V1.0.0

Parity-Protected Superconducting Qubits at 4 K: A Thermodynamic Error-Budget and Cooling-Crossover Analysis

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Scalable superconducting quantum processors are conventionally operated at millikelvin temperatures, where dilution-refrigerator cooling power is capped near tens of microwatts while every added qubit, cable, and amplifier increases the heat load on a stage that cannot be expanded cheaply. The specific gap addressed here is quantitative: no published analysis relates the cooling-power scaling of cryogenic plants, the heat dissipated per qubit, and the thermally activated error rates of superconducting qubits into a single crossover criterion that says when error-protecting a qubit beats cooling it. We derive such a criterion from the BCS gap-edge Boltzmann factor $p_{\mathrm{th}}(T) = e^{-\Delta/k_B T}$, the cooling-power law $\kappa(T) = \kappa_0 (T/T_0)^{\alpha}$, and a per-qubit heat-load assumption. For niobium ($T_c = 9.2$ K, $\Delta/k_B = 16.192$ K) we compute $p_{\mathrm{th}}(4\ \mathrm{K}) = 1.7458 \times 10^{-2}$ against 4.78×10^{-71} at 100 mK, a penalty of 3.65×10^{68}; against a target error $p_{\mathrm{target}} = 10^{-4}$ this demands a parity-selective suppression factor $S \approx 174.6$, and the unsuppressed crossover temperature is $T^{*} \approx 1.758$ K. Aluminum junctions ($T_c = 1.2$ K) are excluded outright at 4 K ($p_{\mathrm{th}} = 0.5897$). With the conservative cooling-power ratio 2×10^4 between the 4 K and millikelvin stages and an assumed 10 nW per-qubit load, the permissible qubit budget rises from 5×10^3 to 10^8. The result reframes refrigeration roadmaps: the binding design parameter becomes parity-selective suppression, not base temperature.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). Parity-Protected Superconducting Qubits at 4 K: A Thermodynamic Error-Budget and Cooling-Crossover Analysis. QNFO Preprint v1.0.0.

1. Introduction

The operators building superconducting quantum processors face a scaling wall that is thermodynamic before it is electronic. A dilution refrigerator's mixing chamber, operating near 10 millikelvin, delivers a cooling power capped in the microwatt regime — typically about 50 microwatts for standard commercial systems — while the 4 kelvin pulse-tube stage of the same cryostat offers roughly 20,000 times more, often exceeding 1 watt [8]. Every physical qubit added to the processor carries wiring that bridges the 300 K room-temperature flange to the 10 mK mixing chamber, and because the electrical conductivity of the required coaxial lines correlates with thermal conductivity through the Wiedemann–Franz law, this wiring imports a parasitic heat load that scales linearly with qubit count whether or not the qubits are active [8]. The consequence is a hard ceiling: beyond a few hundred to a few thousand qubits, the heat generated by control and measurement systems exceeds the available cooling capacity, and additional qubits degrade the system through thermal crosstalk rather than improving it [9]. Today this ceiling blocks the path from hundred-qubit demonstrators toward the million-qubit regime that fault-tolerant algorithms require, and the standard response — build larger and colder dilution refrigerators — attacks a stage whose cooling power is fixed by the thermodynamics of helium-3/helium-4 mixing rather than by engineering maturity [8].

If the analysis in this paper holds, the hardware roadmap decision of where to spend engineering effort shifts from pushing base temperature down and refrigerator capacity up, toward building error protection into the qubit itself so that it can operate at the 4 K stage, where cooling budget, wiring tolerance, and controls infrastructure are all orders of magnitude more generous. Large-scale 4 K cryogenic operation is mature industrial practice: Fermilab operates and tests superconducting 1.3 GHz cryomodules on standalone Tevatron-style refrigerators at the New Muon Lab facility [2], and the PIP-II project's cryogenic distribution and controls system specifies industrial-grade instrumentation and control hardware, including a Siemens PCS7 system with 26 remote-IO cabinets and 48 relay racks for temperature, valve, level, and heater readouts [4]. A qubit platform compatible with that stage inherits an existing engineering ecosystem rather than a research-grade one.

The candidate qubit platform for 4 K operation is the parity-protected superconducting qubit. Two families exist in the literature. The first embeds topological Josephson junctions in 0–π qubits: a pair of 0–π qubits built on two-dimensional topological-insulator Josephson junctions forms a parity-spin qubit whose even and odd parity ground states are protected against both charge noise and flux noise, with the topological junction itself suppressing fabrication variation in the Josephson energy [1]. The second couples a conventional transmon to Majorana bound states: the top-transmon uses the superconducting charge qubit as a fermion-parity meter that couples directly to fermion parity and is insensitive to subgap quasiparticle excitations [3]. More recently, a gate-tunable superconductor–altermagnet–superconductor Josephson junction has been proposed to give full gate-voltage control of a parity-protected qubit, removing the magnetic-flux biasing that introduces flux noise, control overhead, and limited scalability in current implementations [5]. All three designs share the property that the protected degree of freedom is fermion parity rather than an energy eigenstate separated from the thermal bath by the superconducting gap alone, which is precisely the property a 4 K architecture must exploit.

The gap this paper fills is the quantitative crossover criterion. Prior work has established the thermodynamic bottleneck qualitatively [8], stated the cooling-power scaling law with exponent α between 2 and 3 [9], and performed a first reconciled error-budget calculation for niobium-based parity qubits at 4 K [6]. What is missing is a single derivation that combines the gap-edge thermal excitation probability, the required parity-selective suppression, the cooling-power ratio between stages, and a per-qubit heat load into an explicit crossover temperature and qubit-budget comparison, so that an architect can decide numerically — not rhetorically — whether protection or cooling is the better buy at a given temperature. The contributions are:

- A derivation of the gap-edge thermal excitation probabilities for niobium and aluminum at 4 K and at 100 mK, with fully shown arithmetic, giving $p_{\text{th}}^{\text{Nb}}(4\ \text{K}) = 1.7458 \times 10^{-2}$, $p_{\text{th}}^{\text{Nb}}(100\ \text{mK}) = 4.78 \times 10^{-71}$, and $p_{\text{th}}^{\text{Al}}(4\ \text{K}) = 0.5897$ (Section 4, Steps 3–6).
- A derived requirement on parity-selective suppression: against a target physical error $p_{\text{target}} = 10^{-4}$, an unprotected niobium qubit at 4 K requires $S = 174.6$, and the unsuppressed crossover temperature is $T^{*} = 1.758$ K (Section 4, Steps 7–8).
- A derived qubit-budget comparison under two per-qubit heat-load models: with the conservative cooling-power ratio 2×10^4 from the literature and an assumed 10 nW per-qubit load, the permissible qubit count rises from 5×10^3 at the millikelvin stage to 10^8 at the 4 K stage; under an alternative gate-energy model the counts are 3.62×10^{11} and 1.81×10^{14} (Section 4, Steps 9–10). The thermal-photon population of a 5 GHz readout line rises from 0.0998 to 16.17 photons between the two stages (Step 11).
- An identification of the three error channels that remain unquantified — Nb_2O_5 two-level-system loss, quasiparticle poisoning of unprotected ancillary elements, and thermal noise in 4 K parity readout — carried over from the prior analysis and restated as open conditions on the crossover claim (Sections 4 and 6).

The crossover criterion is stated in terms adjacent fields already use: the same gap-edge Boltzmann factor governs quasiparticle poisoning in superconducting detectors, and the same cooling-power scaling governs any cryogenic detector array, so the numerical criterion may transfer wherever superconducting devices meet limited cryogenic budgets.

2. Background and Related Work

The thermodynamic bottleneck argument has two independent lines of support in the literature, a qubit-design line, and a prior quantitative analysis that this paper extends and partially revises.

On the thermodynamic side, the scaling-asymmetry analysis [8] quantifies the mismatch that motivates this paper: the mixing chamber of a dilution refrigerator, operating at approximately 10 millikelvin, is physically constrained to the microwatt regime, typically capping at 50 microwatts for standard commercial systems, while the 4 K pulse-tube stage offers a cooling budget roughly 20,000 times larger, often exceeding 1 watt [8]. The same work identifies the wiring bottleneck as the tangible mechanism: every qubit requires a signal path bridging the 300 K flange to the 10 mK mixing chamber, the cable material's electrical and thermal conductivities are linked by the Wiedemann–Franz law, and the resulting parasitic heat load scales linearly with qubit count regardless of activity [8]. It further grounds the entropy side of the budget in the Landauer limit, under which logically irreversible error-correction cycles must dump entropy into the thermal bath, and it reports a validation of its numerical model against 2024 peer-reviewed coherence data from niobium trilayer and tantalum qubits, calibrating a microscopic TLS coupling constant and an intrinsic $1/f$ noise floor [8]. The cooling-capacity scaling law is stated independently in the Thermodynamic Imperative analysis [9], which writes $\kappa(T) = \kappa_0 (T/T_0)^{\alpha}$ with α typically between 2 and 3, and derives a scalability condition $\lim_{N\to\infty} Q_{\text{Total}}(N,T)/\kappa(T) < 1$ whose violation means that even with perfect qubit coherence and optimal error-correcting codes, heat from control electronics and measurement would exceed available cooling capacity beyond a few hundred qubits [9]. The same work notes a materials route — moiré superlattices in twisted bilayer graphene or TMDs that could host topological superconducting states with critical temperatures exceeding 4 K — while flagging that twist-angle control at the required $\pm 0.05^{\circ}$ precision across wafer scales remains an unsolved fabrication bottleneck [9]. This paper uses the empirical 20,000-fold ratio of [8] as its conservative baseline and treats the $\alpha \in [2,3]$ law of [9] as an upper-range sensitivity case; the choice is stated explicitly in Section 4.

On the qubit-design side, the parity-protected superconducting qubit based on topological-insulator Josephson junctions [1] proposes two 0–π qubits built on 2D topological-insulator Josephson junctions, arguing that the topological junction provides protection against fabrication variations and thus ensures the identical Josephson junctions the 0–π qubit requires; its model section treats the two lowest states of a single 0–π qubit as a spin-$\frac{1}{2}$ degree of freedom and combines two such qubits shunted by a capacitor into a parity-spin qubit designed to improve charge noise and flux noise simultaneously, with a second-order sweet-spot expansion of the Hamiltonian in the noise parameters $\lambda \in \{n_g, \phi_e\}$ [1]. The same work reports Josephson-to-charging-energy ratios of $E_J^{(2)}/E_c \approx 40$ in the protected regime and $E_J^{(2)}/E_c = 9$ for a symmetric configuration, indicating that the design can be tuned to enhance protection. The top-transmon proposal [3] couples a superconducting charge qubit to Majorana bound states so that the qubit acts as a fermion-parity meter insensitive to subgap excitations. The gate-voltage-control proposal [5] removes the flux-bias overhead of current parity-protected implementations by using a superconductor–altermagnet–superconductor Josephson junction, deriving the Andreev levels of the junction from a BdG Hamiltonian in a scattering-matrix treatment and computing the Fourier coefficients $E_{J,1}$, $E_{J,2}$ of the Josephson potential numerically with a KWANT tight-binding model at parameters $d = 50a$, $W = 7a$, $\Delta = 0.007t$, $L = 325a$, with an InSb-type channel ($m^* = 0.024\, m_e$, $n = 2.7 \times 10^{11}$ cm^{-2}, $\mu \approx 23$ meV) and an aluminum superconductor with $\Delta = 0.18$ meV and $\xi \approx 1000$ nm [5]. The aluminum gap value stated in [5] matters directly for this paper: Step 6 below shows that aluminum is thermodynamically excluded as a junction material at 4 K, so the altermagnetic control mechanism must be re-hosted on a higher-T_c superconductor to serve a 4 K architecture.

On the prior-analysis side, the reconciled thermodynamic error-budget paper [6] is the direct predecessor. It computed $k_B T/h = 83.3465$ GHz at 4 K and 2.08366 GHz at 100 mK, found $p_{\mathrm{th}}^{\mathrm{Nb}}(4\,\mathrm{K}) = 1.5006 \times 10^{-2}$ against approximately 1.12×10^{-73} at 100 mK with an operating-point penalty of order $10^{71.13}$, excluded aluminum at 4 K with $p_{\mathrm{th}}^{\mathrm{Al}}(4\,\mathrm{K}) = 0.5898$, and derived a required suppression factor $S = 150$ on one identified channel against a 10^{-4} target, meeting the target only below $T^{*} \approx 1.82$ K [6]. Its discussion states explicitly what its analysis does not show: three error channels remain unquantified — thermally activated TLS loss in $\mathrm{Nb_2O_5}$, 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 [6]. This paper recomputes the niobium thermal factor from the standard BCS gap convention $\Delta = 1.76\,k_B T_c$, obtaining 1.7458×10^{-2} rather than the 1.5006×10^{-2} of [6] — a factor of 1.16 that does not change any conclusion — and extends the analysis with the cooling-crossover and qubit-budget calculations that [6] does not perform.

Finally, the survey-correction paper [7] contributes a caution about how platform claims propagate into decisions: it demonstrates that a factual error in survey literature — there, the premature announcement that non-Abelian anyons were experimentally confirmed in 2020 in fractional quantum Hall systems, where the confirmed $\nu = 1/3$ observations are Abelian with exchange acting by a scalar phase of unit modulus — propagates into real funding and research decisions, and it complements its corrections with derived figures such as an anyon-sector error-correction threshold $p_c \geq 0.037$, a factor of 37 above an assumed operating error rate of 10^{-3} [7]. The lesson applied here is that the 4 K architecture claim must carry its numbers and its unmet conditions openly, as [7] does and as [6] does, rather than as a qualitative posture.

The cryogenic-engineering works [2] and [4] establish what the 4 K stage's industrial ecosystem looks like. The New Muon Lab facility at Fermilab is a superconducting 1.3 GHz cryomodule test facility in which a single Type III+ cryomodule is cooled and tested using a single Tevatron-style standalone refrigerator [2] — evidence that multi-watt-class cryogenic plants with the thermodynamic headroom [8] quantifies are operated routinely. The PIP-II cryogenic controls design specifies a Siemens PCS7 controls system with 26 remote-IO Rittal cabinets and 48 relay racks for temperature readouts, valve positioners, level, and heater control across the cryogenic distribution system and cryomodules [4] — evidence that the instrumentation and control layer for large 4 K-class cryogenic infrastructure is standardized industrial technology, not a research frontier.

3. Methods

This paper is an analytical derivation with fully shown arithmetic. It performs no experiment, no device simulation, and no measurement; every number in Section 5 is computed in Section 4 from stated inputs, or is labeled a projection with stated assumptions.

Model. The physical error channel that temperature controls directly is thermally activated quasiparticle excitation across the superconducting gap. Under the BCS assumption, the probability scale for a quasiparticle to be thermally excited is the gap-edge Boltzmann factor

$$p_{\mathrm{th}}(T) = \exp\!\left(-\frac{\Delta}{k_B T}\right),$$

where Δ is the superconducting gap and k_B is Boltzmann's constant. We take the BCS gap convention $\Delta = 1.76\,k_B T_c$, so that $\Delta/k_B = 1.76\,T_c$ in kelvin. This is the same channel structure the prior analysis used [6]; the only convention difference is the gap prefactor, and where the two conventions give different numbers we state both and use the BCS value (Step 3 note).

The parity-protection mechanism enters as a single phenomenological suppression factor S_{par}: the ratio of the unprotected thermally activated error probability to the error probability of the protected qubit on the same channel. We define S as this ratio, so that the protected error probability is p_{th}/S; this is the inverse of the residual-fraction convention in which the prior analysis also reports $S = 150$ [6], and the two conventions give reciprocal numbers for the same quantity. The literature supports the existence of such suppression in two forms: the top-transmon's parity meter is insensitive to subgap excitations, [3], and the topological-insulator parity qubit suppresses charge and flux noise simultaneously at second order at its sweet spots [1]. No source fixes a numerical value of S_{par} for a realized device, so we treat S_{par} as a design requirement to be met rather than a measured quantity, and we state an illustrative value only as an assumption.

The thermodynamic side uses the cooling-power ratio between stages. Two characterizations exist in the literature: the empirical ratio of roughly 2×10^4 between the 4 K pulse-tube stage (exceeding 1 W) and the 10 mK mixing chamber (capped near 50 μW) [8], and the power law $\kappa(T) = \kappa_0 (T/T_0)^{\alpha}$ with $\alpha \in [2, 3]$ [9]. We take the empirical ratio as the conservative baseline; the divergence between the two characterizations and its consequence are stated in Section 4.

Assumptions. We make the following assumptions, each labeled where used: (A1) the BCS gap convention $\Delta = 1.76\, k_B T_c$; (A2) a target physical error probability per gate operation $p_{\text{target}} = 10^{-4}$, the same target the prior analysis used [6]; (A3) a per-qubit heat load $q = 10$ nW at both stages, a stand-in value no source provides, used only for the qubit-budget projection; (A4) a readout frequency $f = 5$ GHz for the thermal-photon calculation; (A5) that the parity-protected qubit's protected channel is the thermally activated quasiparticle channel, with the three channels the prior analysis left unquantified — Nb_2O_5 TLS loss, ancilla quasiparticle poisoning, and 4 K readout thermal noise — carried as open conditions rather than folded into the budget [6]. For the alternative load model we additionally assume (A6) a gate rate $f_{\text{gate}} = 10^8$ Hz and a gate energy $E_{\text{gate}} = k_B T$ per operation.

4. Analysis

Every number in this section is computed from the stated inputs with fully shown arithmetic. The inputs are: $T_c^{\text{Nb}} = 9.2$ K (standard niobium critical temperature, used with the BCS convention of assumption A1), $T_c^{\text{Al}} = 1.2$ K (the aluminum critical temperature; the altermagnetic-junction work [5] uses aluminum with $\Delta = 0.18$ meV, consistent with this value), $T_{4\text{K}} = 4$ K, $T_{\text{mK}} = 0.1$ K, $p_{\text{target}} = 10^{-4}$ (assumption A2), $\kappa_{\text{mK}} = 50\ \mu$W and the stage ratio 2×10^4 from [8], $q = 10$ nW (assumption A3), $f_{\text{gate}} = 10^8$ Hz and $E_{\text{gate}} = k_B T$ (assumption A6), and $f = 5$ GHz (assumption A4). Physical constants: $k_B = 1.380649 \times 10^{-23}$ J/K, $h = 6.62607015 \times 10^{-34}$ J s, $\ln 10 = 2.302585$, $\log_{10} e = 0.434294$.

Step 1: Thermal energy scale in frequency units. With the BCS convention (A1), $\Delta/k_B = 1.76\, T_c$.

$$\frac{k_B T}{h}\Big|_{4\,\text{K}} = \frac{1.380649 \times 10^{-23} \times 4}{6.62607015 \times 10^{-34}} = \frac{5.522596 \times 10^{-23}}{6.62607015 \times 10^{-34}} = 8.33465 \times 10^{10}\ \text{Hz} = 83.3465\ \text{GHz}.$$

$$\frac{k_B T}{h}\Big|_{0.1\,\text{K}} = \frac{1.380649 \times 10^{-24}}{6.62607015 \times 10^{-34}} = 2.08366\ \text{GHz}.$$

Step 2: Niobium gap in temperature units.

$$\frac{\Delta_{\text{Nb}}}{k_B} = 1.76 \times 9.2 = 16.192\ \text{K}.$$

Step 3: Niobium gap-edge factor at 4 K.

$$p_{\text{th}}^{\text{Nb}}(4\ \text{K}) = e^{-16.192/4} = e^{-4.048}.$$

Computing $e^{-4.048}$: $e^{-4} = 0.0183156$, $e^{-0.048} = 1 - 0.048 + 0.048^2/2 - \cdots = 0.953134$. Product: $0.0183156 \times 0.953134 = 0.0174580$. So

$$p_{\text{th}}^{\text{Nb}}(4\ \text{K}) = 1.7458 \times 10^{-2}.$$

Convention note. The prior analysis [6] used a gap convention that gave $p_{\text{th}}^{\text{Nb}}(4\ \text{K}) = 1.5006 \times 10^{-2}$; this paper uses the standard BCS prefactor $\Delta = 1.76\, k_B T_c$, giving 1.7458×10^{-2}, a factor of $1.7458/1.5006 = 1.1634$ larger. We take the BCS convention because it is the standard relation for weak-coupling superconductors; the factor of 1.16 changes the required suppression factor from 150 to 174.6 and no conclusion.

Step 4: Niobium gap-edge factor at 100 mK.

$$p_{\text{th}}^{\text{Nb}}(0.1\ \text{K}) = e^{-16.192/0.1} = e^{-161.92}.$$

In base 10: $161.92 \times 0.434294 = 70.3209$. So $p_{\text{th}}^{\text{Nb}}(0.1\ \text{K}) = 10^{-70.3209} = 10^{-71} \times 10^{0.6791} = 4.776 \times 10^{-71}$, i.e.

$$p_{\text{th}}^{\text{Nb}}(0.1\ \text{K}) \approx 4.78 \times 10^{-71}.$$

Step 5: Operating-point penalty.

$$\frac{p_{\text{th}}^{\text{Nb}}(4\ \text{K})}{p_{\text{th}}^{\text{Nb}}(0.1\ \text{K})} = \frac{1.7458 \times 10^{-2}}{4.776 \times 10^{-71}} = 3.655 \times 10^{68}.$$

Moving from the millikelvin stage to 4 K costs a factor of 3.65×10^{68} on the thermally activated quasiparticle channel.

Step 6: Aluminum exclusion. With $\Delta_{\text{Al}}/k_B = 1.76 \times 1.2 = 2.112$ K,

$$p_{\text{th}}^{\text{Al}}(4\ \text{K}) = e^{-2.112/4} = e^{-0.528} = 0.5897.$$

($e^{-0.5} = 0.606531$, $e^{-0.028} = 0.972388$; product 0.58975.) At 4 K, 58.97% of the gap-edge Boltzmann weight is thermally populated: aluminum junctions are excluded outright as a 4 K junction material, consistent with the categorical exclusion of [6]. This bears directly on the altermagnetic control proposal [5], whose junction superconductor is aluminum with $\Delta = 0.18$ meV: the control mechanism must be re-hosted on a higher-T_c superconductor to serve a 4 K architecture.

Step 7: Required parity-selective suppression. The protected error probability is p_{th}/S (Section 3). Setting $p_{\text{th}}/S = p_{\text{target}}$:

$$S = \frac{p_{\text{th}}^{\text{Nb}}(4\ \text{K})}{p_{\text{target}}} = \frac{1.7458 \times 10^{-2}}{10^{-4}} = 174.58 \approx 174.6.$$

A parity-selective mechanism must suppress the thermally activated quasiparticle channel by a factor of about 175 for a niobium qubit at 4 K to meet a 10^{-4} physical error target. The top-transmon's parity meter couples directly to fermion parity and is insensitive to subgap excitations [3], and the parity-spin qubit's second-order sweet-spot protection [1] are the candidate mechanisms; no source fixes a realized numerical value, so $S_{\text{par}} \geq 174.6$ is a design requirement, not a measurement.

Step 8: Unsuppressed crossover temperature. Without parity protection, the qubit meets the target only when $p_{\text{th}}(T^*) = p_{\text{target}}$:

$$\frac{\Delta_{\text{Nb}}}{k_B T^*} = \ln\left(\frac{1}{p_{\text{target}}}\right) = \ln(10^4) = 9.21034,$$

$$T^* = \frac{16.192}{9.21034} = 1.7580\ \mathrm{K}.$$

An unprotected niobium qubit reaches $p_{\mathrm{err}} = 10^{-4}$ from this channel only below $T^* \approx 1.76$ K — a temperature between the two stages, reachable by neither a pulse tube nor cheaply by a dilution unit, which is precisely why suppression rather than cooling is the lever.

Step 9: Qubit-budget comparison, fixed-load model. With the empirical stage ratio of [8], $\kappa_{4\mathrm{K}} = 50\ \mu\mathrm{W} \times 2 \times 10^4 = 1$ W. With the per-qubit load $q = 10$ nW (A3):

$$N_{\mathrm{mK}} = \frac{\kappa_{\mathrm{mK}}}{q} = \frac{50 \times 10^{-6}\ \mathrm{W}}{10 \times 10^{-9}\ \mathrm{W}} = 5 \times 10^3, \qquad N_{4\mathrm{K}} = \frac{1\ \mathrm{W}}{10^{-8}\ \mathrm{W}} = 10^8.$$

The permissible qubit budget rises by the stage ratio, 2×10^4, from 5×10^3 to 10^8.

Step 10: Qubit-budget comparison, gate-energy model. With $E_{\mathrm{gate}} = k_B T$ and $f_{\mathrm{gate}} = 10^8$ Hz (A6), the per-qubit power is $P = k_B T f_{\mathrm{gate}}$:

$$P_{4\mathrm{K}} = 1.380649 \times 10^{-23} \times 4 \times 10^8 = 5.5226 \times 10^{-15}\ \mathrm{W}, \qquad N_{4\mathrm{K}} = \frac{1}{5.5226 \times 10^{-15}} = 1.81 \times 10^{14}.$$

$$P_{\mathrm{mK}} = 1.380649 \times 10^{-23} \times 0.1 \times 10^8 = 1.3806 \times 10^{-16}\ \mathrm{W}, \qquad N_{\mathrm{mK}} = \frac{50 \times 10^{-6}}{1.3806 \times 10^{-16}} = 3.62 \times 10^{11}.$$

Both counts are far above any realistic wiring or control limit; the gate-energy model shows that the fixed-load model (Step 9) is the conservative one, and the crossover conclusion is insensitive to which load model is used.

Step 11: Thermal photon population of a readout line. The Bose–Einstein occupation of a mode at frequency f is $n = 1/(e^{hf/k_BT} - 1)$. With $hf/k_B = 4.79924 \times 10^{-11}$ K s and $f = 5$ GHz (A4), $hf/k_B = 0.239962$ K.

At 4 K: $x = 0.239962/4 = 0.0599905$; $e^x = 1.061825$; $n = 1/0.061825 = 16.17$.

At 100 mK: $x = 0.239962/0.1 = 2.39962$; $e^x = 11.021$; $n = 1/10.021 = 0.0998$.

A 5 GHz readout line at 4 K carries 16.17 thermal photons per mode against 0.0998 at the millikelvin stage — a factor of 162 — so 4 K readout requires attenuation and filtering, or a higher-frequency readout scheme, as an open engineering condition.

Step 12: Cooling-power-law sensitivity. The power law of [9], $\kappa(T) = \kappa_0 (T/T_0)^\alpha$ with $\alpha \in [2, 3]$, gives a stage ratio of $(4/0.01)^\alpha = 400^\alpha$: for $\alpha = 2$, 1.6×10^5; for $\alpha = 3$, 6.4×10^7. Both exceed the empirical 2×10^4 of [8] by factors of 8 to 3200, so the empirical ratio is the conservative choice, and using it can only understate the 4 K advantage. We take the empirical ratio as the baseline under this stated convention.

5. Results

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

1. **Thermal scales.** $k_BT/h = 83.3465$ GHz at 4 K and 2.08366 GHz at 100 mK (computed, Steps 1).
2. **Niobium survives but is only moderately protected.** $p_{\mathrm{th}}^{\mathrm{Nb}}(4\ \mathrm{K}) = 1.7458 \times 10^{-2}$ versus 4.78×10^{-71} at 100 mK; the operating-point penalty is 3.65×10^{68} (computed, Steps 2–5). Under the gap convention of [6] the 4 K value is 1.5006×10^{-2}; the BCS convention is used here (Step 3 note).
3. **Aluminum is excluded.** $p_{\mathrm{th}}^{\mathrm{Al}}(4\ \mathrm{K}) = 0.5897$: aluminum junctions are categorically excluded at 4 K, and the alter-magnetic control mechanism of [5], demonstrated in its numerical model on an aluminum superconductor with $\Delta = 0.18$ meV, must be re-hosted on a higher-T_c superconductor (computed, Step 6).

4. **Required suppression.** Against $p_{\text{target}} = 10^{-4}$, a niobium qubit at 4 K requires a parity-selective suppression $S = 174.6$ on the thermally activated quasiparticle channel (computed, Step 7). This is a design requirement; no source reports a realized device value.
5. **Unsuppressed crossover temperature.** $T^* = 1.758$ K: an unprotected niobium qubit meets the target only below 1.76 K, between the two stages (computed, Step 8).
6. **Qubit budget, fixed-load projection.** Under assumptions A3 and the empirical stage ratio of [8], the permissible qubit count rises from 5×10^3 at the millikelvin stage to 10^8 at the 4 K stage (projection on assumed $q = 10$ nW, Step 9).
7. **Qubit budget, gate-energy projection.** Under assumption A6, the counts are 3.62×10^{11} (mK) and 1.81×10^{14} (4 K); the fixed-load model is the conservative bound (projection, Step 10).
8. **Readout thermal photons.** A 5 GHz mode carries 16.17 thermal photons at 4 K versus 0.0998 at 100 mK, a factor of 162 (computed under assumption A4, Step 11).
9. **Cooling-law sensitivity.** The power law of [9] gives stage ratios 1.6×10^5 ($\alpha = 2$) to 6.4×10^7 ($\alpha = 3$), all above the empirical 2×10^4 baseline (computed, Step 12).

6. Discussion

What the result does and does not show. The analysis establishes a quantified asymmetry: conventional aluminum-junction qubits are categorically excluded at 4 K, while niobium-based parity-protected qubits face a finite, numerically stated requirement — suppression $S \geq 174.6$ on one identified channel — that the mechanisms described in [1] and [3] are designed in principle to provide. It does not demonstrate that a full 4 K fault-tolerant processor is achievable, because three error channels remain unquantified, carried over from the prior analysis [6]: thermally activated two-level-system loss in Nb_2O_5, quasiparticle poisoning of unprotected ancillary elements required for gate operations, and thermal noise in parity readout at 4 K, where the 16.17-photon population of a 5 GHz line (Step 11) is non-negligible for any readout chain not heavily attenuated and filtered. The crossover claim holds only for the one channel the temperature controls directly.

Falsification. The central claim — that parity suppression is a smaller engineering burden than millikelvin cooling at the 4 K operating point — would be falsified by any of the following observations: a realized parity-protected device whose measured suppression on the thermally activated quasiparticle channel falls below 174.6 (or whose suppression mechanism is shown to saturate below that value as E_J/E_C or junction quality is increased); a measurement showing that one of the three unquantified channels carries an error probability at 4 K that exceeds 10^{-4} by a factor no realistic suppression can close; or a cryogenic plant whose millikelvin cooling power scales fast enough with investment that the qubit-budget advantage of Step 9 inverts. The first is the sharpest test, because $S = 174.6$ is a modest, directly measurable number.

Limitations. The result rests on the following assumptions, with their sensitivity:

- (A1) The BCS gap convention $\Delta = 1.76\, k_B T_c$. Strong-coupling corrections for niobium would raise Δ and lower p_{th}; the prior convention of [6] lowers the 4 K factor by 16% and raises S by the same factor. The conclusion is insensitive at the order-of-magnitude level.
- (A2) $p_{\text{target}} = 10^{-4}$. S scales linearly in $1/p_{\text{target}}$: a 10^{-5} target demands $S = 1746$, still modest; a 10^{-2} surface-code-style target demands only $S = 1.75$.
- (A3) $q = 10$ nW per qubit is a stand-in with no source. The qubit budget scales as $1/q$; a 1 μW load lowers $N_{4\text{K}}$ to 10^6, still 200 times the millikelvin count. The stage ratio, not the load, sets the advantage.
- (A5) Only the thermally activated quasiparticle channel is budgeted. The three unquantified channels could dominate; the crossover claim is conditional on their eventual suppression.

- The cooling-power baseline is the empirical ratio of [8]; the power law of [9] gives larger ratios (Step 12), so this choice is conservative.
- S_{par} is treated as a design requirement, not a measured quantity; no source fixes it for a realized device.

Open questions. What suppression does the topological-insulator parity-spin qubit [1] achieve on the quasiparticle channel specifically, as opposed to charge and flux noise? Can the altermagnetic gate-voltage control of [5] be implemented with a niobium-class superconductor, given that its numerical model uses aluminum? Does the $\pm 0.05^\circ$ twist-angle bottleneck that [9] flags for moiré topological superconductors relax enough for those materials to serve as intrinsic 4 K hosts? And what readout frequency and filtering chain reduces the 4 K thermal photon population to a tolerable level without importing the wiring heat load that [8] identifies?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are quantum-processor architects and cryogenic-systems engineers making roadmap decisions about where control, readout, and eventually qubits sit in a cryostat. If the crossover criterion holds, the binding design parameter in their architecture documents changes from base temperature to parity-selective suppression: a design review would specify a required $S_{\text{par}} \geq 174.6$ (for $p_{\text{target}} = 10^{-4}$) as a device-level acceptance number, in place of a specification that qubits operate below some millikelvin setpoint. Facilities that already operate multi-watt 4 K plants — the cryomodule test facility at Fermilab's New Muon Lab, where a single Type III+ 1.3 GHz cryomodule is cooled and tested on a single Tevatron-style standalone refrigerator [2], and the PIP-II cryogenic distribution system with its Siemens PCS7 controls spanning 26 remote-IO cabinets and 48 relay racks [4] — are the existing infrastructure a 4 K qubit platform would inherit. The result is also foundational for adjacent problems: any superconducting detector array whose dark rate is set by the same gap-edge Boltzmann factor faces the same crossover between cooling and suppression, and the cooling-power scaling law of [9] applies to any cryogenic instrument whose heat load grows with channel count.

(b) Current readiness. This work is analytical and not peer reviewed; against the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and formulated, analytical treatment without experimental validation). The qubit designs it draws on are themselves proposals or early-stage theory in the cited literature: the parity-spin qubit of [1] is a circuit-model proposal with numerical coherence analysis, the top-transmon of [3] is a proposal whose conclusion states the parity-meter property, and the altermagnetic junction of [5] is demonstrated through a tight-binding numerical model. No claim of device, manufacturing, or deployment readiness is made or supportable.

(c) Smallest next test. The smallest de-risking measurement is a single-device determination of S_{par}: fabricate one top-transmon-class device [3], hold it at 4 K, and measure the ratio of thermally activated quasiparticle-induced error events with the parity-meter coupling on versus off, using an on/off coupling modulation as a design assumption. The outcome that refutes the crossover claim is a measured suppression $S_{\text{par}} < 174.6$ at the maximum achievable E_J/E_C, or a suppression that saturates below that value as E_J/E_C is increased. A companion calculation — a master-equation estimate of the Nb_2O_5 TLS channel at 4 K using literature loss tangents — would close or confirm the first of the three open channels.

(d) Binding constraints. In the order they likely bind: first, materials — the aluminum exclusion (Step 6) forces every junction, including the altermagnetic control junction of [5], onto a higher-T_c superconductor, and niobium junction fabrication with the uniformity the parity-spin qubit requires is itself a fabrication question [1]; second, readout — the 16.17 thermal photons per 5 GHz mode at 4 K (Step 11) demand an attenuated, filtered, or higher-frequency readout chain whose added components must fit inside the per-qubit heat budget of assumption A3; third, the three unquantified error channels of Section 6, which no current source bounds; fourth, wiring, which the Wiedemann–Franz analysis of [8] shows scales linearly with qubit count at any stage, though the 4 K stage tolerates 2×10^4 times more of it. Regulatory or cost constraints are not quantifiable from the sources and remain open questions.

(e) Order-of-magnitude bounds (projections). Under the stated assumptions A3 and the empirical stage ratio of [8], the 4 K stage supports 10^8 qubits at 10 nW each against 5×10^3 at the millikelvin stage; under the gate-energy model A6 the counts are 1.81×10^{14} and 3.62×10^{11}. These are projections, not measurements: they assume the per-qubit loads stated, ignore wiring and readout overhead, and bound only the thermodynamic term. The suppression requirement itself, $S = 174.6$, is a computed number, not a projection.

(f) Risks and misuse paths. The result could be misused to justify skipping millikelvin investment before the three unquantified channels are closed, on the strength of a one-channel budget; the survey-correction lesson of [7] — that a qualitative platform claim propagating into funding decisions without its unmet conditions stated openly distorts allocation — applies directly. Any roadmap document citing this analysis should carry the three open channels and the unmeasured status of S_{par} alongside the crossover numbers, as this paper does.

8. Conclusion

We derived a thermodynamic error-budget and cooling-crossover criterion for operating parity-protected superconducting qubits at 4 K, with fully shown arithmetic. The niobium gap-edge Boltzmann factor is $p_{\mathrm{th}}^{\mathrm{Nb}}(4\ \mathrm{K}) = 1.7458 \times 10^{-2}$ against 4.78×10^{-71} at 100 mK, an operating-point penalty of 3.65×10^{68}; aluminum is excluded outright at 4 K with $p_{\mathrm{th}}^{\mathrm{Al}} = 0.5897$. Against a 10^{-4} target, the parity-selective suppression requirement is $S = 174.6$, and the unsuppressed crossover temperature is $T^* = 1.758$ K — between the stages, reachable by neither cheaply. Under the conservative empirical cooling-power ratio of 2×10^4 [8] and an assumed 10 nW per-qubit load, the permissible qubit budget rises from 5×10^3 to 10^8; the cooling-power law of [9] would only widen this gap. The binding design parameter for a 4 K superconducting architecture is therefore parity-selective suppression, not base temperature — a requirement that is modest, numerically stated, and directly testable on a single device, while three error channels (Nb_2O_5 TLS loss, ancilla quasiparticle poisoning, and 4 K readout thermal noise, where a 5 GHz line carries 16.17 thermal photons) remain as open conditions on the full architecture claim.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.0). 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, openai/gpt-oss-120b, moonshotai/kimi-k2.6, 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: the execution sandbox was unavailable, so a second language model recomputed 5 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 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] Parity-protected superconducting qubit based on topological insulators. arXiv:2308.12027v1. https://arxiv.org/abs/2308.12027v1 [2] Design and testing of the New Muon Lab cryogenic system at Fermilab. arXiv:1206.6796v1. https://arxiv.org/abs/1206.6796v1 [3] Top-transmon: hybrid superconducting qubit for parity-protected quantum computation. arXiv:1105.0315v1. https://arxiv.org/abs/1105.0315v1 [4] Fermilab PIP-II CDS and CM Cryogenic Controls System. arXiv:2509.15369v1. https://arxiv.org/abs/2509.15369v1 [5] Full Gate-Voltage Control of a Parity-Protected Superconducting Qubit with an Altermagnetic Josephson Junction. arXiv:2607.04097v1. https://arxiv.org/abs/2607.04097v1 [6] https://papers.qnfo.org/papers/parity-protected-superconducting-qubits-at-4-k-a-reconciled-thermodynamic-error-. QNFO: Parity-Protected Superconducting Qubits at 4 K: A Reconciled Thermodynamic Error-Budget Analysis. [7] https://papers.qnfo.org/papers/paper-beyond-the-qubit. QNFO: Beyond the Qubit: Constructive Paradigms for Post-Particle Computation. [8] https://papers.qnfo.org/papers/thermodynamic-constraints-and-architectural-inversions-in-scalable-quantum-information-systems. QNFO: Thermodynamic and Quantum Constraints on Scalable Quantum Computing. [9] https://papers.qnfo.org/papers/thermodynamic-imperative. QNFO: Thermodynamic Imperative.