

Planar Bosonic Error Correction with Heavy Fluxonium: A Quantitative Reassessment of Hardware-Efficient GKP Stabilization

Planar Bosonic Error Correction with Heavy Fluxonium: A Quantitative Reassessment of Hardware-Efficient GKP Stabilization

Abstract

Bosonic quantum error correction (QEC) encodes one logical qubit into an oscillator mode, promising fault tolerance with far fewer physical components than qubit-based surface codes. Until recently, however, all superconducting demonstrations of bosonic QEC relied on centimeter-scale three-dimensional microwave cavities controlled by fixed-frequency transmon qubits, whose bit-flip errors dominated the logical error budget. The recent demonstration of bosonic error correction with fluxonium [1],[3] changes this picture: a heavy fluxonium qubit with a bit-flip lifetime of $T_1^{\text{bf}} = 451 \pm 70 \mu\text{s}$ controls an on-chip Archimedean spiral resonator, and measurement-free (autonomous-style) error correction with rapid fluxonium reset extends the logical GKP lifetime by a factor of 1.59 ± 0.05 . In this paper we (i) place the result in the context of qubit-level and oscillator-level QEC theory, including continuous-time correction and beyond-qubit encodings; (ii) derive explicitly, from the reported numbers alone, the bit-flip rate budget, the corrected logical lifetime, the per-cycle error probability, and the residual (uncorrected) error channel strength; (iii) connect the result to a resource-commensurable comparison framework in which bosonic codes require $5\text{--}40\times$ fewer photons and roughly $100\times$ fewer modes than surface codes at a target logical error rate of $p_L = 10^{-6}$ [10]; and (iv) identify the failure modes — fluxonium dephasing, finite-energy GKP squeezing, reset latency, and spiral-resonator loss — that will govern whether planar bosonic QEC can scale. We argue that the fluxonium-GKP platform is currently the most credible route to planar, hardware-efficient bosonic QEC, while flagging the specific measurements that would falsify this assessment.

1. Introduction

The central promise of bosonic QEC is economy of hardware: instead of encoding one logical qubit into many two-level systems, one encodes into a single high-quality harmonic oscillator whose large Hilbert space (d -dimensional, with d set by the average photon number $\bar{n}$) absorbs errors in a structured, correctable way. The Gottesman–Kitaev–Preskill (GKP) code is the canonical example for protection against both position and momentum shifts, and it maps naturally onto the harmonic oscillator phase space [6]. Yet the practical history of bosonic QEC in superconducting circuits has been shaped by an awkward asymmetry: the oscillator (a 3D cavity) is long-lived, while the ancillary control qubit (a transmon) is the weak link, with bit-flip times of order tens of microseconds that cap the logical lifetime regardless of how good the bosonic code is.

The fluxonium-GKP experiment [1],[3] attacks exactly this asymmetry. Heavy fluxonium — a superconducting artificial atom with a large shunt capacitance and an inductive loop, whose low-frequency transitions suppress matrix elements for relaxation — achieves a bit-flip lifetime of $451 \pm 70 \mu\text{s}$, exceeding that of any control qubit used in prior bosonic-QEC demonstrations. Paired with an on-chip Archimedean spiral resonator (a planar structure with a mode volume several orders of magnitude smaller than a 3D cavity), it enables fully planar bosonic QEC: finite-energy GKP state preparation, stabilization via measurement-free error correction, and rapid fluxonium reset. The headline result is a logical lifetime extension factor of 1.59 ± 0.05 .

This paper is a quantitative reassessment, not a replication. Our contributions are:

1. **Error-budget arithmetic.** From the reported T_1^{bf} and extension factor, we derive the bit-flip rate, the corrected and uncorrected logical lifetimes, the per-cycle error probabilities, and the residual error-channel strength, with every arithmetic step shown (Section 4).
2. **Architectural context.** We situate the result within continuous-time QEC [2], beyond-qubit encodings [6], entanglement-assisted codes [4], and a resource-commensurable comparison of bosonic versus surface codes [10].
3. **Cross-platform analogy.** We draw a structured analogy to multi-spin molecular qubits, where a $T_2 \ll T_1$ hierarchy similarly dictates which error channel a correction cycle must target first [11].
4. **Falsifiability.** We state explicitly what measurements would undermine the claim that planar fluxonium-GKP is the leading route to scalable bosonic QEC (Section 6).

A methodological note on scope: this analysis uses only numbers reported in [1],[3] and publicly stated figures from [10],[11],[13]; we perform no simulations and report no new measurements. All derived quantities in Section 5 trace to Section 4 arithmetic.

2. Background and Related Work

The source experiment. Reference [1] (identical content to [3], the v2 record of the same work) reports the first fully planar superconducting implementation of bosonic QEC. The key hardware elements are (a) a heavy fluxonium qubit with $T_1^{\text{bf}} = 451 \pm 70 \mu\text{s}$ — an order of magnitude or more beyond the transmon bit-flip times that limited earlier bosonic demonstrations — and (b) an Archimedean spiral resonator, a planar spiral inductor-capacitor structure whose mode volume is several orders of magnitude smaller than the centimeter-scale 3D cavities used previously. The protocol prepares finite-energy GKP states, stabilizes them with measurement-free (feedback-free) error correction, and exploits rapid fluxonium reset so that the ancilla returns to its ground state quickly enough to act repeatedly on the oscillator. The measured logical lifetime extension is 1.59 ± 0.05 . This work is the direct object of our analysis and supplies every input number used in Section 4.

Beyond-qubit encodings. Reference [6], "Quantum error correction beyond qubits," provides the theoretical foundation for treating oscillators as native error-correcting substrates rather than as mere buses. It frames QEC for systems whose natural error model is continuous (small shifts in position or momentum, photon loss) rather than discrete bit flips, which is precisely the setting of GKP codes. Our argument in Section 6 — that bosonic codes are architecturally "native" for oscillator hardware — rests on this framing.

Continuous-time QEC. Reference [2] develops continuous-time quantum error correction (CTQEC), in which both noise and correction are modeled as processes continuous in time, using continuous weak measurement and feedback, analyzed through the subsystem principle (the idea that protected information lives in a tensor factor of the system's Hilbert space). The measurement-free stabilization used in [1] is a discrete-time cousin: rather than continuously weakly measuring error syndromes, the fluxonium performs repeated coherent, dissipative correction cycles. The CTQEC formalism supplies the right language for the limit in which cycle period T_{cyc} becomes short compared with the oscillator error timescales — a limit the planar platform is structurally positioned to approach because of its fast reset.

Entanglement-assisted codes. Reference [4] introduces entanglement-assisted quantum error-correcting codes (EAQECC), which allow the encoder to consume pre-shared entanglement with a reference to construct codes even when the stabilizer formalism alone would fail (e.g., for codes with non-commuting generator requirements). While [1] does not use entanglement assistance, EAQECC theory is relevant to a scaling question we raise in Section 6: whether multiple spiral-resonator/fluxonium cells could share entangled ancilla resources to boost the correctable error set beyond what standalone bosonic correction achieves.

Resource comparisons. Reference [10] (QNFO corpus) performs a resource-commensurable comparison of cat, GKP, binomial, and surface codes using a photons-per-logical-qubit metric at a target logical error rate $p_L = 10^{-6}$, finding that bosonic codes require $5\text{--}40\times$ fewer photons and roughly $100\times$ fewer modes than surface codes. This is the quantitative backbone of the "hardware-efficiency" claim in [1]'s abstract. We use [10]'s metric in Section 4 to convert the measured lifetime extension into a statement about photon budgets.

Cross-platform error hierarchies. Reference [11] analyzes two-level error correction for multi-spin molecular qubits, where fast pure dephasing (T_2) coexists with slower energy relaxation (T_1), so that a correction scheme must prioritize the dominant channel. The fluxonium-GKP system exhibits the mirror-image hierarchy: the dominant ancilla error is the *slow* bit flip ($T_1^{\text{bf}} = 451\ \mu\text{s}$), while oscillator errors (shifts, photon loss) are handled by the code itself. The structural lesson from [11] — that correction-cycle design should be commensurate with the slowest uncorrected channel — motivates our per-cycle error analysis in Section 4.

Auditing methodology. Reference [13] demonstrates a nine-stage audit pipeline for quantum-hardware advantage claims, including due diligence and external literature triangulation. We adopt its spirit here: every quantitative claim in this paper is either arithmetic derived from [1]'s reported numbers or an explicitly labeled projection with stated assumptions, and Section 6 is devoted to adversarial self-criticism.

Excluded sources. The bibliography supplied for this preprint contains four records ([5],[7],[8],[9]) that are administratively withdrawn or removed arXiv submissions unrelated to the topic (fictitious-content withdrawal, duplicate withdrawal, pseudonymous removal, and plagiarism withdrawal, respectively). We cite none of them substantively and note this as a corpus-quality limitation in Section 6. This leaves eight substantive sources, all of which we use.

3. Methods

Our method is *derived-metric analysis*: we take the small set of numbers reported in [1],[3] and propagate them through standard QEC error-budget relations to obtain quantities the experiment implies but does not state directly. No simulation is performed; every result is either exact arithmetic on reported inputs or a labeled projection with stated assumptions.

Inputs (all from [1],[3] unless noted):

- I_1 : ancilla bit-flip lifetime $T_1^{\text{bf}} = 451\ \mu\text{s}$, with uncertainty $\pm 70\ \mu\text{s}$.
- I_2 : logical lifetime extension factor $F = 1.59$, with uncertainty ± 0.05 .
- I_3 (from [10]): bosonic codes need $5\text{--}40\times$ fewer photons and $\sim 100\times$ fewer modes than surface codes at $p_L = 10^{-6}$.

Derived quantities and their defining relations:

1. *Ancilla bit-flip rate.* $\Gamma_{\text{bf}} = 1/T_1^{\text{bf}}$.
2. *Corrected logical lifetime.* If the uncorrected logical lifetime is $T_L^{(0)}$ and the extension factor is F , then $T_L^{(1)} = F T_L^{(0)}$. Since [1] reports F but not $T_L^{(0)}$ in the fetched abstract, we treat $T_L^{(0)}$ as a free scale and report all results as functions of it, or as ratios — this is the honest treatment given our input block.
3. *Per-cycle bit-flip probability.* For a correction cycle of duration T_{cyc} , the probability that the ancilla flips during one cycle, under a Poisson model with rate Γ_{bf} , is

$$p_{\text{bf}}(T_{\text{cyc}}) = 1 - e^{-\Gamma_{\text{bf}} T_{\text{cyc}}} \approx \Gamma_{\text{bf}} T_{\text{cyc}} \quad \text{for } \Gamma_{\text{bf}} T_{\text{cyc}} \ll 1.$$

1. *Residual error-channel strength.* The lifetime extension factor relates the uncorrected and corrected logical decay rates:

$$\Gamma_L^{(1)} = \frac{\Gamma_L^{(0)}}{F}, \quad \Gamma_L^{(i)} = \frac{1}{T_L^{(i)}}.$$

The fractional reduction in logical decay rate is $1 - 1/F$.

1. *Photon-budget projection.* Using [10]'s metric, the photon count per logical qubit at $p_L = 10^{-6}$ for GKP scales as $N_\gamma^{\text{GKP}} \in [N_\gamma^{\text{surf}}/40, N_\gamma^{\text{surf}}/5]$; we combine this with the measured F to project how much additional headroom the fluxonium ancilla buys relative to a transmon-ancilla baseline, under the explicit assumption (stated in Section 5) that the logical error budget is ancilla-bit-flip-dominated.

Uncertainty propagation. For a derived quantity $q = f(I_1, I_2)$ of independent inputs with uncertainties $\sigma_{I_1}, \sigma_{I_2}$, we use first-order propagation:

$$\sigma_q^2 = \left(\frac{\partial f}{\partial I_1} \right)^2 \sigma_{I_1}^2 + \left(\frac{\partial f}{\partial I_2} \right)^2 \sigma_{I_2}^2.$$

All uncertainties quoted in Section 5 come from this formula with the inputs' stated uncertainties.

4. Analysis

We now carry out every derivation explicitly.

4.1 Ancilla bit-flip rate

Input I_1 (source: [1],[3] abstract): $T_1^{\text{bf}} = 451 \mu\text{s} = 451 \times 10^{-6} \text{ s} = 4.51 \times 10^{-4} \text{ s}$.

$$\Gamma_{\text{bf}} = \frac{1}{T_1^{\text{bf}}} = \frac{1}{4.51 \times 10^{-4} \text{ s}} = 2.2173 \times 10^3 \text{ s}^{-1}.$$

So the ancilla bit-flip rate is $\Gamma_{\text{bf}} \approx 2.2173 \times 10^3 \text{ s}^{-1}$, i.e., about 2.2173 bit flips per millisecond of exposure.

Uncertainty. With $\sigma_T = 70 \mu\text{s} = 7.0 \times 10^{-5} \text{ s}$ and $f(T) = 1/T$, we have $\partial f / \partial T = -1/T^2$, so

$$\sigma_\Gamma = \frac{\sigma_T}{T^2} = \frac{7.0 \times 10^{-5}}{(4.51 \times 10^{-4})^2} = \frac{7.0 \times 10^{-5}}{2.034 \times 10^{-7}} = 3.441 \times 10^2 \text{ s}^{-1}.$$

Thus $\Gamma_{\text{bf}} = (2.22 \pm 0.34) \times 10^3 \text{ s}^{-1}$, a relative uncertainty of $0.344/2.217 \approx 15.5\%$, matching the relative uncertainty of T_1^{bf} itself ($70/451 = 0.15521064301552107$), as expected for a reciprocal quantity.

4.2 Fractional reduction in logical decay rate

Input I_2 (source: [1],[3]): $F = 1.59 \pm 0.05$.

The corrected logical decay rate is

$$\Gamma_L^{(1)} = \frac{\Gamma_L^{(0)}}{F} = \frac{\Gamma_L^{(0)}}{1.59} = 0.6289 \Gamma_L^{(0)}.$$

The fractional reduction in logical decay rate is

$$1 - \frac{1}{F} = 1 - 0.6289 = 0.3711,$$

i.e., error correction removes approximately 37.1% of the logical decay rate. Equivalently, the residual (uncorrected) fraction of the logical error budget is $1/F = 0.6289 \approx 62.9\%$.

Uncertainty. With $f(F) = 1/F$, $\partial f/\partial F = -1/F^2 = -1/(1.59)^2 = -1/2.5281 = -0.39555$, so

$$\sigma_{1/F} = 0.39555 \times 0.05 = 0.01978 \approx 0.020.$$

The residual fraction is therefore 0.629 ± 0.020 , and the removed fraction is 0.371 ± 0.020 .

4.3 Per-cycle ancilla bit-flip probability

The measurement-free protocol in [1] relies on the fluxonium returning to its ground state via rapid reset between correction interactions. The relevant question is: during one correction cycle of duration T_{cyc} , what is the probability that the ancilla has flipped?

Using the Poisson model with the rate from Section 4.1:

$$p_{\text{bf}}(T_{\text{cyc}}) = 1 - e^{-\Gamma_{\text{bf}} T_{\text{cyc}}}.$$

For the small-argument expansion to hold we need $\Gamma_{\text{bf}} T_{\text{cyc}} \ll 1$, i.e., $T_{\text{cyc}} \ll 1/\Gamma_{\text{bf}} = 4.51 \times 10^{-4} \text{ s} = 451 \mu\text{s}$.

Worked example at $T_{\text{cyc}} = 10 \mu\text{s}$ (a representative fast-reset cycle time; we state clearly that this cycle time is an assumption, not a reported number):

$$\Gamma_{\text{bf}} T_{\text{cyc}} = (2.2173 \times 10^3 \text{ s}^{-1})(1.0 \times 10^{-5} \text{ s}) = 2.2173 \times 10^{-2}.$$

$$p_{\text{bf}} = 1 - e^{-0.022173} = 1 - 0.97807 = 0.02193 \approx 2.19\%.$$

(Here $e^{-0.022173} \approx 1 - 0.022173 + 0.022173^2/2 = 1 - 0.022173 + 0.000246 = 0.978073$.)

For comparison, a transmon ancilla with a typical bit-flip time of $T_1^{\text{tm}} = 30 \mu\text{s}$ (a representative value for the fixed-frequency transmons used in prior 3D-cavity bosonic demonstrations, per the context in [1]; we label this an assumption) would give

$$\Gamma_{\text{bf}}^{\text{tm}} = \frac{1}{3.0 \times 10^{-5} \text{ s}} = 3.333 \times 10^4 \text{ s}^{-1},$$

$$\Gamma_{\text{bf}}^{\text{tm}} T_{\text{cyc}} = (3.333 \times 10^4)(1.0 \times 10^{-5}) = 0.3333,$$

$$p_{\text{bf}}^{\text{tm}} = 1 - e^{-0.3333} = 1 - 0.71653 = 0.28347 \approx 28.3\%.$$

(Here $e^{-0.3333} \approx 1 - 0.3333 + 0.05555 - 0.006173 + 0.000514 = 0.71653$.)

The ratio of per-cycle ancilla error probabilities is

$$\frac{p_{\text{bf}}^{\text{tm}}}{p_{\text{bf}}^{\text{flx}}} = \frac{0.28347}{0.02193} = 12.93 \approx 12.9.$$

Equivalently, the ratio of rates is simply the ratio of lifetimes:

$$\frac{\Gamma_{\text{bf}}^{\text{tm}}}{\Gamma_{\text{bf}}^{\text{flx}}} = \frac{T_1^{\text{bf}}}{T_1^{\text{tm}}} = \frac{451}{30} = 15.03 \approx 15.0,$$

with the small-argument expansions diverging slightly (12.9 vs 15.0) because the transmon case is no longer deep in the $\Gamma T \ll 1$ regime.

Uncertainty on $p_{\text{bf}}^{\text{flx}}$ at $T_{\text{cyc}} = 10 \mu\text{s}$. With $f(T) = 1 - e^{-T_{\text{cyc}}/T}$,
 $\partial f / \partial T = (T_{\text{cyc}}/T^2) e^{-T_{\text{cyc}}/T} = (10^{-5}/(4.51 \times 10^{-4})^2)(0.97807)$. Compute:
 $T^2 = 2.034 \times 10^{-7} \text{ s}^2$; $T_{\text{cyc}}/T^2 = 10^{-5}/2.034 \times 10^{-7} = 49.16 \text{ s}^{-1}$; times 0.97807 gives
 48.09 s^{-1} . Then

$$\sigma_p = 48.09 \times 7.0 \times 10^{-5} = 3.37 \times 10^{-3}.$$

So $p_{\text{bf}}^{\text{flx}} = 0.0219 \pm 0.0034$ at $T_{\text{cyc}} = 10 \mu\text{s}$.

4.4 Implied uncorrected logical lifetime scale

The experiment reports $F = 1.59$ but (in our fetched source material) not the absolute $T_L^{(0)}$. We can, however, bound the consistency of the result with the ancilla budget. If the corrected logical lifetime were entirely limited by residual ancilla bit flips, and if each ancilla flip during a cycle maps to a logical error with probability η (the "flip-to-logical" conversion factor, a property of the GKP code's finite-energy envelope), then

$$\Gamma_L^{(1)} \approx \eta \Gamma_{\text{bf}} p_{\text{bf}}(T_{\text{cyc}}) \cdot N_{\text{cyc}},$$

where $N_{\text{cyc}} = T_{\text{obs}}/T_{\text{cyc}}$ is the number of cycles in an observation time T_{obs} , so that per unit time the ancilla-induced logical rate is $\eta \Gamma_{\text{bf}}$ (the p_{bf} and N_{cyc} factors cancel: $p_{\text{bf}} \cdot N_{\text{cyc}}/T_{\text{obs}} = \Gamma_{\text{bf}}$). This cancellation is the key structural insight: in the steady state, the ancilla-induced logical error rate is independent of cycle speed and equals $\eta \Gamma_{\text{bf}}$ — cycle speed only matters for keeping up with oscillator errors, not for suppressing ancilla errors.

Setting $\Gamma_L^{(1)} = \eta \Gamma_{\text{bf}}$ and using $\Gamma_{\text{bf}} = 2.2173 \times 10^3 \text{ s}^{-1}$:

$$T_L^{(1)} = \frac{1}{\eta \Gamma_{\text{bf}}} = \frac{1}{\eta \times 2.2173 \times 10^3} \text{ s} = \frac{4.51 \times 10^{-4}}{\eta} \text{ s}.$$

For $\eta = 0.1$ (i.e., only 10% of ancilla flips corrupt the logical state — plausible for a GKP code with good shift-error tolerance, though we emphasize η is an assumption):

$$T_L^{(1)} = \frac{4.51 \times 10^{-4}}{0.1} = 4.51 \times 10^{-3} \text{ s} = 4.51 \text{ ms}.$$

This is a *projection*, clearly labeled as such: it assumes ancilla-bit-flip dominance of the corrected logical budget and $\eta = 0.1$.

4.5 Photon-budget connection

Input I_3 (source: [10]): at $p_L = 10^{-6}$, bosonic codes need 5–40 $\times$ fewer photons and $\sim 100\times$ fewer modes than surface codes.

The measured extension factor $F = 1.59$ multiplies the usable logical lifetime per photon budget. Under the assumption that logical lifetime scales linearly with the photon (i.e., energy) budget allocated — an assumption we state explicitly — the fluxonium ancilla effectively delivers a 1.59 $\times$ improvement in the photons-per-logical-qubit metric relative to a transmon-ancilla baseline at fixed p_L . Combined with [10]'s range, the effective advantage versus surface codes becomes

$$5 \times 1.59 = 7.95 \approx 8.0 \quad \text{to} \quad 40 \times 1.59 = 63.6 \approx 64$$

fewer photons per logical qubit, under the stated linearity assumption. This is a projection, not a measurement; its uncertainty is dominated by the $[5, 40]$ range from [10] and by the validity of the linearity assumption, which we critique in Section 6.

5. Results

All quantities below are computed in Section 4 from inputs I_1 and I_2 ([1],[3]) and I_3 ([10]), or are explicitly labeled projections.

R1 (computed). Ancilla bit-flip rate: $\Gamma_{\text{bf}} = (2.22 \pm 0.34) \times 10^3 \text{ s}^{-1}$, from $T_1^{\text{bf}} = 451 \pm 70 \mu\text{s}$.

R2 (computed). Error correction removes a fraction 0.371 ± 0.020 of the logical decay rate; the residual uncorrected fraction is 0.629 ± 0.020 (from $F = 1.59 \pm 0.05$).

R3 (computed, with one stated assumption). Per-cycle ancilla bit-flip probability at an assumed cycle time $T_{\text{cyc}} = 10 \mu\text{s}$: $p_{\text{bf}}^{\text{fx}} = 0.0219 \pm 0.0034$ ($2.19\% \pm 0.34\%$). Under the further assumption of a transmon baseline with $T_1^{\text{tm}} = 30 \mu\text{s}$, the per-cycle probability would be $p_{\text{bf}}^{\text{tm}} \approx 28.3\%$, a factor of ≈ 12.9 worse (rate ratio 15.0).

R4 (computed, structural). In the steady state, the ancilla-induced logical error rate is $\eta \Gamma_{\text{bf}}$, independent of cycle speed; faster reset improves the *transient* and the ability to track oscillator errors, but the floor is set by Γ_{bf} itself. With $\eta = 0.1$ (assumed), the ancilla-limited corrected logical lifetime would be $T_L^{(1)} \approx 4.51 \text{ ms}$ — a projection.

R5 (projection). Under the assumption that logical lifetime scales linearly with photon budget at fixed $p_L = 10^{-6}$, the fluxonium ancilla's $1.59\times$ lifetime extension translates the [10] photon-advantage range of $5\text{--}40\times$ into an effective $8.0\text{--}64\times$ fewer photons per logical qubit versus surface codes. Uncertainty: dominated by the $[5, 40]$ input range and the unquantified linearity assumption; we assign no numerical error bar to this projection.

R6 (computed). The reciprocal consistency check: the relative uncertainty of Γ_{bf} (15.5%) equals that of T_1^{bf} ($70/451 = 15.5\%$), confirming internal consistency of the propagation.

No simulation results, no empirical measurements beyond those reported in [1],[3], and no numbers outside Sections 4–5 appear in this paper.

6. Discussion

Strengths of the platform. The decisive architectural fact is that the ancilla is no longer the bottleneck. With $\Gamma_{\text{bf}} \approx 2.22 \times 10^3 \text{ s}^{-1}$, the fluxonium's bit-flip channel is weak enough that the logical budget can, in principle, be dominated by oscillator-side errors (photon loss, dephasing-induced shifts, finite-energy GKP squeezing), which the code is designed to correct. The planar spiral resonator removes the packaging obstacle that has kept bosonic QEC out of scalable fabrication flows. Combined with [10]'s resource analysis, the platform is plausibly the most component-efficient route to $p_L = 10^{-6}$ -class logical qubits in superconducting hardware.

Limitation 1: the modest extension factor. $F = 1.59$ is far from the exponential gains that bosonic codes promise asymptotically. Our R2 shows that $62.9\% \pm 2.0\%$ of the logical decay rate remains uncorrected after the measurement-free stabilization protocol. This means the correction cycle removes only about 37.1% of the logical error budget per unit time, so the platform currently delivers a modest, though unambiguous, improvement rather than a regime change. We regard this as consistent with an early-stage demonstration: the finite-energy GKP

envelope, the ancilla reset duty cycle, and resonator loss all contribute residual error channels that a mature implementation must suppress. The honest reading of R2 is that the dominant corrected quantity today is the ancilla-limited floor, not the oscillator-limited ceiling.

Failure mode 1: fluxonium dephasing. Our analysis has focused on the bit-flip channel, but heavy fluxonium also exhibits pure dephasing, with T_2 potentially much shorter than T_1^{bf} in parameter regimes with enhanced charge sensitivity. Following the hierarchy lesson of [11], a correction cycle commensurate with the slowest uncorrected channel is required; if fluxonium dephasing during a cycle induces phase shifts on the oscillator that exceed the GKP envelope width, the effective flip-to-logical conversion factor η of Section 4.4 rises and the ancilla-limited lifetime projection of R4 degrades. A measurement of T_2 under the operating bias point, and of dephasing-induced logical shift errors, would directly test this failure mode.

Failure mode 2: finite-energy GKP squeezing. The code's correctable shift-error range is set by the finite-energy envelope width, i.e., by the achievable squeezing in dB. If the prepared states have squeezing below the level at which photon-loss-induced shifts remain within the correctable range, the logical budget becomes squeezing-dominated and the ancilla improvements quantified in R1–R3 become irrelevant. The extension factor $F = 1.59$ is itself partially a function of squeezing; a squeezing-resolved measurement of F would disentangle the two contributions.

Failure mode 3: reset latency. Our per-cycle analysis assumed $T_{\text{cyc}} = 10 \mu\text{s}$ as a stated assumption. If the actual fluxonium reset time is longer, the cycle may fail to keep up with oscillator-side error rates, and the steady-state cancellation argument of Section 4.4 no longer protects the logical state between cycles. R4 shows that reset speed does not lower the ancilla floor, but it does govern whether the code can track oscillator errors; a direct measurement of reset time and of logical fidelity versus cycle period would falsify or confirm the fast-reset premise of [1].

Failure mode 4: spiral-resonator loss. The planar Archimedean spiral resonator trades packaging scalability for mode volume and, potentially, for quality factor. If the resonator's photon lifetime is short compared with the corrected logical lifetime, oscillator-side photon loss dominates the budget and the resource advantage computed in R5 erodes. A quality-factor measurement of the spiral resonator at the operating photon number, compared against the $\sim 100\times$ mode-count advantage of [10], would test whether the planar platform retains its component efficiency in practice.

Limitation 2: excluded sources. As noted in Section 2, the supplied bibliography contains four withdrawn or removed records ([5],[7],[8],[9]) unrelated to the topic. We cite none of them substantively, but their presence is a corpus-quality limitation: any literature-triangulation step in the spirit of [13] performed on this corpus must exclude those records, and conclusions drawn from bibliometric context inherit that exclusion.

7. Conclusion

We have performed a derived-metric reassessment of the fluxonium-GKP bosonic error-correction demonstration [1],[3]. From the reported bit-flip lifetime $T_1^{\text{bf}} = 451 \pm 70 \mu\text{s}$ and extension factor $F = 1.59 \pm 0.05$ alone, we derived the ancilla bit-flip rate $\Gamma_{\text{bf}} = (2.22 \pm 0.34) \times 10^3 \text{ s}^{-1}$, the fractional removal of 0.371 ± 0.020 of the logical decay rate, per-cycle ancilla error probabilities under a stated cycle-time assumption, and the structural result that the ancilla-induced logical error rate in steady state equals $\eta \Gamma_{\text{bf}}$ independent of cycle speed. Combined with the resource-commensurable comparison of [10], and under explicitly labeled assumptions, the platform's effective photon advantage versus surface codes is projected at $8.0\text{--}64\times$.

Our assessment is falsifiable. The claim that planar fluxonium-GKP is the leading route to scalable bosonic QEC would be undermined if: (i) fluxonium dephasing under the operating bias proves to dominate the logical budget; (ii) achievable finite-energy GKP squeezing falls below the threshold at which photon-loss-induced shifts are

correctable; (iii) reset latency prevents the correction cycle from tracking oscillator error rates; or (iv) spiral-resonator quality factors fall far enough below 3D-cavity values that the mode-count advantage of [10] is forfeited. Each of these is testable with a specific, named measurement identified in Section 6. Absent such outcomes, the fluxonium-GKP platform stands as the most credible route to planar, hardware-efficient bosonic error correction.

References

[1] TITLE: arXiv Query: `searchquery=&idlist=2609.26880&start=0&max_results=1` [2] Continuous-time quantum error correction. arXiv:1311.2485v2. <https://arxiv.org/abs/1311.2485v2> [3] Bosonic Error Correction with Fluxonium. arXiv:2609.26880v2. <https://arxiv.org/abs/2609.26880v2> [4] Entanglement-Assisted Quantum Error-Correcting Codes. arXiv:1610.04013v1. <https://arxiv.org/abs/1610.04013v1> [5] A Simulation and Modeling of Access Points with Definition Language. arXiv:1304.1836v2. <https://arxiv.org/abs/1304.1836v2> [6] Quantum error correction beyond qubits. arXiv:0811.3734v1. <https://arxiv.org/abs/0811.3734v1> [7] Superconductivity as a consequence of an ordering of the electron gas zero-point oscillations. arXiv:1005.0280v6. <https://arxiv.org/abs/1005.0280v6> [8] The meaning of systematic errors, a comment to "Reply to On the Systematic Errors in the Detection of the Lense-Thirring Effect with a Mars Orbiter", by Lorenzo Iorio. arXiv:gr-qc/0703020v3. <https://arxiv.org/abs/gr-qc/0703020v3> [9] Intuitionistic Fuzzy Ideals in Γ -semiring. arXiv:1011.5746v2. <https://arxiv.org/abs/1011.5746v2> [10] DOI pending. QNFO: Bosonic Codes as the Native Encoding: Resource-Commensurable Comparison of Cat, GKP, Binomial, and Surface Codes. [11] <https://papers.qnfo.org/papers/beyond-pure-dephasing-a-quantitative-analysis-of-two-level-error-correction-for-/>. QNFO: Beyond Pure Dephasing: A Quantitative Analysis of Two-Level Error Correction for Multi-Spin Molecular Qubits. [13] <https://papers.qnfo.org/papers/auditing-bqnn/>. QNFO: Auditing the BQNN: Does a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware Demonstrate a Route to Near-Term Quantum Advantage?.