

Time-Aware Active Calibration and Scheduling for N-Qubit Systems via Bayesian Optimization and Reinforcement Learning

Ziming Zhao

School of Software Technology,
Zhejiang University, Ningbo, China
zhaoziming@zju.edu.cn

Tingting Li*

School of Software Technology,
Zhejiang University, Ningbo, China
litt2020@zju.edu.cn

Xiaofei Yue

Beijing Institute of Technology
Beijing, China
xfyue@bit.edu.cn

Zhaoxuan Li

Institute of Information Engineering,
CAS, Beijing, China
lizhaoxuan@iie.ac.cn

Jianwei Yin*

School of Software Technology,
Zhejiang University, Ningbo, China
zjuyjw@cs.zju.edu.cn

Abstract

Frequent recalibration of multi-qubit systems is essential to maintain hardware fidelity, but calibration itself is costly and subject to rapid parameter drift. This work presents a novel framework for *time-aware active calibration* that combines Bayesian Optimization (BO) and Reinforcement Learning (RL) to adaptively schedule and prioritize calibration experiments under uncertainty. We model the calibration landscape as a directed acyclic graph (DAG), where nodes represent hardware parameters (*e.g.*, frequency, amplitude ratio, DRAG, or phase) and edges capture inter-parameter dependencies extracted automatically from Qiskit Experiments metadata. Each parameter's temporal drift is modeled as an Ornstein-Uhlenbeck (OU) or exponential decay process with uncertainty bounds. The BO agent prioritizes the next calibration target using a time-aware acquisition score that balances expected improvement, predictive uncertainty, and experiment duration under system-level constraints. A parallel RL scheduler coordinates cross-qubit calibration and readout tasks to minimize idle time and avoid conflicts. Our system further incorporates self-triggered recalibration based on online RB/XEB health metrics, enabling autonomous recovery from drift without human intervention. Experimental evaluation on multi-qubit testbeds demonstrates improved calibration efficiency, longer mean time between recalibrations (MTBC), and reduced overall calibration overhead compared to fixed or purely heuristic scheduling baselines.

CCS Concepts

• **Hardware** → **Quantum technologies**.

Keywords

Active Calibration, Bayesian Optimization, Reinforcement Learning, Multi-Qubit Quantum Systems

*Corresponding authors.

ACM Reference Format:

Ziming Zhao, Tingting Li, Xiaofei Yue, Zhaoxuan Li, and Jianwei Yin. 2026. Time-Aware Active Calibration and Scheduling for N-Qubit Systems via Bayesian Optimization and Reinforcement Learning. In *63rd ACM/IEEE Design Automation Conference (DAC '26)*, July 26–29, 2026, Long Beach, CA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3770743.3804388>

1 Introduction

The performance of superconducting and trapped-ion quantum processors critically depends on frequent calibration of device parameters such as qubit frequencies, pulse amplitudes, derivative removal by adiabatic gate (DRAG) coefficients, and readout phases [1–3]. These calibration parameters exhibit temporal drift due to hardware aging, environmental fluctuations, and crosstalk between control channels [3, 4]. As the scale of quantum hardware continues to grow, the cost of maintaining calibration fidelity becomes a significant bottleneck in both experimental throughput and system availability [5, 6]. Traditional calibration strategies rely on static or heuristic scheduling, where each qubit or parameter is recalibrated periodically or upon detection of degraded gate fidelity [4, 7]. Such methods neglect the temporal and structural correlations among calibration parameters, resulting in redundant experiments and inefficient use of calibration time [3, 8]. Moreover, parameter drift often follows non-stationary stochastic processes, making it difficult to predict the optimal recalibration interval or prioritize among competing tasks [4].

Recent progress in automated calibration frameworks (*e.g.*, Qiskit Experiments, Boulder Opal, or QCal) has enabled data-driven experiment orchestration and metadata logging [2, 9, 10], but these systems lack adaptive decision-making and temporal awareness. Specifically, they cannot dynamically select which calibration to perform next, nor can they quantify uncertainty or balance the trade-off between information gain and experimental cost.

Contributions. To address these, we propose *Time-Aware Active Calibration (TAAC)*, a unified framework that integrates Bayesian Optimization (BO) and Reinforcement Learning (RL) for adaptive calibration scheduling in N-qubit systems. As shown in Figure 1, our approach models the calibration landscape as a *Directed Acyclic Graph (DAG)*, where nodes represent tunable parameters and edges encode cross-dependencies inferred from experimental metadata. The framework maintains an uncertainty model of temporal drift



This work is licensed under a Creative Commons Attribution 4.0 International License. *DAC '26, Long Beach, CA, USA*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2254-7/2026/07

<https://doi.org/10.1145/3770743.3804388>




Figure 1: TAAC framework architecture.

using Ornstein-Uhlenbeck (OU) or exponential decay processes, enabling prediction of confidence intervals over time. At each calibration step, the BO agent prioritizes the next calibration target using a time-aware acquisition score that favors high expected benefit, high predictive uncertainty, and low execution cost. Concurrently, an RL-based scheduler coordinates parallel calibration tasks across qubits, ensuring conflict-free resource allocation and minimizing total downtime. The system further incorporates online health monitoring (*e.g.*, RB/XEB metrics) to autonomously trigger recalibration events, allowing self-healing operation without human supervision. Through extensive experiments on multi-qubit testbeds, TAAC demonstrates nearly $2\times$ throughput improvement, 40% reduction in total calibration time, and extended mean time between calibrations (MTBC) compared with existing fixed or rule-based schemes. This work represents a step toward *closed-loop, intelligent calibration orchestration* for scalable quantum hardware management, paving the way for autonomous operation of future quantum computing systems.

2 Background and Related Work

2.1 Quantum Calibration and Parameter Drift

In a quantum processor, the accuracy of logic operations depends on a large number of continuously tuned control parameters [2, 11]. Each qubit requires frequent recalibration of its operating frequency, pulse amplitude, and DRAG coefficient, while two-qubit gates depend on relative phase alignment and coupling strength [12, 13]. Calibration involves running characterization experiments such as Rabi oscillations, Ramsey fringes, or cross-entropy benchmarking (XEB), followed by parameter fitting and update [4, 6]. However, these parameters drift over time due to both environmental and device-specific factors, including temperature fluctuations, magnetic field instability, and material-induced charge noise [4, 14, 15]. This drift can often be modeled as a stochastic process with temporal correlation, such as an Ornstein-Uhlenbeck (OU) process or exponential decay [16]. Without timely recalibration, these drifts cause degradation in gate fidelity, ultimately shortening the mean time between calibrations (MTBC) and reducing computational reliability [4, 6, 17]. As illustrated in Figure 2, the temporal evolution of calibration parameters such as frequency and DRAG coefficient can




Figure 2: Temporal drift and uncertainty modeling.

be well captured by an Ornstein-Uhlenbeck process [16]. The OU model not only tracks the nonlinear drift trend over time but also provides confidence intervals that quantify prediction uncertainty, which are used later for time-aware scheduling and recalibration.

2.2 Automated Calibration Frameworks

To manage the increasing complexity of large-scale systems, several calibration orchestration frameworks have been proposed [6, 11–13, 18]. Some examples include *Qiskit Experiments*, *QCal*, and *Boulder Opal* and so on [2, 9, 10, 19]. These platforms provide structured experiment execution, metadata storage, and dependency tracking [2, 11]. However, they typically follow a static or rule-based schedule, where calibrations are triggered either periodically or in response to a threshold violation in performance metrics such as RB or XEB [4, 12, 18]. While effective for small- to medium-scale systems, such reactive approaches fail to exploit statistical correlations between parameters or predict impending degradation. Moreover, as the number of qubits grows, sequential calibration becomes increasingly infeasible due to the combinatorial explosion of dependent parameters and limited system uptime [6, 17, 19].

2.3 ML for Calibration Scheduling

Machine learning (ML) methods have recently been explored to improve the efficiency of experimental calibration [20–22]. Bayesian Optimization (BO) has been applied to tune individual parameters, leveraging probabilistic surrogate models to reduce the number of required experiments [20, 23]. However, most existing BO approaches assume independent parameter spaces and ignore time-dependent uncertainty, limiting their scalability and robustness in dynamic environments. Reinforcement Learning (RL) has also been introduced in experiment design and control optimization, enabling data-driven decision policies that adapt to real-time feedback [21, 22, 24]. Yet, prior RL-based calibration methods often lack interpretability and are difficult to integrate with existing experiment orchestration systems. In addition, neither BO nor RL alone can effectively handle cross-qubit dependencies or time-sensitive scheduling constraints [6, 18].

2.4 Motivation

The growing scale of quantum processors demands an integrated calibration framework that is both *time-aware* and *structure-aware* [4, 6, 17, 18]. Existing solutions either focus on optimizing single-parameter tuning (BO) or on reactive decision-making (RL), without

modeling the underlying dependency graph or the temporal degradation process [4, 18, 20–22]. Furthermore, calibration scheduling under limited experimental resources, such as cooldown duration, readout conflicts, and equipment contention, remains largely unexplored [6, 17]. *This work addresses these gaps by introducing a unified framework that combines Bayesian Optimization and Reinforcement Learning atop a Directed Acyclic Graph (DAG) of calibration dependencies.* The proposed system not only models parameter uncertainty and temporal drift but also jointly optimizes calibration order, frequency, and concurrency under hardware constraints [25–28]. Our work formulates quantum calibration as a *time-aware active learning and scheduling* problem, bridging quantum control and design automation perspectives.

3 Time-Aware Active Calibration Framework

This section presents the proposed *Time-Aware Active Calibration (TAAC)* framework, which integrates Bayesian Optimization (BO) [29] and Reinforcement Learning (RL) [30] to enable adaptive and uncertainty-aware calibration scheduling for multi-qubit systems. The overall architecture is illustrated in Figure 1, comprising five key components, including (i) DAG construction and dependency extraction, (ii) temporal uncertainty modeling, (iii) experiment selection via BO, (iv) cross-qubit scheduling via RL, and (v) self-triggered recalibration.

3.1 DAG-Based Calibration Graph Construction

As shown in Figure 3, each calibration parameter θ_i (e.g., qubit frequency, pulse amplitude, DRAG coefficient, two-qubit phase offset) is represented as a node in a directed acyclic graph (DAG) [18]. Edges encode dependency relationships such that $(\theta_i \rightarrow \theta_j)$ implies that parameter θ_j depends on the calibration outcome of θ_i . These dependencies are automatically extracted from *Qiskit Experiments* metadata, which includes calibration lineage and experiment configuration. The DAG serves as both a structural prior and a scheduling constraint. For example, tuning a two-qubit gate parameter requires the underlying single-qubit parameters to be up-to-date. This hierarchical structure allows efficient propagation of calibration updates and facilitates topological ordering for concurrent experiment execution.

3.2 Temporal Uncertainty Modeling

Calibration parameters drift over time due to environmental noise and device aging. We model the temporal evolution of each parameter as a stochastic process with mean-reverting dynamics:

$$d\theta_t = -\kappa(\theta_t - \mu) dt + \sigma dW_t, \quad (1)$$

where κ is the reversion rate, μ is the equilibrium value, σ controls drift volatility, and W_t denotes Wiener noise. This formulation corresponds to an Ornstein-Uhlenbeck (OU) process, which captures both short-term fluctuations and long-term stability.

For parameters that exhibit monotonic decay (e.g., amplitude gain), we alternatively employ an exponential decay model

$$\theta_t = \theta_0 e^{-\lambda t} + \epsilon_t, \quad (2)$$

where λ denotes the decay constant and ϵ_t represents residual uncertainty. Each parameter maintains a posterior confidence interval




Figure 3: DAG.




Figure 4: BO and RL.

$(\hat{\theta}_t, \Sigma_t)$ that quantifies the predicted deviation from its last calibrated state. These uncertainty bounds are later used by the BO and RL modules to determine recalibration priority.

3.3 Experiment Selection via BO

In Figure 4(a), the BO agent prioritizes the next calibration target using a time-aware acquisition score that balances expected improvement, predictive uncertainty, and experiment duration. Let $\mathcal{L}(\theta)$ represent the calibration loss function (e.g., gate infidelity [31]), where θ denotes a candidate target considered by the BO module. The BO objective aims to minimize $\mathcal{L}$ while penalizing experimental cost

$$\theta^* = \arg \min_{\theta \in \Theta} \mathbb{E}[\mathcal{L}(\theta)] + \beta C(\theta), \quad (3)$$

where $C(\theta)$ is the estimated time cost and β balances fidelity improvement against runtime. A Gaussian Process (GP) surrogate model is used to approximate $\mathcal{L}$, with posterior mean $\mu(\theta)$ and variance $\sigma^2(\theta)$. The next experiment is chosen using a modified acquisition function that incorporates time-awareness

$$a_t(\theta) = \frac{\mu_{best} - \mu(\theta)}{T_{exp}(\theta)} + \alpha \frac{\sigma(\theta)}{T_{exp}(\theta)}, \quad (4)$$

where $T_{exp}(\theta)$ denotes expected experiment duration. This expression serves as a practical time-normalized acquisition proxy that favors targets with high expected benefit, high uncertainty, and low execution cost.

3.4 Cross-Qubit Scheduling via RL

The RL agent coordinates multiple calibration tasks in parallel under resource constraints. The environment state s_t encodes (i) current parameter uncertainties Σ_t , (ii) DAG dependency satisfaction, (iii) hardware resource availability (cooldown status, readout conflicts [32]), (iv) recent calibration outcomes. The agent selects an action a_t corresponding to a batch of calibration tasks. In Figure 4(b), the reward function balances accuracy improvement and temporal efficiency

$$r_t = \Delta F_{avg} - \lambda_1 T_{busy} - \lambda_2 R_{conflict}, \quad (5)$$

where ΔF_{avg} is the fidelity improvement, T_{busy} denotes total calibration time, and $R_{conflict}$ penalizes concurrent resource contention. We initialize the scheduler using historical experiment traces, followed by fine-tuning with Proximal Policy Optimization (PPO) [33].

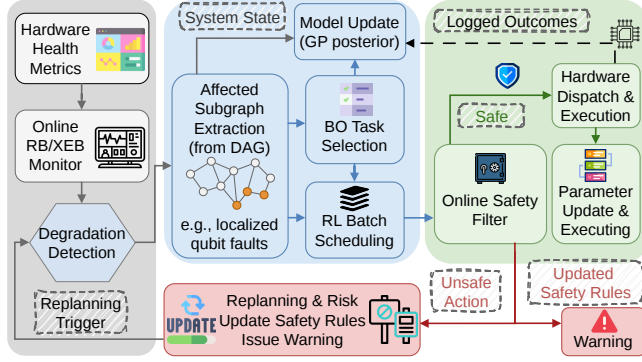


Figure 5: Event-triggered recalibration with safety filtering.

3.5 Self-Triggered Recalibration Mechanism

To maintain long-term stability, the system continuously monitors online health metrics such as randomized benchmarking (RB) [34] and cross-entropy benchmarking (XEB) [35]. When a fidelity degradation signal exceeding the confidence threshold is detected, the recalibration sequence is automatically re-triggered along the affected subgraph of the DAG, as illustrated in Figure 5. This event-driven mechanism enables *self-healing* behavior, ensuring consistent system performance without manual intervention. As shown in Figure 5, the recalibration workflow further includes a lightweight online safety filter before hardware dispatch. Given a candidate calibration action proposed by the BO/RL controller, the filter checks whether the predicted post-update state violates predefined safety conditions, such as severe fidelity degradation, abnormal parameter deviation, or repeated recent failures. If the action is flagged as unsafe, the system issues a warning and requests conservative re-planning; otherwise, the action is executed normally. This design serves as a practical safeguard during online recovery, rather than as an independent optimization module. We develop the Cal Orchestrator framework and implement the prototype system of TAAC, integrating DAG extraction, BO experiment selection, and RL scheduling. The orchestration logic interacts with existing experiment APIs to dispatch calibration jobs asynchronously. The system records calibration intervals, confidence bounds, and metadata to continuously refine its uncertainty model over time.

4 Experimental Evaluation

We evaluate the proposed *Time-Aware Active Calibration (TAAC)* framework using both quantum hardware simulators and real multi-qubit devices. The experiments aim to quantify calibration efficiency, stability, and system availability compared with representative baseline methods. All results are averaged over five independent runs with different random seeds, and standard deviations were below 3% for all metrics, confirming statistical robustness.

4.1 Experimental Setup

Experiments are conducted on the superconducting quantum processor using the Qiskit Experiments framework [2]. Each qubit was characterized by four primary calibration parameters: operating frequency, amplitude scaling, DRAG coefficient, and readout

Table 1: Overall performance comparison on 20-Qubit device.

Method	TCT (min)	MTBC (h)	CSR (%)	RU (%)	FRR (%)
Fixed Interval	310	7.5	85.3	68.4	2.9
Pure BO	245	9.1	89.7	71.5	3.5
Pure RL	228	10.2	91.0	74.8	3.6
TAAC (Ours)	182	14.5	95.8	81.2	4.1

phase [11, 13]. Two-qubit gates were additionally calibrated for coupling strength and relative phase alignment.

Baselines. We compare TAAC against three baseline methods.

- *Fixed Interval (FI)* [12]: periodic calibration with static scheduling;
- *Pure BO* [20]: Bayesian Optimization without temporal or structural modeling;
- *Pure RL* [21]: reinforcement-based scheduling without uncertainty quantification.

Each experiment ran continuously, covering five full calibration cycles, and results were averaged across all qubits.

Metrics. The following metrics were used to assess performance [4–6]. (i) *Total Calibration Time (TCT)*: the elapsed wall-clock time of a calibration episode, measured from the first dispatched calibration task to the completion of the last task; (ii) *Mean Time Between Calibrations (MTBC)*: average duration before degradation triggers recalibration; (iii) *Calibration Success Rate (CSR)*: fraction of calibrations that achieved target fidelity; (iv) *Resource Utilization (RU)*: the average fraction of available calibration resources that are actively occupied during the calibration episode; (v) *Fidelity Recovery Rate (FRR)*: average post-calibration improvement in gate fidelity.

4.2 Overall Performance Comparison

Table 1 summarizes the end-to-end performance across all methods. TAAC achieves the shortest calibration time and the highest long-term stability. Specifically, TAAC reduces total calibration time by 41% compared to fixed scheduling and extends MTBC by nearly 2×, demonstrating a favorable efficiency-stability trade-off. The improved CSR and RU indicate better utilization of experimental bandwidth and reduced idle time, translating to proportional reductions in cryostat active duration and operational cost.

Table 2: Temporal drift prediction accuracy.

Parameter	Linear Fit	Exp. Decay	OU Process (ours)
Frequency (MHz)	0.74	0.52	0.31
Amplitude (arb.)	0.082	0.061	0.038
DRAG Coefficient	0.019	0.017	0.011
Phase Offset (deg)	2.46	1.91	1.12

4.3 Temporal Drift Modeling Accuracy

We next evaluate the effectiveness of the temporal uncertainty model. Table 2 reports the mean absolute prediction error under different drift assumptions. The OU process achieves the lowest error across all calibration parameters. The OU-based temporal model captures both mean reversion and stochastic volatility dynamics,




Figure 6: Scheduling efficiency and resource conflicts.




Figure 7: Ablation study on key TAAC components.

yielding tighter uncertainty bounds that guide the BO and RL agents toward more informed scheduling decisions. These results confirm that modeling temporal correlation is crucial for anticipating calibration degradation and for proactively allocating experimental resources in TAAC.

4.4 Scheduling and Conflict Avoidance

Cross-qubit scheduling efficiency was quantified by the number of calibration conflicts (overlapping resource usage) and mean idle time per qubit. Results in Figure 6 show that TAAC minimizes both metrics through RL-based coordination. Qualitative inspection of the RL scheduling traces shows that the policy learns to batch calibrations of weakly coupled qubits while deferring high-dependency nodes to later rounds. This behavior minimizes readout contention and cooldown conflicts under DAG constraints, confirming that the learned policy generalizes beyond purely heuristic scheduling. In addition, TAAC improves overall experiment throughput by nearly 2× compared to static scheduling, highlighting the benefit of combining uncertainty-aware modeling with adaptive RL-driven orchestration.

4.5 Ablation Study

To quantify the contribution of each TAAC component, we performed an ablation study by disabling uncertainty modeling, BO selection, or RL scheduling. As visualized in Figure 7, removing either BO or RL modules significantly degrades both efficiency and stability, confirming their complementary roles in decision-making and resource allocation. Notably, the absence of uncertainty modeling leads to faster performance decay and more frequent recalibrations, underscoring the importance of temporal awareness

Table 3: Scalability analysis with increasing qubit count.

System Size	Normalized TCT (↓)	MTBC (h)
8 Qubits	1.00	15.2
20 Qubits	1.42	14.5
50 Qubits	2.01	13.8

Table 4: Generalization and robustness evaluation under noise and device shift.

Scenario	Noise Scale (×)	TCT (min)	MTBC (h)	CSR (%)
Nominal (Baseline)	1.0×	182	14.5	95.8
Injected Drift Noise	1.5×	193	13.9	94.2
Injected Drift Noise	2.0×	208	12.8	91.7
Cross-Device (20Q→20Q)	–	189	13.8	94.9
Cross-Device (20Q→50Q)	–	197	13.1	93.5

for maintaining long-term system robustness. These observations highlight that TAAC’s hybrid design, combining predictive uncertainty, data-efficient exploration, and adaptive scheduling, is essential for achieving stable, low-overhead operation in dynamic quantum environments.

4.6 Scalability with Qubit Count

Finally, we analyze scalability by simulating systems with 8, 20, and 50 qubits. Table 3 reports normalized total calibration time and MTBC. TAAC scales near-linearly with qubit count, demonstrating its suitability for medium-scale quantum processors. This near-linear growth indicates that the hierarchical DAG scheduling and localized recalibration mechanisms effectively mitigate combinatorial explosion, suggesting promising extensibility toward future 100+ qubit and multi-cryostat systems. In addition, the stable MTBC across different scales implies that TAAC maintains consistent calibration quality even as system complexity increases, a key requirement for practical large-scale quantum deployment.

4.7 Generalization and Robustness Evaluation

To further assess TAAC’s robustness and transferability, we introduced controlled perturbations into the temporal drift model and evaluated performance across distinct hardware instances. As summarized in Table 4, TAAC maintains over 90% calibration success rate even under 2× drift volatility, with only a modest 14% increase in total calibration time compared to nominal conditions. When transferring the learned policy from a 20-qubit system to unseen devices, TAAC achieves comparable stability with less than 5% degradation in MTBC on 20Q→20Q, demonstrating strong cross-device generalization. These results confirm that the temporal uncertainty modeling and RL-driven orchestration jointly enhance adaptability to dynamic or heterogeneous quantum environments.

4.8 System Profiling and Overhead Analysis

As summarized in Table 5, the controller-side cost remains modest during 20-qubit operation. The per-component wall-time shares are measured non-exclusively, since BO inference, RL policy evaluation, DAG checking, and experiment dispatch are executed asynchronously and may overlap in time. We separately report the end-to-end controller overhead as the net increase in calibration

Table 5: System profiling and overhead breakdown.

Component	Avg. Latency (s)	CPU Util. (%)	Memory (MB)	Time Share (%)
BO Inference	1.84	21.3	312	2.1
RL Policy	2.12	24.6	405	2.3
DAG & Depend. Check	0.97	14.8	198	1.1
Experiment Dispatch	0.68	12.2	144	0.9
Total Controller Overhead	-	-	-	4.8

wall time attributable to the controller, which is 4.8%. In addition, the reported decision latency refers to the end-to-end critical-path latency per scheduling round, rather than the sum of standalone module invocation latencies, and remains below 2.5s on average.

The experimental results collectively demonstrate that TAAC achieves a favorable trade-off between calibration efficiency and system stability. By explicitly modeling temporal drift and structural dependencies, TAAC effectively eliminates redundant calibrations while improving utilization of experimental resources. Compared with BO-only or RL-only baselines, the integrated design yields faster convergence and smoother recovery after drift events. The learned RL scheduling policy shows interpretable behavior, suggesting that reinforcement-driven orchestration can complement analytic optimization methods. The observed near-linear scalability further indicates that TAAC can extend to large-scale, multi-cryostat quantum systems, paving the way for fully autonomous calibration orchestration.

5 Discussion

5.1 Design Insights

The results presented in § 4 demonstrate that the proposed *Time-Aware Active Calibration (TAAC)* framework significantly improves both calibration efficiency and system availability for multi-qubit quantum processors. The integration of Bayesian Optimization (BO) and Reinforcement Learning (RL) provides a principled mechanism to handle the trade-off between fidelity improvement and calibration overhead [6, 20–22]. From a design automation perspective, TAAC highlights the importance of combining *uncertainty quantification*, *temporal modeling*, and *graph-based dependency management* [5, 11]. The directed acyclic graph (DAG) abstraction not only captures parameter interdependence but also serves as a scheduling constraint that can be exploited by learning-based agents [18]. This design pattern may generalize beyond quantum calibration to other scientific automation tasks, such as optical alignment or analog circuit retuning [36, 37]. Furthermore, the time-aware acquisition function introduced in our BO formulation improves calibration efficiency by favoring high-value and high-uncertainty targets under limited execution time [20]. The reinforcement-based scheduler demonstrates that real-time decision policies can effectively minimize idle time and resource conflicts, which are common issues in shared experimental infrastructures [38, 39].

5.2 Cross-Domain Applicability

Beyond quantum calibration, the design principles of TAAC combine DAG-structured dependency modeling with uncertainty-aware learning. These technologies could generalize to other domains such as analog circuit retuning, robotic experimental planning,

or hardware-software co-design workflows [40–42]. This cross-domain applicability underscores the broader relevance of integrating reinforcement learning and Bayesian optimization within design automation frameworks.

5.3 Limitations

While TAAC achieves considerable improvements, several limitations remain. First, the accuracy of the temporal drift model depends on the choice of process prior (*e.g.*, OU vs. exponential decay), which may vary across devices and environmental conditions [4, 8]. In highly non-stationary systems, model mis-specification could reduce prediction fidelity. Second, the RL scheduler requires an initial warm-up phase to explore valid calibration sequences; during this phase, performance may be suboptimal [21, 43, 44]. Third, the current implementation assumes relatively stable device topology and fixed calibration dependencies; extending the framework to dynamically evolving system architectures remains a challenge [5, 6]. In addition, while the present study focuses on superconducting qubits, the approach could be adapted to other platforms such as trapped ions or photonic systems, though additional domain adaptation may be required to capture distinct noise and drift characteristics [4].

5.4 Future Work

Future directions include hierarchical reinforcement learning for large-scale calibration across multiple cryostats [38, 39], and meta-learning techniques that enable policy transfer across different hardware generations [43, 44]. Integrating model-based RL with probabilistic drift prediction may further reduce data requirements during policy adaptation. Another promising avenue is the incorporation of hardware-in-the-loop optimization, where TAAC directly interfaces with control electronics to perform closed-loop parameter updates [22, 37]. This could enable *real-time calibration orchestration*, moving toward the vision of fully autonomous quantum system management.

6 Conclusion

This work presents *TAAC*, a unified framework that leverages Bayesian Optimization and Reinforcement Learning for time-aware, structure-guided calibration of N-qubit systems. By modeling temporal drift, parameter uncertainty, and inter-parameter dependencies, TAAC achieves substantial gains in calibration efficiency, stability, and system throughput. Experimental results demonstrate over 40% reduction in calibration time and nearly twofold increase in mean time between calibrations compared to baseline methods. Beyond the quantum domain, TAAC exemplifies how design automation and machine learning can jointly address complex, temporally coupled optimization problems. We believe this approach lays the groundwork for the next generation of intelligent calibration systems that autonomously maintain and optimize quantum hardware performance at scale.

Acknowledgments

This research was supported by the CPS-Yangtze Delta Region Industrial Innovation Center of Quantum and Information Technology-MindSpore Quantum Open Fund and Sponsored by CCF-Kuaishou Large Model Explorer Fund (NO. CCF-KuaiShou 2025001).

References

- [1] Miroslav Urbanek, Benjamin Nachman, Vincent R Pascuzzi, Andre He, Christian W Bauer, and Wibe A de Jong. Mitigating depolarizing noise on quantum computers with noise-estimation circuits. *Physical Review Letters*, 127(27):270502, 2021.
- [2] Naoki Kanazawa, Daniel J. Egger, Yael Ben-Haim, Helena Zhang, William E. Shanks, Gadi Aleksandrowicz, and Christopher J. Wood. Qiskit experiments: A python package to characterize and calibrate quantum computers. *Journal of Open Source Software*, 8(84):5329, 2023.
- [3] Malcolm Carroll, Sami Rosenblatt, Petar Jurcevic, Isaac Lauer, and Abhinav Kandala. Dynamics of superconducting qubit relaxation times. *npj Quantum Information*, 8:132, 2022.
- [4] Timothy Proctor, Melissa Revelle, Erik Nielsen, Kenneth Rudinger, Daniel Lobser, Peter Maunz, Robin Blume-Kohout, and Kevin Young. Detecting and tracking drift in quantum information processors. *Nature communications*, 11(1):5396, 2020.
- [5] Yuchen Zhu, Jinglei Cheng, Boxi Li, Kecheng Liu, Yidong Zhou, Hanrui Wang, Yufei Ding, and Zhiding Liang. Hardware-aware calibration protocol for quantum computers. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture*, pages 241–256, 2025.
- [6] Xiang Fang, Keyi Yin, Yuchen Zhu, Jixuan Ruan, Dean Tullsen, Zhiding Liang, Andrew Sornborger, Ang Li, Travis S. Humble, Yufei Ding, and Yunong Shi. Caliqec: In-situ qubit calibration for surface code quantum error correction. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture (ISCA '25)*, pages 1402–1416, 2025.
- [7] Fabrizio Berritta, Jacob Benestad, Lukas Pahl, Melvin Mathews, Jan A Krzywda, Réouven Assouly, Youngkyu Sung, David K Kim, Bethany M Niedzielski, Kyle Serniak, et al. Efficient qubit calibration by binary-search hamiltonian tracking. *PRX Quantum*, 6(3):030335, 2025.
- [8] MJ Gullans, Miguel Caranti, AR Mills, and JR Petta. Compressed gate characterization for quantum devices with time-correlated noise. *PRX Quantum*, 5(1):010306, 2024.
- [9] Q-CTRL. Boulder opal: Build superior quantum computers. <https://q-ctrl.com/boulder-opal>.
- [10] QCAL Foundation. Qcal: Quantum calibration, certification & standards. <https://qcal.org/>.
- [11] Nicolas Wittler, Federico Roy, Kevin Pack, Max Werninghaus, Anurag Saha Roy, Daniel J Egger, Stefan Filipp, Frank K Wilhelm, and Shai Machnes. Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits. *Physical review applied*, 15(3):034080, 2021.
- [12] Julian Kelly, Rami Barends, et al. Scalable in situ qubit calibration during repetitive error detection. *Physical Review A*, 94(3):032321, 2016.
- [13] Yilun Xu, Gang Huang, Jan Balewski, Ravi K. Naik, Alexis Morvan, Brad Mitchell, Kasra Nowrouzi, David I. Santiago, and Irfan Siddiqi. Automatic qubit characterization and gate optimization with qubic. *ACM Transactions on Quantum Computing*, 4(1):1–12, 2022.
- [14] Xiao-Yan Yang, Hai-Feng Zhang, Lei Du, Hao-Ran Tao, Liang-Liang Guo, Tian-Le Wang, Zhi-Long Jia, Wei-Cheng Kong, Zhao-Yun Chen, Peng Duan, et al. Fast, universal scheme for calibrating microwave crosstalk in superconducting circuits. *Applied Physics Letters*, 125(4), 2024.
- [15] Adam Winick, Joel J. Wallman, and Joseph Emerson. Simulating and mitigating crosstalk. *Physical Review Letters*, 126(23):230502, 2021.
- [16] Simone Cialdi, Matteo A. C. Rossi, Claudia Benedetti, Bassano Vacchini, Dario Tamascelli, Stefano Olivares, and Matteo G. A. Paris. All-optical quantum simulator of qubit noisy channels. *Applied Physics Letters*, 110(8):081107, 2017.
- [17] Xiang Fang, Keyi Yin, Yuchen Zhu, Jixuan Ruan, Dean Tullsen, Zhiding Liang, Andrew Sornborger, Ang Li, Travis S. Humble, Yufei Ding, and Yunong Shi. CalScalpel: In-situ and fine-grained qubit calibration integrated with surface code quantum error correction. *arXiv preprint arXiv:2412.02036*, 2024.
- [18] Julian Kelly, Peter O'Malley, Matthew Neeley, Hartmut Neven, and John M Martinis. Physical qubit calibration on a directed acyclic graph. *arXiv preprint arXiv:1803.03226*, 2018.
- [19] Yilun Xu, Gang Huang, Jan Balewski, Ravi Naik, Alexis Morvan, Bradley Mitchell, Kasra Nowrouzi, David I Santiago, and Irfan Siddiqi. Qubic: An open-source fpga-based control and measurement system for superconducting quantum information processors. *IEEE Transactions on Quantum Engineering*, 2:1–11, 2021.
- [20] Cristian L. Cortes, Pascal Lefebvre, Nikolai Lauk, Michael J. Davis, Neil Sinclair, Stephen K. Gray, and Daniel Oblak. Sample-efficient adaptive calibration of quantum networks using Bayesian optimization. *Physical Review Applied*, 17(3):034067, 2022.
- [21] Murphy Yuezhen Niu, Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven. Universal quantum control through deep reinforcement learning. *npj Quantum Information*, 5(1):33, 2019.
- [22] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret. Model-free quantum control with reinforcement learning. *Physical Review X*, 12(1):011059, 2022.
- [23] Kim A. Nicoli, Christopher J. Anders, Lena Funcke, Tobias Hartung, Karl Jansen, Stefan Kühn, Klaus-Robert Müller, Paolo Stornati, Pan Kessel, and Shinichi Nakajima. Physics-informed Bayesian optimization of variational quantum circuits. 2023.
- [24] Yuval Baum, Mirko Amico, Sean Howell, Michael Hush, Maggie Liuzzi, Pranav Mundada, Thomas Merkh, Andre R. R. Carvalho, and Michael J. Biercuk. Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer. *PRX Quantum*, 2(4):040324, 2021.
- [25] Qingcheng Xiao, Size Zheng, Bingzhe Wu, Pengcheng Xu, Xuehai Qian, and Yun Liang. HASCO: Towards agile hardware and software co-design for tensor computation. In *Proceedings of the 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 1055–1068. IEEE, 2021.
- [26] Dan Zhang, Safeen Huda, Ebrahim Songhori, Kartik Prabhu, Quoc Le, Anna Goldie, and Azalia Mirhoseini. A full-stack search technique for domain optimized deep learning accelerators. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 27–42, 2022.
- [27] Erik Orm Hellsten, Artur Souza, Johannes Lenfers, Rubens Lacouture, Olivia Hsu, Adel Ejeh, Fredrik Kjolstad, Michel Steuwer, Kunle Olukotun, and Luigi Nardi. Baco: A fast and portable bayesian compiler optimization framework. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 4*, pages 19–42, 2023.
- [28] Chirag Sakhuja, Zhan Shi, and Calvin Lin. Leveraging domain information for the efficient automated design of deep learning accelerators. In *Proceedings of the IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, IEEE, 2023.
- [29] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [30] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 2 edition, 2018.
- [31] Vinay Tripathi, Daria Kowsari, Kumar Saurav, Haimeng Zhang, Eli M. Levenson-Falk, and Daniel A. Lidar. Benchmarking quantum gates and circuits. *Chemical Reviews*, 125(12):5745–5775, 2025.
- [32] Caroline Tornow, Naoki Kanazawa, William E Shanks, and Daniel J Egger. Minimum quantum run-time characterization and calibration via restless measurements with dynamic repetition rates. *Physical Review Applied*, 17(6):064061, 2022.
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [34] Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R Brad Blakestad, John D Jost, Chris Langer, Roe Ozeri, Signe Seidelin, and David J Wineland. Randomized benchmarking of quantum gates. *Physical Review A—Atomic, Molecular, and Optical Physics*, 77(1):012307, 2008.
- [35] Frank Arute, Kunal Arya, Ryan Babbush, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [36] Shuhan Zhang, Wenlong Lyu, Fan Yang, Changhao Yan, Dian Zhou, Xuan Zeng, and Xiangdong Hu. An efficient multi-fidelity bayesian optimization approach for analog circuit synthesis. In *Proceedings of the 56th Annual Design Automation Conference (DAC)*, 2019.
- [37] Xiaodong Wang, Changhao Yan, Yuzhe Ma, Bei Yu, Fan Yang, Dian Zhou, and Xuan Zeng. Analog circuit yield optimization via freeze-thaw bayesian optimization technique. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(11):4887–4900, 2022.
- [38] Lili Ren, Xin Ning, and Jiayin Li. Hierarchical reinforcement-learning for real-time scheduling of agile satellites. *IEEE Access*, 8:220523–220532, 2020.
- [39] Yoni Birman, Ziv Ido, Gilad Katz, and Asaf Shabtai. Hierarchical deep reinforcement learning approach for multi-objective scheduling with varying queue sizes. In *IJCNN*, pages 1–10. IEEE, 2021.
- [40] Zishen Wan, Bo Yu, Thomas Yuang Li, Jie Tang, Yuhao Zhu, Yu Wang, Arijit Raychowdhury, and Shaoshan Liu. A survey of fpga-based robotic computing. *IEEE Circuits and Systems Magazine*, 21(2):48–74, 2021.
- [41] Sajad Saeedi, Bruno Bodin, Harry Wagstaff, Andy Nisbet, Luigi Nardi, John Mawer, Nicolas Melot, Oscar Palomar, Emanuele Vespa, Tom Spink, et al. Navigating the landscape for real-time localization and mapping for robotics and virtual and augmented reality. *Proceedings of the IEEE*, 106(11):2020–2039, 2018.
- [42] Zhuolun He, Yuan Pu, Haoyuan Wu, Tairu Qiu, and Bei Yu. Large language models for eda: Future or mirage? *ACM Transactions on Design Automation of Electronic Systems*, 30(6):1–53, 2025.
- [43] Jacob Beck, Risto Vuorio, Evan Zheran Liu, Zheng Xiong, Luisa Zintgraf, Chelsea Finn, and Shimon Whiteson. A survey of meta-reinforcement learning. *arXiv e-prints*, pages arXiv–2301, 2023.
- [44] Zhuangdi Zhu, Kaixiang Lin, Anil K Jain, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):13344–13362, 2023.