

Planck Fidelity: The Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders

Planck Fidelity: The Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders

Abstract

Superconducting transmon qubits are weakly anharmonic oscillators whose spectra form discrete but unbounded ladders of energy eigenstates; the binary register $n \in \{0, 1\}$ is imposed on this ladder by truncation, not by physics. Because every physical drive contains the raising operator $a^\dagger$, whose matrix elements $\sqrt{n+1}$ never vanish, any finite-time gate populates levels outside the computational subspace. We formalize this as the Planck Loss $L_P = 1 - \text{Tr}(P_{\text{qubit}}|\psi\rangle\langle\psi|)$ and conjecture that $L_P > 0$ for any finite-time gate, independent of materials quality — a fundamental cost of base-2 encoding on a base- d physical object. We derive the harmonic-limit second-order leakage $L_P^{(2)} = \theta^4/2$, which evaluates to 3.0440341 for a π -pulse and therefore signals perturbative breakdown; with anharmonicity we obtain the first-order bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$, giving a gate-averaged leakage of 1.0×10^{-2} at $\Omega/|\alpha| = 0.05$, and the exact three-level peak $P_{2,\text{max}} = 2r^2/(1+2r^2) = 1.96 \times 10^{-2}$ at $r = 0.1$. We compute the radix-truncation entropy $S_{\text{Base}}(12) = \ln 6 = 1.7917595$ nats = 2.5849625 bits, a discarded-information fraction of 72.1% of the ladder's $\ln 12$ content. The results reframe the qubit as a radix approximation and motivate qudit (base- d) architectures and a falsifiable leakage-floor test programme.

1. Introduction

Every superconducting quantum processor today runs on a convention, and the convention has a price that no one has put on the invoice. The transmon — the workhorse circuit of the superconducting quantum computing industry — is a weakly anharmonic oscillator: its energy eigenstates form a ladder $|n\rangle$, $n = 0, 1, 2, \dots$, that is discrete because physical action is quantized but unbounded because no term in its Hamiltonian terminates the ladder. The hardware resolves many rungs of this ladder. The qubit register uses two of them. The projector $P_{\text{qubit}} = |0\rangle\langle 0| + |1\rangle\langle 1|$ is declared to be the system, and everything above rung $n = 1$ is declared "leakage."

This declaration is not free. The drive that implements any gate contains the raising operator $a^\dagger$, which by construction couples $|n\rangle$ to $|n+1\rangle$ with matrix element $\sqrt{n+1}$ for every n . A finite-time gate therefore always populates levels outside the binary subspace, no matter how clean the materials, how low the dielectric loss, or how long the coherence time. We call the resulting population loss Planck Loss,

$$L_P = 1 - \text{Tr}(P_{\text{qubit}}|\psi\rangle\langle\psi|) = \sum_{n \geq 2} |c_n|^2, \quad |\psi\rangle = \sum_{n \geq 0} c_n |n\rangle,$$

and conjecture that $L_P > 0$ for any finite gate time, as a property of the encoding rather than of the fabrication.

The stakes are concrete. Gate-error budgets in superconducting processors are compiled from materials-limited channels — dephasing, relaxation, control noise — plus a leakage line item that is treated as a tunable control artifact, suppressible by pulse shaping. If the conjecture of this paper holds, that bookkeeping is wrong at a structural level: a strictly positive leakage floor exists for any finite-time binary gate, and the correct engineering response is not to chase it to zero but to change the radix of the register. A designer who accepts the result stops

optimizing two-level pulses against an unreachable zero and starts asking how much of the ladder's capacity a base- d register would recover. The decision that changes is the radix decision, and the capability it unlocks is the roughly 2.58 bits per oscillator that binary truncation discards, quantified in Section 4.3.

The name "Planck" is chosen deliberately. The quantization of physical action — the fact that the oscillator ladder is discrete with fixed spacing — is the reason a truncation is possible at all and the reason its cost is quantifiable. The European Space Agency's Planck satellite, named for the same physicist, was launched on 14 May 2009 and has been surveying the sky stably and continuously since 13 August 2009 [2]; its scientific programme was designed to extract essentially all of the information in the cosmic microwave background temperature anisotropies, because only a small fraction of the information available in the CMB had been extracted by prior missions [1]. That programme — extract all of the information a quantized physical channel carries, rather than a conventionally selected subset of it — is precisely the attitude we argue quantum hardware should adopt toward the bosonic ladder.

Our contributions are:

1. A perturbative derivation of the harmonic-limit leakage $L_P^{(2)} = \theta^4/2$ for a driven transmon Hamiltonian, with all arithmetic shown, showing perturbative breakdown for a π -pulse (Section 4.1).
2. A first-order leakage bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$ and an exact three-level leakage formula, evaluated with explicit arithmetic at stated operating points (Section 4.2).
3. An information-theoretic accounting of the binary truncation as an approximation entropy $S_{\text{Base}} = \ln(d/2)$, evaluated for a qudit of dimension $d = 12$ as $\ln 6 = 1.7917595$ nats = 2.5849625 bits, together with the discarded-information fraction $f_{\text{disc}} = 0.721057$ (Section 4.3).
4. A definition of the Planck Fidelity benchmark metric F_P and a falsifiability analysis of the conjecture $L_P > 0$, including the escape routes that would defeat it (Sections 3.5 and 4.4).

Throughout, "leakage" means population outside the computational subspace; "radix" means the base of the logical register (2 for a qubit, d for a qudit); "nats" means logarithms in base e .

2. Background and Related Work

The literature base for this paper divides into two groups: the Planck CMB mission papers [1]–[8], which we use as methodological and conceptual analogies (flagged explicitly each time), and the QNFO line of work [9]–[11], which is the direct scientific lineage of the truncation-cost thesis.

[1] The Scientific Programme of Planck states that for 40 years the cosmic microwave background has been the most important source of information about the geometry and contents of the Universe, that only a small fraction of the information available in the CMB has been extracted to date, and that Planck, the third space CMB mission after COBE and WMAP, is designed to extract essentially all of the information in the CMB temperature anisotropies. The design philosophy — extract all the information the quantized field carries, not a conventionally truncated subset — is the exact analogue of our thesis that a register should match the physical Hilbert-space dimension rather than a chosen binary slice of it.

[2] Planck Early Results: The Planck mission records that the Planck satellite was launched on 14 May 2009, has been surveying the sky stably and continuously since 13 August 2009, carries 74 detectors sensitive to frequencies between approximately 25 and 1000 GHz, and that after its first year the major elements of the satellite's performance exceeded their original technical requirements. We use it as the factual anchor for the mission whose name we borrow, and for the demonstration that a quantized physical channel can be operated continuously against its full design capacity rather than a nominal subset.

[3] The Planck Catalogue of Compact Sources (PCCS) describes nine single-frequency catalogues of compact Galactic and extragalactic sources detected over the entire sky in the first 15 months of operations, covering 30–857 GHz, with 90% completeness at 180 mJy in the best channel. The catalogue is itself an act of truncation: sources below a sensitivity threshold are declared outside the registered set, exactly as levels above $n = 1$ are declared outside the qubit. The stated completeness figure is a quantitative model for what a truncation threshold costs in coverage — a number the catalogue states rather than hides, which is the reporting practice we demand for transmon leakage.

[4] The Planck catalogue of Sunyaev-Zeldovich sources presents an all-sky catalogue of clusters and cluster candidates derived over the first 15.5 months. The catalogue's distinction between confirmed clusters and candidates illustrates the general accounting problem of separating registered signal from registered-but-unvalidated excess, structurally the same bookkeeping as separating computational population from leakage population: both require that the unvalidated excess be counted and reported, not discarded.

[5] The statistical study of extragalactic radio sources uses the Early Release Compact Source Catalogue to measure number counts dN/dS at 30, 44, 70, 100, 143 and 217 GHz, finding counts at 30, 44 and 70 GHz in good agreement with WMAP at nearby frequencies. Counting population above a flux threshold is the direct statistical analogue of counting population above the truncation rung of a ladder; the paper's practice of measuring the thresholded number counts dN/dS directly across six frequency bands, and validating them against WMAP at 30, 44 and 70 GHz, is the caution we adopt: a truncation-dependent observable must be measured against the untruncated object, not against a model of it.

[6] The HFI spectral response paper reports that the High Frequency Instrument spectral response was determined through ground-based tests of the focal plane in a cryogenic environment prior to launch, the goal being to measure the relative spectral response including out-of-band signal rejection, and that out-of-band rejection was verified through comparisons with HFI zodiacal light observations, where any out-of-band sky signal received would be evident in the data. Out-of-band rejection is the instrumental counterpart of leakage suppression: a detector nominally assigned to one frequency band inevitably responds outside it, and the honest engineering response is to measure and quantify that response rather than declare it zero. This is precisely the epistemic stance we demand for transmon leakage.

[7] The Second Planck Catalogue of Compact Sources is a catalogue of compact sources detected in single-frequency maps over the full mission duration and supersedes previous versions of the Planck compact source catalogues; its validation section reports that no clearly significant discrepancies in positions or flux densities were found. The act of superseding an earlier truncated catalogue with a longer-integration, more complete one — while quantifying the residual discrepancies against external data — models our central methodological recommendation: re-baseline the register against the full physical object, then quantify what the earlier truncation missed.

[8] The LFI calibration paper describes the pipeline used to calibrate Low Frequency Instrument timelines into thermodynamic temperatures over four years of uninterrupted operations, using the spin-synchronous modulation of the CMB dipole as calibrator. A calibration pipeline that converts raw physical signal into a declared unit system, with a known reference and a stated accuracy, is the structural template for our Planck Fidelity metric: a declared fidelity figure is meaningful only relative to a stated reference subspace, and the accuracy of the reference itself must be quantified.

[9] QNFO: Planck Fidelity: Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders is the most technically developed antecedent of this paper. It derives the first-order leakage bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$ and the exact three-level results $P_{2,\max} = 2r^2/(1 + 2r^2)$ and $\overline{P}_2 = r^2/(1 + 2r^2)$, and discusses the model-class scope of the conjecture.

[10] QNFO: Planck Fidelity: The Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders is the immediate predecessor of the present text. It states the harmonic-limit result $c_2^{(2)} = -\theta^2/\sqrt{2}$ and $L_P^{(2)} = \theta^4/2$, evaluating to **3.0440341** for a π -pulse and signaling perturbative breakdown; it states the anharmonic projection $L_P^{(\text{est})} \approx (\Omega/\alpha)^2$ with values 1.0×10^{-2} , 2.5×10^{-3} , and 4.0×10^{-2} at $\Omega/\alpha = 0.1$, 0.05 , and 0.2 ; and its discussion concedes that the prefactor is pulse-shape dependent, that no rigorous operator lower bound on L_P is derived, and that a skeptic could accept the structural claim while denying practical significance if engineered counter-sequences suppress leakage below materials-limited error rates. We carry forward its harmonic-limit derivation (Section 4.1), treat its $(\Omega/\alpha)^2$ estimate as superseded by the more explicit bound of [9] (Section 4.2), and preserve its self-critical framing in Section 6.

[11] QNFO: Project Rosetta v2.0 is the conceptual source of the radix-error framing. With its "Planck-Radix Correction" it corrects an earlier version by affirming that the transmon spectrum IS discrete (Planck quantization), so that binary truncation is a radix error, not a physics error; it asserts that the transmon is a qudit of dimension $d \sim 12$, not a qubit; and it defines $S_{\text{Base}} = \ln(d/2) \sim 1.79$ nats as the base-encoding entropy cost. Its statement of results tabulates a Translation Calorimeter as a falsifiable measurement of the approximation entropy S_A , a Rosetta Constant $R = k_B T \ln(1/\alpha_r)$, and a decomposition $S_{\text{Total}} = S_{\text{Obs}} + S_{\text{Base}} + S_{\text{Arch}} + S_{\text{Alg}} + S_{\text{Trot}}$; its discussion proposes three paths forward — analog bosonic simulation, fermionic native-discrete platforms, and hybrid architectures with bounded approximation entropy at the interface. We adopt the quantity S_{Base} , re-derive its value with explicit arithmetic in Section 4.3, and design our falsifiability analysis (Section 4.4) to be compatible in spirit with its calorimetric protocol.

3. Methods

3.1 Model

We model the transmon as a weakly anharmonic oscillator,

$$H_0 = \hbar \omega a^\dagger a + \frac{\hbar \alpha}{2} a^\dagger a^\dagger a a,$$

where ω is the ladder spacing, α the anharmonicity (the detuning of the $|1\rangle \leftrightarrow |2\rangle$ transition from ω), and $a, a^\dagger$ the ladder operators satisfying $a^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle$. A classical drive of envelope $\varepsilon(t)$ couples all adjacent levels:

$$H_{\text{drv}}(t) = \hbar \varepsilon(t) (a e^{i\omega_d t} + a^\dagger e^{-i\omega_d t}),$$

with drive frequency ω_d . The computational (qubit) subspace is $\mathcal{H}_q = \text{span}\{|0\rangle, |1\rangle\}$ with projector $P_{\text{qubit}} = |0\rangle\langle 0| + |1\rangle\langle 1|$, and Planck Loss for a state $|\psi\rangle = \sum_{n \geq 0} c_n |n\rangle$ is

$$L_P = 1 - \text{Tr}(P_{\text{qubit}} |\psi\rangle\langle\psi|) = \sum_{n \geq 2} |c_n|^2.$$

The mechanism is kinematic: $a^\dagger$ has nonzero matrix elements $\langle n+1 | a^\dagger | n \rangle = \sqrt{n+1}$ for all n , so the drive is not block-diagonal with respect to the qubit subspace, and no choice of materials enters H_{drv} .

3.2 Metric and conjecture

For a gate acting for time T_g starting from $|\psi_0\rangle \in \mathcal{H}_q$, define $L_P(T_g) = 1 - \langle \psi_0 | U^\dagger(T_g) P_{\text{qubit}} U(T_g) | \psi_0 \rangle$.

Planck Fidelity conjecture. For the Hamiltonian above with any drive $\varepsilon(t) \neq 0$ of finite duration T_g , and in the absence of any engineered counter-leakage mechanism, $L_P(T_g) > 0$.

Following [9], we state the scope of this conjecture explicitly: it is a model-class result, holding for finite-time, finite-amplitude drives of the stated $a^\dagger$ -form under the rotating-wave approximation, and shaped envelopes can null the first-order leakage at isolated, measure-zero parameter points. The falsifiability analysis of Section 4.4 is built on this scope.

3.3 Leakage bounds

Two leakage formulas are adopted from [9] and re-derived with full arithmetic in Section 4.2. The first is a first-order perturbative bound for the detuned three-level system,

$$L_2 = 8 \left(\frac{\Omega}{|\alpha|} \right)^2 \sin^2 \left(\frac{\alpha T_g}{2} \right),$$

where Ω is the drive Rabi amplitude and α the anharmonic detuning. The second is the exact three-level solution for constant resonant drive, parametrized by $r = \Omega/|\alpha|$:

$$P_{2,\max} = \frac{2r^2}{1 + 2r^2}, \quad \overline{P}_2 = \frac{r^2}{1 + 2r^2},$$

the peak and time-averaged populations on $|2\rangle$. The harmonic-limit second-order amplitude of Section 4.1 is carried from [10].

3.4 Approximation entropy of the radix

Following [11], we define the base-encoding entropy cost of imposing a binary register on a d -level ladder as

$$S_{\text{Base}}(d) = \ln \left(\frac{d}{2} \right) \text{ nats},$$

the information by which the physical ladder exceeds the binary register. For $d = 2$, $S_{\text{Base}} = 0$; for the transmon qudit with $d \approx 12$ as asserted in [11], we compute S_{Base} explicitly in Section 4.3. We also define the discarded-information fraction

$$f_{\text{disc}}(d) = \frac{\ln(d/2)}{\ln d},$$

the share of the full ladder's $\ln d$ information content that the binary register discards.

3.5 Planck Fidelity benchmark metric

We define the Planck Fidelity of a gate acting on input $|\psi_k\rangle$ as

$$F_P = 1 - \overline{L}_P, \quad \overline{L}_P = \frac{1}{N} \sum_{k=1}^N \sum_{n \geq 2} |\langle n | U(T) | \psi_k \rangle|^2,$$

where T is the gate time, $U(T)$ the full (untruncated) evolution operator, and the average runs over a set of N input states spanning $\mathcal{H}_q$. The metric differs from a standard qubit gate fidelity in that the reference evolution $U(T)$ is evaluated on the full ladder, so that leakage is counted as error by construction rather than absorbed into a truncated simulation. The calibration-pipeline analogy of [8] applies: F_P is meaningful only relative to the declared reference subspace P_{qubit} , which must be stated with the metric.

4. Analysis

Every input number is stated with its source; every arithmetic step is shown.

4.1 Second-order leakage amplitude in the harmonic limit

Inputs. Pulse area $\theta = \int \varepsilon(t) dt$ (model parameter; a π -pulse has $\theta = \pi/2$ in the convention where excitation probability is $\sin^2 \theta$); ladder matrix elements $\langle n+1 | a^\dagger | n \rangle = \sqrt{n+1}$ (from the definition of $a^\dagger$).

Derivation. Starting from $|0\rangle$, two applications of the raising operator reach $|2\rangle$. In second-order perturbation theory the amplitude is

$$c_2^{(2)} = \frac{(-i\theta)^2}{2!} \langle 2 | (a^\dagger)^2 | 0 \rangle.$$

Now $\langle 2 | (a^\dagger)^2 | 0 \rangle = \langle 2 | a^\dagger | 1 \rangle \langle 1 | a^\dagger | 0 \rangle = \sqrt{2} \cdot 1 = \sqrt{2}$, and $2! = 2$, so

$$c_2^{(2)} = \frac{-\theta^2}{2} \cdot \sqrt{2} = -\frac{\theta^2}{\sqrt{2}}.$$

Leakage probability.

$$L_P^{(2)} = |c_2^{(2)}|^2 = \frac{\theta^4}{2}.$$

Numerical evaluation for a π -pulse. With $\theta = \pi/2$:

$$\theta = 1.5707963 \dots, \quad \theta^2 = 2.4674011 \dots, \quad \theta^4 = 2.4674011^2 = 6.0880682 \dots,$$

$$L_P^{(2)} = \frac{6.0880682}{2} = 3.0440341.$$

This exceeds unity, which signals the breakdown of second-order perturbation theory in the harmonic limit: with zero anharmonicity, a resonant drive transfers population up the ladder without bound. The correct reading is qualitative but sharp: in the harmonic limit there is no finite- θ pulse that keeps L_P small; the binary register is incompatible with the uncorrected harmonic oscillator. Perturbation theory is valid only when anharmonicity detunes the higher rungs, to which we now turn.

4.2 Anharmonic leakage: first-order bound and exact three-level result

Inputs (illustrative assumptions, stated explicitly, following [9]). (A1) Drive amplitude Ω and anharmonic detuning α , same angular-frequency units. (A2) The first-order bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$ holds for the detuned three-level subsystem. (A3) Illustrative operating point $\Omega/|\alpha| = 0.05$ (perturbative check: $8 \times 0.05^2 = 0.02 \ll 1$, satisfied), with gate timing $\alpha T_g/2 = \pi/2$, the worst case within the first half-period of the anharmonicity beat. (A4) For the exact three-level result, $r = f_\Omega/f_\alpha = 20 \text{ MHz}/200 \text{ MHz} = 0.1$, so $r^2 = 0.01$; these frequencies are illustrative assumptions, not measurements.

First-order bound. At the stated operating point,

$$L_2 = 8 \times (0.05)^2 \times \sin^2\left(\frac{\pi}{2}\right) = 8 \times 0.0025 \times 1 = 0.02,$$

giving $F_P = 1 - 0.02 = 0.98$. Gate-averaging $\sin^2$ over uniformly distributed gate times (mean $1/2$):

$$\overline{L}_2 = 8 \times 0.0025 \times \frac{1}{2} = 0.01, \quad \overline{F}_P = 0.99.$$

Scaling check. Doubling $\Omega/|\alpha|$ from 0.05 to 0.10 gives $\overline{L}_2 = 4 \times 0.01 = 0.04$; the ratio $0.04/0.01 = 4 = 2^2$ confirms the quadratic scaling numerically.

Exact three-level peak. With $r^2 = 0.01$:

$$P_{2,\max} = \frac{2 \times 0.01}{1 + 2 \times 0.01} = \frac{0.02}{1.02} = 0.019608 \approx 1.96 \times 10^{-2},$$

$$\overline{P}_2 = \frac{0.01}{1.02} = 0.009804 \approx 9.80 \times 10^{-3}.$$

π -pulse duration. The resonant π -pulse duration at $f_\Omega = 20$ MHz is

$$t_\pi = \frac{1}{2f_\Omega} = \frac{1}{2 \times 20 \times 10^6 \text{ s}^{-1}} = 2.5 \times 10^{-8} \text{ s} = 25 \text{ ns}.$$

Consistency with the cruder projection. The estimate $L_P^{(\text{est})} \approx (\Omega/\alpha)^2$ carried from [10] gives 1.0×10^{-2} at $\Omega/\alpha = 0.1$, 2.5×10^{-3} at 0.05, and 4.0×10^{-2} at 0.2 (arithmetic: $0.1^2 = 0.01$, $0.05^2 = 0.0025$, $0.2^2 = 0.04$). The gate-averaged bound $\overline{L}_2 = 0.01$ at $\Omega/|\alpha| = 0.05$ is a factor of 4 above the crude estimate at the same ratio; we attribute this to the explicit prefactor 8 and the gate-averaging factor $1/2$ in the bound, and we treat the bound of [9] as the more conservative and more explicit figure. Both are projections under stated assumptions, not measurements. The realism of the illustrative frequency ratio $f_\Omega/f_\alpha = 0.1$ is an open empirical question; the derived leakage values scale quadratically in r , so any correction to the assumed ratio rescales them by its square.

4.3 Approximation entropy for the transmon qudit

Inputs. $d = 12$ (the transmon qudit dimension asserted in [11]); $\ln$ denotes the natural logarithm.

Derivation.

$$S_{\text{Base}}(12) = \ln \frac{12}{2} = \ln 6 = \ln 2 + \ln 3 = 0.6931472 + 1.0986123 = 1.7917595 \text{ nats}.$$

Converted to bits:

$$S_{\text{Base}}(12)|_{\text{bits}} = \frac{1.7917595}{\ln 2} = \frac{1.7917595}{0.6931472} = 2.5849625 \text{ bits}.$$

Discarded-information fraction. The full ladder content is $\ln 12 = \ln 4 + \ln 3 = 1.3862944 + 1.0986123 = 2.4849067$ nats. Then

$$f_{\text{disc}} = \frac{\ln 6}{\ln 12} = \frac{1.7917595}{2.4849067} = 0.721057,$$

i.e. 72.1% of the ladder's information content is discarded by binary truncation; the radix efficiency is $\eta = 1 - 0.721057 = 0.278943 \approx 27.9\%$. Equivalently, one $d = 12$ oscillator carries $\log_2 12 = \ln 12 / \ln 2 = 2.4849067 / 0.6931472 = 3.5849625$ bits, so a binary register captures $2/3.5849625 = 0.5578859$ of the oscillator's log-capacity while requiring $\lceil \log_2 12 \rceil = 4$ qubits to represent its index space.

Scaling table. The scaling of the quantity S_{Base} defined in Section 3.4 is:

d	$S_{\text{Base}}(d) = \ln(d/2)$ (nats)	(bits)
2	$\ln 1 = 0$	0
4	$\ln 2 = 0.6931472$	1.0000000
6	$\ln 3 = 1.0986123$	1.5849625
12	$\ln 6 = 1.7917595$	2.5849625

(arithmetic: bits via division by $\ln 2 = 0.6931472$; e.g. $1.0986123/0.6931472 = 1.5849625$.)

4.4 Falsifiability threshold for the conjecture $L_P > 0$

The conjecture of Section 3.2 is falsified if a finite-time gate on a physical transmon achieves $L_P = 0$ exactly. Operationally, since any measurement has finite resolution δL , the conjecture is tested by the chain:

Under assumptions (A1)–(A2), the detuned three-level subsystem sustains strictly positive leakage $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2) > 0$ for any nonzero drive over a finite time, except at the measure-zero gate timings where $\sin^2(\alpha T_g/2) = 0$, which [9] notes can be nulled by shaped envelopes only at isolated parameter points. We emphasize that L_2 is a first-order estimate of leakage in a truncated three-level subsystem, not a proven lower bound on the full-ladder leakage L_P ; the conjecture $L_P > 0$ is consistent with this estimate but is not bounded below by it, and a rigorous lower bound remains an open question (Section 6). The escape routes are: (i) $\Omega \rightarrow 0$, which for a fixed pulse area $\theta = \Omega T_g$ forces $T_g \rightarrow \infty$ — the adiabatic limit, not a finite-time gate; (ii) a drive engineered to exactly cancel the $a^\dagger$ coupling on all rungs simultaneously, which the structure $[a^\dagger, P_{\text{qubit}}] \neq 0$ forbids for a subspace-preserving pulse within the model class. We state this as a derivation of the *structure* of the conjecture, with the numerical magnitude carried by the projections of Section 4.2.

5. Results

We report only (i) quantities computed in Section 4 and (ii) projections with stated assumptions.

R1 (computed). In the harmonic limit ($\alpha \rightarrow 0$), the second-order leakage amplitude for a pulse of area θ is $c_2^{(2)} = -\theta^2/\sqrt{2}$, giving $L_P^{(2)} = \theta^4/2$; for $\theta = \pi/2$ this evaluates to 3.0440341, exceeding unity and signaling perturbative breakdown — the harmonic oscillator admits no perturbatively small-leakage binary gate at all.

R2 (projection; assumptions (A1)–(A3)). With anharmonic detuning at $\Omega/|\alpha| = 0.05$ and worst-case timing, the first-order bound gives $L_2 = 0.02$ ($F_P = 0.98$) and a gate-averaged $\overline{L}_2 = 0.01$ ($\overline{F}_P = 0.99$); the scaling is quadratic, confirmed numerically by the factor-of-4 increase to 0.04 at $\Omega/|\alpha| = 0.10$.

R3 (computed; assumptions (A4)). The exact three-level solution at $r = 0.1$ gives $P_{2,\text{max}} = 0.02/1.02 = 1.96 \times 10^{-2}$ and $\overline{P}_2 = 9.80 \times 10^{-3}$, with $t_\pi = 25$ ns at the illustrative $f_\Omega = 20$ MHz.

R4 (computed). The base-encoding entropy of a $d = 12$ transmon qudit is $S_{\text{Base}} = \ln 6 = 1.7917595$ nats = 2.5849625 bits, matching the value stated in [9] and [11].

R5 (computed). The discarded-information fraction is $f_{\text{disc}} = 1.7917595/2.4849067 = 0.721057$: binary truncation discards 72.1% of the $d = 12$ ladder's $\ln 12$ information content, with radix efficiency $\approx 27.9\%$ and full log-capacity $\log_2 12 = 3.5849625$ bits.

R6 (computed scaling law). $S_{\text{Base}}(d) = \ln(d/2)$ vanishes only at $d = 2$ and grows logarithmically: 0, 0.6931472, 1.0986123, 1.7917595 nats at $d = 2, 4, 6, 12$ respectively. Any base-2 register on a ladder with $d > 2$ carries a strictly positive encoding cost under this definition.

R7 (structural result). The conjecture $L_P > 0$ for finite-time gates holds under assumptions (A1)–(A2) because $L_2 \propto (\Omega/|\alpha|)^2 > 0$ whenever $\Omega > 0$ and $|\alpha| < \infty$ away from measure-zero timings; the adiabatic route ($T_g \rightarrow \infty$) removes the leakage only by removing the gate.

R8 (programme, not result). The Planck Fidelity metric F_P of Section 3.5 is specified and ready for execution; no simulation or experimental numbers are reported in this paper. Benchmarking F_P against experimental leakage data is the designated empirical test.

6. Discussion

Limitations of the leakage estimates. The first-order bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$ is first order in Ω and neglects levels $n \geq 3$; the exact three-level solution is itself a truncation. The prefactor depends on pulse shape, rise time, and DRAG (Derivative Removal by Adiabatic Gate) pulse shaping, a standard counter-leakage pulse family; the true coefficient could differ by an order of magnitude. We have not derived a rigorous operator lower bound on L_P for a general pulse — only a structural argument plus scaling estimates. A skeptic could accept the structural claim ($a^\dagger$ couples all rungs) while denying that the resulting leakage is practically significant, if engineered counter-sequences suppress it below materials-limited error rates.

Limitations of S_{Base} . The quantity $S_{\text{Base}} = \ln(d/2)$ is adopted from [11] and is a bookkeeping entropy, not a thermodynamic or operational entropy; we have not shown that it corresponds to any measurable resource cost (time, energy, error budget). The choice $d = 12$ is taken from [11] without independent derivation; since S_{Base} depends logarithmically on d , its value is conditional on that input — at $d = 6$ the discarded fraction would be $\ln 3 / \ln 6 = 1.0986123 / 1.7917595 = 0.613147$, still large but different. It is also basis- and projector-dependent: a different choice of computational subspace changes L_P , so the conjecture should be read as a statement about the binary projector specifically.

Limitations of the literature base. Eight of the eleven references [1]–[8] concern the Planck CMB mission and come from a disjoint domain (cosmology instrumentation), used here only as methodological analogies; every statement we make about them is drawn from their own texts, and their role here is analogical. The substantive scientific lineage is [9]–[11], whose leakage formulas, entropy definition, and self-critical framing we adopt and re-derive.

Failure modes and falsification. The single observation that would falsify the main result is a finite-time gate on a well-characterized transmon that achieves leakage strictly zero up to arbitrarily small materials-limited error floors — i.e., if pulse shaping can drive L_P to zero faster than all other error channels and hold it there across a continuum of gate timings, the conjecture fails. Conversely, observation of a leakage floor that persists as materials-limited errors (dephasing, dielectric loss) are independently suppressed would support it. A second falsification route: if the transmon's usable ladder were actually bounded near $d \approx 2$ by physics rather than control, the radix-error framing would lose force.

Against ourselves. A skeptic may rightly object that leakage in practice is dominated by ordinary control error, that randomized compiling and DRAG-style pulse shaping reduce leakage to levels far below other error sources, and that calling the residual "fundamental" is rhetoric until a proven lower bound exists. We agree that this is the weakest point of the paper: our 1.0×10^{-2} gate-averaged figure is a model estimate under assumptions (A1)–(A3), and the honest claim is the structural one — the coupling $\sqrt{n+1} \varepsilon(t)$ never vanishes — not the magnitude.

Open questions. (i) A rigorous lower bound on $L_P(T_g)$ for bounded-amplitude drives of fixed pulse area. (ii) The scaling of gate infidelity with resolvable d and its relation to $S_{\text{Base}} = \ln(d/2)$. (iii) Experimental benchmarking of the Planck Fidelity metric against leakage data. (iv) Architecture implications: whether base- d qudit encodings recover the ≈ 2.58 bits per transmon that binary truncation discards.

7. Practical Implications and Engineering Pathways

(a) Users and decisions. The users are superconducting-processor architects and pulse-design engineers who compile gate-error budgets. If the result holds, the leakage line item in those budgets changes status: from a tunable control artifact to be minimized toward zero, to a bounded structural cost to be measured, reported, and traded against other channels — and the radix decision (qubit versus d -level qudit register per oscillator) enters the architecture trade space, with the ≈ 2.58 bits per $d = 12$ oscillator quantified in Section 4.3 as the capacity at stake.

(b) Readiness. 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 claim above TRL 3 is made, and no manufacturing readiness is implied.

(c) Smallest next test. The smallest de-risking experiment is a three-level numerical check plus a single-qubit experiment: simulate the exact three-level system of Section 4.2 for a calibrated pulse and measure P_2 by leakage-detection (e.g. a mapping pulse from $|2\rangle$ to a distinguishable state) on one transmon at $r \approx 0.1$. The outcome that refutes the structural claim is a measured P_2 consistent with zero across a continuum of gate timings and pulse shapes at fixed pulse area; the outcome that supports it is a residual P_2 floor that tracks the predicted $(\Omega/|\alpha|)^2$ scaling as Ω is varied.

(d) Binding constraints. Within the model class, the first binding constraint is control bandwidth: suppressing leakage by reducing $\Omega/|\alpha|$ lengthens gates as $t_\pi = 1/(2f_\Omega)$, so at the illustrative $f_\Omega = 20$ MHz the 25 ns π -pulse sets the speed–leakage trade-off. Whether coherence time, control electronics linearity, or measurement fidelity binds first in a qudit implementation is an open question this paper does not have the grounding to rank.

(e) Order-of-magnitude bounds (projections). Under assumptions (A1)–(A3), the gate-averaged leakage floor scales as $\overline{L}_2 = 4(\Omega/|\alpha|)^2$: at $\Omega/|\alpha| = 0.05, 0.01$; at $0.025, 2.5 \times 10^{-3}$; at $0.01, 4.0 \times 10^{-4}$. These are projections from the derived formula, not measurements, and carry the prefactor uncertainty stated in Section 6.

(f) Risks and misuse. The main misuse is rhetorical: quoting the 1.0×10^{-2} figure as a measured floor, or citing the $\backslash$ \$-scale of any industry, neither of which this paper establishes. The structural claim is about a model class; presenting it as a proven universal bound would overstate the result.

8. Conclusion

We have formalized the Planck Fidelity conjecture: binary truncation of a discrete-but-unbounded bosonic ladder incurs a materials-independent leakage $L_P > 0$ for any finite-time gate, because the drive couples all levels through $a^\dagger$ with matrix elements $\sqrt{n+1}$ that never vanish. We showed that in the harmonic limit the second-order leakage probability $L_P^{(2)} = \theta^4/2$ exceeds unity for a π -pulse (3.0440341), so the binary register is incompatible with an uncorrected harmonic oscillator; that finite anharmonicity suppresses the leakage to a gate-averaged $\overline{L}_2 = 0.01$ at $\Omega/|\alpha| = 0.05$ under the first-order bound, with the exact three-level peak $P_{2,\text{max}} = 1.96 \times 10^{-2}$ at $r = 0.1$; and that the radix-truncation entropy is exactly $S_{\text{Base}}(12) = \ln 6 = 1.7917595$ nats = 2.5849625 bits, a discarded-information fraction of 72.1%. The

transmon is best regarded as a qudit; the qubit is an approximation whose entropy cost we have made explicit. The remaining gap — a rigorous lower bound — is the natural next step, and the conjecture stands or falls on the falsifiable leakage-floor test of Section 4.4.

Appendix A. Divergence report

D1 (leakage bound). Draft A claimed a universal lower bound $L_P \geq 1/d$ derived from a uniform-amplitude argument, and reported simulated leakage values (0.25, 0.17, 0.13, 0.10, 0.083 for $d = 4, 6, 8, 10, 12$) from numerical diagonalisation. Drafts B and C instead derived a perturbative scaling estimate and explicitly reported no simulation results. The disagreement stems from a methodological convention: Draft A treated the $1/d$ bound as derivable and its simulations as executed, while Drafts B and C held that no rigorous bound or simulation was available. Resolution: the main text follows Drafts B and C, because the $1/d$ "bound" is not actually a lower bound (a state concentrated in $\{|0\rangle, |1\rangle\}$ has $L_P = 0$ in a d -level space; the uniform-distribution argument bounds nothing), and because Draft A's simulation numbers cannot be verified. Draft A's simulation table is dropped entirely.

D2 (bits conversion). Draft B reported $S_{\text{Base}}(12) \approx 2.5840$ bits; Draft C reported 2.5850 bits. The conflict is a rounding convention: the exact conversion is $1.7917595/0.6931472 = 2.5849625$ bits, so Draft B's 2.5840 was a rounding error rather than a different calculation. Resolution: the main text adopts 2.5849625 bits.

D3 (leakage formula choice). Draft B used the crude projection $L_P^{(\text{est})} \approx (\Omega/\alpha)^2$ with its value 1.0×10^{-2} at $\Omega/\alpha = 0.1$; Draft C used the explicit first-order bound $L_2 = 8(\Omega/|\alpha|)^2 \sin^2(\alpha T_g/2)$ of [9] with its gate-averaged value 0.01 at $\Omega/|\alpha| = 0.05$. The conventions differ in which operating point and which formula is headline. Resolution: the main text presents both, treats the explicit bound of [9] as the more conservative figure, and reports the crude projection as a consistency check (Section 4.2).

Appendix B. Claim attribution

ID	Claim	Source drafts	Agreement status
C1	The transmon is a weakly anharmonic oscillator with a discrete but unbounded ladder; the binary register is imposed by truncation, and leakage L_P is population outside $\text{span}\{ 0\rangle, 1\rangle\}$	A, B, C	CONVERGENT
C2	Planck Fidelity conjecture: $L_P > 0$ for any finite-time gate with nonzero drive, as a property of the encoding rather than of fabrication	A, B, C	CONVERGENT
C3	Harmonic-limit second-order leakage $L_P^{(2)} = \theta^4/2$, evaluating to 3.0440341 for a π -pulse and signaling perturbative breakdown	B, C	CONVERGENT
C4	First-order leakage bound $L_2 = 8(\Omega/ \alpha)^2 \sin^2(\alpha T_g/2)$ with gate-averaged value 0.01 at $\Omega/ \alpha = 0.05$, adopted from [9] and re-derived with shown arithmetic	B, C	CONVERGENT
C5	Universal lower bound $L_P \geq 1/d$ plus simulated leakage values (0.25, 0.17, 0.13, 0.10, 0.083 for $d = 4, 6, 8, 10, 12$) from numerical diagonalisation	A	DIVERGENT (rejected; see Appendix A, D1)
C6	Exact three-level leakage $P_{2,\text{max}} = 2r^2/(1 + 2r^2)$ and $\overline{P}_2 = r^2/(1 + 2r^2)$, evaluated at $r = 0.1$ giving 1.96×10^{-2} and 9.80×10^{-3} , with $t_\pi = 25$ ns at $f_\Omega = 20$ MHz	B, C	CONVERGENT
C7	Approximation entropy $S_{\text{Base}} = \ln(d/2)$ with $d = 12$ from [11], evaluated as $\ln 6 = 1.7917595$ nats	A, B, C	CONVERGENT

ID	Claim	Source drafts	Agreement status
C8	Bits conversion of $S_{\text{Base}}(12)$	B, C	DIVERGENT (resolved as 2.5849625 bits; see Appendix A, D2)
C9	Headline leakage figure: crude projection $L_P^{(\text{est})} \approx (\Omega/\alpha)^2$ at $\Omega/\alpha = 0.1$ versus explicit bound of [9] at $\Omega/ \alpha = 0.05$	B, C	DIVERGENT (both retained; explicit bound treated as more conservative; see Appendix A, D3)
C10	Discarded-information fraction $f_{\text{disc}} = \ln 6 / \ln 12 = 0.721057$, i.e. 72.1% of the ladder's $\ln 12$ content discarded by binary truncation	B, C	CONVERGENT
C11	Planck Fidelity metric $F_P = 1 - \overline{L_P}$ referenced to the full untruncated evolution, with the reference subspace stated alongside the metric, following the calibration-reference practice of [8]	B, C	CONVERGENT
C12	Analogical use of the Planck CMB mission literature [1]–[8], with every statement about them drawn from their own texts and their role flagged as analogical each time	B, C	CONVERGENT
C13	Falsifiability of the conjecture via persistence of a leakage floor as materials-limited errors are suppressed, compatible in spirit with the calorimetric protocol of [11]	A, B, C	CONVERGENT
C14	Model-class scoping of the conjecture: finite-time, finite-amplitude drives of the stated $a^\dagger$ -form under the rotating-wave approximation, with shaped envelopes nulling first-order leakage only at isolated, measure-zero parameter points	B, C	CONVERGENT
C15	Radix decision (qubit versus base- d qudit register per oscillator) enters the architecture trade space, with the ≈ 2.58 bits per $d = 12$ oscillator as the capacity at stake	B, C	CONVERGENT
C16	No simulation or experimental results reported; the Planck Fidelity benchmark and the three-level numerical check are specified as the designated next tests	B, C	CONVERGENT

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.3-codeagent). 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; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- 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 20 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 the retrieved full-text excerpts; 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] The Scientific Programme of Planck. arXiv:astro-ph/0604069v1. <https://arxiv.org/abs/astro-ph/0604069v1>
- [2] Planck Early Results: The Planck mission. arXiv:1101.2022v2. <https://arxiv.org/abs/1101.2022v2>
- [3] Planck 2013 results. XXVIII. The Planck Catalogue of Compact Sources. arXiv:1303.5088v2. <https://arxiv.org/abs/1303.5088v2>
- [4] Planck 2013 results. XXIX. Planck catalogue of Sunyaev-Zeldovich sources. arXiv:1303.5089v2. <https://arxiv.org/abs/1303.5089v2>
- [5] Planck Early Results: Statistical properties of extragalactic radio sources in the Planck Early Release Compact Source Catalogue. arXiv:1101.2044v2. <https://arxiv.org/abs/1101.2044v2>
- [6] Planck 2013 results. IX. HFI spectral response. arXiv:1303.5070v2. <https://arxiv.org/abs/1303.5070v2>
- [7] Planck 2015 results. XXVI. The Second Planck Catalogue of Compact Sources. arXiv:1507.02058v2. <https://arxiv.org/abs/1507.02058v2>
- [8] Planck 2015 results. V. LFI calibration. arXiv:1505.08022v2. <https://arxiv.org/abs/1505.08022v2>
- [9] QNFO: Planck Fidelity: Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders
- [10] QNFO: Planck Fidelity: The Fundamental Leakage Cost of Binary Truncation of Bosonic Ladders
- [11] QNFO: Project Rosetta: The Approximation Entropy & The Fractal Limits of Digital Physics — v2.0