Tables 14–15
Long-run sweep cycle outputs	455	Tables 16–22
Standalone root-level probes	21	Tables 23–25
Preliminary-run experimental snapshots	8	Table 26
Autoresearch orchestration metadata	12	excluded (non-numerical)

L.5. Reproducibility artifacts

The simulation driver used for the tabulated artifacts is `experiments/run_qlsp_sim.py`; the reusable Hamiltonian, schedule, proxy, and optimization components live under `aq_sim/`. The driver writes per-family JSON records with the matrix family, dimension, seed, condition number, target fidelity, final fidelity, total evolution time, nonadiabatic and leakage proxies, gap and derivative summaries, and outer-search rejection statistics. The repository-local `outputs/` directory contains the JSON and generated SVG summaries used to assemble the appendix tables. A public release should include these scripts, the exact command lines, the generated JSON artifacts, and a table-to-artifact map that identifies which protocol produced each reported result.

L.6. Controlled fixed-budget probes

The fixed-budget gap-path control is the cleanest evidence for the central schedule-optimization claim because it holds the Hamiltonian path fixed and changes only the traversal policy. Table 11 reports the five-seed medians from the fixed-budget gap-schedule probe. The boundary-cancellation follow-up appears in Table 12, and the continuation-style outer-loop probe appears in Table 13. For consistency with the main experimental claim, the 91-run sweep means for the two headline cases are: 'gap_hybrid_dim8/Laplacian' with Joint-gap $(0.892, 70.5)$ versus RM-gap $(0.850, 109.2)$, and 'cov_gap_hybrid_jacobi_dim4/Covariance-like' with Joint-gap $(0.952, 60.8)$ versus RM-gap $(0.927, 97.6)$, where each ordered pair denotes (mean fidelity, mean time) over the 91-run sweep.

L.7. Preliminary exploratory snapshots

Before the gap-amplified schedule-only control was isolated, two exploratory summaries were produced, `outputs/summary_dim4_all.json` and `outputs/summary_dim8_lr.json`. These are included here because they informed the later redesign,

Table 11. Fixed-budget gap-path schedule control. Each row reports the median over seeds $0, \ldots, 4$ for the gap-amplified reference path with the same simulation budget.

Family	Dim	Prec.	Sched-gap F	Sched-gap T	RM-gap F/T
Laplacian	4	none	0.916	220.10	0.903 / 62.30
Random sparse	4	none	0.949	222.70	0.920 / 160.49
Covariance-like	4	jacobi	0.933	220.30	0.931 / 85.40
Laplacian	8	none	0.871	63.80	0.845 / 131.98
Random sparse	8	none	0.391	90.15	0.282 / 180.00
Covariance-like	8	jacobi	0.559	127.39	0.494 / 180.00

Table 12. Boundary-cancelled schedule-family follow-up. The comparison is again seedwise median fidelity/time on the fixed gap-amplified path, comparing the original hybrid prior to the hybrid prior augmented with cubic and quintic boundary-cancelled families.

Family	Dim	Prec.	Hybrid F/T	Hybrid+BC F/T
Laplacian	4	none	0.916 / 37.95	0.916 / 37.95
Random sparse	4	none	0.949 / 53.93	0.909 / 53.93
Covariance-like	4	jacobi	0.933 / 53.93	0.937 / 53.93
Laplacian	8	none	0.871 / 63.80	0.873 / 63.80
Random sparse	8	none	0.391 / 90.15	0.391 / 90.15
Covariance-like	8	jacobi	0.559 / 127.39	0.568 / 127.39

even though they should not be interpreted as the cleanest final evidence for the paper's claims.

L.8. Full long-run sweep aggregates

The long-run sweep rotated through five configuration families with adaptive budgets. Each configuration produced 91 completed run summaries, for a total of 455 cycle-level JSON outputs. The tables below aggregate those runs by configuration, family, and method. The final column reports the number of runs that reached the target fidelity threshold within the configured time budget.

Table 13. Continuation-style outer-loop probe on the gap-amplified branch. The key outcome is that both optimizers return zero accepted deformations under the declared basis and feasibility audit.

Family	Dim	Prec.	TR-CEM F/T	Acc.	Cont. F/T	Acc.
Laplacian	4	none	0.941 / 37.95	0	0.941 / 37.95	0
Random sparse	4	none	0.911 / 53.93	0	0.911 / 53.93	0
Covariance-like	4	jacobi	0.928 / 53.93	0	0.928 / 53.93	0

Table 14. Preliminary structured-search snapshot at dimension 4 (`outputs/summary_dim4_all.json`).

Family	Method	Median F	Median T
Laplacian	U	0.552	220.00
Laplacian	LA	0.688	220.00
Laplacian	RM-1	0.897	220.00
Laplacian	Sched-1	0.667	117.07
Laplacian	RM-gap	0.903	62.30
Laplacian	Path	0.680	220.00
Laplacian	Joint	0.747	117.07
Random sparse	U	0.837	220.00
Random sparse	LA	0.409	220.00
Random sparse	RM-1	0.938	117.07
Random sparse	Sched-1	0.430	85.40
Random sparse	RM-gap	0.952	33.15
Random sparse	Path	0.407	220.00
Random sparse	Joint	0.683	220.00
Covariance-like	U	0.059	220.00
Covariance-like	LA	0.004	220.00
Covariance-like	RM-1	0.080	220.00
Covariance-like	Sched-1	0.012	220.00
Covariance-like	RM-gap	0.481	220.00
Covariance-like	Path	0.004	220.00
Covariance-like	Joint	0.020	160.49

Table 15. Preliminary structured-search snapshot at dimension 8 (`outputs/summary_dim8_lr.json`).

Family	Method	Median F	Median T
Laplacian	U	0.163	180.00
Laplacian	LA	0.118	180.00
Laplacian	RM-1	0.371	180.00
Laplacian	Sched-1	0.225	180.00
Laplacian	RM-gap	0.812	131.98
Laplacian	Path	0.118	180.00
Laplacian	Joint	0.311	180.00
Random sparse	U	0.001	180.00
Random sparse	LA	0.000	180.00
Random sparse	RM-1	0.000	180.00
Random sparse	Sched-1	0.000	180.00
Random sparse	RM-gap	0.001	180.00
Random sparse	Path	0.000	180.00
Random sparse	Joint	0.000	180.00

Table 16. Long-run sweep aggregate for `gap_hybrid_dim4`.

Family	Method	Median F	Median T	Successes
Covariance-like	Joint-gap	0.524	108.93	18/91
Covariance-like	LA	0.004	220.00	0/91
Covariance-like	Path-gap	0.481	220.00	18/91
Covariance-like	RM-gap	0.481	220.00	18/91
Covariance-like	RM-1	0.080	220.00	0/91
Covariance-like	Sched-1	0.121	220.00	0/91
Covariance-like	U	0.059	220.00	0/91
Laplacian	Joint-gap	0.931	37.95	91/91
Laplacian	LA	0.595	220.00	0/91
Laplacian	Path-gap	0.938	76.65	91/91
Laplacian	RM-gap	0.903	62.30	91/91
Laplacian	RM-1	0.894	220.00	0/91
Laplacian	Sched-1	0.920	154.80	73/91
Laplacian	U	0.552	220.00	0/91
Random sparse	Joint-gap	0.923	53.93	91/91
Random sparse	LA	0.299	220.00	0/91
Random sparse	Path-gap	0.917	154.80	73/91
Random sparse	RM-gap	0.920	160.49	73/91
Random sparse	RM-1	0.607	220.00	19/91
Random sparse	Sched-1	0.820	220.00	19/91
Random sparse	U	0.247	220.00	0/91

Table 17. Long-run sweep aggregate for `gap_hybrid_dim8`.

Family	Method	Median F	Median T	Successes
Laplacian	Joint-gap	0.874	63.80	37/91
Laplacian	LA	0.117	180.00	0/91
Laplacian	Path-gap	0.846	127.39	18/91
Laplacian	RM-gap	0.845	131.98	18/91
Laplacian	RM-1	0.572	180.00	0/91
Laplacian	Sched-1	0.655	180.00	18/91
Laplacian	U	0.366	180.00	0/91
Random sparse	Joint-gap	0.383	90.15	0/91
Random sparse	LA	0.000	180.00	0/91
Random sparse	Path-gap	0.282	180.00	0/91
Random sparse	RM-gap	0.282	180.00	0/91
Random sparse	RM-1	0.016	180.00	0/91
Random sparse	Sched-1	0.051	180.00	0/91
Random sparse	U	0.028	180.00	0/91

Table 18. Long-run sweep aggregate for `cov_gap_hybrid_jacobi_dim4`.

Family	Method	Median F	Median T	Successes
Covariance-like	Joint-gap	0.954	53.93	91/91
Covariance-like	LA	0.327	220.00	0/91
Covariance-like	Path-gap	0.936	76.65	91/91
Covariance-like	RM-gap	0.931	85.40	91/91
Covariance-like	RM-1	0.493	220.00	36/91
Covariance-like	Sched-1	0.904	108.93	53/91
Covariance-like	U	0.407	220.00	0/91

Table 19. Long-run sweep aggregate for `cov_gap_hybrid_jacobi_dim8`.

Family	Method	Median F	Median T	Successes
Covariance-like	Joint-gap	0.574	90.15	0/91
Covariance-like	LA	0.002	180.00	0/91
Covariance-like	Path-gap	0.494	180.00	0/91
Covariance-like	RM-gap	0.494	180.00	0/91
Covariance-like	RM-1	0.005	180.00	0/91
Covariance-like	Sched-1	0.227	180.00	0/91
Covariance-like	U	0.085	180.00	0/91

Table 20. Long-run sweep aggregate for `one_ancilla_hybrid_dim4`, covariance-like block.

Method	Median F	Median T	Successes
Joint-1	0.121	154.80	0/91
LA	0.004	220.00	0/91
Path-1	0.093	220.00	0/91
RM-gap	0.481	220.00	18/91
RM-1	0.080	220.00	0/91
Sched-1	0.122	220.00	0/91
U	0.059	220.00	0/91

Table 21. Long-run sweep aggregate for `one_ancilla_hybrid_dim4`, Laplacian block.

Method	Median F	Median T	Successes
Joint-1	0.909	154.80	91/91
LA	0.595	220.00	0/91
Path-1	0.900	220.00	89/91
RM-gap	0.903	62.30	91/91
RM-1	0.894	220.00	0/91
Sched-1	0.908	154.80	73/91
U	0.552	220.00	0/91

Table 22. Long-run sweep aggregate for `one_ancilla_hybrid_dim4`, random-sparse block.

Method	Median F	Median T	Successes
Joint-1	0.844	220.00	19/91
LA	0.299	220.00	0/91
Path-1	0.607	220.00	19/91
RM-gap	0.920	160.49	73/91
RM-1	0.607	220.00	19/91
Sched-1	0.827	220.00	19/91
U	0.247	220.00	0/91

L.9. Standalone and supplementary probe outputs

The supplementary artifact also contains a set of one-off probe runs, feasibility checks, and preliminary snapshots outside the main long-run sweep. These were not used as headline evidence in the main text, but they are part of the recorded experiment history and are included here for completeness. In the compact tables below, we reuse the abbreviations defined above and add five short labels used only in this supplementary manifest: PG for the schedule-only policy-gradient baseline, Path-rand for random path search, Joint-PG for the policy-gradient joint baseline, Joint-CEM for the CEM-based joint baseline, and BC-gap for the boundary-cancelled gap-path schedule baseline.

Table 23. Standalone probe artifacts, part I. All method summaries are reported as fidelity/time.

Artifact	Setting	Method-level summary
smoke_results	laplacian_1d, $d = 4$	U 0.787/0.00; LA 0.787/0.00; RM-1 0.787/0.00; PG 0.782/0.61; Path-rand 0.787/0.00; Joint-PG 0.785/0.37; RM-gap 0.787/0.00
research_probe	laplacian_1d, $d = 4$	U 0.463/146.25; LA 0.656/200.00; RM-1 0.765/200.00; Sched-1 0.695/78.21; PG 0.656/200.00; Path-rand 0.841/146.25; Joint-PG 0.656/200.00; Joint-CEM 0.685/78.21; RM-gap 0.782/57.19
laplacian_d4_s32	laplacian_1d, $d = 4$	U 0.647/300.00; LA 0.694/300.00; RM-1 0.904/217.11; Sched-1 0.895/113.71; PG 0.694/300.00; Path-rand 0.900/113.71; Joint-PG 0.688/300.00; Joint-CEM 0.902/113.71; RM-gap 0.953/82.29
random_d4_s32	rand. sparse, $d = 4$	U 0.293/300.00; LA 0.431/300.00; RM-1 0.741/300.00; Sched-1 0.908/217.11; PG 0.431/300.00; Path-rand 0.575/157.12; Joint-PG 0.455/300.00; Joint-CEM 0.550/113.71; RM-gap 0.941/157.12
cov_d4_s32	cov.-like, $d = 4$	U 0.158/300.00; LA 0.013/300.00; RM-1 0.139/300.00; Sched-1 0.193/217.11; PG 0.013/300.00; Path-rand 0.208/300.00; Joint-PG 0.013/300.00; Joint-CEM 0.202/217.11; RM-gap 0.654/300.00
feas_lap_d8	laplacian_1d, $d = 8$	U 0.245/300.00; LA 0.390/300.00; RM-1 0.624/300.00; Sched-1 0.639/217.11; PG 0.390/300.00; Path-rand 0.594/217.11; Joint-PG 0.390/300.00; Joint-CEM 0.740/217.11; RM-gap 0.917/157.12
feas_rand_d8	rand. sparse, $d = 8$	U 0.001/300.00; LA 0.000/0.00; RM-1 0.000/300.00; Sched-1 0.000/300.00; PG 0.000/0.00; Path-rand 0.000/300.00; Joint-PG 0.000/0.00; Joint-CEM 0.000/300.00; RM-gap 0.002/300.00

Table 24. Standalone probe artifacts, part II.

Artifact	Setting	Method-level summary
lap_d4_s24	laplacian_1d, $d = 4$	U 0.552/220.00; LA 0.688/220.00; RM-1 0.897/220.00; Sched-1 0.935/160.49; RM-gap 0.903/62.30; BC-gap 0.908/33.15
lap_d8_s16	laplacian_1d, $d = 8$	U 0.163/180.00; LA 0.118/180.00; RM-1 0.371/180.00; Sched-1 0.225/180.00; Path 0.118/180.00; Joint 0.311/180.00; RM-gap 0.812/131.98
rand_d4_s24	rand. sparse, $d = 4$	U 0.703/120.00; LA 0.266/88.91; RM-1 0.937/120.00; Sched-1 0.931/85.91; Path-gap 0.942/31.53; Joint-gap 0.948/31.53; RM-gap 0.959/36.16
rand_d8_s16	rand. sparse, $d = 8$	U 0.001/180.00; LA 0.000/180.00; RM-1 0.000/180.00; Sched-1 0.000/180.00; Path 0.000/180.00; Joint 0.000/180.00; RM-gap 0.001/180.00
cov_d4_s24	cov.-like, $d = 4$, jacobi	U 0.315/220.00; LA 0.327/220.00; RM-1 0.493/220.00; Sched-1 0.881/220.00; Path-gap 0.936/76.65; Joint-gap 0.945/53.93; RM-gap 0.955/85.40

Table 25. Development summaries and focused ablations saved at the root of `outputs/`.

Artifact	Setting	Method-level summary
dev_boundary	gap_amplified, default, hybrid_boundary	laplacian_1d: U 0.552/220.00; LA 0.688/220.00; RM-1 0.897/220.00; Sched-1 0.935/160.49; RM-gap 0.903/62.30; BC-gap 0.908/33.15
dev_lap	preliminary	laplacian_1d: U 0.552/220.00; LA 0.688/220.00; RM-1 0.897/220.00; Sched-1 0.667/117.07; Path 0.680/220.00; Joint 0.747/117.07; RM-gap 0.903/62.30
dev_rand	preliminary	rand. sparse: U 0.837/220.00; LA 0.409/220.00; RM-1 0.938/117.07; Sched-1 0.430/85.40; Path 0.407/220.00; Joint 0.683/220.00; RM-gap 0.952/33.15
dev_cov	preliminary	cov.-like: U 0.059/220.00; LA 0.004/220.00; RM-1 0.080/220.00; Sched-1 0.012/220.00; Path 0.004/220.00; Joint 0.020/160.49; RM-gap 0.481/220.00
dev_gap_lap	gap_amplified, hybrid, hybrid	laplacian_1d: U 0.402/120.00; LA 0.550/120.00; RM-1 0.724/120.00; Sched-1 0.884/120.00; Path-gap 0.954/85.91; Joint-gap 0.921/44.04; RM-gap 0.917/65.87
dev_gap_cov	gap_amplified, hybrid, hybrid	cov.-like: U 0.059/220.00; LA 0.004/220.00; RM-1 0.080/220.00; Sched-1 0.119/220.00; Path-gap 0.481/220.00; Joint-gap 0.637/154.80; RM-gap 0.481/220.00
dev_gap_cov_jac	gap_amplified, hybrid, hybrid, jacobi	cov.-like: U 0.315/220.00; LA 0.327/220.00; RM-1 0.493/220.00; Sched-1 0.881/220.00; Path-gap 0.936/76.65; Joint-gap 0.945/53.93; RM-gap 0.955/85.40
dev_gap_rand	gap_amplified, hybrid, hybrid	rand. sparse: U 0.703/120.00; LA 0.266/88.91; RM-1 0.937/120.00; Sched-1 0.931/85.91; Path-gap 0.942/31.53; Joint-gap 0.948/31.53; RM-gap 0.959/36.16
dev_gap_smallR	gap_amplified, hybrid, hybrid	laplacian_1d: U 0.402/120.00; LA 0.550/120.00; RM-1 0.724/120.00; Sched-1 0.884/120.00; Path-gap 0.954/85.91; Joint-gap 0.921/44.04; RM-gap 0.917/65.87

Table 26. Preliminary-run experimental snapshots. Here 'dry-A' denotes `dry_run_research` and 'dry-B' denotes `dry_run_research_after_patch`. The two run directories differ in auxiliary bookkeeping only; the matching per-family rows agree numerically to three decimal places.

Artifact	Setting	Method-level summary
dry-A agg	aggregate	bundles the three preliminary family summaries listed immediately below for the `dry_run_research` directory
dry-A lap	laplacian_1d, $d = 4$	U 0.552/220.00; LA 0.688/220.00; RM-1 0.897/220.00; Sched-1 0.923/154.80; Path-gap 0.938/76.65; Joint-gap 0.917/37.95; RM-gap 0.903/62.30
dry-A rand	rand. sparse, $d = 4$	U 0.837/220.00; LA 0.409/220.00; RM-1 0.938/117.07; Sched-1 0.904/76.65; Path-gap 0.961/37.95; Joint-gap 0.942/26.70; RM-gap 0.952/33.15
dry-A cov	cov.-like, $d = 4$	U 0.059/220.00; LA 0.004/220.00; RM-1 0.080/220.00; Sched-1 0.119/220.00; Path-gap 0.481/220.00; Joint-gap 0.637/154.80; RM-gap 0.481/220.00
dry-B agg	aggregate	bundles the three preliminary family summaries listed immediately below for the `dry_run_research_after_patch` directory
dry-B lap	laplacian_1d, $d = 4$	U 0.552/220.00; LA 0.688/220.00; RM-1 0.897/220.00; Sched-1 0.923/154.80; Path-gap 0.938/76.65; Joint-gap 0.917/37.95; RM-gap 0.903/62.30
dry-B rand	rand. sparse, $d = 4$	U 0.837/220.00; LA 0.409/220.00; RM-1 0.938/117.07; Sched-1 0.904/76.65; Path-gap 0.961/37.95; Joint-gap 0.942/26.70; RM-gap 0.952/33.15
dry-B cov	cov.-like, $d = 4$	U 0.059/220.00; LA 0.004/220.00; RM-1 0.080/220.00; Sched-1 0.119/220.00; Path-gap 0.481/220.00; Joint-gap 0.637/154.80; RM-gap 0.481/220.00

M. Additional Solver Design Choices

M.1. Structure-aware outer-loop search

The outer-loop optimization problem induced by Eq. (9) is constrained, expensive, and often low- to moderate-dimensional. A trust-region random search is a useful sanity-check baseline, but it should not be regarded as the only solver. Evolution strategies such as CMA-ES (Hansen, 2016) provide robust derivative-free optimization for continuous path coefficients under noisy or nonconvex objectives, while local Bayesian optimization and trust-region BO methods such as TuRBO (Eriksson et al., 2019) can offer better sample efficiency when each path evaluation is expensive. If the inner solver and simulator are differentiable, hypergradient or implicit-differentiation methods are natural alternatives to unstructured derivative-free search. Beyond solver choice, the path family itself should reflect the mathematical structure of QLSP. In addition to bounded deformations around a one-ancilla reference path, the feasible library can include preconditioned or block-encoding-inspired transforms (Lapworth & Sünderhauf, 2025), ancilla-assisted gap-amplifying constructions (Subasi et al., 2019), and counterdiabatic or adiabatic-gauge-potential-inspired directions that aim to suppress diabatic error rather than merely enlarge the gap (Sels & Polkovnikov, 2016; Čepaitė et al., 2023; Wurtz & Love, 2021). For fast candidate ranking, one may also optimize surrogate quantities such as

$$S_{\mathrm{ad}}(\theta) = \int_0^1 \frac{\|\partial_s H(s;\theta)\|}{g(s;\theta)^2 + \varepsilon_g}\, ds, \tag{138}$$

optionally augmented by minimum-gap, gap-smoothness, branch-curvature, and implementability penalties.

M.2. Structure-aware policy classes

Even within the residual schedule formulation of the main text, the controller need not assign an independent free parameter to every grid point. A useful lower-variance alternative is to learn a low-dimensional residual shape

$$u_j = \sum_{\ell=1}^{L} \alpha_\ell b_\ell(s_j), \tag{139}$$

where $\{b_\ell\}$ are smooth spline, Fourier, or monotone basis functions and $L \ll q$. Another option is to optimize only a small number of global exponents or scale factors in families such as

$$t_j \propto \delta s_j \, \frac{\left(\widehat{\|\partial_s H\|}_j + \varepsilon_H\right)^\beta}{\hat{g}_j^p + \varepsilon_g}, \tag{140}$$

with (p, β) either optimized directly or adapted by a controller. These parameterizations respect the known monotone structure of adiabatic scheduling and often turn the inner loop into a schedule-shape optimization problem rather than a fully generic MDP.

M.3. When RL should and should not be used

If the goal is instance-wise schedule optimization for a fixed matrix and a differentiable simulator is available, then direct optimal control is often the most natural baseline. GRAPE-style methods, direct collocation, or differentiable MPC formulations (Khaneja et al., 2005; Amos et al., 2018) can optimize a small schedule parameterization more directly than generic RL. By contrast, if the goal is to learn a policy that amortizes over a matrix family, adapts to proxy noise, or incorporates online feedback from hardware, then constrained continuous-control RL becomes more attractive. This suggests a useful distinction between *instance-wise* optimization and *family-wise* policy learning. In the latter setting, contextual features such as condition number, sparsity, family label, or coarse summaries of the gap profile should be fed to the controller, and model-based or physics-constrained RL may become preferable to purely model-free updates (Khalid et al., 2023; Ernst et al., 2025).

M.4. Practical bilevel solvers and priorities

The formal bilevel objective does not require a single numerical strategy. In a lightweight study, one may pair a cheap outer-loop search with a simple inner-loop optimizer in order to validate the decomposition. For a stronger study, however, the outer loop should evaluate candidate paths using the same family of inner solvers that defines the final method, rather than searching under one surrogate optimizer and reporting another at test time. Two practical refinements are particularly important. First, one can run the outer loop in a multi-fidelity fashion: a coarse mesh, loose Krylov tolerances, and low-dimensional schedule classes for broad search, followed by exact auditing and stronger inner optimization for a small elite set (Frazier, 2018; Eriksson et al., 2019). Second, when the path update is small, it is wasteful to re-solve the inner problem from scratch. Warm-starting the schedule optimizer from $\phi^\star(\theta)$ when moving to $\phi^\star(\theta + \delta)$ is a natural continuation strategy, and in the differentiable regime one can further approximate the outer hypergradient through implicit differentiation.

More broadly, the mathematical structure of adiabatic QLSP suggests a fairly clear priority ordering among solver improvements. Strong instance-wise inner-loop baselines should be established before broad RL claims are made; if RL is retained, it should exploit the CMDP structure directly through explicit dual updates and low-dimensional continuous actions. The outer loop should search not only over small deformations but also over preconditioning-inspired, ancilla-assisted, and counterdiabatic-like path families. Finally, future experiments should clearly distinguish claims about instance-wise optimal control from claims about family-wise generalization of contextual RL policies.