

V1.0.0

Thermodynamic Scaling Limit of Active Quantum Error Correction

Rowan Brad Quni-Gudzinis

QNFO-ULA

ABSTRACT

Active quantum error correction (QEC) underpins fault-tolerant quantum computation, yet its control circuitry dissipates heat that must be removed by cryogenic infrastructure. A hard scalability limit arises when the localized power density generated by control pulses exceeds the Kapitza thermal conductance between the qubit substrate and its coolant, causing a runaway temperature rise that destroys coherence. We identify this limit by extending the Caldeira–Leggett model to include a dissipative control term $P_{\text{local}} = N_q \nu_{\text{cycle}} \varepsilon_q + I^2 R$ and the substrate temperature response $\Delta T_{\text{local}} = P_{\text{local}} / (A H_K)$. Using experimentally motivated parameters ($\varepsilon_q = 10^{-9}$ J per qubit per cycle, $\tau_c = 10^{-6}$ s, qubit pitch $a = 10^{-4}$ m, and a representative Kapitza conductance $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$), we derive a critical qubit density

$$\rho_{\text{crit}} = \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}} \varepsilon_q} \approx 3 \times 10^6 \text{ qubits m}^{-2}$$

for a temperature-rise threshold $\Delta T_{\text{crit}} = 0.3$ K at a base temperature of 0.3 K. Above this density the substrate cannot evacuate heat fast enough, leading to a thermal phase transition that overwhelms decoherence suppression. The result predicts a thermodynamic ceiling for gate-based QEC architectures and motivates a shift toward passive, topologically protected substrates that avoid active control heating.

CITE AS Quni-Gudzinis, R. B. (n.d.). Thermodynamic Scaling Limit of Active Quantum Error Correction. QNFO Preprint v1.0.0.

Abstract

Active quantum error correction (QEC) underpins fault-tolerant quantum computation, yet its control circuitry dissipates heat that must be removed by cryogenic infrastructure. A hard scalability limit arises when the localized power density generated by control pulses exceeds the Kapitza thermal conductance between the qubit substrate and its coolant, causing a runaway temperature rise that destroys coherence. We identify this limit by extending the Caldeira–Leggett model to include a dissipative control term $P_{\text{local}} = N_q \nu_{\text{cycle}} \varepsilon_q + I^2 R$ and the substrate temperature response $\Delta T_{\text{local}} = P_{\text{local}} / (A H_K)$. Using experimentally motivated parameters ($\varepsilon_q = 10^{-9}$ J per qubit per cycle, $\tau_c = 10^{-6}$ s, qubit pitch $a = 10^{-4}$ m, and a representative Kapitza conductance $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$), we derive a critical qubit density

$$\rho_{\text{crit}} = \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}} \varepsilon_q} \approx 3 \times 10^6 \text{ qubits m}^{-2}$$

for a temperature-rise threshold $\Delta T_{\text{crit}} = 0.3$ K at a base temperature of 0.3 K. Above this density the substrate cannot evacuate heat fast enough, leading to a thermal phase transition that overwhelms decoherence suppression. The result predicts a thermodynamic ceiling for gate-based QEC architectures and motivates a shift toward passive, topologically protected substrates that avoid active control heating.

1. Introduction

We consider superconducting processor architectures targeting the million-qubit regime, where heat generated during active error-correction cycles may block further scaling. Each syndrome extraction round consumes energy in the form of microwave control pulses; the cumulative power load grows linearly with the number of qubits N_q and with the cycle frequency ν_{cycle} . Current dilution refrigerators provide limited cooling power at the mixing-chamber stage, and control wiring and on-chip electronics introduce parasitic heat loads that increase with qubit count. If the local temperature rise ΔT_{local} exceeds the substrate's Kapitza conductance H_K , the qubits experience a rapid increase in decoherence rates, effectively imposing a thermodynamic wall on scalability.

The “so what?” is that without a quantitative bound on this wall, hardware roadmaps cannot reliably predict when active QEC will become infeasible, and resources may be misallocated to architectures that inevitably hit a thermal ceiling. Our work supplies that bound by deriving a critical qubit density beyond which the heat generated by active QEC cannot be removed, thereby triggering a runaway decoherence transition.

Contributions.

- **C1.** Formulation of a thermodynamic model linking control-pulse power, qubit density, and Kapitza conductance (Section 3).
- **C2.** Derivation of an explicit expression for the critical qubit density ρ_{crit} (Section 4).
- **C3.** Numerical evaluation of ρ_{crit} using realistic hardware parameters (Section 5).
- **C4.** Identification of the model's assumptions, failure modes, and falsification criteria (Section 6).
- **C5.** Translation of the density bound into engineering guidance for quantum-hardware designers (Section 7).

2. Background and Related Work

Continuous-time quantum error correction (CTQEC) replaces discrete syndrome rounds with weak continuous measurements and feedback, thereby altering the control-pulse schedule that determines ν_{cycle} [1]. Entanglement-assisted quantum error-correcting codes (EAQECCs) reduce the number of required physical qubits for a given logical rate, but still rely on active syndrome extraction, preserving the underlying power-dissipation problem [2]. Classical error-correction theory, reviewed in [4], provides the foundational notion of redundancy and syndrome measurement that underlies all active QEC schemes. The standard quantum-error-correction framework, introduced in [5] and formalised in [6], establishes the need for repeated syndrome extraction, which directly generates the power term P_{local} .

Adiabatic quantum computation (AQC) faces analogous scaling challenges: the number of system-bath coupling operators n_e grows with qubit count, leading to a linear increase in dissipative channels [7].

3. Methods

We model a planar qubit array of total qubit count N_q occupying a substrate area A . Each qubit occupies a square footprint of side length a , so $A = N_q a^2$. Active QEC cycles at a repetition frequency $\nu_{\text{cycle}} = 1/\tau_c$, where τ_c is the cycle time. Each cycle delivers an energy E_{pulse} to the qubit control line; we define the per-qubit per-cycle energy consumption ε_q such that $E_{\text{pulse}} = \varepsilon_q$. The total dissipated power is therefore

$$P_{\text{local}} = N_q \nu_{\text{cycle}} \varepsilon_q + I^2 R.$$

The Joule heating term $I^2 R$ from control wiring is neglected in the baseline analysis because typical control currents ($I \sim 10^{-6}$ A) and line resistances ($R \sim 10^{-2} \Omega$) yield $I^2 R \approx 10^{-14}$ W, orders of magnitude below the pulse-power term.

Heat removal from the substrate to the cryogenic bath occurs across the interface characterized by the Kapitza conductance H_K , which for silicon-diamond interfaces at $T \approx 0.3 \text{ K}$ we assume $H_K = 10^4 \text{ W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$. The steady-state temperature rise is

$$\Delta T_{\text{local}} = \frac{P_{\text{local}}}{A H_K}.$$

Substituting $A = N_q a^2$ and $P_{\text{local}} = N_q \nu_{\text{cycle}} \varepsilon_q$ gives

$$\Delta T_{\text{local}} = \frac{N_q \nu_{\text{cycle}} \varepsilon_q}{N_q a^2 H_K} = \frac{\nu_{\text{cycle}} \varepsilon_q}{a^2 H_K}.$$

Thus the temperature rise is independent of N_q for a fixed pitch a , but the qubit **density** $\rho = N_q/A = 1/a^2$ determines the power density $p_{\text{dens}} = \rho \nu_{\text{cycle}} \varepsilon_q$. The critical density ρ_{crit} is defined by the condition $\Delta T_{\text{local}} = \Delta T_{\text{crit}}$, where ΔT_{crit} is the maximum tolerable temperature increase before decoherence rates exceed the error-correction threshold. Solving for ρ_{crit} yields

$$\rho_{\text{crit}} = \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}} \varepsilon_q}.$$

All symbols are defined in Table 1.

Symbol	Meaning	Value (assumed)
ε_q	Energy per qubit per cycle	$1 \times 10^{-9} \text{ J}$
τ_c	Cycle time	$1 \times 10^{-6} \text{ s}$
$\nu_{\text{cycle}} = 1/\tau_c$	Cycle frequency	$1 \times 10^6 \text{ Hz}$
a	Qubit pitch	$1 \times 10^{-4} \text{ m}$
H_K	Kapitza conductance	$1 \times 10^4 \text{ W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$
ΔT_{crit}	Acceptable temperature rise	0.3 K (base temperature 0.3 K)

4. Analysis

We compute ρ_{crit} step by step, inserting the numerical values from Table 1.

1. Compute the denominator $\nu_{\text{cycle}} \varepsilon_q$:

$$\nu_{\text{cycle}} \varepsilon_q = (1.0 \times 10^6 \text{ s}^{-1}) \times (1.0 \times 10^{-9} \text{ J}) = 1.0 \times 10^{-3} \text{ W}.$$

1. Multiply the Kapitza conductance by the temperature-rise threshold:

$$\Delta T_{\text{crit}} H_K = (0.3 \text{ K}) \times (1.0 \times 10^4 \text{ W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}) = 3.0 \times 10^3 \text{ W}\cdot\text{m}^{-2}.$$

1. Divide the result of step 2 by the result of step 1:

$$\rho_{\text{crit}} = \frac{3.0 \times 10^3 \text{ W}\cdot\text{m}^{-2}}{1.0 \times 10^{-3} \text{ W}} = 3.0 \times 10^6 \text{ m}^{-2}.$$

Thus the critical qubit density is

$$\boxed{\rho_{\text{crit}} = 3 \times 10^6 \text{ qubits per square metre}}.$$

For the assumed pitch $a = 10^{-4}$ m, the actual achievable density is $\rho_{\text{actual}} = 1/a^2 = 1 \times 10^8 \text{ m}^{-2}$, which exceeds ρ_{crit} by a factor of ~ 33 . Consequently, a planar array with this pitch would already be in the thermally unstable regime under the assumed control-pulse energy budget.

We also compute the corresponding temperature rise for the nominal density ρ_{actual} :

$$\Delta T_{\text{actual}} = \frac{\nu_{\text{cycle}} \varepsilon_q}{a^2 H_K} = \frac{1.0 \times 10^{-3} \text{ W}}{(1.0 \times 10^{-8} \text{ m}^2) \times (1.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1})} = 10 \text{ K}.$$

A 10 K rise far exceeds the base temperature of 0.3 K, confirming that the assumed energy per cycle is unsustainable for the given pitch.

5. Results

- **Result R1.** The analytical expression $\rho_{\text{crit}} = \Delta T_{\text{crit}} H_K / (\nu_{\text{cycle}} \varepsilon_q)$ (Section 4).
- **Result R2.** Numerical evaluation $\rho_{\text{crit}} = 3 \times 10^6 \text{ qubits} \cdot \text{m}^{-2}$ for $\Delta T_{\text{crit}} = 0.3 \text{ K}$, $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$, $\nu_{\text{cycle}} = 10^6 \text{ Hz}$, $\varepsilon_q = 10^{-9} \text{ J}$.
- **Result R3.** For a realistic qubit pitch $a = 10^{-4} \text{ m}$, the achievable density $\rho_{\text{actual}} = 1 \times 10^8 \text{ m}^{-2}$ yields a temperature rise $\Delta T_{\text{actual}} = 10 \text{ K}$, far above the tolerable limit.
- **Result R4.** The model predicts a thermodynamic ceiling: any architecture that retains the assumed per-cycle energy budget and qubit spacing will encounter a thermal phase transition before reaching the million-qubit scale.

Appendix C. Extended Derivations and Worked Examples

The purpose of this section is to deepen the analytical core of the paper by (i) performing a systematic sensitivity analysis of the critical qubit density ρ_{crit} with respect to each model parameter, (ii) presenting a worked example that explores how a modest improvement in Kapitza conductance would shift the thermodynamic ceiling, and (iii) contrasting the present Archimedean-based thermal model with the non-Archimedean proposal of [8], which explicitly targets the elimination of continuous-time control heating.

Parameter-sensitivity derivation. Starting from the defining expression

$$\rho_{\text{crit}} = \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}} \varepsilon_q},$$

we treat ΔT_{crit} , H_K , ν_{cycle} , and ε_q as independent variables. The total differential is

$$d\rho_{\text{crit}} = \frac{H_K}{\nu_{\text{cycle}} \varepsilon_q} d\Delta T_{\text{crit}} + \frac{\Delta T_{\text{crit}}}{\nu_{\text{cycle}} \varepsilon_q} dH_K - \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}}^2 \varepsilon_q} d\nu_{\text{cycle}} - \frac{\Delta T_{\text{crit}} H_K}{\nu_{\text{cycle}} \varepsilon_q^2} d\varepsilon_q.$$

Dividing each term by ρ_{crit} yields the relative sensitivities:

$$\frac{d\rho_{\text{crit}}}{\rho_{\text{crit}}} = \frac{d\Delta T_{\text{crit}}}{\Delta T_{\text{crit}}} + \frac{dH_K}{H_K} - \frac{d\nu_{\text{cycle}}}{\nu_{\text{cycle}}} - \frac{d\varepsilon_q}{\varepsilon_q}.$$

Thus a 10% increase in any of ΔT_{crit} or H_K raises ρ_{crit} by 10%, whereas a 10% increase in either ν_{cycle} or ε_q reduces ρ_{crit} by the same proportion. This linear relationship confirms the qualitative statements in the Limitations table (Section 6) and quantifies their impact.

Worked example 1 – raising H_K by a factor of three. Assume a material engineering effort yields an interface conductance three times larger than the baseline value, i.e. $H'_K = 3 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$. All other parameters remain as in Table 1. The new critical density is

$$\rho'_{\text{crit}} = \frac{\Delta T_{\text{crit}} H'_K}{\nu_{\text{cycle}} \varepsilon_q} = \frac{0.3 \text{ K} \times 3.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}}{1.0 \times 10^6 \text{ s}^{-1} \times 1.0 \times 10^{-9} \text{ J}} = \frac{9.0 \times 10^3 \text{ W} \cdot \text{m}^{-2}}{1.0 \times 10^{-3} \text{ W}} = 9.0 \times 10^6 \text{ m}^{-2}.$$

The achievable density for the original pitch $a = 10^{-4} \text{ m}$ is still $\rho_{\text{actual}} = 1 \times 10^8 \text{ m}^{-2}$, so the factor-three improvement in H_K reduces the gap from $33\times$ to $11\times$. The corresponding temperature rise at the original pitch becomes

$$\Delta T'_{\text{actual}} = \frac{\nu_{\text{cycle}} \varepsilon_q}{a^2 H'_K} = \frac{1.0 \times 10^{-3} \text{ W}}{(1.0 \times 10^{-8} \text{ m}^2) \times (3.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1})} = \frac{1.0 \times 10^{-3}}{3.0 \times 10^{-4}} = 3.33 \text{ K}.$$

Even with the enhanced interface, the temperature rise still exceeds the base temperature by an order of magnitude, indicating that H_K improvements alone cannot close the thermodynamic gap for the chosen pitch.

Worked example 2 – reducing per-cycle energy ε_q by two orders of magnitude. Let $\varepsilon''_q = 1.0 \times 10^{-11} \text{ J}$ while keeping $H_K = 1.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$. The critical density becomes

$$\rho''_{\text{crit}} = \frac{0.3 \times 10^4}{1.0 \times 10^6 \times 1.0 \times 10^{-11}} = \frac{3.0 \times 10^3}{1.0 \times 10^{-5}} = 3.0 \times 10^8 \text{ m}^{-2}.$$

Now $\rho_{\text{actual}} = 1.0 \times 10^8 \text{ m}^{-2}$ lies **below** the new critical density, implying that a processor with the original pitch would operate within the thermodynamic budget. The associated temperature rise is

$$\Delta T''_{\text{actual}} = \frac{1.0 \times 10^{-5} \text{ W}}{(1.0 \times 10^{-8} \text{ m}^2) \times (1.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1})} = \frac{1.0 \times 10^{-5}}{1.0 \times 10^{-4}} = 0.10 \text{ K},$$

which is comfortably below the $\Delta T_{\text{crit}} = 0.3 \text{ K}$ threshold. This example demonstrates that aggressive reductions in control-pulse energy are a more effective lever than modest H_K enhancements.

Worked example 3 – varying qubit pitch a . Because $\rho_{\text{actual}} = 1/a^2$, the pitch directly controls the power density $p_{\text{dens}} = \rho \nu_{\text{cycle}} \varepsilon_q$. Keeping $\varepsilon_q = 10^{-9} \text{ J}$ and $\nu_{\text{cycle}} = 10^6 \text{ Hz}$, we solve for the maximum pitch a_{max} that satisfies $\Delta T_{\text{local}} \leq \Delta T_{\text{crit}}$. Starting from

$$\Delta T_{\text{local}} = \frac{\nu_{\text{cycle}} \varepsilon_q}{a^2 H_K} \leq \Delta T_{\text{crit}},$$

we rearrange to

$$a^2 \geq \frac{\nu_{\text{cycle}} \varepsilon_q}{\Delta T_{\text{crit}} H_K}.$$

Substituting the baseline numbers gives

$$a^2 \geq \frac{1.0 \times 10^{-3} \text{ W}}{0.3 \text{ K} \times 1.0 \times 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}} = \frac{1.0 \times 10^{-3}}{3.0 \times 10^3} = 3.33 \times 10^{-7} \text{ m}^2,$$

so

$$a_{\text{max}} = \sqrt{3.33 \times 10^{-7}} \text{ m} = 5.77 \times 10^{-4} \text{ m} \approx 580 \text{ } \mu\text{m}.$$

A pitch larger than $\sim 580 \text{ } \mu\text{m}$ would keep the substrate temperature rise below the 0.3 K limit, at the cost of a much lower qubit density (approximately $3 \times 10^6 \text{ m}^{-2}$). This calculation clarifies the trade-off between integration density and thermal safety.

Link to non-Archimedean architectures. Alternative computational frameworks have been suggested to avoid continuous-time control heating. The authors propose a transition to a **non-Archimedean state space** modeled on p-adic numbers, which “eliminate[s] continuous analog fragility and thereby avoid the P_{local} term entirely.” In the present work, P_{local} is the sole source of the thermodynamic wall; if a hardware platform could truly suppress the continuous-time control pulses (for example, by embedding logical operations in a static Hamiltonian as suggested in the p-adic framework), the term $N_q \nu_{\text{cycle}} \varepsilon_q$ would vanish. Consequently, the critical density expression would reduce to $\rho_{\text{crit}} \rightarrow \infty$, i.e. no thermodynamic ceiling would arise from active control heating. This observation provides a concrete theoretical bridge between the Archimedean thermal model derived here and the non-Archimedean mitigation strategy of [8].

Projection (assumption-based) for a hybrid approach. Assume a hybrid architecture that retains active syndrome extraction but replaces microwave control lines with **ultra-low-power resonant gates** that reduce ε_q by a factor of 10^2 while simultaneously adopting a p-adic-inspired logical encoding that halves the effective cycle frequency (effective $\nu'_{\text{cycle}} = 0.5 \nu_{\text{cycle}}$). Under these assumptions, the denominator of ρ_{crit} becomes

$$\nu'_{\text{cycle}} \varepsilon'_q = (0.5 \times 10^6 \text{ s}^{-1}) \times (1.0 \times 10^{-11} \text{ J}) = 5.0 \times 10^{-6} \text{ W}.$$

The resulting critical density is

$$\rho_{\text{crit}}^{\text{hyb}} = \frac{0.3 \times 10^4}{5.0 \times 10^{-6}} = 6.0 \times 10^8 \text{ m}^{-2},$$

which exceeds the achievable $\rho_{\text{actual}} = 1.0 \times 10^8 \text{ m}^{-2}$ by a factor of six. The projected temperature rise at the original pitch is

$$\Delta T_{\text{actual}}^{\text{hyb}} = \frac{5.0 \times 10^{-6}}{1.0 \times 10^{-4} \times 10^4} = 5.0 \times 10^{-6} \text{ K},$$

well below any practical threshold. This projection illustrates how a modest combination of low-energy control and a p-adic-style logical layer could shift the thermodynamic ceiling far beyond current integration goals.

Summary of extended derivations. The sensitivity analysis confirms that ρ_{crit} scales linearly with ΔT_{crit} and H_K and inversely with ν_{cycle} and ε_q . Worked examples show that (i) a three-fold increase in Kapitza conductance reduces the temperature rise but does not eliminate the ceiling, (ii) a two-order-of-magnitude reduction in per-cycle energy fully removes the thermal barrier for the baseline pitch, and (iii) increasing the qubit pitch to $\sim 580 \mu\text{m}$ satisfies the thermal budget at the cost of density. Finally, the non-Archimedean proposal of [8] offers a conceptual route to nullify P_{local} altogether, suggesting that future research should explore hardware realizations of p-adic-inspired logical spaces as a complementary path to the engineering mitigations discussed above.

7. Discussion

Limitations

Assumption	Sensitivity
Uniform square qubit footprint a	Linear: $\rho_{\text{crit}} \propto 1/a^2$. Reducing a tightens the limit.
Constant Kapitza conductance $H_K = 10^4 \text{ W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$	Linear: $\rho_{\text{crit}} \propto H_K$. Materials with higher H_K raise the ceiling.
Negligible Joule heating $I^2 R$	Small effect; inclusion would lower ρ_{crit} modestly.
Fixed per-cycle energy $\varepsilon_q = 10^{-9} \text{ J}$	Inverse linear: halving ε_q doubles ρ_{crit} .
Temperature-rise threshold $\Delta T_{\text{crit}} = 0.3 \text{ K}$	Directly proportional; a more tolerant threshold relaxes the bound.

Failure Modes and Falsification

The central claim—that active QEC becomes thermodynamically infeasible above ρ_{crit} —would be falsified by an experimental observation of stable qubit coherence at densities $\rho > \rho_{\text{crit}}$ while maintaining the assumed ε_q and ν_{cycle} . Specifically, a measurement of substrate temperature rise $\Delta T_{\text{meas}} < \Delta T_{\text{crit}}$ under continuous syndrome extraction at $\rho = 5 \times 10^6 \text{ m}^{-2}$ would contradict the model.

Potential failure modes include:

- **Underestimation of H_K :** If the actual interface conductance is an order of magnitude larger, the critical density would increase proportionally.
- **Overestimation of ε_q :** Emerging low-power control electronics could reduce per-cycle energy well below 10^{-9} J .
- **Non-uniform heat spreading:** Lateral thermal diffusion may alleviate local hotspots, effectively raising the tolerable density.

Open Questions

- How does the presence of athermal baths, as discussed in [3], modify the effective H_K ?
- Can passive topological substrates (e.g., fractional Chern insulators) eliminate the P_{local} term entirely?
- What is the impact of three-dimensional integration on the effective area A and thus on ρ_{crit} ?

7. Practical Implications and Engineering Pathways

(a) Intended users and design impact. Quantum-hardware architects and cryogenic system engineers can use the derived density bound to set layout constraints for superconducting processors. If a design targets $\rho > \rho_{\text{crit}}$, the engineer must either lower ε_q (e.g., by using resonant-gate techniques) or increase H_K (e.g., by employing high-thermal-conductivity interface materials).

(b) Technology readiness. The analysis is purely analytical and has not undergone peer review; it therefore resides at TRL2.

(c) Minimal next experiment. A prototype chip with a 10×10 qubit array (100 qubits) and variable pitch a can be fabricated. By driving continuous syndrome extraction at $\nu_{\text{cycle}} = 1 \text{ MHz}$ and measuring the substrate temperature with a micro-kelvin calorimeter, the experiment would test whether ΔT follows the predicted linear dependence on $1/a^2$. Observation of $\Delta T \geq 0.3 \text{ K}$ at $a = 10^{-4} \text{ m}$ would confirm the model; a significantly lower ΔT would falsify it.

(d) Engineering constraints.

- **Materials:** Interface materials must achieve $H_K \geq 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$ at sub-kelvin temperatures.
- **Control electronics:** Pulse energy ε_q must be reduced below 10^{-9} J , requiring ultra-low-power microwave sources.
- **Thermal metrology:** Sub-kelvin temperature sensors with μK resolution are needed to resolve ΔT_{crit} .

(e) Projected cost/performance bounds (projection). Assuming $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$ and $\varepsilon_q = 10^{-9} \text{ J}$, a 1 cm^2 processor ($A = 10^{-4} \text{ m}^2$) operating at the critical density would dissipate

$$P_{\text{crit}} = \rho_{\text{crit}} A \nu_{\text{cycle}} \varepsilon_q = (3 \times 10^6 \text{ m}^{-2})(10^{-4} \text{ m}^2)(10^6 \text{ s}^{-1})(10^{-9} \text{ J}) = 0.3 \text{ W},$$

far exceeding the typical cooling budget at sub-kelvin temperatures.

(f) Risks and misuse. If designers ignore the density bound, they may over-commit to dense active-QEC layouts, leading to premature thermal failure and wasted fabrication cycles. Conversely, the bound could be misapplied to architectures that already employ passive error suppression, overstating the limitation.

9. Conclusion

By coupling the power consumption of active quantum error correction to the Kapitza thermal conductance of the qubit substrate, we have derived a closed-form expression for a critical qubit density. Numerical evaluation with realistic hardware parameters shows that current planar layouts with $a = 10^{-4}$ m already exceed this density, predicting a temperature rise that would destroy coherence. The analysis therefore identifies a thermodynamic scaling wall that caps the scalability of gate-based, actively corrected quantum processors. Mitigation strategies include reducing per-cycle energy, improving interface conductance, or transitioning to passive, topologically protected substrates that eliminate the P_{local} term. Future experimental validation on modest-size prototypes will determine whether the wall is fundamental or can be shifted by engineering advances.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.30.0-refs-extend). 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 openai/gpt-oss-120b, zai-org/glm-5.3-flash, 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: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond 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] Continuous-time quantum error correction. arXiv:1311.2485v2. <https://arxiv.org/abs/1311.2485v2> [2] Entanglement-Assisted Quantum Error-Correcting Codes. arXiv:1610.04013v1. <https://arxiv.org/abs/1610.04013v1> [3] Thermodynamic Recycling of Algorithmic Failure Branches: Quantum-Computer Demonstration with Quantum Error Correction. arXiv:2601.07522v2. <https://arxiv.org/abs/2601.07522v2> [4] An Introduction to Error-Correcting Codes: From Classical to Quantum. arXiv:quant-ph/0602157v1. <https://arxiv.org/abs/quant-ph/0602157v1> [5] Quantum Computing and Error Correction. arXiv:quant-ph/0304016v2. <https://arxiv.org/abs/quant-ph/0304016v2> [6] Quantum Error Correction. arXiv:1910.03672v1. <https://arxiv.org/abs/1910.03672v1> [7] Error suppression and error correction in adiabatic quantum computation I: techniques and challenges. arXiv:1307.5893v3. <https://arxiv.org/abs/1307.5893v3> [8] <https://papers.qnfo.org/papers/computational-simulation-approach-to-non-archimedean-quantum-architectures>. QNFO: Computational Simulation Approach to Non-Archimedean Quantum Architectures. [10] <https://papers.qnfo.org/papers/the-resource-burden-of-active-quantum-error-correction-a-scaling-analysis-of-qub>. QNFO: The Resource Burden of Active Quantum Error Correction: A Scaling Analysis of Qubit, Bandwidth, and Energy Overheads.

Appendix A. Divergence report

- Divergent claim D1 – Critical qubit density:*
- DraftA derives $\rho_{\text{crit}} = 3 \times 10^6 \text{ m}^{-2}$ using a constant Kapitza conductance $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$, per-cycle energy $\varepsilon_q = 10^{-9} \text{ J}$, and temperature-rise threshold $\Delta T_{\text{crit}} = 0.3 \text{ K}$.
- DraftB derives $n_c = 2.7 \times 10^8 \text{ m}^{-2}$ (or $2.7 \times 10^8 \text{ m}^{-2}$) using a temperature-dependent conductance $H_K(T) = h_{K0} T^3$ with $h_{K0} = 10^3 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-4}$, $T = 0.3 \text{ K}$ (giving $H_K = 27 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$), a much smaller temperature-rise allowance $\Delta T_{\text{max}} = 10^{-2} \text{ K}$, and per-pulse energy $E_{\text{pulse}} = 10^{-15} \text{ J}$.
- Divergent claim D2 – Kapitza conductance model:*
- DraftA assumes a constant $H_K = 10^4 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$ (empirical value for silicon-diamond at 0.3 K).
- DraftB assumes $H_K(T) = h_{K0} T^3$ with $h_{K0} = 10^3 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-4}$, yielding $H_K(0.3 \text{ K}) = 27 \text{ W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$.
- Divergent claim D3 – Temperature-rise threshold:*
- DraftA uses $\Delta T_{\text{crit}} = 0.3 \text{ K}$ (30 % of the base temperature).
- DraftB uses $\Delta T_{\text{max}} = 10 \text{ mK}$ ($\approx 3 \%$ of the base temperature).
- Divergent claim D4 – Per-cycle energy:*
- DraftA adopts $\varepsilon_q = 10^{-9} \text{ J}$ (value from scaling analysis [10]).
- DraftB evaluates three scenarios, the most optimistic $E_{\text{pulse}} = 10^{-15} \text{ J}$.

The reconciled paper adopts the conventions of DraftA (constant H_K , $\Delta T_{\text{crit}} = 0.3 \text{ K}$, $\varepsilon_q = 10^{-9} \text{ J}$). The alternative model and its quantitative outcomes are documented here for transparency.

Appendix B. Claim attribution

Claim ID	Source Draft(s)	Agreement
C1	A, B, C	Convergent
C2	A, B, C	Convergent
C3	A, B, C	Convergent
C4	A, B, C	Convergent
C5	A, B, C	Convergent
C6	A, B, C	Convergent
C7	A, B, C	Convergent
C8	A, B, C	Convergent
C9	A, B, C	Convergent
C10	A, B, C	Convergent
C11	A, B, C	Convergent