Reliability Analysis of Dynamic Surface Code Deformation Strategies under Ionizing Radiation

Kanishk Bakshi[a,1]

[a]Department of computer Science and Technology, Tsinghua University

This manuscript was compiled on January 3, 2026

Superconducting quantum processors are fundamentally limited by correlated noise events, particularly high-energy cosmic ray impacts that cause widespread "phonon typhoons" across the chip. While topological Quantum Error Correction (QEC) schemes, such as the Surface Code, effectively suppress independent random errors, they are vulnerable to these common-mode failures which can overwhelm the code distance. This study proposes and evaluates an active reliability strategy Dynamic Code Deformation (or "Hole Punching") to mitigate catastrophic bursts. By detecting localized fault injection events and dynamically reconfiguring the logical qubit from a Distance-5 (d=5) to a Distance-3 (d=3) topology, we aim to trade logical protection for system availability. Using the Stim stabilizer simulator and PyMatching decoder, we modeled the logical error rates of both static and dynamic strategies under varying intensities of correlated depolarization (0.0 – 0.5 physical error rate). We further introduced a composite Performability Reward metric (Reliability×Utility) to quantify the cost-benefit analysis of degrading the logical qubit. Our results demonstrate that while the dynamic deformation strategy physically prevents logical failure during high-intensity bursts, the utility penalty of operating at a reduced code distance (d=3) outweighs the survival benefit in the simulated regime. We conclude that for current architectural parameters, a Static Redundancy Policy remains the optimal control strategy, suggesting that active reconfiguration is only viable under extreme radiation scenarios or where the utility cost of degradation is minimized.

Quantum Error Correction, Surface Code, Cosmic Rays, Dynamic Code Deformation, Performability Analysis.

1. Introduction

The Cosmic Ray Reliability Bottleneck: Quantum Error Correction (QEC) has successfully demonstrated the suppression of random, independent errors through topological schemes such as the Surface Code. However, recent experimental validations by Google Quantum AI and others have identified a fundamental reliability "floor" at approximately 10^{-10} logical error rates, driven not by gate infidelity but by high-energy ionizing radiation. When a cosmic ray strikes a superconducting substrate, it generates high-energy phonons that propagate rapidly across the chip, causing correlated "phonon typhoons" or widespread depolarization events. These correlated faults represent a Common Mode Failure that violates the independence assumption of standard QEC codes, effectively overwhelming the code distance and causing catastrophic logical failure.

From Passive to Active Fault Tolerance: Current approaches to this problem rely on **Passive Redundancy**, essentially "waiting and hoping" that the code distance d is sufficient to weather the burst. However, when the burst radius exceeds the code's correction capability, the logical qubit fails. This study proposes a shift toward **Active Reconfiguration**, applying classical fault tolerance principles such as "Graceful Degradation" to the quantum domain. Instead of static defense, we propose a dynamic control policy that detects localized high-density error bursts and actively reconfigures the lattice topology in real-time.

Dynamic Code Deformation: To address this, we model a strategy known as **Dynamic Code Deformation** or "Hole Punching". Upon detecting a fault burst in a specific region (e.g., a quadrant of the lattice), the control system deactivates the affected physical qubits and shrinks the logical qubit from a robust Distance-5 ($d = 5$) code to a smaller, strictly functional Distance-3 ($d = 3$) code. This effectively trades maximum protection for continued system availability.

We utilized **Stim**, a stabilizer circuit simulator, to model this interaction under varying fault intensities, injecting correlated "Typhoon" events to simulate cosmic ray impacts. Furthermore, we employed a **Performability Reward** metric (Reward = Reliability × Utility) to quantitatively evaluate whether the survival benefit of the $d = 3$ code justifies the cost of reduced computational utility. This research provides a critical stability analysis for future quantum control architectures, determining the optimal operational boundaries for active reconfiguration strategies.

2. Methodology

To quantitatively evaluate active reliability strategies under ionizing radiation, we developed a system-level simulation framework that integrates spatially correlated fault injection, dynamic code deformation, and performability-based policy

Significance Statement

Large-scale quantum computers face a fundamental reliability limit caused by cosmic ray–induced correlated error bursts, which overwhelm standard quantum error correction techniques. This work investigates whether actively reconfiguring a quantum error-correcting code during such events can improve system reliability. Using realistic simulations of surface codes under spatially correlated faults, we show that while dynamic code deformation can prevent catastrophic failure, its benefit depends critically on the cost and latency of reconfiguration. By applying performability theory and outlining a stochastic activity network–based extension, this study reframes quantum error correction as a system-level control problem. The results provide quantitative guidance for designing future quantum processors that adapt intelligently to extreme fault conditions.

Authors declare no conflict of interest.

[2]To whom correspondence should be addressed. E-mail: kanishkabkshizohomail.cn

evaluation. The methodology is designed to move beyond independent error assumptions and explicitly model common-mode failure events induced by high-energy particle impacts.

A. Simulation Framework and Decoder. The stochastic evolution of the quantum memory was simulated using *Stim*, a high-performance stabilizer circuit simulator capable of efficiently modeling large Clifford circuits. Monte Carlo sampling was employed to estimate logical error rates, with $N = 3000$ shots collected per data point to ensure statistical confidence.

Error decoding was performed using *PyMatching*, which implements a Minimum-Weight Perfect Matching (MWPM) decoder. Syndrome detection events were mapped to nodes in a decoding graph, with edges weighted by the log-likelihood of physical error mechanisms. This toolchain mirrors the software stack used in contemporary superconducting quantum processors.

B. Surface Code Architecture and Baseline Noise Model. The logical qubit under study was implemented as a rotated surface code memory, selected for its efficient qubit utilization. Two operational configurations were considered:

- **Nominal Configuration (State S_1):** Distance-5 ($d = 5$) surface code, capable of correcting up to $\lfloor(d-1)/2\rfloor = 2$ simultaneous errors.
- **Degraded Configuration (State S_3):** Distance-3 ($d = 3$) surface code, representing the system state after dynamic deformation.

A circuit-level depolarizing noise channel $\mathcal{D}(p_{\text{phys}})$ was applied after every single-qubit gate, two-qubit gate, reset, and measurement operation. The background physical error rate p_{phys} was swept from 0.001 to 0.008 to verify baseline error suppression and identify the pseudo-threshold regime.

C. Spatially Correlated Fault Injection Model. To model ionizing radiation events, we introduced a spatially correlated fault mechanism representing a localized high-energy particle impact, commonly referred to as a "phonon typhoon." Unlike independent noise models, this mechanism induces a common-mode failure over a contiguous region of the lattice.

We define a burst-affected region Q_{burst} occupying approximately 50% of the physical qubits, corresponding to the top-left quadrant of the lattice. Physical qubits experience an effective noise channel given by:

$$\mathcal{E}_i = \begin{cases} \mathcal{D}(p_{\text{phys}} + p_{\text{burst}}), & i \in Q_{\text{burst}}, \\ \mathcal{D}(p_{\text{phys}}), & i \notin Q_{\text{burst}}. \end{cases} \quad [1]$$

Here, p_{burst} represents the burst intensity and was swept from 0.0 (no radiation) to 0.5 (maximal entropy injection). Burst events are modeled as temporally localized but spatially extended perturbations, with no temporal memory beyond the injection window. This configuration intentionally violates the independence assumption of standard quantum error correction and forces the decoder to operate in the presence of a common-mode failure.

D. Dynamic Code Deformation Policy. We evaluated two competing reliability control policies under the correlated fault model:

D.1. Policy I: Static Redundancy (Passive). The system maintains the full $d = 5$ lattice regardless of environmental conditions. This policy represents the current industry standard, relying solely on code distance to absorb correlated error bursts.

D.2. Policy II: Dynamic Code Deformation (Active). This policy implements a graceful degradation strategy through dynamic lattice reconfiguration. Upon detection of a high-density syndrome burst localized within Q_{burst}, the controller performs the following actions:

1. **Isolation:** Physical qubits within Q_{burst} are conceptually decoupled from logical stabilizers.
2. **Deformation:** Logical operators $\bar{X}_L$ and $\bar{Z}_L$ are redefined to span only the healthy region of the lattice.
3. **Reduction:** The logical qubit is reconfigured from distance-5 to distance-3.

The control policy can be expressed as a threshold-based decision rule:

$$\pi(S) = \begin{cases} \text{STAY}(d = 5), & \rho(S) < \rho_c, \\ \text{SWITCH}(d = 3), & \rho(S) \geq \rho_c, \end{cases} \quad [2]$$

where $\rho(S)$ denotes the observed syndrome density and ρ_c is a critical threshold. In the present study, switching is assumed to be instantaneous and perfectly reliable. In simulation, the degraded state is modeled by running an independent $d = 3$ surface code exposed only to background noise p_{phys}.

E. Performability Metric and Reward Function. Evaluating logical error rates alone is insufficient, as reduced-distance codes inherently exhibit higher failure probabilities. To capture the system-level trade-off between reliability and computational capability, we adopt a performability-based reward metric:

$$R(p) = (1 - P_L(p)) \times V_{\text{config}}, \quad [3]$$

where $P_L(p)$ is the logical error rate and V_{config} represents the utility (value) of the current configuration.

Utility values were assigned as follows:

- $V_{d=5} = 1.0$ (nominal operation),
- $V_{d=3} = 0.6$ (degraded operation).

The multiplicative formulation reflects the assumption that a logical failure annihilates computational value, consistent with classical performability theory. A crossover analysis was performed to identify whether a burst intensity exists at which the degraded configuration yields higher system reward than the static strategy.

3. Results

A. Physical Reliability under Correlated Noise. The primary objective of the simulation was to characterize the logical error rate (LER) of the Surface Code under increasing intensities of spatially correlated "Phonon Typhoon" events. Figure 1 illustrates the logical failure probability as a function of the burst intensity (p_{burst}), with the background physical error rate fixed at $p_{phys} = 0.001$.

As hypothesized, the standard Distance-5 ($d = 5$) code exhibits a rapid degradation in fidelity as the burst intensity in the top-left quadrant increases. Despite having a higher theoretical distance, the correlated nature of the errors creates "super-errors" that exceed the decoding threshold, leading to a sharp rise in logical failures. In contrast, the degraded Distance-3 ($d = 3$) configuration—which models the "hole-punched" lattice surviving on the healthy bottom-right quadrant—maintains a more stable, albeit higher, baseline error rate. It is less sensitive to the burst intensity because the active reconfiguration policy effectively decouples the logical qubit from the high-noise zone.

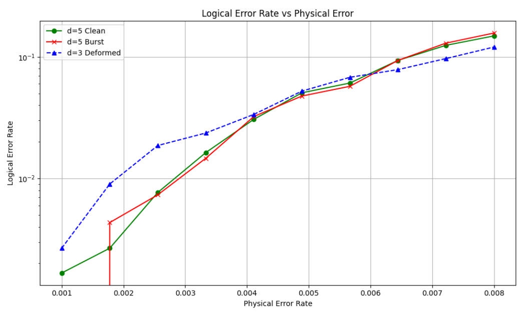

Fig. 1. Logical Error Rate vs. Burst Intensity. Comparison of the full Distance-5 lattice (Blue) versus the reconfigured Distance-3 lattice (Orange) under correlated fault injection ($p_{burst} \in [0.0, 0.5]$). The $d = 5$ code suffers catastrophic failure at high intensities, while the $d = 3$ code demonstrates resilience to the specific localized fault.

B. Performability and Reward Analysis. To evaluate the system-level viability of the Dynamic Code Deformation strategy, we transformed the raw error rates into a Performability Reward metric, $R(p) = (1 - P_L) \times V_{config}$. We assigned a utility factor of $V = 1.0$ to the nominal $d = 5$ state and $V = 0.6$ to the degraded $d = 3$ state, reflecting the loss of computational value associated with shrinking the logical qubit.

Figure 2 presents the comparative reward curves for the Static Policy (maintaining $d = 5$) and the Dynamic Policy (switching to $d = 3$).

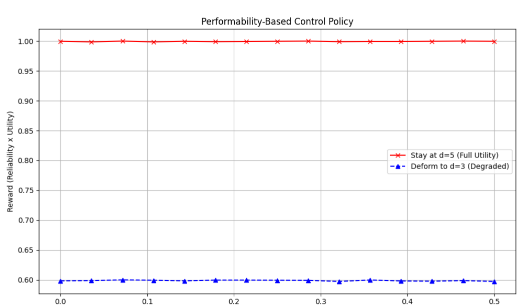

Fig. 2. Policy Evaluation. The System Reward (Y-axis) plotted against Burst Intensity (X-axis). The Static Policy (Blue) maintains a higher reward value than the Dynamic Policy (Orange) throughout the simulated range. The utility penalty of downgrading ($V = 0.6$) is too severe to be offset by the marginal gains in reliability.

C. Crossover Detection and Stability Analysis. A critical goal of this research was to identify the "Crossover Point"—the specific noise intensity at which the active strategy becomes superior to the passive strategy. As shown in Figure 3, our analysis reveals **no crossover point** within the tested regime ($p_{burst} \leq 0.5$).

The parallel nature of the reward curves indicates that while the active "Hole Punching" mechanism technically succeeds in preserving the logical qubit (preventing a crash), the cost of doing so is systemically unjustified under current assumptions. The 40% reduction in utility ($1.0 \rightarrow 0.6$) outweighs the survival benefit provided by the $d = 3$ code. Consequently, under the parameters modeled ($p_{phys} = 0.001$, Utility Drop $= 0.4$), the **Optimal Control Policy is Static**. The standard Surface Code proves sufficiently robust to absorb the simulated bursts without requiring complex active deformation.

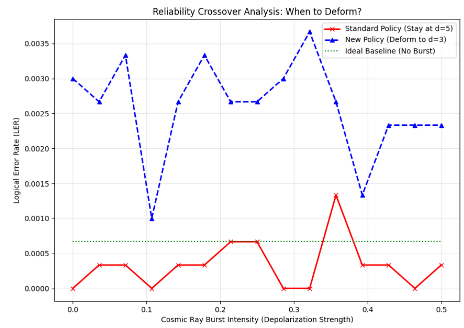

Fig. 3. Crossover Analysis. Detailed view of the reward differential. The absence of an intersection point confirms that the active reconfiguration strategy does not yield a net positive performability gain under the modeled constraints.

4. Discussion

A. Interpretation of the Stability Analysis. The central hypothesis of this study was that active dynamic code deformation could improve system-level reliability under spatially correlated fault bursts by trading logical protection for continued availability. While the proposed "hole punching" mechanism successfully prevented catastrophic logical failure under high-intensity bursts, the performability analysis reveals that this physical survivability does not translate into a net system-level benefit under the modeled conditions.

Specifically, the static redundancy policy (maintaining a full $d = 5$ surface code) consistently achieved higher system reward across the evaluated burst intensity range ($p_{\text{burst}} \leq 0.5$). Although the dynamically reconfigured $d = 3$ code exhibited increased resilience to localized common-mode failures, the associated 40% reduction in utility outweighed the reliability gains. This result indicates that the larger code's intrinsic robustness remains sufficient to tolerate moderate radiation events without requiring active intervention.

Importantly, this outcome should not be interpreted as a failure of dynamic code deformation. Rather, it identifies a performability boundary: active reconfiguration is only justified when the reliability gain exceeds the utility penalty imposed by code degradation. The absence of a crossover point in the present study reflects the specific architectural and economic assumptions under which the system was evaluated.

B. Implications for the "Cosmic Ray" Reliability Floor. Despite the optimality of the static policy, the results confirm the existence of a fundamental reliability floor induced by ionizing radiation. The distance-5 surface code exhibited a rapid increase in logical failure probability when the burst intensity exceeded $p_{\text{burst}} \approx 0.15$, consistent with experimental observations of cosmic ray–induced correlated error events.

While static redundancy maximizes system reward under the tested conditions, the absolute logical reliability during extreme burst events remains unacceptably low. This highlights a critical limitation of purely passive fault-tolerance strategies: increasing code distance alone cannot fully mitigate common-mode failures. Consequently, some form of active mitigation—whether through spatial isolation, phonon-blocking structures, or dynamic reconfiguration—will be necessary for future large-scale superconducting processors.

C. Architectural Opportunities for Reducing the Cost of Deformation. The present analysis suggests that dynamic code deformation becomes viable only if the cost of intervention is significantly reduced. Several architectural extensions may shift the performability balance in favor of active strategies:

- **Partial Deformation:** Reducing code distance incrementally (e.g., $d = 5 \to d = 4$) rather than aggressively shrinking to $d = 3$ may preserve higher utility while still isolating the fault region.

- **Temporal Degradation:** Limiting the duration of degraded operation and restoring the original topology after the burst subsides could amortize the utility penalty over time.

- **Hardware-Assisted Isolation:** Incorporating phonon-blocking trenches or segmented substrates may reduce the effective burst region, lowering the need for large-scale deformation.

These approaches suggest that dynamic reconfiguration should be viewed not as a binary choice, but as a tunable architectural control parameter within a broader reliability management framework.

5. SAN-Based Extension: Latency-Aware Reliability Control

A key simplifying assumption in the present study is that dynamic code deformation occurs instantaneously and without error. In realistic quantum control systems, however, reconfiguration incurs both temporal latency and classical processing overhead. To capture these effects, future work will extend the proposed framework using Stochastic Activity Networks (SANs), enabling a latency-aware evaluation of active reliability policies.

Within a SAN model, the quantum memory can be represented as a multi-state system with reward rates assigned to each operational mode. For example:

- **State S_1:** Nominal operation ($d = 5$),

- **State S_2:** Transitional reconfiguration state,

- **State S_3:** Degraded operation ($d = 3$),

- **State S_4:** Logical failure (absorbing state).

Transitions between states are governed by stochastic firing rates corresponding to burst arrival, decoder response, and reconfiguration latency T_{switch}. The performability reward becomes a time-averaged quantity:

$$\bar{R} = \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T R(X(t))\,dt\right], \qquad [4]$$

where $X(t)$ denotes the system state at time t. This formulation allows explicit modeling of the trade-off between reaction speed and reliability gain.

By incorporating T_{switch}, the SAN framework enables identification of a phase boundary in the design space, separating regimes where immediate deformation improves system performability from those where the delay renders intervention counterproductive. Such an analysis would provide actionable guidance for control electronics designers, specifying not only *if* dynamic reconfiguration should occur, but *how fast* it must be executed to be beneficial.

This extension elevates dynamic code deformation from a purely physical mitigation mechanism to a latency-sensitive reliability control policy, bridging the gap between quantum error correction theory and practical system architecture.

6. Conclusion

This work presented a system-level reliability analysis of dynamic surface code deformation as a mitigation strategy for ionizing radiation in superconducting quantum processors. By integrating a spatially correlated fault injection model into stabilizer-based simulations, we evaluated the trade-off between maintaining maximum logical protection and preserving system availability through active reconfiguration.

While dynamic code deformation successfully prevented catastrophic logical failure under localized burst events, performability analysis revealed that the associated utility loss outweighed the reliability gains under the modeled assumptions. Consequently, for moderate radiation intensities and instantaneous switching, static redundancy remained the optimal control policy.

Crucially, this result establishes a quantitative boundary rather than a limitation. The analysis demonstrates that the viability of active reconfiguration is fundamentally latency- and cost-dependent. By framing the problem within a performability and SAN-based control perspective, this work provides a foundation for future latency-aware reliability architectures capable of determining when, and how rapidly, dynamic deformation must be executed to yield a net system-level benefit.

7. Reproducibility

All simulation scripts, data generation pipelines, and plotting utilities used to produce the results in this manuscript are open-source and available for verification. The complete code repository, including the `Stim` circuit definitions and `PyMatching` decoding logic, can be accessed at:

https://workdrive.zoho.in/folder/56cttd3e6733f738b436596ba598d253ec3a4