FADE: Fidelity-Aware Drift-Embedded Gate Optimizer

Full Name:
Sashreek Satapathy

Email Address:
sashreeksatapathydpsnt15@gmail.com

Research Topic:
Quantum Gate Optimization Under Time-Dependent Hardware Drift

Institution:
Delhi Public School, Vadodara, India

FADE: Fidelity-Aware Drift-Embedded Gate Optimizer

Sashreek Satapathy | Delhi Public School, Vadodara, India

1. Research Problem & Introduction

Quantum gate optimization rests on an assumption that does not hold in practice. GRAPE (Khaneja et al., 2005) and CRAB (Doria et al., 2011) both treat the hardware Hamiltonian as fixed at calibration. Physical qubits do not cooperate: superconducting systems exhibit $1/f$ flux noise that shifts qubit frequencies on the timescale of minutes; trapped-ion systems experience laser intensity variation and motional heating over the same window. A gate that performs well at calibration will, in general, perform differently thirty minutes into an experiment, causing calibrated gates to degrade unpredictably during execution. No existing optimizer addresses this — temporal drift is treated as an external nuisance managed through recalibration, not as a variable the optimizer should reason about.

FADE (Fidelity-Aware Drift-Embedded optimization) incorporates stochastic hardware drift directly into the gate optimization objective, targeting fidelity sustained across a full experimental session rather than at a single calibration point. The motivation is concrete: NISQ devices operate without quantum error correction, so gate errors accumulate without algorithmic recourse. Addressing drift at the level of pulse design — rather than through repeated recalibration — would improve reliability on every current hardware platform. In short, FADE aims to design quantum gates that remain accurate throughout a computation, not merely at the moment of calibration.

2. Literature Review

GRAPE computes analytic gradients of single-point gate fidelity over piecewise-constant pulse segments; CRAB reparameterizes pulses through a truncated functional basis for bandwidth robustness. Both converge to high fidelity under controlled conditions, and both embed the same structural assumption: optimization is performed against a static Hamiltonian. DRAG and second-order GRAPE refine specific imperfections without disturbing this. Ensemble averaging over discrete Hamiltonian offsets adds static robustness but treats uncertainty as a fixed distribution rather than a time-correlated process, and it provides no incentive to minimize recalibration frequency across a session.

The Lindblad master equation (Breuer & Petruccione, 2002) captures coherent qubit evolution alongside energy relaxation ($1/T_1$) and dephasing ($1/T_2$). Extending it to account for drift requires treating the Hamiltonian as time-varying: $H(t) = H_0 + \delta H(t, \xi)$, where δH is a stochastic perturbation parameterized by platform-specific noise. To the best of our knowledge, no published framework uses this structure to define a session-integrated optimization objective. Bayesian optimization (Shahriari et al., 2016) and reinforcement learning (Fosel et al., 2018) have each been applied to gate-related problems, but neither redefines what the optimizer maximizes — both operate within the same single-point fidelity paradigm. The gap is structural: every existing method asks how to maximize fidelity at one moment. None asks how to sustain it across a session.

3. Research Question

Can an objective function that explicitly models time-dependent hardware drift produce control pulses that sustain higher gate fidelity across a full experimental session than conventional single-shot optimizers?

The independent variable is the structure of the optimization objective: a session-integrated, drift-aware formulation versus the standard single-point fidelity. The dependent variable is gate accuracy over time, measured across multiple simulated sessions on physically motivated hardware models. An affirmative result would establish that temporal fidelity degradation is partly a consequence of how the optimization problem is posed, not solely of hardware physics.

4. Methodology

Phase 1 — Gap analysis: Each major quantum control method (GRAPE, CRAB, DRAG, gradient-free) and ML optimization method (Bayesian optimization, reinforcement learning, evolutionary strategies, neural controllers) is characterized by its objective class, Hamiltonian assumptions, and treatment of parametric uncertainty. A systematic comparison identifies unexplored pairings. No such mapping exists in the literature; this analysis is an independent contribution from which FADE emerges.

Phase 2 — Framework development: Gate evolution under drift is modeled via the Lindblad equation with stochastic Hamiltonian $\delta H(t, \xi)$, where ξ is drawn from platform-specific noise models calibrated against published data. The session-level objective is: $\mathbf{J(\theta) = E_\xi[\int_0^T F(\rho_\theta(t,\xi), \rho_{target})dt] - \lambda C(\theta)}$, where $C(\theta)$ penalizes pulses requiring recalibration to stay above a 0.99 fidelity threshold and λ governs the fidelity–recalibration trade-off. In effect, $J(\theta)$ rewards pulses that remain accurate throughout a session, not just at its start. Pulse parameters are refined via GRAPE-style gradients against $J(\theta)$; drift model hyperparameters are tuned by Bayesian optimization, suited to this task because each objective evaluation requires Monte Carlo integration over drift trajectories.

Phase 3 — Benchmarking: FADE is evaluated against GRAPE and CRAB using QuTiP on two hardware models: a superconducting transmon qubit ($1/f$ flux noise from published spectra) and a trapped-ion qubit (laser noise and motional heating from the literature). Each session simulates a target X-rotation over thirty minutes with drift trajectories sampled per the respective noise model. Performance metrics: session-average fidelity, fidelity variance, and recalibration frequency. Success criterion: statistically significant improvement on ≥ 2 of 3 metrics versus both baselines on both platforms. Each simulation is repeated across multiple independent drift realizations to ensure statistical robustness.

5. Significance and Novelty

FADE corrects a problem formulation, not merely an algorithm. Optimizing instantaneous fidelity was tractable as a starting point, but drift was never absent — the field chose to handle it post-hoc through recalibration. The Lindblad equation, stochastic Hamiltonian modeling, GRAPE-style gradients, and Bayesian hyperparameter tuning each appear in prior work; their integration within a session-level objective for quantum gate optimization has, to the best of our knowledge, no precedent, and the Phase 1 gap analysis provides the evidentiary basis for that claim. As national quantum programs (US, EU, India) work to bring NISQ hardware to practical utility, addressing drift at the level of the optimization objective is a specific, tractable contribution to that effort.

6. Conclusion and Feasibility

All tools required — QuTiP, Python, NumPy, SciPy — are freely available, and all simulations are designed to run on standard personal computing hardware within reasonable time constraints. Initial validation will use reduced-dimensional simulations to ensure tractability before scaling to full session-length benchmarks. Drift parameters are sourced from published experimental characterizations; no hardware access is needed. The three-phase structure is robust to partial results: Phases 1 and 2 are publishable contributions independent of simulation outcomes. Validation is simulation-based; the paper will present results as theoretical proof-of-concept and specify the hardware measurements that would constitute empirical confirmation. Computational cost of Monte Carlo fidelity estimation will be managed via first-order Taylor expansion around the drift-free trajectory and variance reduction techniques — identified mitigations, not open defeaters. If the framework succeeds, FADE shifts quantum gate design from static snapshot optimization to session-aware pulse control: hardware that is accurate not just at calibration, but throughout computation.

References

Breuer, H.-P., & Petruccione, F. (2002). *The Theory of Open Quantum Systems.* Oxford University Press.

Doria, P., Calarco, T., & Montangero, S. (2011). Optimal control technique for many-body quantum dynamics. *Physical Review Letters, 106*(19), 190501.

Fosel, T., Tighineanu, P., Weiss, T., & Marquardt, F. (2018). Reinforcement learning with neural networks for quantum feedback. *Physical Review X, 8*(3), 031084.

Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbruggen, T., & Glaser, S. J. (2005). Optimal control of coupled spin dynamics. *Journal of Magnetic Resonance, 172*(2), 296–305.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & de Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE, 104*(1), 148–175.