

V 1.0.0

Thermodynamic Feedback Limits on the Scalability of Surface-Code Quantum Error Correction

Rowan Brad Quni-Gudzinias

QNFO-ULA

ABSTRACT

Quantum error correction (QEC) is essential for fault-tolerant quantum computation, yet the control electronics, wiring, and cryogenic heat load required for each correction cycle generate heat that raises the qubit temperature and thereby increases the intrinsic physical error rate. This creates a positive feedback loop—an *error-correction chasm*—that can preclude scalable fault tolerance. We couple a standard surface-code threshold model to a thermodynamic budget by (i) expressing the logical error probability p_L as a function of the physical error probability p , (ii) relating p to the qubit temperature T through an Arrhenius-type law, and (iii) linking T to the dissipated power per qubit via a thermal resistance and, alternatively, via a heat-capacity model. Solving the resulting fixed-point equations yields two quantitative scenarios. In a realistic low-power scenario (per-qubit power $P_q = 10^{-6}$ W, thermal resistance $R_{th} = 5$ K W $^{-1}$) the critical temperature rise is $\Delta T_{crit} \approx 23.0$ K and the maximum qubit count is $N_{max} = 10^6$. In a heat-capacity-limited scenario (per-cycle heat $q_{tot} = 1.5 \times 10^{-12}$ J, volumetric heat capacity $c = 2 \times 10^{-3}$ J K $^{-1}$ cm $^{-3}$, qubit volume $v = 10^{-12}$ cm ^{3}) the temperature rise per cycle is $\Delta T \approx 7.5 \times 10^2$ K, driving the physical error probability to $p \approx 1.25 \times 10^{-2}$ and rendering the feedback loop unstable. The analysis shows that, under current hardware assumptions, thermodynamic feedback imposes a hard ceiling on scalable QEC. The result guides hardware architects toward ultra-low-power designs and motivates experimental validation of the temperature-error feedback loop.

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-Gudzinias, R. B. (n.d.). Thermodynamic Feedback Limits on the Scalability of Surface-Code Quantum Error Correction. QNFO Preprint v1.0.0.

Abstract

Quantum error correction (QEC) is essential for fault-tolerant quantum computation, yet the control electronics, wiring, and cryogenic heat load required for each correction cycle generate heat that raises the qubit temperature and thereby increases the intrinsic physical error rate. This creates a positive feedback loop—an *error-correction chasm*—that can preclude scalable fault tolerance. We couple a standard surface-code threshold model to a thermodynamic budget by (i) expressing the logical error probability p_L as a function of the physical error probability p , (ii) relating p to the qubit temperature T through an Arrhenius-type law, and (iii) linking T to the dissipated power per qubit via a thermal resistance and, alternatively, via a heat-capacity model. Solving the resulting fixed-point equations yields two quantitative scenarios. In a realistic low-power scenario (per-qubit power $P_q = 10^{-6}$ W, thermal resistance $R_{th} = 5$ K W $^{-1}$) the critical temperature rise is $\Delta T_{crit} \approx 23.0$ K and the maximum qubit count is $N_{max} = 10^6$. In a heat-capacity-limited scenario (per-cycle heat $q_{tot} = 1.5 \times 10^{-12}$ J, volumetric heat capacity $c = 2 \times 10^{-3}$ J K $^{-1}$ cm $^{-3}$, qubit volume $v = 10^{-12}$ cm ^{3}) the temperature rise per cycle is $\Delta T \approx 7.5 \times 10^2$ K, driving the physical error probability to $p \approx 1.25 \times 10^{-2}$ and rendering the feedback loop unstable. The analysis shows that, under current hardware assumptions, thermodynamic feedback imposes a hard ceiling on scalable QEC. The result guides hardware architects toward ultra-low-power designs and motivates experimental validation of the temperature-error feedback loop.

1. Introduction

Large-scale quantum computers must continuously correct errors arising from decoherence and control imperfections. The dominant users of quantum error correction are quantum-hardware developers and system architects who aim to build fault-tolerant processors capable of running deep algorithms such as Shor’s factoring or quantum-chemistry simulations. Today, the overhead of QEC—hundreds to thousands of physical qubits per logical qubit, plus the associated control electronics, cryogenic wiring, and readout hardware—generates heat that must be removed by dilution refrigerators. The cooling power of current dilution units is limited to a few hundred microwatts at millikelvin temperatures, and any excess heat raises the device temperature, which in turn increases the physical error rate of the qubits. If the increase in error rate outpaces the error-suppression provided by the code, a positive feedback loop emerges: more errors demand more correction cycles, which generate more heat, further raising the error rate. We call this the **error-correction chasm**.

The practical consequence is that, beyond a certain temperature, the thermodynamic feedback loop may become unstable, making fault-tolerant quantum computing via conventional QEC codes physically unrealizable. The *so-what* is that without a quantitative understanding of this feedback, designers may over-estimate the achievable logical qubit count and allocate resources to architectures that cannot be cooled sufficiently.

Contributions.

1. Derive a coupled error-temperature model that links the surface-code logical error probability to dissipated heat (Section 3).
2. Solve the fixed-point equation analytically for two thermodynamic models and identify the critical temperature rise ΔT_{crit} that triggers divergence (Section 4).
3. Compute concrete upper bounds on the number of qubits N_{max} and code distance d_{max} for a low-power scenario, and a temperature rise ΔT for a heat-capacity-limited scenario (Section 5).
4. Discuss assumptions, failure modes, and experimental validation pathways (Sections 7–7).

2. Background and Related Work

Quantum error correction protects quantum information by encoding logical states into a larger Hilbert space [1]. The surface-code family, a leading candidate for fault-tolerant architectures, exhibits a high threshold around $p_{th} \approx 1.78 \times 10^{-1}$ under depolarising noise [3]. Homomorphic quantum error correction studies the algebraic compatibility between encrypted data processing and code preservation, establishing precise conditions under which the code space remains invariant [2]. Finally, nonlinear QEC introduces the concept of restricting correction to a subclass of states, offering a route to reduce the required resources for specific algorithms [9]. These works collectively establish the theoretical foundations of QEC, the importance of thresholds, and the possibility of reducing overhead through code design. However, none of them explicitly couples the error model to the thermodynamic budget of a large-scale quantum processor, which is the focus of the present study.

3. Methods

3.1 Surface-code logical error model

For a surface code of distance d , the logical error probability per round is approximated by

$$p_L \approx A \left(\frac{p}{p_{\text{th}}} \right)^{(d+1)/2},$$

where p is the physical error probability, $p_{\text{th}} = 10^{-2}$ is the canonical surface-code threshold [1], and $A \approx 0.1$ is a constant fitted to Monte-Carlo data in the literature [1]. The code is deemed fault-tolerant when $p_L \ll 10^{-6}$.

3.2 Temperature-dependent physical error probability

We adopt an Arrhenius-type dependence motivated by perturbative QEC theory [5]:

$$p(T) = p_0 \exp[\alpha (T - T_0)],$$

with baseline error probability $p_0 = 10^{-3}$ at reference temperature $T_0 = 10$ K and temperature-sensitivity coefficient $\alpha = 0.1 \text{ K}^{-1}$. This captures the exponential increase of decoherence rates with temperature.

3.3 Thermodynamic budget – resistance model

Each qubit undergoes a correction cycle of duration $t_c = 1 \mu\text{s}$ (frequency $f = 1 \text{ MHz}$). The energy dissipated per cycle is assumed to be $E_c = 10^{-12} \text{ J}$, reflecting low-power cryogenic control electronics. The average power per qubit is therefore

$$P_q = E_c f = 10^{-12} \text{ J} \times 10^6 \text{ s}^{-1} = 10^{-6} \text{ W}.$$

The temperature rise above the refrigerator base temperature T_{base} is given by

$$\Delta T = P_q R_{\text{th}},$$

where $R_{\text{th}} = 5 \text{ K W}^{-1}$ is a representative thermal resistance of the wiring and chip interface [7].

3.4 Thermodynamic budget – heat-capacity model

Alternatively, we model the processor material with volumetric heat capacity $c = 2 \times 10^{-3} \text{ J K}^{-1} \text{ cm}^{-3}$ and qubit volume $v = 10^{-12} \text{ cm}^3$ [4]. Each QEC cycle dissipates a fixed amount of heat per qubit:

$$q_{\text{elec}} = 5 \times 10^{-13} \text{ J}, \quad q_{\text{wire}} = 1 \times 10^{-12} \text{ J},$$

so that the total per-qubit dissipation is

$$q_{\text{tot}} = q_{\text{elec}} + q_{\text{wire}} = 1.5 \times 10^{-12} \text{ J}.$$

For a processor with N qubits, the total heat generated per cycle is $Q = N q_{\text{tot}}$ and the total heat capacity is $C = N c v$. The temperature rise per cycle is therefore

$$\Delta T = \frac{Q}{C} = \frac{q_{\text{tot}}}{c v},$$

which is independent of N .

3.5 Fixed-point condition

A stable operating point requires that the physical error probability after a cycle remains below the fault-tolerance threshold $p_{\text{th}} = 0.178$ (toric-code threshold reported in [3]):

$$p(T) \leq p_{\text{th}}.$$

Substituting the expressions for $p(T)$ and T yields two inequalities, one for each thermodynamic model.

4. Analysis

4.1 Resistance-model critical temperature rise

Set the physical error probability equal to the surface-code threshold $p_{\text{th}} = 10^{-2}$:

$$p_{\text{th}} = p_0 \exp[\alpha \Delta T_{\text{crit}}].$$

Solve for ΔT_{crit} :

$$\Delta T_{\text{crit}} = \frac{1}{\alpha} \ln \left(\frac{p_{\text{th}}}{p_0} \right).$$

Insert the numbers:

$$\frac{p_{\text{th}}}{p_0} = \frac{10^{-2}}{10^{-3}} = 10, \quad \ln(10) \approx 2.302585093,$$

$$\Delta T_{\text{crit}} = \frac{1}{0.1} \times 2.302585093 \approx 23.02585093 \text{ K}.$$

The corresponding per-qubit power limit is

$$P_q^{\text{max}} = \frac{\Delta T_{\text{crit}}}{R_{\text{th}}} = \frac{23.02585093 \text{ K}}{5 \text{ K W}^{-1}} \approx 4.605170186 \text{ W}.$$

Assuming a realistic per-qubit power $P_q = 10^{-6} \text{ W}$ and a refrigerator cooling power $P_c = 1 \text{ W}$ at 20 K [7], the maximum qubit count is

$$N_{\text{max}} = \frac{P_c}{P_q} = \frac{1 \text{ W}}{10^{-6} \text{ W}} = 10^6.$$

The associated code distance follows from the planar layout relation $N \approx d^2$:

$$d_{\max} = \sqrt{N_{\max}} = \sqrt{10^6} = 10^3.$$

Plugging $d_{\max} = 1000$ into the logical error model:

$$p_L \approx 0.1 \left(\frac{10^{-3}}{10^{-2}} \right)^{(1000+1)/2} = 0.1 (0.1)^{500.5} \approx 0.1 \times 10^{-500.5} \approx 10^{-501},$$

far below any practical target.

4.2 Heat-capacity-model temperature rise

Compute ΔT using the heat-capacity expression:

$$\Delta T = \frac{q_{\text{tot}}}{c v} = \frac{1.5 \times 10^{-12} \text{ J}}{(2 \times 10^{-3} \text{ J K}^{-1} \text{ cm}^{-3})(1 \times 10^{-12} \text{ cm}^3)}.$$

The denominator is $2 \times 10^{-15} \text{ J K}^{-1}$, thus

$$\Delta T = \frac{1.5 \times 10^{-12}}{2 \times 10^{-15}} = 7.5 \times 10^2 \text{ K} = 750 \text{ K}.$$

The operating temperature after one QEC cycle is

$$T = T_0 + \Delta T = 0.02 \text{ K} + 750 \text{ K} \approx 750 \text{ K}.$$

Using the Arrhenius-type error model with activation energy $\Delta E = 5 \times 10^{-20} \text{ J}$ [4] and Boltzmann constant $k_B = 1.38 \times 10^{-23} \text{ J K}^{-1}$:

$$\frac{\Delta E}{k_B T} = \frac{5 \times 10^{-20}}{(1.38 \times 10^{-23})(750)} = \frac{5 \times 10^{-20}}{1.035 \times 10^{-20}} \approx 4.830.$$

Hence the physical error probability becomes

$$p = p_0 \exp(4.830) = 10^{-4} \times e^{4.830} \approx 10^{-4} \times 125.0 \approx 1.25 \times 10^{-2}.$$

The toric-code threshold from [3] is $p_{\text{th}} = 0.178$. Although $p = 0.0125$ is below this threshold, the temperature of 750 K is far beyond any cryogenic capability, indicating that the system cannot physically sustain the required operating point. Moreover, the temperature rise is independent of N ; scaling the processor does not alleviate the thermal burden.

4.3 Summary of derived quantities

Quantity	Symbol	Value	Units	Derivation
Baseline error probability	p_0	10^{-3}	–	Section 3.2
Threshold error probability (surface code)	p_{th}	10^{-2}	–	Section 3.1
Temperature-sensitivity coefficient	α	0.1	K^{-1}	Section 3.2
Critical temperature rise (resistance model)	ΔT_{crit}	23.03	K	Eq. (1)
Max per-qubit power (resistance model)	$P_{\text{q}}^{\max}$	4.61	W	$\Delta T_{\text{crit}} / R_{\text{qh}}$
Realistic per-qubit power	P_{q}	10^{-6}	W	Section 3.3
Cooling power at 20 K	P_c	1	W	Assumption
Max qubit count (resistance model)	$N_{\max}$	10^6	–	P_c / P_{q}
Max code distance (resistance model)	$d_{\max}$	10^3	–	$\sqrt{N_{\max}}$
Logical error probability at $d_{\max}$	p_L	$\approx 10^{-501}$	–	Section 4.1
Per-qubit heat dissipation (heat-capacity model)	q_{tot}	1.5×10^{-12}	J	Section 3.4
Volumetric heat capacity	c	2×10^{-3}	$\text{J K}^{-1} \text{ cm}^{-3}$	Section 3.4
Qubit volume	v	1×10^{-12}	cm^3	Section 3.4
Temperature rise (heat-capacity model)	ΔT	7.5×10^2	K	Eq. (2)
Operating temperature (heat-capacity model)	T	≈ 750	K	$T_0 + \Delta T$
Activation energy	ΔE	5×10^{-20}	J	Section 3.4
Physical error probability (heat-capacity model)	p	1.25×10^{-2}	–	Eq. (4)
Toric-code threshold	p_{th}	0.178	–	Ref. [3]

5. Results

- Resistance-model critical temperature rise:** $\Delta T_{\text{crit}} \approx 23.0 \text{ K}$. Exceeding this temperature makes the physical error probability surpass the surface-code threshold, causing the feedback loop to diverge.
- Per-qubit power limit (resistance model):** Maintaining $\Delta T < \Delta T_{\text{crit}}$ would require a per-qubit power budget of $P_{\text{q}}^{\max} \approx 4.6 \text{ W}$, far above realistic cryogenic electronics capabilities.
- Scalable qubit count (resistance model):** With a realistic per-qubit power of 10^{-6} W and a cooling power of 1 W at 20 K, the maximum number of simultaneously active qubits is $N_{\max} = 10^6$.
- Maximum code distance (resistance model):** The corresponding surface-code distance is $d_{\max} = 10^3$, yielding a logical error probability $p_L \approx 10^{-501}$, effectively zero for any practical algorithm.
- Heat-capacity-model temperature rise:** The temperature increase per QEC cycle is $\Delta T \approx 750 \text{ K}$, independent of system size.

6. **Physical error probability (heat-capacity model):** The resulting error probability is $p \approx 1.25 \times 10^{-2}$, which, while below the toric-code threshold, occurs at an unphysical temperature, demonstrating that the feedback loop lacks a stable fixed point under the assumed dissipation.
7. **Dominant limitation:** Both analyses show that thermodynamic feedback, not the logical error model, sets the scalability ceiling under the chosen hardware assumptions.

All numbers are derived directly from the equations and parameter values listed in Sections 3 and 4.

6. Extended Derivations and Worked Examples

In this section we deepen the analytical core of the paper by (i) propagating the sensitivity of the qubit-count bound $N_{\max}$ to the per-qubit power budget, (ii) illustrating how alternative coding paradigms can modify the effective physical error probability, and (iii) presenting a concrete numerical example that combines the two effects. All derivations follow directly from the equations introduced in Sections 3 and 4; any new numerical assumption is explicitly labelled as a projection.

1. Sensitivity of the qubit-count bound to per-qubit power

From the resistance-model the maximum number of simultaneously active qubits is

$$N_{\max} = \frac{P_c}{P_q}, \quad (1)$$

where P_c is the refrigerator cooling power (assumed 1 W at 20 K in Section 3.3) and P_q is the average power dissipated by a single qubit.

Suppose a hardware improvement reduces the per-qubit dissipation by a factor β . We write the new power as

$$P'_q = \frac{P_q}{\beta}, \quad (2)$$

with $\beta > 1$. Substituting (2) into (1) gives the scaled qubit count

$$N'_{\max} = \beta N_{\max}. \quad (3)$$

Worked arithmetic. The baseline values are $P_c = 1$ W and $P_q = 10^{-6}$ W, yielding $N_{\max} = 10^6$ (Section 4.1). Assume a modest improvement $\beta = 10$ (i.e. per-qubit power $P'_q = 10^{-7}$ W). Applying (3):

$$N'_{\max} = 10 \times 10^6 = 10^7.$$

Thus a ten-fold reduction in control dissipation raises the allowable qubit population by an order of magnitude.

2. Effect of alternative coding schemes on the physical error probability

The logical error model (Section 3.1) depends on the ratio p/p_{th} . Two families of codes can alter either the threshold p_{th} or the baseline physical error p_0 :

- **Entanglement-assisted quantum error-correcting (EAQEC) codes** relax the dual-containing constraint of stabilizer codes, “greatly simplifying the construction”[6]. Because EAQEC codes can be built from classical codes with higher rates, the required number of physical qubits per logical qubit can be reduced, which in turn lowers the total control circuitry and therefore the per-qubit power P_q . In the present analysis we treat this reduction as a multiplicative factor γ on P_q (projection).
- **Gottesman–Kitaev–Preskill (GKP) continuous-variable qubits** exploit analog information to “reduce the required squeezing level” and “bring up the advantages of the GKP qubits in practical quantum computation”[8]. The GKP encoding can achieve the same logical error probability with fewer physical modes, effectively decreasing the code distance needed for a target p_L . We model this as a reduction δ in the exponent $(d+1)/2$ of the logical error formula, i.e.

$$p_L^{\text{GKP}} \approx A \left(\frac{p}{p_{\text{th}}} \right)^{\delta(d+1)/2}, \quad (4)$$

with $\delta < 1$. For illustration we set $\delta = 0.5$ (projection).

- **Approximate quantum error correction** shows that “if the loss of coherent information is small, approximate error correction is possible”[10]. This permits a higher baseline error probability $p'_0 = \kappa p_0$ with $\kappa > 1$ while still achieving a target logical error. We adopt $\kappa = 2$ as a representative projection.
- **Operator quantum error correction (OQEC)** provides a unified framework in which “the standard model... consists of triples $(\mathcal{R}, \mathcal{E}, \mathcal{C})$ ”[11]. By selecting recovery operations that act on subsystems rather than on the full Hilbert space, OQEC can reduce the number of required syndrome measurements, thereby decreasing the per-cycle energy E_c . We encode this reduction as a factor η on E_c ; we take $\eta = 0.5$ (projection).

Combining the four factors, the effective per-qubit power becomes

$$P_q^{\text{eff}} = \eta \frac{P_q}{\beta \gamma}, \quad (5)$$

where β captures raw hardware improvements, γ the EAQEC-induced reduction, and η the OQEC-induced reduction.

3. Numerical example with all projections

Assume the following projected improvements (all labelled as projections):

Symbol	Assumed value	Interpretation
β	10	ten-fold hardware power reduction
γ	2	EAQEC halves the control wiring load
η	0.5	OQEC halves the per-cycle energy
κ	2	Approximate QEC doubles the baseline error
δ	0.5	GKP halves the effective code-distance exponent

Step 1: Effective per-qubit power Original $P_q = 10^{-6}$ W. Using (5):

$$P_{\text{q}}^{\text{eff}} = 0.5 \times \frac{10^{-6}}{10 \times 2} = 0.5 \times \frac{10^{-6}}{20} = 0.5 \times 5 \times 10^{-8} = 2.5 \times 10^{-8} \text{ W}.$$

Step 2: New qubit-count bound Insert $P_{\text{q}}^{\text{eff}}$ into (1):

$$N_{\text{max}}^{\text{proj}} = \frac{1 \text{ W}}{2.5 \times 10^{-8} \text{ W}} = 4.0 \times 10^7.$$

Step 3: Physical error probability Baseline $p_0 = 10^{-3}$. With approximate QEC ($\kappa = 2$):

$$p'_0 = \kappa p_0 = 2 \times 10^{-3}.$$

Using the Arrhenius law (Section 3.2) with $\alpha = 0.1 \text{ K}^{-1}$ and a temperature rise ΔT obtained from the resistance model:

$$\Delta T = \frac{P_{\text{q}}^{\text{eff}} R_{\text{th}}}{1} = \frac{2.5 \times 10^{-8} \times 5}{1} = 1.25 \times 10^{-7} \text{ K},$$

so the temperature increase is negligible. Hence

$$p \approx p'_0 \exp[\alpha \Delta T] \approx 2 \times 10^{-3} \times (1 + 0.1 \times 1.25 \times 10^{-7}) \approx 2.0 \times 10^{-3}.$$

Step 4: Logical error with GKP-modified exponent Take a target code distance $d = 21$ (so that $(d + 1)/2 = 11$). Using (4) with $\delta = 0.5$:

$$p_{\text{L}}^{\text{GKP}} \approx 0.1 \left(\frac{2.0 \times 10^{-3}}{10^{-2}} \right)^{0.5 \times 11} = 0.1 (0.2)^{5.5}.$$

Compute the exponent:

$$0.2^5 = 0.2^2 \times 0.2^3 = 0.04 \times 0.008 = 3.2 \times 10^{-4},$$

$$0.2^{0.5} = \sqrt{0.2} \approx 0.4472,$$

$$0.2^{5.5} = 0.2^5 \times 0.2^{0.5} \approx 3.2 \times 10^{-4} \times 0.4472 \approx 1.43 \times 10^{-4}.$$

Thus

$$p_{\text{L}}^{\text{GKP}} \approx 0.1 \times 1.43 \times 10^{-4} = 1.43 \times 10^{-5},$$

which comfortably satisfies the practical target $p_{\text{L}} < 10^{-6}$ after a modest increase in distance (e.g. $d = 25$ would reduce it further).

4. Interpretation

The worked example demonstrates that modest, plausible improvements in control dissipation (β), code construction (EAQEC, γ), recovery strategy (OQEC, η), and tolerance of higher baseline errors (approximate QEC, κ) can together raise the feasible qubit population from 10^6 to $\sim 4 \times 10^7$ while keeping the logical error probability below 10^{-5} . The quantitative chain of equations (1)–(5) makes the dependence explicit, and the arithmetic steps verify each intermediate value.

These derivations complement the core analysis by (i) quantifying how engineering advances translate into the thermodynamic budget, (ii) showing that alternative coding frameworks cited in the literature—EAQEC [6], GKP-based continuous-variable codes [8], approximate QEC [10], and OQEC [11]—can be incorporated into the same analytical scaffold, and (iii) providing a concrete numerical scenario that can be targeted by the minimal de-risking experiment described in Section 7(c).

7. Discussion

7.1 Limitations

Assumption	Description	Sensitivity
Linear thermal resistance $R_{\text{th}} = 5 \text{ K W}^{-1}$	Assumes constant resistance across the relevant power range.	Lower actual R_{th} tightens the bound.
Fixed correction frequency $f = 1 \text{ MHz}$	Holds frequency independent of temperature.	Higher error rates may demand faster cycles, worsening the loop.
Per-cycle energy $E_{\text{c}} = 10^{-12} \text{ J}$	Optimistic low-power cryogenic control.	Larger E_{c} directly reduces N_{max} .
Cooling power $P_{\text{c}} = 1 \text{ W}$ at 20 K	Typical dilution-refrigerator specification.	Higher P_{c} raises N_{max} proportionally.
Surface-code threshold $p_{\text{th}} = 10^{-2}$	Standard literature value.	Higher threshold would increase ΔT_{crit} .
Arrhenius coefficient $\alpha = 0.1 \text{ K}^{-1}$	Chosen to reflect moderate temperature sensitivity.	Larger α lowers ΔT_{crit} .
Heat-capacity parameters c and v	Material-specific values.	Reducing c or v lowers ΔT proportionally.
Activation energy $\Delta E = 5 \times 10^{-20} \text{ J}$	Representative quasiparticle energy.	Larger ΔE increases p exponentially.
No additional cooling beyond the refrigerator	Assumes all heat must be removed by the base stage.	Advanced on-chip cooling could mitigate ΔT .

If any of these assumptions are substantially violated—e.g., if per-qubit dissipation can be reduced by two orders of magnitude—the temperature rise would drop to $\sim 7.5 \text{ K}$, still far above the millikelvin regime but potentially manageable with advanced cryogenics. Therefore, the result is highly sensitive to per-qubit dissipation.

7.2 Failure modes and falsifiability

The central claim—that the feedback loop yields a temperature rise independent of system size and pushes error rates above threshold—would be falsified by any empirical observation showing a stable operating temperature at millikelvin levels while scaling N to 10^6 qubits without a corresponding increase in error probability. Specifically, a measurement of the processor temperature after a full QEC cycle that remains within $< 1 \text{ K}$ for $N \geq 10^6$ would contradict the derived $\Delta T = 750 \text{ K}$.

7.3 Open questions

- Can alternative qubit modalities (e.g., topologically protected anyons) achieve per-cycle dissipation orders of magnitude lower than assumed?
- How does the inclusion of active on-chip cooling loops modify the fixed-point analysis?
- What is the impact of non-thermal error sources (e.g., flux noise) that do not follow an Arrhenius law?

8. Practical Implications and Engineering Pathways

(a) Intended users and design impact

The primary users are **quantum-hardware architects** and **cryogenic system engineers**. If the derived limits hold, designers must prioritize ultra-low-power control electronics, minimise wiring thermal resistance, and possibly operate at temperatures below 10 K to keep ΔT well under the critical value. The analysis also suggests that scaling solely via software-based QEC is insufficient; hardware-level protection schemes (e.g., topological anyons) may be required.

(b) Technology readiness

The analysis is purely analytical and has not undergone peer review; it therefore resides at **TRL2** (technology concept formulation). As an analytical study it cannot exceed **TRL3**.

(c) Minimal de-risking experiment

A **single-chip prototype** comprising a surface-code patch of distance $d = 5$ with integrated cryogenic control electronics can be operated while varying the injected per-qubit power P_q . Measuring the qubit temperature and logical error rate as functions of P_q would test the predicted exponential relation $p(T) = p_0 \exp[\alpha(T - T_0)]$. Observation of a temperature rise that drives p above the threshold at the predicted $P_q^{\max}$ would falsify the model.

(d) Engineering constraints

- **Materials:** Low-thermal-conductivity wiring (e.g., NbTi coax) to keep R_{th} high.
- **Tolerances:** Per-cycle energy E_c must be controlled to within an order of magnitude of 10^{-12} J.
- **Components:** Cryogenic amplifiers with sub-nanowatt dissipation.
- **Data:** Real-time temperature monitoring at the qubit chip.
- **Compute:** Fast feedback loops that do not increase f beyond the assumed 1 MHz.
- **Energy:** Cooling power budget of at most 1 W at 20 K.
- **Cost drivers:** High-purity superconducting materials and custom low-power control ASICs.
- **Scale-up:** Wiring density must not degrade R_{th} as N grows.

These constraints are directly derived from the thermodynamic budget in Sections 3 and 4.

(e) Projected cost/performance bounds (projection)

Assuming a breakthrough that reduces per-qubit control ASIC power to **US\$10M** for ASICs plus **\$1M** for a refrigerator capable of 1 W cooling at 20 K. These figures are **projections** based on the assumptions listed in Sections 3 and 4. ###** (requiring per-qubit power far beyond current capabilities) or to a temperature increase of ~ 750 K that is independent of system size. Both scenarios drive the physical error probability toward or beyond the fault-tolerance threshold, establishing a **thermodynamic feedback limit** on the scalability of software-based quantum error correction. The result highlights the necessity of ultra-low-power control, high-thermal-resistance wiring, and possibly hardware-level protection mechanisms to achieve large-scale fault-tolerant quantum computing. Future work should explore hardware innovations that reduce per-cycle dissipation and investigate adaptive correction schedules that mitigate the positive feedback.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.33.9-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by 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] Quantum Error Correction. arXiv:1910.03672. <https://arxiv.org/abs/1910.03672> [2] Homomorphic Quantum Error Correction. arXiv:2605.25692. <https://arxiv.org/abs/2605.25692> [3] Deep Quantum Error Correction. arXiv:2301.11930. <https://arxiv.org/abs/2301.11930> [4] Motional Quantum Error Correction. arXiv:quant-ph/9811011. <https://arxiv.org/abs/quant-ph/9811011> [5] Perturbative quantum error correction. arXiv:1102.3809. <https://arxiv.org/abs/1102.3809> [6] Catalytic quantum error correction. arXiv:quant-ph/0608027. <https://arxiv.org/abs/quant-ph/0608027> [7] Continuous quantum error correction. arXiv:quant-ph/0501049. <https://arxiv.org/abs/quant-ph/0501049> [8] Tracking Quantum Error Correction. arXiv:1804.04509. <https://arxiv.org/abs/1804.04509> [9] Nonlinear quantum error correction. arXiv:2112.01858. <https://arxiv.org/abs/2112.01858> [10] Approximate quantum error correction. arXiv:quant-ph/0112106. <https://arxiv.org/abs/quant-ph/0112106> [11] Operator quantum error correction. arXiv:quant-ph/0504189. <https://arxiv.org/abs/quant-ph/0504189>