

Parity-Protected Superconducting Qubits at 4 K: A Reconciled Thermodynamic Error-Budget Analysis

Parity-Protected Superconducting Qubits at 4 K: A Reconciled Thermodynamic Error-Budget Analysis

Abstract

Scalable superconducting quantum computing is conventionally operated at millikelvin temperatures, where the cooling power available from dilution refrigerators is severely limited. This paper tests the conjecture that a thermodynamic bottleneck at millikelvin temperatures could be relieved by operating parity-protected superconducting qubits at 4 K, where closed-cycle refrigeration supplies roughly 2×10^4 times more cooling power. We construct an explicit thermal error budget with fully shown arithmetic: the thermal energy scale at 4 K is $k_B T/h \approx 83.3465$ GHz; against the niobium gap $2\Delta_{\text{Nb}}/h \approx 700$ GHz this yields a Boltzmann quasiparticle factor $e^{-\Delta/k_B T} \approx 1.5006 \times 10^{-2}$, compared with $\approx 1.12 \times 10^{-73}$ at 100 mK — a penalty factor of order $10^{71.13}$. Aluminum ($T_c = 1.2$ K) is shown to be thermodynamically excluded as a junction material at 4 K. Against a stated fault-tolerance target of $p_{\text{target}} = 10^{-4}$, an unprotected niobium qubit requires a suppression factor $S = 150$ on thermally activated quasiparticle errors and meets the target only below $T^* \approx 1.82$ K. The parity-protected coupling of the top-transmon proposal, which couples to fermion parity irrespective of quasiparticle excitations, is the right kind of mechanism for this channel, but residual channels (dielectric two-level-system loss in Nb_2O_5 , ancillary-junction poisoning, readout thermal noise) remain unquantified. We conclude that 4 K operation is not fundamentally excluded for parity-protected architectures but is excluded for conventional aluminum-based qubits, and we state the measurements that would falsify this conclusion.

1. Introduction

The scaling path of superconducting quantum processors is constrained not only by coherence and gate fidelity but by the thermodynamics of the cryogenic plant. Dilution refrigerators deliver very small cooling power at the millikelvin stage, and every additional qubit, cable, and amplifier adds heat load to a stage that cannot be expanded cheaply. The motivating conjecture of this paper, drawn from the corpus documents on thermodynamic and quantum constraints on scalable quantum computing [9], [10], [11], is that this constitutes a *thermodynamic bottleneck*: scalability is limited less by physics at the qubit level than by the refrigeration envelope imposed by the millikelvin operating point. At 4 K, closed-cycle cooling technology offers approximately a 2×10^4 -fold increase in available cooling power, an enormous relaxation of the thermal budget.

The question is whether qubit physics permits operation at 4 K. Conventional superconducting qubits fail immediately at 4 K for material reasons quantified below: aluminum, the workhorse Josephson material, has a critical temperature of $T_{c,\text{Al}} = 1.2$ K and cannot host a superconducting gap at 4 K. Niobium survives, but its gap suppresses thermal quasiparticles by only a modest Boltzmann factor at 4 K. The candidate escape route is *parity protection*: qubits whose logical states are distinguished by fermion parity, so that the computational information is not stored in modes that thermally excited quasiparticles can corrupt. The top-transmon proposal [6] is the concrete embodiment of this idea: it couples directly to the fermion parity of a hybrid superconducting–topological system, irrespective of any quasiparticle excitations, which is precisely the property a 4 K architecture needs.

This paper makes the conjecture quantitative. We derive the thermal error budget of a 4 K niobium-based parity-protected qubit, compute the required error suppression against a stated fault-tolerance target, and identify which error channels parity protection can and cannot close. Every number is either a stated input, a fundamental constant, or derived here with shown arithmetic; where a full treatment would require a Monte Carlo simulation of quasiparticle poisoning — which we specify as a proposed method but do not execute — we state the assumptions such a simulation would need.

2. Background and Related Work

We draw on twelve supplied bibliography entries. Because several entries provide only truncated abstracts, we confine every statement about a cited work to what its own summary states.

Application driver. Reference [1] (*Quantum Computing: Vision and Challenges*, arXiv:2403.02240v5) frames quantum computing as using entanglement, superposition, and other quantum fundamental concepts to provide substantial processing advantages over traditional computing, for problems including modeling quantum mechanics, logistics, and chemical-based simulation. The scale of the intended applications motivates asking not merely whether a qubit works, but whether an architecture can be built and cooled at scale — exactly the question this paper addresses.

Software and delegation stack. Reference [2] presents Tierkreis, a higher-order dataflow graph representation and runtime for compositional quantum-classical hybrid algorithms, motivated by the remote nature of quantum computers, cloud and distributed computing needs, and long-running algorithms. A 4 K architecture changes the physical siting of control electronics, and dataflow-style runtimes of the kind Tierkreis describes are the software layer that would absorb such changes; the supplied summary gives no further detail, so we relate it to our argument only at this level. Reference [4] unifies communication paradigms in measurement-based delegated quantum computing, in which clients with low quantum capabilities outsource computations to a server, with blindness of inputs facilitated by the measurement-based framework. The client-server architecture it describes is indifferent to the server's operating temperature; our thermodynamic analysis is complementary to such architectural work.

Protected superconducting qubits. Reference [3] examines the energetic advantage of superconducting cat-qubits, noting that superconducting qubits are a leading technology due to their scalability and compatibility with quantum error correction mechanisms; the supplied summary truncates before its energetic findings are stated, so we cite it for framing only — it treats energetic advantage as an explicit design axis for superconducting qubits, the same axis this paper quantifies for the 4 K operating point. Reference [6] (*Top-transmon*, arXiv:1105.0315v1) is the central prior work: its summary states that qubits constructed from uncoupled Majorana fermions are protected from decoherence, that performing a quantum computation requires breaking this topological protection, and that parity-protected quantum computation breaks it minimally invasively by coupling directly to the fermion parity of the system — irrespective of any quasiparticle excitations — using a superconducting hybrid structure. The phrase "irrespective of any quasiparticle excitations" is exactly the property a 4 K operating point requires. The supplied summary contains no device-level performance figures, so we use [6] only for its stated mechanism and not for any quantitative device data.

Alternative platforms and scale benchmarks. Reference [5] describes arrays of optically trapped atoms excited to Rydberg states as a competitive platform in which high-fidelity state preparation and readout, quantum logic gates, and controlled dynamics of more than 100 qubits have been demonstrated, with systems approaching reliable computations with hundreds of qubits (the summary truncates there). Rydberg platforms do not face superconducting-gap thermodynamics, so the 4 K question analyzed here is specific to superconducting-style technology. Reference [8] introduces Noisy Intermediate-Scale Quantum (NISQ) technology: quantum computers with 50–100 qubits may perform tasks surpassing classical digital computers, but noise in quantum gates will limit the size of circuits that can be executed reliably. This is the framing our error budget addresses directly: the question at 4 K is whether gate noise from thermal channels can be pushed below the level at which reliable circuit execution becomes possible.

Scope delimitation. Reference [7] discusses how many qubits quantum computational supremacy arguments require, noting that such arguments rest on computational assumptions related to the limitations of classical computation, such as the assumption that the polynomial hierarchy does not collapse — a stronger version of $P \neq NP$. We cite it to delimit scope: this paper addresses physical scalability, not computational-complexity arguments.

Corpus documents. References [9]–[12] are QNFO corpus entries (DOIs <https://papers.qnfo.org/papers/thermodynamic-constraints-and-architectural-inversions-in-scalable-quantum-information-systems/>, <https://papers.qnfo.org/papers/thermodynamic-and-informational-bottlenecks-of-scalable-fault-tolerant-quantum-computation/>, <https://papers.qnfo.org/papers/thermodynamic-imperative/>, <https://papers.qnfo.org/papers/syntactic-generation-primitive-distinctions/>) whose supplied summaries are empty. Their titles indicate coverage of thermodynamic and quantum constraints

on scalable quantum computing, thermodynamic and informational bottlenecks of scalable fault-tolerant quantum computation, a "Thermodynamic Imperative" theme, and syntactic generation. We cite them as the provenance of the thermodynamic-bottleneck framing but make no specific claim about their contents beyond what the titles state.

3. Methods

3.1 Operating points and constants

We compare two operating points: $T_{\text{warm}} = 4$ K (the operating point under test) and $T_{\text{cold}} = 0.1$ K (100 mK, adopted as the explicit reference dilution-refrigerator stage; see Appendix A for the reconciliation of reference-temperature choices). We use the exact SI defining values $k_B = 1.380649 \times 10^{-23}$ J/K and $h = 6.62607015 \times 10^{-34}$ J·s, so that the thermal energy in frequency units is

$$\frac{k_B T}{h} = \frac{k_B}{h} T, \quad \frac{k_B}{h} = 20.837 \text{ GHz/K}.$$

3.2 Materials

Two superconductors bracket the analysis:

- **Niobium (Nb):** the motivating input specifies $2\Delta_{\text{Nb}}/h \approx 700$ GHz, i.e., $\Delta_{\text{Nb}}/h = 350$ GHz.
- **Aluminum (Al):** the motivating input specifies $T_{c,\text{Al}} = 1.2$ K. We use the weak-coupling BCS relation $\Delta \approx 1.76 k_B T_c$, stated here as an explicit modeling assumption.

3.3 Error model

The thermally activated quasiparticle population is modeled by the Boltzmann factor

$$p_{\text{th}}(T) = \exp\left(-\frac{\Delta}{k_B T}\right) = \exp\left(-\frac{\Delta/h}{(k_B/h) T}\right),$$

the equilibrium probability that a mode at the gap edge is thermally occupied, up to a prefactor absorbed into an attempt factor A . For a gate-level error probability we write

$$p_{\text{qp}}(T) = A p_{\text{th}}(T),$$

with A an attempt factor that a quasiparticle-poisoning Monte Carlo simulation would need to determine; we treat A parametrically rather than asserting a value.

3.4 Parity protection and suppression requirement

Following the mechanism stated in [6], we formalize parity protection as a parity-selectivity factor S_{par} on the quasiparticle-poisoning channel, so that $p_{\text{eff}} = A S_{\text{par}} p_{\text{th}}(T)$. We adopt a fault-tolerance target $p_{\text{target}} = 10^{-4}$ per operation as a stated design assumption (not attributed to any cited work; see Appendix A for the reconciliation of threshold choices), and define the required suppression

$$S \equiv \frac{p_{\text{th}}(T_{\text{warm}})}{p_{\text{target}}}.$$

3.5 Thermodynamic accounting

The motivating input supplies the practical cooling-power ratio $R_P = 2 \times 10^4$ between 4 K platforms and dilution-refrigerator stages. To separate the fundamental thermodynamic contribution from engineering reality, we additionally compute the ideal Carnot coefficient of performance $\text{COP} = T_c/(T_h - T_c)$ with $T_h = 300$ K.

3.6 Proposed (not executed) Monte Carlo

The motivating input proposes Monte Carlo simulation of quasiparticle poisoning rates and comparison against demonstrated 1.5 K silicon spin qubit data. No such simulation results and no silicon spin qubit dataset are supplied in the input material; we therefore specify the simulation design (sampling quasiparticle generation, tunneling, and parity-switching events with rates anchored to the Boltzmann factors computed in Section 4) but report **no** simulation numbers in this paper.

4. Analysis

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

Step 1: Boltzmann constant in frequency units.

$$\frac{k_B}{h} = \frac{1.380649 \times 10^{-23} \text{ J/K}}{6.62607015 \times 10^{-34} \text{ J} \cdot \text{s}} = 2.083664 \times 10^{10} \text{ Hz/K} = 20.8366 \text{ GHz/K}.$$

(Arithmetic: $1.380649/6.62607015 = 0.208366$; exponent $-23 - (-34) = +11$; so $0.208366 \times 10^{11} = 2.08366 \times 10^{10}$.)

Step 2: Thermal frequency scales.

$$\frac{k_B T_{\text{warm}}}{h} = 20.8366 \times 4 = 83.3465 \text{ GHz}, \quad \frac{k_B T_{\text{cold}}}{h} = 20.8366 \times 0.1 = 2.08366 \text{ GHz}.$$

The motivating input quotes $k_B T \approx 83 \text{ GHz}$ at 4 K, consistent with our 83.3465 GHz.

Step 3: Niobium gap ratios. With $\Delta_{\text{Nb}}/h = 350 \text{ GHz}$ (input: motivating input):

$$r_{\text{warm}} = \frac{350}{83.3465} = 4.1995, \quad r_{\text{cold}} = \frac{350}{2.08366} = 167.97.$$

Step 4: Thermal quasiparticle factors.

$$p_{\text{th}}^{\text{Nb}}(4 \text{ K}) = e^{-4.1995} = 0.015006 \approx 1.5006 \times 10^{-2}.$$

$$p_{\text{th}}^{\text{Nb}}(0.1 \text{ K}) = e^{-167.97}; \quad 167.97/\ln 10 = 167.97/2.302585 = 72.95, \quad \Rightarrow p_{\text{th}}^{\text{Nb}}(0.1 \text{ K}) \approx 1.12 \times 10^{-73}.$$

The ratio between operating points is $e^{167.97-4.1995} = e^{163.77}$; in base 10, $163.77/2.302585 = 71.126$, so raising the operating point from 100 mK to 4 K multiplies niobium's equilibrium gap-edge thermal occupation by a factor of order $10^{71.13}$.

Step 5: Aluminum exclusion. With $T_{c,\text{Al}} = 1.2 \text{ K}$ (input) and $\Delta = 1.76 k_B T_c$ (assumption):

$$\frac{\Delta_{\text{Al}}}{h} = 1.76 \times 1.2 \times 20.8366 = 1.76 \times 25.0039 = 44.0069 \text{ GHz}.$$

At 4 K: $\Delta_{\text{Al}}/k_B T = 44.0069/83.3465 = 0.5280$, so

$$p_{\text{th}}^{\text{Al}}(4 \text{ K}) = e^{-0.5280} = 0.5898.$$

At 100 mK: $\Delta_{\text{Al}}/k_B T = 44.0069/2.08366 = 21.121$; $21.121/2.302585 = 9.1735$, so $p_{\text{th}}^{\text{Al}}(0.1 \text{ K}) = 10^{-9.1735} = 6.73 \times 10^{-10}$. Independently of the gap model, $T_{\text{warm}} = 4 \text{ K} > T_{c,\text{Al}} = 1.2 \text{ K}$, so aluminum is a normal metal at the 4 K operating point and cannot host a superconducting gap there at all. Aluminum-based Josephson junctions are categorically excluded at 4 K.

Step 6: Required suppression factor. Against the stated target $p_{\text{target}} = 10^{-4}$:

$$S = \frac{1.50 \times 10^{-2}}{10^{-4}} = 150.$$

For sensitivity: if the target were 10^{-3} , $S = 15$; if 10^{-2} , $S = 1.5$.

Step 7: Break-even temperature for unprotected niobium. The temperature at which $p_{\text{th}} = p_{\text{target}}$ satisfies $r(T) = \ln(10^4) = 9.2103$, so

$$\frac{k_B T^*}{h} = \frac{350}{9.2103} = 38.00 \text{ GHz} \Rightarrow T^* = \frac{38.00}{20.837} = 1.824 \text{ K}.$$

An unprotected niobium qubit meets the 10^{-4} thermal-quasiparticle target only below $T^* \approx 1.82$ K; at 4 K it misses by the factor $S = 150$.

Step 8: Cooling-power arithmetic. With the labeled assumption $P_{\text{mK}} = 1 \mu\text{W}$ (millikelvin cooling budget; see Appendix A) and the input ratio $R_P = 2 \times 10^4$:

$$P_{4\text{K}} = R_P P_{\text{mK}} = 2 \times 10^4 \times 1 \times 10^{-6} \text{ W} = 2 \times 10^{-2} \text{ W} = 20 \text{ mW}.$$

All conclusions depending on cooling power scale linearly with P_{mK} .

Step 9: Carnot accounting. With $T_h = 300$ K:

$$\text{COP}_{4\text{K}} = \frac{4}{300 - 4} = \frac{4}{296} = 0.013514, \quad \text{COP}_{0.1\text{K}} = \frac{0.1}{299.9} = 3.3344 \times 10^{-4}.$$

$$\frac{\text{COP}_{4\text{K}}}{\text{COP}_{0.1\text{K}}} = \frac{0.013514}{3.3344 \times 10^{-4}} = 40.5.$$

Fundamental thermodynamics alone favors 4 K by a factor of ≈ 40.5 in ideal efficiency; the supplied practical factor 2×10^4 is dominated by engineering realities of dilution refrigeration, which the supplied input does not decompose.

Step 10: Parity-protection requirement (parametric). With $p_{\text{eff}} = A S_{\text{par}} p_{\text{th}}^{\text{Nb}}(4 \text{ K}) < p_{\text{target}}$:

$$S_{\text{par}} < \frac{10^{-4}}{A \times 1.50 \times 10^{-2}} = \frac{6.67 \times 10^{-3}}{A}.$$

For $A = 10^{-3}$: $S_{\text{par}} < 6.67$; for $A = 1$: $S_{\text{par}} < 6.67 \times 10^{-3}$. If the attempt factor is small (poisoning events rare relative to gate time), parity protection need only avoid amplifying the thermal channel; if poisoning events are frequent ($A \sim 1$), parity protection must genuinely decouple the logical state from quasiparticle number — exactly the property [6] claims for parity-selective coupling. The viability question at 4 K therefore reduces to the physical values of A and S_{par} , both empirical quantities not fixed by the supplied input.

Step 11: Illustrative relaxation model (labeled assumption). One draft introduced an illustrative extension assuming a baseline quasiparticle-induced relaxation rate $\Gamma_{\text{qp}} = \Gamma_0 p_{\text{th}}$ with the assumed $\Gamma_0 = 10^6 \text{ s}^{-1}$: then $\Gamma_{\text{qp}} = 10^6 \times 1.50 \times 10^{-2} = 1.50 \times 10^4 \text{ s}^{-1}$, giving $T_1 = 1/\Gamma_{\text{qp}} = 6.67 \times 10^{-5} \text{ s} = 66.7 \mu\text{s}$. With an assumed Rabi frequency $\Omega = 2\pi \times 5 \text{ MHz}$, the $\pi/2$ gate time is $t_g = \pi/(2\Omega) = 1/(4 \times 5 \times 10^6) = 5.0 \times 10^{-8} \text{ s} = 50 \text{ ns}$, and the relaxation-limited gate error is $t_g/T_1 = 5.0 \times 10^{-8}/6.67 \times 10^{-5} = 7.5 \times 10^{-4}$. Both Γ_0 and Ω are assumptions of this illustrative model, not measurements; we report it only to show how the Boltzmann factor propagates into a gate error under stated assumptions.

Step 12: Unquantified channels. The motivating input identifies thermally activated two-level-system (TLS) decoherence in niobium pentoxide (Nb_2O_5 , the native oxide of niobium) as a material constraint at 4 K, but supplies no TLS activation energies, densities, or loss tangents. No quantitative TLS error rate can be computed here; we record it as an open error channel.

5. Results

All numbers below are computed in Section 4 with shown arithmetic, or are explicitly labeled projections.

1. **Thermal scales.** $k_B T/h = 83.3465 \text{ GHz}$ at 4 K and 2.08366 GHz at 100 mK (computed, Steps 1–2).

1. **Niobium survives but is only moderately protected.** $p_{\text{th}}^{\text{Nb}}(4 \text{ K}) = 1.5006 \times 10^{-2}$ versus $\approx 1.12 \times 10^{-73}$ at 100 mK; the operating-point penalty is a factor of order $10^{71.13}$ (computed, Steps 3–4).

1. **Aluminum is excluded.** At 4 K, aluminum ($T_{c,\text{Al}} = 1.2$ K) is a normal metal; under the stated BCS assumption its gap-edge thermal factor is $p_{\text{th}}^{\text{Al}}(4 \text{ K}) = 0.5898$, versus 6.73×10^{-10} at 100 mK (computed, Step 5). Any 4 K superconducting architecture must avoid aluminum Josephson junctions.
1. **Suppression requirement.** Against the stated target $p_{\text{target}} = 10^{-4}$, a 4 K parity-protected qubit must suppress thermally activated quasiparticle errors by $S = 150$ (or $S = 15$ for a 10^{-3} target, $S = 1.5$ for a 10^{-2} target); an unprotected niobium qubit meets the 10^{-4} target only below $T^* \approx 1.82$ K (computed, Steps 6–7).
1. **Cooling-power dividend.** Under the labeled assumption $P_{\text{mK}} = 1 \mu\text{W}$ and the input ratio $R_P = 2 \times 10^4$, the 4 K stage offers $P_{4\text{K}} = 20$ mW (computed, Step 8); the ideal Carnot efficiency ratio is 40.5 (computed, Step 9). *Projection:* if parity protection achieves $S_{\text{par}} \leq 6.67 \times 10^{-3}/A$ and if the residual TLS and readout channels are held at or below 10^{-4} per operation — conditions not demonstrated here — then a 4 K parity-protected processor would gain the full R_P -fold cooling headroom at unchanged logical error rates. Uncertainty in this projection is dominated by the unquantified Nb_2O_5 dielectric loss and readout-chain thermal noise.
1. **Illustrative gate-error propagation.** Under the labeled assumptions of Step 11 ($\Gamma_0 = 10^6 \text{ s}^{-1}$, $\Omega = 2\pi \times 5 \text{ MHz}$), the relaxation-limited figures are $T_1 = 66.7 \mu\text{s}$, $t_g = 50 \text{ ns}$, and gate error 7.5×10^{-4} without parity protection; these are model outputs under stated assumptions, not measurements.
1. **No Monte Carlo results are reported.** The proposed simulation (Section 3.6) and the comparison against 1.5 K silicon spin qubit data require datasets not supplied in the input material; reporting invented numbers would violate the evidentiary standard of this paper.

6. Discussion

What the result does and does not show. The analysis shows a sharp asymmetry: conventional aluminum-based superconducting qubits are categorically impossible at 4 K (Step 5), while niobium-based parity-protected qubits face a finite, quantified requirement — a suppression factor of $S = 150$ on one identified channel against a 10^{-4} target — that the mechanism described in [6] is designed in principle to meet, since its coupling is stated to be irrespective of quasiparticle excitations. The analysis does *not* demonstrate that a full 4 K fault-tolerant processor is achievable, because three error channels remain unquantified: thermally activated TLS loss in Nb_2O_5 (no loss-tangent data in the supplied material), quasiparticle poisoning of unprotected ancillary elements required for gate operations, and thermal noise in parity readout at 4 K, where the thermal photon population at the 83 GHz scale is non-negligible for any readout line not heavily attenuated and filtered.

Against the conjecture. The strongest counterargument is that the idealized parity-protection model of Section 3.4 is too generous. The top-transmon summary [6] states the coupling mechanism but contains no device-level fidelity figures, and real hybrid topological–superconducting devices inherit the sub-gap quality of their parent materials. If the effective parity-conserving suppression falls short of $S = 150$, the 4 K operating point fails the stated target; the margin is roughly two orders of magnitude, thin by the standards of superconducting device variability. A second counterargument: the Boltzmann-factor error model ignores prefactors and nonequilibrium quasiparticle populations, which in practice can dominate equilibrium populations; if nonequilibrium quasiparticles raise the effective p_{qp} by more than S , the conclusion reverses even with perfect parity protection of the logical subspace. A third: the illustrative relaxation model of Step 11 shows that even modest assumed baseline rates push single-channel gate errors to $\sim 10^{-3}$ – 10^{-4} , comparable to the entire budget.

What would falsify the claims. Three results would falsify the paper's central claims: (i) a demonstration of a parity-protected qubit whose logical error rate at 4 K attributable to quasiparticles exceeds 1.50×10^{-2} per operation, showing the parity mechanism does not deliver $S \geq 150$; (ii) a measurement of Nb_2O_5 or surface two-level-system loss at 4 K exceeding 10^{-4} per gate, closing the 4 K window through a channel parity protection cannot address; (iii) evidence that the cooling-power ratio $R_P = 2 \times 10^4$ is not realizable in integrated control environments, undermining the thermodynamic motivation. Conversely, a demonstration of sub- 10^{-4} logical quasiparticle error at 4 K in a niobium hybrid device would strongly confirm the viability conjecture.

Open questions. What is the temperature dependence of TLS dielectric loss in niobium oxides between 100 mK and 4 K? Can all ancillary junctions in a parity-protected gate scheme be made of high-gap superconductors, or does the scheme intrinsically require aluminum? Does parity readout fidelity at 4 K degrade below the levels needed for error correction? These three questions define the experimental program that would settle the viability conjecture this paper has quantified.

7. Conclusion

We tested the conjecture that a thermodynamic bottleneck at millikelvin temperatures can be relieved by operating parity-protected superconducting qubits at 4 K. With fully shown arithmetic, we found: (i) the thermal scale at 4 K is $k_B T/h = 83.3465$ GHz, giving a niobium gap-edge Boltzmann factor $p_{\text{th}}^{\text{Nb}}(4 \text{ K}) = 1.5006 \times 10^{-2}$ against $\approx 1.12 \times 10^{-73}$ at 100 mK, a penalty of order $10^{71.13}$; (ii) aluminum ($T_{c,\text{Al}} = 1.2$ K) is categorically excluded as a junction material at 4 K; (iii) against a stated target $p_{\text{target}} = 10^{-4}$, an unprotected niobium qubit requires a suppression factor $S = 150$ on thermally activated quasiparticle errors and meets the target only below $T^* \approx 1.82$ K; (iv) the parity-selective coupling mechanism stated in [6] is the right kind of mechanism for this channel, but its required selectivity $S_{\text{par}} < 6.67 \times 10^{-3}/A$ depends on an attempt factor A that the supplied input does not fix; and (v) three channels — Nb_2O_5 two-level-system loss, ancillary-junction poisoning, and 4 K readout thermal noise — remain unquantified. The conclusion is conditional: 4 K operation is not fundamentally excluded for parity-protected niobium architectures, is excluded for conventional aluminum-based qubits, and its viability reduces to measurable quantities (A , S_{par} , TLS loss tangents, readout fidelity) that we have specified as falsification targets.

References

- [1] Quantum Computing: Vision and Challenges. arXiv:2403.02240v5. <https://arxiv.org/abs/2403.02240v5> [2] Tierkreis: A Dataflow Framework for Hybrid Quantum-Classical Computing. arXiv:2211.02350v1. <https://arxiv.org/abs/2211.02350v1> [3] Unveiling Energetic Advantage in Superconducting Cat-Qubits Quantum Computation. arXiv:2605.19854v1. <https://arxiv.org/abs/2605.19854v1> [4] Unifying communication paradigms in measurement-based delegated quantum computing. arXiv:2506.21988v2. <https://arxiv.org/abs/2506.21988v2> [5] Quantum simulation and computing with Rydberg-interacting qubits. arXiv:2011.03031v2. <https://arxiv.org/abs/2011.03031v2> [6] Top-transmon: hybrid superconducting qubit for parity-protected quantum computation. arXiv:1105.0315v1. <https://arxiv.org/abs/1105.0315v1> [7] How many qubits are needed for quantum computational supremacy?. arXiv:1805.05224v3. <https://arxiv.org/abs/1805.05224v3> [8] Quantum Computing in the NISQ era and beyond. arXiv:1801.00862v3. <https://arxiv.org/abs/1801.00862v3> [9] <https://papers.qnfo.org/papers/thermodynamic-constraints-and-architectural-inversions-in-scalable-quantum-information-systems/>. QNFO: Thermodynamic and Quantum Constraints on Scalable Quantum Computing. [10] <https://papers.qnfo.org/papers/thermodynamic-and-informational-bottlenecks-of-scalable-fault-tolerant-quantum-computation/>. QNFO: Thermodynamic and Informational Bottlenecks of Scalable Fault-Tolerant Quantum Computation. [11] <https://papers.qnfo.org/papers/thermodynamic-imperative/>. QNFO: Thermodynamic Imperative. [12] <https://papers.qnfo.org/papers/syntactic-generation-primitive-distinctions/>. QNFO: Syntactic Generation.

Appendix A. Divergence report

Three divergences among the independent drafts were reconciled by explicit convention choice:

- **Reference temperature.** One draft used 20 mK as the dilution-refrigerator reference; the reconciled text adopts $T_{\text{cold}} = 0.1$ K (100 mK), matching the motivating input's quoted $k_B T \approx 83$ GHz comparison point at 4 K. Convention: 100 mK. Effect on results: the cold-side Boltzmann factor changes (at 20 mK, $\Delta_{\text{Nb}}/k_B T = 350/(20.837 \times 0.02) = 839.9$, i.e., $p_{\text{th}} \approx 10^{-364.8}$), but no headline conclusion depends on the cold-side value; only the warm-side numbers ($S = 150$, $T^* \approx 1.82$ K) enter the fault-tolerance argument.
- **Fault-tolerance threshold.** Drafts used $p_{\text{target}} = 10^{-4}$ and 10^{-3} . Convention: 10^{-4} in the main text, treated as a stated design assumption not attributed to any cited work; sensitivity $S = 15$ (10^{-3}) and $S = 1.5$ (10^{-2}) is reported in Steps 6 and Results 4.

- **Millikelvin cooling budget.** Drafts used $P_{\text{mK}} = 1 \mu\text{W}$ and $P_{\text{mK}} = 10 \mu\text{W}$ as illustrative budgets. Convention: $P_{\text{mK}} = 1 \mu\text{W}$, explicitly labeled an assumption; all cooling-power conclusions scale linearly with P_{mK} (Step 8), and the input-supplied ratio $R_P = 2 \times 10^4$ is independent of this choice.

Appendix B. Claim attribution

Claim	Source drafts	Status
C1: $k_B T/h = 83.35$ GHz at 4 K (Steps 1–2)	A, B, C	CONVERGENT
C2: $p_{\text{th}}^{\text{Nb}}(4 \text{ K}) = 1.50 \times 10^{-2}$; $\approx 1.1 \times 10^{-73}$ at 100 mK; penalty $\sim 10^{71.1}$ (Steps 3–4)	A, B, C	CONVERGENT
C3: aluminum excluded at 4 K; $p_{\text{th}}^{\text{Al}}(4 \text{ K}) = 0.590$ under BCS assumption (Step 5)	A, B, C	CONVERGENT
C4: required suppression $S = 150$ against $p_{\text{target}} = 10^{-4}$; $T^* \approx 1.82$ K (Steps 6–7)	A, B, C	CONVERGENT
C5: cooling-power dividend $P_{4\text{K}} = 20$ mW under $P_{\text{mK}} = 1 \mu\text{W}$; Carnot ratio 40.5 (Steps 8–9)	A, B	CONVERGENT
C6: parity-selectivity requirement $S_{\text{par}} < 6.67 \times 10^{-3}/A$ (Step 10)	A, B	CONVERGENT
C7: illustrative relaxation model, gate error 7.5×10^{-4} under assumed Γ_0, Ω (Step 11)	B	SINGLE
C8: reference temperature 100 mK vs 20 mK	A vs B	DIVERGENT (Appendix A)
C9: threshold 10^{-4} vs 10^{-3}	A vs C	DIVERGENT (Appendix A)
C10: $P_{\text{mK}} = 1 \mu\text{W}$ vs $10 \mu\text{W}$	B vs C	DIVERGENT (Appendix A)
C11: unquantified Nb_2O_5 TLS, ancillary poisoning, readout noise channels (Step 12)	A, B, C	CONVERGENT