

Two-Level System Defects as a Random Ensemble: Counting and Dephasing Estimates for Superconducting Qubits

Two-Level System Defects as a Random Ensemble: Counting and Dephasing Estimates for Superconducting Qubits

Abstract

Parasitic two-level system (TLS) defects in the amorphous tunnel barriers of superconducting qubits are a major identified source of decoherence, yet device designers lack closed-form estimates linking defect density and junction geometry to the probability that a device is degraded and to the resulting dephasing rate. This paper models the defect population of a junction as a random ensemble — a random dynamical system in which defect positions and frequencies are random variables and the qubit is a deterministic observer — and derives two complementary closed-form results. A Poisson counting argument gives the probability that a junction hosts a strongly coupled defect: for an assumed density of 10^{24} m^{-3} , a barrier volume of $4 \times 10^{-23} \text{ m}^3$, and a fractional coupling window of 2×10^{-3} , the mean count is $\lambda = 0.08$ and the single-junction probability is $P_1 = 0.0769$; on a 100-junction chip the expected count is 8 and the affected-chip probability is 0.99966. A dephasing calculation for off-resonant defects gives $\Gamma_\varphi = 5 \text{ Hz}$ ($T_\varphi = 0.2 \text{ s}$) for a defect-rich five-defect configuration and $800 \mu\text{s}$ under a telegraph-noise model for a switching defect. The results convert defect decoherence from a device-by-device pathology into an ensemble-level design parameter with an explicit falsification register.

1. Introduction

Superconducting quantum processors lose coherence faster than their control electronics or circuit designs require, and a large share of the loss is traced to the materials of the device itself. Atomic-scale defects in the amorphous oxides and tunnel barriers of superconducting circuits behave as parasitic two-level systems (TLS): quantum systems with two accessible states whose transition frequencies lie in the gigahertz band occupied by the qubit. When a defect resides in the tunnel barrier of a Josephson junction and couples strongly to the qubit mode, it produces readout failures and spectral features visible in multi-photon spectroscopy and shiftable by strain tuning; the experimental study of a strongly coupled TLS in a transmon tunnel barrier documents exactly this three-partite interaction among TLS, qubit, and readout resonator [6]. The people who face this problem are qubit device engineers, process-integration teams, and cryogenic measurement groups: every fabrication run produces junctions whose coherence is limited by an uncontrolled defect population, and the yield of high-coherence devices is set in part by how often a defect lands close enough in frequency and space to matter.

Today that likelihood is treated empirically, device by device. A defect that ruins qubit j on chip A is characterized there, but the same fabrication process places different defects on chip B , so the fix does not transfer. What is blocked is predictive design: a fabrication change cannot be evaluated by its effect on the defect ensemble until the ensemble's statistical consequences are computable before the device is made. If the defect population is treated statistically — as a random ensemble with measurable density, frequency distribution, and coupling distribution — then three decisions change. Fabrication teams can set a target defect density from a required device yield. Architects can budget coherence margins per junction rather than per chip. Spectroscopy campaigns can be sized in advance, because the expected number of strongly coupled defects per chip follows from the same parameters.

The core idea is a change of statistical frame. The defect ensemble is a random dynamical system in the technical sense: probability theory supplies the description of the unknown microscopic configuration, dynamical systems supplies the evolution rules, and the combination describes systems known only inaccurately, whether from physical uncertainty or computational limitation [2]. The qubit is a deterministic probe coupled to this random background. The same reframing has precedent in disordered materials: in amorphous solids, structural disorder long hindered first-principles identification of the microscopic carriers of plasticity, until well-defined topological defects — hedgehog defects in three-dimensional amorphous solids — were shown to be identifiable and correlatable with soft spots [4]. The TLS defects of superconducting junctions are the analogous microscopic carriers of decoherence, and this paper imports the same move: name the carriers, then treat their spatial and spectral arrangement as a random ensemble.

The contributions are:

- A formulation of the junction defect population as a random dynamical system, with defect spectral positions as the random state and the qubit as the deterministic observer (Section 3).
- A closed-form derivation of the near-resonance probability $P \geq 1 = 1 - e^{-\lambda}$ under two parameterizations — spectral counting and spatio-spectral counting — evaluated with all arithmetic shown to $P \geq 1 = 0.0999$ and $P_1 = 0.0769$ respectively, with a chip-level extension giving $P_{\text{chip}} = 0.99966$ for 100 junctions (Section 4).
- A derivation of the ensemble-averaged dispersive dephasing rate $\Gamma_\varphi = \sum_i (g_i/\Delta_i)^2 \gamma_i$, evaluated to $T_\varphi = 0.2$ s for five sub-band defects, and a telegraph-noise dephasing rate $\Gamma_\varphi = \varepsilon^2/(8\gamma)$, evaluated to $T_\varphi = 800$ μs (Section 4).
- A sensitivity analysis showing that the counting result moves by an order of magnitude per tenfold change in defect density, and that the dephasing rate is quadratic in coupling (Section 4).
- A falsification register: the specific spectroscopy observations that would refute the counting model and the noise model, and a list of the assumptions on which each number rests (Sections 6 and 7).

2. Background and Related Work

The literature relevant to this paper falls into three groups: the physics of TLS defects in superconducting circuits, the mathematics of random systems, and a body of work that shares the TLS acronym but concerns the Transport Layer Security protocol. All three are discussed here, because the acronym collision is itself a practical obstacle to literature search in this area and because the protocol-side works illustrate a general point about statistical treatment of large configuration spaces.

On the physics side, the central experimental anchor is the study of readout failures in superconducting qubits due to TLS defects in tunnel junctions [6]. That work identifies material defects as the origin of parasitic two-level systems that constitute a major source of decoherence in superconducting qubits, and it studies a strongly coupled TLS residing in the tunnel barrier of a transmon qubit using multi-photon spectroscopy and TLS strain tuning to explore the spectrum of the interacting three-partite system of TLS, qubit, and readout resonator. Its summary states that the interacting spectrum is rich, which is precisely the regime where single-defect intuition fails and ensemble statistics become necessary; the present paper supplies that statistical layer. The summary supplied for [6] gives no numerical defect densities, so no density value is imported from it here; all densities used below are stated assumptions of this paper.

The statistical machinery comes from two traditions. The theory of random dynamical systems [2] is described as a comparatively recent development combining ideas and methods from the well-developed areas of probability theory and dynamical systems, motivated by inaccurate knowledge of the particular physical system or by

computational and theoretical limitations such as insufficient computational power, inefficient algorithms, or insufficiently developed mathematical tools. This motivation matches the defect problem exactly: the microscopic defect configuration of a given junction is unknowable in detail, so the configuration must be treated as a random state evolved under known dynamics. On the fluctuation side, the review of one-dimensional stochastic growth and its connection to Gaussian ensembles of random matrices [8] establishes that surface fluctuations of the polynuclear growth model coincide, for both curved and flat growth, with limit distribution functions from Gaussian random-matrix ensembles, with the connection explained through associated point processes. The methodological lesson imported here is that spatial or spectral point processes of disordered systems can carry non-trivial fluctuation statistics, which motivates treating the defect frequency positions as a point process rather than as independent fixed quantities.

Two further mathematical results shape the count statistics used in Section 4. The theorem on random perfect matchings in regular graphs [7] proves that in all regular robust expanders every edge is asymptotically equally likely to be contained in a uniformly chosen perfect matching, and that for any fixed matching or spanning regular graph the overlap $|M \cap E(N)|$ is approximately Poisson distributed. The Poisson limit for an overlap count of a random object against a fixed background is the same limit structure that Section 4 uses for the count of defects falling inside a fixed spatial-and-frequency window, although the physical setting of independent point defects is far simpler than the expander-graph setting and needs no expansion properties. And in amorphous solids, the hedgehog topological defect work [4] states that structural disorder renders the concept of topological defects difficult to apply and hinders first-principles identification of the microscopic carriers of plasticity and of soft spots, while recently proposed well-defined topological defects in glasses can be identified and correlated with rearrangement-prone regions. This is the closest structural analogue to the present program: defects that were individually invisible become statistically and even topologically characterizable once the right observable is chosen. If TLS defects are eventually identified with such structural objects, the defect density ρ used in this paper stops being a fit parameter and becomes predictable from glass structure; until then it remains an input to be measured.

Three works share the TLS acronym with the physics subject but concern the Transport Layer Security protocol, and they are cited for completeness of the literature landscape and for one structural parallel. The Helol tunnel work [1] targets the combinatorial properties of Client Hello packets in the ubiquitous Transport Layer Security protocol to build a covert channel for data exfiltration and command-and-control, on the premise that such channels are integral parts of modern cyberattacks. The active TLS stack fingerprinting work [3] uses active measurements to collect server characteristics at scale, so that metadata can reveal hidden relations and commonalities among deployments — including, as its summary notes, identification of previously unknown cybercriminal infrastructure as cyber-threat intelligence. The post-quantum signature placement study [5] examines how signature-algorithm choice in TLS 1.3 couples with certificate-hierarchy structure, chain exposure during the handshake, and role-dependent cryptographic cost. The parallel worth noting is methodological: all three treat a large combinatorial space of protocol configurations statistically, asking what a random draw from that space looks like and what its aggregate consequences are. The present paper asks the same question of a defect ensemble. No technical content flows between the protocol works and the physics model; the connection is the shared statistical posture toward a large configuration space, and the disambiguation itself matters because a defect-density literature search that mixes the two senses of TLS will retrieve the wrong corpus entirely.

Finally, four QNFO corpus entries frame the broader context. The trapped-ion ultrametric testbed [11] organizes sixteen published records from a single research program, spanning December 2025 to August 2026, into one testable claim: that trapped-ion quantum simulators are the first near-term platform on which ultrametric (p-adic) structure in quantum dynamics can be accepted or rejected by measurement. Its significance here is as a model of how a coherence-related research program can be organized around an explicitly falsifiable register, a discipline

Section 6 adopts; its supplied summary contains no material on superconducting qubits. The entry on structural versus driven quantum coherence [9] is supplied with no summary text, so no specific claim from it can be used; it is relevant only as a marker that the distinction between coherence arising from structure and coherence induced by driving is an active organizing question, which is the distinction this paper draws between defect-intrinsic dynamics and qubit-driven defect response. The toy model of non-local information storage in a quantum cellular automaton [10] and the work on operationalizing generalized symmetries [12] are likewise supplied without summary text; their summaries give no further detail, and they are related to this paper only through their titles' indication of quantum-information and symmetry formalisms, which the present paper does not use.

3. Methods

3.1 The random dynamical system formulation

A random dynamical system, in the sense of [2], consists of a state space, a dynamical flow on that state space, and a probability measure describing inaccurate knowledge of the state or of the system itself. The defect problem maps onto this structure as follows. The state space is the configuration space of the junction: a configuration is the list of transition frequencies $\omega = (\omega_1, \dots, \omega_{N_d})$ of the N_d TLS defects in the junction volume, together with their couplings $g = (g_1, \dots, g_{N_d})$ to the qubit's electric field. Conditional on the configuration, the qubit and readout resonator dynamics are deterministic, in the three-partite structure documented in [6]. The randomness enters through the measure on configuration space: defect frequencies, positions, and couplings are not known for a given device and are modeled as draws from ensemble distributions.

Three ensemble distributions are needed as inputs:

1. The defect count N_d in the junction volume, treated as a random variable with mean $\bar{N}_d$ (or, equivalently, a spatial density ρ).
2. The frequency distribution of defect transitions, modeled as uniform over a band of width Δ_0 . A uniform spectral density is the standard null model for an uncharacterized amorphous defect bath; any measured deviation from uniformity would itself be a finding about the material.
3. The coupling distribution, characterized by a typical value g_0 and treated as independent of frequency. This independence is an assumption; strain tuning of the kind used in [6] could test it, since strain shifts defect frequencies and any resulting change in coupling would reveal correlation.

The qubit is the deterministic observer. Its frequency ω_q is fixed by design, and all ensemble averages below are taken over the defect configuration measure with ω_q held fixed. This asymmetry is what makes the framework predictive: the fabrication process controls the ensemble parameters (ρ, Δ_0, g_0) , and the observable decoherence is an ensemble average at fixed ω_q .

3.2 Counting model

Defects are modeled as point entities distributed independently and uniformly at random with number density ρ inside a barrier of area A and thickness d , so the barrier volume is

$$V = A d.$$

Each defect has a transition frequency drawn uniformly from a band of width Δ_0 . A defect is called *strongly coupled* if two conditions hold: it lies inside the barrier volume V , where electric-field participation and therefore coupling are largest, and its transition frequency lies within $\pm\Gamma$ of the qubit frequency. The fractional frequency window is

$$f_{\Gamma} = \frac{2\Gamma}{\Delta_0}.$$

Because positions and frequencies are independent and uniform, the expected number of strongly coupled defects in one junction is

$$\lambda = \rho V f_{\Gamma} = \rho A d \frac{2\Gamma}{\Delta_0},$$

and the count N is Poisson distributed with mean λ :

$$\Pr(N = k) = \frac{\lambda^k e^{-\lambda}}{k!}, \quad \Pr(N \geq 1) = 1 - e^{-\lambda}.$$

The Poisson assignment is the structural analogue of the approximately Poisson overlap counts proved for random matchings against fixed regular subgraphs [7]; here independence of defect placement makes the argument elementary. An equivalent purely spectral parameterization uses a total defect number $\bar{N}_d$ spread over a band B : the expected count in a two-sided frequency window of width 2δ is $\lambda = \bar{N}_d \cdot 2\delta/B$, with the same Poisson law. For a chip of N_j junctions with independent defect populations, the total count is Poisson with mean $N_j \lambda$, and

$$P_{\text{chip}} = 1 - e^{-N_j \lambda}.$$

3.3 Dephasing models

Two dephasing mechanisms are treated, corresponding to two regimes.

Dispersive dephasing (off-resonant defects). For a defect detuned by Δ_i from the qubit with coupling g_i and intrinsic defect relaxation rate γ_i , the qubit acquires a small admixture of the lossy defect; to leading order in g_i/Δ_i the defect contributes a dephasing rate

$$\Gamma_{\varphi}^{(\text{disp})} = \sum_{i=1}^n \left(\frac{g_i}{\Delta_i} \right)^2 \gamma_i.$$

Rates from distinct defects add. This form is valid when $g_i \ll |\Delta_i|$ for all counted defects; near-resonant defects are handled separately by the counting model.

Telegraph-noise dephasing (switching defect). A TLS that switches stochastically between its two states shifts the qubit frequency by a state-dependent amount. In the classical-noise approximation, the qubit frequency is perturbed by a random telegraph process $\xi(t)$ taking values $\pm\epsilon/2$ with switching rate γ per direction, so the autocorrelation function is

$$C(\tau) = \langle \xi(t) \xi(t + \tau) \rangle = \frac{\epsilon^2}{4} e^{-2\gamma|\tau|}.$$

For a qubit at the noise's zero-frequency crossing, the pure-dephasing rate in the long-time limit is the time integral of the correlation function:

$$\Gamma_{\varphi}^{(\text{tn})} = \int_0^{\infty} C(\tau) d\tau = \frac{\epsilon^2}{4} \cdot \frac{1}{2\gamma} = \frac{\epsilon^2}{8\gamma},$$

with dephasing time $T_{\varphi} = 1/\Gamma_{\varphi}$. This is the standard filter-function result at zero detuning for classical two-state noise; it is exact within the classical-noise and weak-coupling approximations stated here.

All frequencies in the derivations are ordinary frequencies (cycles per second), not angular frequencies; a device using angular frequencies must insert the corresponding factors of 2π consistently.

3.4 Parameter values

All numerical inputs are stated assumptions, chosen to be representative of the model's use rather than measured values; each is carried explicitly through the arithmetic of Section 4.

Parameter	Symbol	Assumed value
Defect density	ρ	10^{24} m^{-3}
Junction barrier area	A	$2 \times 10^{-14} \text{ m}^2$
Barrier thickness	d	$2 \times 10^{-9} \text{ m}$
Coupling window (half-width)	Γ	10^6 s^{-1}
TLS frequency band	Δ_0	10^9 s^{-1}
Junctions per chip	N_j	100
Telegraph coupling amplitude	ε	10^5 s^{-1}
Telegraph switching rate	γ	10^6 s^{-1}
Total defects (spectral model)	$\bar{N}_d$	100
Spectral band (spectral model)	B	19 GHz
Qubit frequency	ω_q	5 GHz
Resonance window (spectral model)	δ	10 MHz
Typical coupling (dispersive)	g_0	1 MHz
Defect relaxation rate (dispersive)	γ_0	10 kHz
Sub-band defect count (dispersive)	n	5

4. Analysis

4.1 Barrier volume

Inputs: $A = 2 \times 10^{-14} \text{ m}^2$ and $d = 2 \times 10^{-9} \text{ m}$ (Section 3.4 assumptions; the area corresponds to a junction of order $0.1 \mu\text{m} \times 0.2 \mu\text{m}$).

$$V = A d = (2 \times 10^{-14}) \times (2 \times 10^{-9}) \text{ m}^3 = 4 \times 10^{-23} \text{ m}^3.$$

4.2 Fractional coupling window

Inputs: $\Gamma = 10^6 \text{ s}^{-1}$, $\Delta_0 = 10^9 \text{ s}^{-1}$ (Section 3.4).

$$f_\Gamma = \frac{2\Gamma}{\Delta_0} = \frac{2 \times 10^6}{10^9} = 2 \times 10^{-3}.$$

4.3 Single-junction defect probability (spatio-spectral model)

Inputs: $\rho = 10^{24} \text{ m}^{-3}$, $V = 4 \times 10^{-23} \text{ m}^3$ (Section 4.1), $f_\Gamma = 2 \times 10^{-3}$ (Section 4.2).

Step 1 — mean count:

$$\lambda = \rho V f_\Gamma = 10^{24} \times 4 \times 10^{-23} \times 2 \times 10^{-3}.$$

Step 2 — arithmetic: $10^{24} \times 4 \times 10^{-23} = 40$; then $40 \times 2 \times 10^{-3} = 0.08$.

$$\lambda = 0.08.$$

Step 3 — probability of at least one strongly coupled defect:

$$P_1 = 1 - e^{-\lambda} = 1 - e^{-0.08}.$$

Step 4 — arithmetic: expanding,
 $e^{-0.08} = 1 - 0.08 + 0.0032 - 0.00008533 + 0.00000171 = 0.923116$, so

$$P_1 = 1 - 0.923116 = 0.076884 \approx 0.0769.$$

Under these assumptions, roughly one junction in thirteen hosts a strongly coupled TLS.

4.4 Chip-level extension

Inputs: $N_j = 100$ (Section 3.4), $\lambda = 0.08$ (Section 4.3).

Step 1 — chip mean count: $\lambda_{\text{chip}} = N_j \lambda = 100 \times 0.08 = 8$.

Step 2 — probability at least one junction is affected: $P_{\text{chip}} = 1 - e^{-8}$.

Step 3 — arithmetic: $e^{-8} = e^{-4} \times e^{-4} \approx 0.0183156 \times 0.0183156 \approx 0.00033546$, so

$$P_{\text{chip}} = 1 - 0.00033546 = 0.99966454 \approx 0.99966.$$

Step 4 — standard deviation of the chip count: for a Poisson variable the standard deviation is the square root of the mean, $\sigma_{\text{chip}} = \sqrt{8} \approx 2.8284$.

A 100-junction chip under these assumptions is essentially certain to contain strongly coupled defects, with chip-to-chip count fluctuation of order ± 3 .

4.5 Spectral counting model and near-resonance probability

Inputs: $\bar{N}_d = 100$ defects uniformly distributed over $B = 19$ GHz; two-sided resonance window $2\delta = 20$ MHz = 0.02 GHz (Section 3.4).

Step 1 — expected near-resonant count:

$$\lambda = \frac{\bar{N}_d \cdot 2\delta}{B} = \frac{100 \times 0.02 \text{ GHz}}{19 \text{ GHz}} = \frac{2}{19} = 0.105263.$$

Step 2 — Poisson approximation: $P_{\geq 1} = 1 - e^{-0.105263}$. Arithmetic:
 $e^{-0.105263} = 1 - 0.105263 + 0.005540 - 0.000194 + 0.0000051 = 0.900088$, so

$$P_{\geq 1} = 1 - 0.900088 = 0.099912 \approx 0.0999.$$

Step 3 — exact binomial check: $P_{\geq 1} = 1 - (1 - 0.02/19)^{100}$. Arithmetic: $1 - 0.02/19 = 0.9989474$; $\ln(0.9989474) = -0.00105319$; $100 \times (-0.00105319) = -0.105319$; $e^{-0.105319} = 0.900038$; so $P_{\geq 1} = 0.099962$, confirming the Poisson approximation to within approximately 5×10^{-5} .

Under these assumptions, approximately one qubit in ten hosts a defect within 10 MHz of its frequency.

Step 4 — mean defect spacing:

$$\bar{s} = \frac{B}{\bar{N}_d} = \frac{19 \text{ GHz}}{100} = 190 \text{ MHz}.$$

The spacing exceeds the resonance window $\delta = 10 \text{ MHz}$ by a factor of 19 and the assumed defect linewidth $\gamma_0 = 10 \text{ kHz}$ by a factor of 1.9×10^4 , so adjacent defects do not overlap spectrally through their linewidths.

4.6 Dispersive dephasing rate

Inputs: $n = 5$ defects in the $\pm 100 \text{ MHz}$ sub-band, each with $g_i = g_0 = 1 \text{ MHz}$, $\gamma_i = \gamma_0 = 10 \text{ kHz}$, and detuning magnitudes averaging $|\bar{\Delta}| = 100 \text{ MHz}$ (Section 3.4).

Step 1 — per-defect ratio: $(g_0/|\bar{\Delta}|)^2 = (1 \text{ MHz}/100 \text{ MHz})^2 = (0.01)^2 = 10^{-4}$.

Step 2 — per-defect dephasing rate:

$$\Gamma_{\varphi}^{(1)} = \left(\frac{g_0}{|\bar{\Delta}|} \right)^2 \gamma_0 = 10^{-4} \times 10^4 \text{ s}^{-1} = 1 \text{ s}^{-1}.$$

Step 3 — sum over $n = 5$ defects:

$$\Gamma_{\varphi}^{(\text{disp})} = 5 \times 1 \text{ s}^{-1} = 5 \text{ s}^{-1}.$$

Step 4 — dephasing time:

$$T_{\varphi} = \frac{1}{\Gamma_{\varphi}} = \frac{1}{5 \text{ s}^{-1}} = 0.2 \text{ s}.$$

Under these assumptions, five off-resonant defects in the $\pm 100 \text{ MHz}$ sub-band produce a pure-dephasing time of 0.2 s , far above typical transmon coherence times; the dispersive channel from sub-band defects is therefore not the dominant limitation in this parameterization, and the near-resonant defects counted in Section 4.3–4.5 are the ones that matter.

4.7 Telegraph-noise dephasing rate

Inputs: $\varepsilon = 10^5 \text{ s}^{-1}$, $\gamma = 10^6 \text{ s}^{-1}$ (Section 3.4).

Step 1 — rate:

$$\Gamma_{\varphi}^{(\text{tn})} = \frac{\varepsilon^2}{8\gamma} = \frac{(10^5)^2}{8 \times 10^6} \text{ s}^{-1} = \frac{10^{10}}{8 \times 10^6} \text{ s}^{-1} = 1.25 \times 10^3 \text{ s}^{-1}.$$

Step 2 — dephasing time:

$$T_{\varphi} = \frac{1}{1.25 \times 10^3 \text{ s}^{-1}} = 8 \times 10^{-4} \text{ s} = 800 \mu\text{s}.$$

A single switching defect with these parameters produces $T_{\varphi} = 800 \mu\text{s}$, which is at the level of good transmon coherence times and therefore observable. The rate is quadratic in the frequency shift amplitude ε : halving ε to $5 \times 10^4 \text{ s}^{-1}$ gives $\Gamma_{\varphi} = 312.5 \text{ s}^{-1}$ and $T_{\varphi} = 3.2 \text{ ms}$, a fourfold improvement.

4.8 Sensitivity analysis

Counting result versus defect density. The mean count is $\lambda = \rho V f_{\Gamma}$, linear in ρ . A tenfold reduction in density, $\rho = 10^{23} \text{ m}^{-3}$, gives $\lambda = 0.008$ and

$$P_1 = 1 - e^{-0.008} \approx 0.00797,$$

so the single-junction probability falls by a factor of about ten, from 0.0769 to 0.0080. At the chip level, $\lambda_{\text{chip}} = 0.8$ and $P_{\text{chip}} = 1 - e^{-0.8} = 0.5507$: even a tenfold-denser-or-cleaner process changes the chip-level statement from near-certainty to a coin flip. Conversely, a tenfold increase, $\rho = 10^{25} \text{ m}^{-3}$, gives $\lambda = 0.8$, $P_1 = 0.5507$, and $P_{\text{chip}} = 1 - e^{-80} \approx 1$. The counting result moves by one order of magnitude in P_1 per tenfold change in ρ ; the chip-level probability saturates and is insensitive to density improvements once $N_j \lambda \gg 1$.

Counting result versus window width. The same linearity holds in $f_\Gamma = 2\Gamma/\Delta_0$. Halving the accepted detuning window to $\Gamma = 5 \times 10^5 \text{ s}^{-1}$ halves λ to 0.04 and gives $P_1 = 0.0392$.

Dephasing versus coupling. The dispersive rate scales as $(g_i/\Delta_i)^2$: a defect ten times closer in frequency, at $|\Delta| = 10 \text{ MHz}$ with the same $g_0 = 1 \text{ MHz}$, contributes $\Gamma_\varphi^{(1)} = (0.1)^2 \times 10^4 \text{ s}^{-1} = 100 \text{ s}^{-1}$, a hundredfold increase, and five such defects give $T_\varphi = 1/500 \text{ s} = 2 \text{ ms}$. The quadratic sensitivity means the dispersive dephasing budget is dominated by the closest defects, not the most numerous ones.

Telegraph rate versus switching speed. $\Gamma_\varphi^{(\text{tn})} = \varepsilon^2/(8\gamma)$ is inversely proportional to γ : a slower-switching defect ($\gamma = 10^5 \text{ s}^{-1}$) with the same ε gives $\Gamma_\varphi = 1.25 \times 10^4 \text{ s}^{-1}$ and $T_\varphi = 80 \mu\text{s}$. Slow switchers are more damaging than fast ones at equal amplitude, because the qubit tracks the quasi-static frequency shift for longer.

5. Results

All numbers in this section are computed in Section 4 from the stated assumptions of Section 3.4; none are measured values.

R1 (Single-junction near-resonance probability). For the spatio-spectral model with $\rho = 10^{24} \text{ m}^{-3}$, $V = 4 \times 10^{-23} \text{ m}^3$, and $f_\Gamma = 2 \times 10^{-3}$: mean strongly coupled count $\lambda = 0.08$ and $P_1 = 0.0769$ (Section 4.3). Roughly one junction in thirteen hosts a strongly coupled TLS.

R2 (Chip-level probability). For $N_j = 100$ independent junctions: chip mean count $\lambda_{\text{chip}} = 8$, $P_{\text{chip}} = 0.99966$, count standard deviation $\sigma_{\text{chip}} = 2.83$ (Section 4.4). A 100-junction chip is essentially certain to contain strongly coupled defects, with chip-to-chip fluctuation of order ± 3 .

R3 (Spectral-model cross-check). With $\bar{N}_d = 100$ defects over $B = 19 \text{ GHz}$ and a 20 MHz two-sided window: $\lambda = 0.105263$, Poisson $P_{\geq 1} = 0.0999$, exact binomial $P_{\geq 1} = 0.099962$; mean defect spacing $\bar{s} = 190 \text{ MHz}$ (Section 4.5). Approximately one qubit in ten is near-resonant; the two counting parameterizations agree in order of magnitude.

R4 (Dispersive dephasing). Five off-resonant defects at $|\Delta| = 100 \text{ MHz}$, $g_0 = 1 \text{ MHz}$, $\gamma_0 = 10 \text{ kHz}$: $\Gamma_\varphi = 5 \text{ s}^{-1}$, $T_\varphi = 0.2 \text{ s}$ (Section 4.6). Sub-band dispersive dephasing is subdominant in this parameterization.

R5 (Telegraph dephasing). One switching defect with $\varepsilon = 10^5 \text{ s}^{-1}$, $\gamma = 10^6 \text{ s}^{-1}$: $\Gamma_\varphi = 1.25 \times 10^3 \text{ s}^{-1}$, $T_\varphi = 800 \mu\text{s}$ (Section 4.7). A single slow switcher at these parameters produces coherence-scale dephasing.

R6 (Sensitivities). P_1 is linear in ρ and in f_Γ : a tenfold density reduction moves P_1 from 0.0769 to 0.0080 and P_{chip} from 0.99966 to 0.5507 (Section 4.8). The dispersive rate is quadratic in g/Δ ; the telegraph rate is proportional to ε^2 and inversely proportional to γ .

6. Discussion

The principal result is the conversion of two assumed densities and two assumed windows into device-level probabilities and times: $P_1 = 0.0769$, $P_{\text{chip}} = 0.99966$, $T_\varphi = 0.2$ s (dispersive) and $800 \mu\text{s}$ (telegraph). The value of these numbers is conditional on the assumptions, and the conditionality is the point of this section.

What would falsify the counting model. The model predicts that strongly coupled defects appear as a Poisson process across nominally identical junctions, with mean $\lambda = \rho V f_\Gamma$. A spectroscopy campaign over many junctions that found the number of strongly coupled defects per junction systematically far from Poisson — for example, strongly bimodal, with most junctions defect-free and a few carrying many defects — would refute independent uniform placement. Equally, a measured density of strongly coupled defects inconsistent with $\rho V f_\Gamma$ by more than the Poisson fluctuation $\sqrt{N_j \lambda} \approx 2.8$ across a chip population would falsify the assumed $\rho = 10^{24} \text{ m}^{-3}$ or the uniform spectral distribution. The strain-tuning protocol of [6] supplies the identification method: each strongly coupled TLS found by multi-photon spectroscopy is one count of N .

What would falsify the dephasing models. The dispersive prediction $T_\varphi = 0.2$ s for the assumed sub-band population implies that sub-band defects contribute negligibly to observed dephasing; if measured T_φ in defect-characterized devices tracked the sub-band defect count rather than the near-resonant count, the dispersive formula or its validity conditions ($g_i \ll |\Delta_i|$) would be implicated. The telegraph prediction is sharper: $T_\varphi = 800 \mu\text{s}$ with quadratic dependence on ε and inverse dependence on γ . A measured dephasing time for a device with an identified switching defect that was neither quadratic in the measured frequency shift nor inverse in the measured switching rate would falsify the classical telegraph model as applied here.

Limitations. The results rest on the following assumptions, each with its sensitivity:

- **Defect density $\rho = 10^{24} \text{ m}^{-3}$ is an assumption, not a measurement.** P_1 is linear in ρ ; the entire counting result scales with it. No density value is available from the supplied summary of [6], and the paper imports none.
- **Uniform spectral distribution.** If defect frequencies cluster near qubit frequencies (a known concern in the TLS literature, though no clustering value is supplied by the grounding), the effective f_Γ rises and P_1 with it; the uniform model is then optimistic.
- **Independence of positions and frequencies.** Correlated placement (for example, density variations across the barrier) breaks the Poisson law; the expander-graph result of [7] shows Poisson statistics can survive weaker independence assumptions, but no such relaxation is proved here.
- **Coupling–frequency independence.** Assumed in Section 3.1; testable by the strain tuning of [6], untested here.
- **Classical-noise and weak-coupling approximations.** The telegraph result is exact only within these approximations; a defect with ε comparable to γ or to the qubit–defect coupling requires the full quantum treatment, which is not derived here.
- **Ordinary-frequency convention.** All rates are in cycles per second; angular-frequency conventions shift numerical prefactors by powers of 2π .
- **No resonant (strong-coupling) dynamics.** Defects that land within a linewidth of the qubit are counted but their dephasing is not modeled; the counting and dephasing results cover disjoint regimes.

Open questions. Three stand out. First, what is the actual ρ in state-of-the-art tunnel barriers, and is it a controllable process parameter? Second, do defect frequencies cluster, and if so with what distribution? Third, can TLS defects be identified with structural objects of the glass, as the hedgehog-defect program proposes for plasticity carriers in

amorphous solids [4]; if so, ρ becomes predictable rather than fitted. The falsifiability-register discipline of the trapped-ion ultrametric program [11] — organizing a research program around explicitly accept-or-reject claims — is the model for how these questions should be retired.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The users are superconducting-qubit fabrication and process-integration teams and qubit chip architects. If the counting model holds, a fabrication team can invert the relation: to hold the affected-junction fraction below a target P_1^{target} , the required density is $\rho^{\text{target}} = -\ln(1 - P_1^{\text{target}})/(V f_\Gamma)$. For example, holding $P_1 \leq 0.01$ under the same V and f_Γ requires $\lambda \leq 0.01005$, i.e. $\rho \leq 0.01005/(4 \times 10^{-23} \times 2 \times 10^{-3}) \text{ m}^{-3} = 1.26 \times 10^{23} \text{ m}^{-3}$, an eightfold reduction from the assumed 10^{24} m^{-3} . Architects can budget coherence margin per junction (one draw from the P_1 distribution) rather than per chip, and spectroscopy campaigns can be sized in advance from the expected count $N_j \lambda = 8$ strongly coupled defects per 100-junction chip.

(b) Readiness level. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept formulated). No device has been built or tested against these formulas here; analytical work of this kind does not exceed TRL 3, and nothing in this paper supports a claim beyond TRL 2.

(c) Smallest next test. The smallest de-risking calculation is a retrospective one: take an existing multi-junction spectroscopy dataset of the kind produced by the protocol of [6], count strongly coupled TLS per junction, and compare the histogram against the Poisson prediction with ρ fitted from the data. The model is refuted if the per-junction count distribution rejects Poisson (for example, by a variance substantially exceeding the mean across a uniform device population) or if a single fitted ρ cannot reconcile counts across junction geometries of different V at fixed f_Γ .

(d) Binding constraints. The first binding constraint is data, not fabrication: the model's inputs (ρ , the spectral distribution, the coupling distribution) require a systematic per-junction defect census, and no such census value is supplied by the grounding. The second is spectroscopy throughput, since identifying each strongly coupled defect requires the multi-photon spectroscopy and strain-tuning protocol of [6]. Whether barrier-defect density is actually controllable by process changes is an open question the grounding does not answer. Materials and tolerance constraints on the barrier itself (thickness uniformity, area control) enter only through V and are second-order compared with the unknown ρ .

(e) Order-of-magnitude projections. Labeled as projections, with all inputs from Section 3.4 and Section 4: a process improvement that reduces ρ tenfold moves the affected-junction fraction from 7.7% to 0.8% and the affected-chip probability from 0.99966 to 0.55; a further tenfold reduction moves P_{chip} to $1 - e^{-0.08} = 0.0769$. These are consequences of the assumed baseline density and carry its full uncertainty; they are design-space maps, not process forecasts.

(f) Risks and misuse paths. None arise from the physics content. The only misuse risk identified is interpretive: presenting the assumed density 10^{24} m^{-3} as a measured value would lend unearned precision to yield projections. The paper's own register — every number traced to a stated assumption — is the mitigation.

8. Conclusion

Treating the TLS defect population of a Josephson junction as a random ensemble converts defect decoherence from a device-by-device pathology into a statistical design parameter. The Poisson counting argument gives, under stated assumptions, a single-junction strongly coupled defect probability of $P_1 = 0.0769$, a chip-level probability of $P_{\text{chip}} = 0.99966$ for 100 junctions with count fluctuation ± 2.8 , and a spectral-model cross-check at $P_{\geq 1} = 0.0999$. The dephasing calculations give $T_\varphi = 0.2$ s for the dispersive channel of five sub-band defects and $T_\varphi = 800$ μ s for a single telegraph switcher, with the sensitivities that the counting probability is linear in density and window width, the dispersive rate quadratic in g/Δ , and the telegraph rate quadratic in shift amplitude and inverse in switching rate. Every number traces to an explicitly stated assumption, and the falsification register specifies the spectroscopy observations — non-Poisson count distributions, or dephasing scalings inconsistent with the derived dependences — that would refute the model. The framework's usefulness now rests on an empirical question the analysis cannot answer: whether defect density and spectral distribution in real tunnel barriers match the assumptions well enough, or can be steered well enough, for the ensemble statistics to become a working yield tool.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Helol Tunnel: Covert Channel Exploitation of TLS Extensibility & Privacy Features. arXiv:2610.01009v1. <https://arxiv.org/abs/2610.01009v1> [2] Random Dynamical Systems. arXiv:math/0608162v1. <https://arxiv.org/abs/math/0608162v1> [3] Active TLS Stack Fingerprinting: Characterizing TLS Server Deployments at Scale. arXiv:2206.13230v3. <https://arxiv.org/abs/2206.13230v3> [4] Hedgehog topological defects in 3D amorphous solids. arXiv:2407.20631v2. <https://arxiv.org/abs/2407.20631v2> [5] Signature Placement in Post-Quantum TLS Certificate Hierarchies: An Experimental Study of ML-DSA and SLH-DSA in TLS 1.3 Authentication. arXiv:2604.06100v3. <https://arxiv.org/abs/2604.06100v3> [6] Readout failures in superconducting qubits due to TLS-defects in tunnel junctions. arXiv:2605.02755v2. <https://arxiv.org/abs/2605.02755v2> [7] Random perfect matchings in regular graphs. arXiv:2301.10131v2. <https://arxiv.org/abs/2301.10131v2> [8] One-dimensional

stochastic growth and Gaussian ensembles of random matrices. arXiv:math-ph/0505038v2. <https://arxiv.org/abs/math-ph/0505038v2> [9] QNFO: Structural vs Driven Quantum Coherence [10] QNFO: Computational Toy Model of Non-Local Information Storage in a Quantum Cellular Automaton [11] QNFO: The Trapped-Ion Ultrametric Testbed: A Falsifiability Register for Testing p-Adic Structure in Quantum Dynamics [12] QNFO: Operationalizing Generalized Symmetries