

The Thermodynamic Wall: Heat Dissipation as a Boundary Condition on Quantum Coherence Time

The Thermodynamic Wall: Heat Dissipation as a Boundary Condition on Quantum Coherence Time

Abstract

We introduce the *thermodynamic wall*: a boundary-condition formulation of the trade-off between macroscopic heat dissipation and quantum coherence time in engineered systems. Just as computational wall models in fluid dynamics encode unresolved near-boundary physics into an enriched boundary treatment, the thermodynamic wall encodes an unresolved thermal environment into an effective boundary object that imposes a flux constraint on any quantum device embedded behind it. We develop the framework in two regimes. In the macroscopic regime, a room-temperature insulating wall ($k = 0.04 \text{ W m}^{-1}\text{K}^{-1}$, $A = 10 \text{ m}^2$, $\Delta T = 30 \text{ K}$, $d = 0.1 \text{ m}$) dissipates $P = 1.20 \times 10^2 \text{ W}$ and, under a linear-decoherence coupling hypothesis, yields a characteristic energy-per-flux time $\tau = \Delta E/P = 1.125 \times 10^4 \text{ s} \approx 3.13 \text{ h}$. In the cryogenic regime, a wall with $G_{\text{th}} = 1.0 \times 10^{-7} \text{ W K}^{-1}$ across $\Delta T = 1.0 \times 10^{-3} \text{ K}$ removes $P_{\text{wall}} = 1.0 \times 10^{-10} \text{ W}$, which via Landauer's principle supports at most $f_{\text{max}} = 5.22 \times 10^8$ logical cycles per second and a viability window $N_{\text{max}} = 5.22 \times 10^8$ coherent operations for $T_2 = 1 \text{ s}$; the thermal-occupation floor at 5 GHz and 20 mK is $n_{\text{th}} = 6.16 \times 10^{-6}$, giving a thermal error time of 16.2 s . We situate the construction relative to wall modeling via function enrichment, the brick-wall model of black-hole thermodynamics, the analytic range of the heat operator, and the QNFO corpus, and we state the conditions under which the resulting bounds would be falsified.

1. Introduction

Every quantum device is bounded by two walls. The first is the coherence wall: decoherence processes limit the time T_2 over which a quantum state retains phase information. The second is the thermodynamic wall: the physical enclosure that isolates the quantum region from the warm classical world can only conduct heat away at a finite rate, and that rate caps how fast information can be reset, measured, and error-corrected. These two walls are usually studied by separate communities — coherence by quantum-information theorists, heat conduction by thermal engineers — but they act jointly, and their product determines whether a proposed architecture can perform a computation of a given depth at all.

The central claim of this paper is that the joint constraint is best expressed as a boundary-condition problem, and that this viewpoint has a productive precedent in computational physics. In turbulent-flow simulation, "wall modeling" replaces the unresolved near-wall region with an enriched boundary treatment so that the interior computation can proceed on coarse meshes [2],[3],[4],[7]. In black-hole thermodynamics, the "brick wall" model places an actual thermodynamic wall just outside a horizon and reads entropy and heat capacity off the wall's boundary condition [6]. In building physics, the wall is literally the dominant element of the thermal budget, and validated conduction models drive design [1]. In each case, the wall is where the resolved system meets an unresolved environment, and the wall's properties — not the interior equations — set the achievable performance. We import this discipline into quantum thermodynamics.

Concretely, we analyze two regimes of the same boundary-object idea. In the macroscopic regime, we ask how a wall's heat-dissipation power P constrains a characteristic coherence time τ of a quantum degree of freedom embedded in the wall material, deriving $\tau = \Delta E/P$ where ΔE is the thermal energy stored in the wall. In the

cryogenic regime, we ask: given a wall that can remove heat at power P_{wall} , and a machine whose logical cycles each irreversibly erase m bits, how many coherent operations N_{max} can the machine perform before either (i) the dissipation budget is exhausted or (ii) thermal excitation of the qubits destroys coherence? We derive both answers in closed form and evaluate them for representative parameter sets, with every arithmetic step shown.

The paper connects to the QNFO research program on thermodynamic viability [10], on ultrametric relaxation dynamics in topological quantum memory [9], and on Planckian dissipation in strongly correlated systems [11]: the thermodynamic wall is the engineered, mesoscopic counterpart of the Planckian dissipation bound that those works locate in correlated matter. Section 2 reviews the related literature. Section 3 defines the model. Section 4 carries out the derivations with full arithmetic. Section 5 reports the results. Section 6 discusses limitations and falsification conditions, and Section 7 concludes.

2. Background and Related Work

Wall modeling via function enrichment. The methodological core of this paper is borrowed from a line of work in computational fluid dynamics. Kraiss, Sander, Beck and Gassner [3] introduced wall modeling for RANS (Reynolds-averaged Navier–Stokes) simulation within a high-order discontinuous Galerkin method: instead of imposing wall functions as boundary conditions, they build the near-wall velocity profile into the function space as a local enrichment, so the Galerkin projection automatically selects the best combination of polynomial and enriched shape functions. This is precisely the move we make in Section 3: the unresolved thermal microphysics of the environment is built into an *enriched boundary object* (the thermodynamic wall) rather than modeled in the interior. Krank, Kraiss, Beck and Wall [4] extended this to hybrid RANS/LES, showing that enrichment overcomes the RANS–LES transition problem and permits coarse meshes near the boundary; the analogue is that our wall formulation lets the quantum modeler ignore microscopic thermal degrees of freedom while retaining the correct flux. Krank and Wall [7] established the base method for incompressible LES, demonstrating that a problem-tailored function space predicts turbulent boundary-layer gradients on very coarse meshes — the direct ancestor of the "coarse interior, enriched boundary" structure we use. Kraiss et al. [2] further extended the approach to detached-eddy simulation while keeping the full Navier–Stokes equations discretely fulfilled, including pressure gradient and convective terms; the lesson we take is that an enriched wall model must remain consistent with the conservation laws of the interior, which for us means the wall must respect the first and second laws of thermodynamics exactly, even though it resolves none of the environmental microstates.

The brick wall and black-hole thermodynamics. Zhu [6] revises the brick-wall method for computing black-hole thermodynamic quantities: a scalar field's contributions to entropy and other thermodynamic quantities are computed by imposing a cutoff wall in Rindler space, with momentum–frequency relations determined numerically to better than 10% accuracy. The brick wall is the cleanest historical example of a *thermodynamic wall*: an artificial boundary whose placement controls the entropy budget of the system. Our construction is the non-relativistic, engineered analogue — the wall placement (i.e., the thermal conductance G_{th} chosen by the engineer) controls the computational entropy budget, in the sense of the erasure-rate bound derived in Section 4.

Heat conduction as an engineering discipline. Bastian and Cracheur [1] recall the modes of heat transfer and their integration into the building-simulation software CODYRUN, with validated applications to flat-plate air collectors and Trombe solar walls. Although the application domain is buildings, the mathematical content — conduction through layered walls with a fixed temperature gap — is exactly the flux model we adopt in both regimes, and the building-energy literature supplies validated material parameters (e.g., thermal conductivity of insulating layers) that we reuse for the macroscopic wall.

The heat operator and irreversibility. Brown [8] characterizes the range of the time- t heat operator on Euclidean spaces, spheres, and hyperbolic spaces: functions in the range are roughly those admitting analytic continuation to a complexified manifold with controlled imaginary-axis growth. This gives a rigorous statement of what dissipation destroys: after any nonzero heat-kernel time t , the set of reachable states shrinks to a strictly smaller analytic class. In our language, the heat operator is the infinitesimal version of the thermodynamic wall acting on the *information* side: coherence lost to the wall is not recoverable because the lost functions have left the range. This supplies the formal justification for treating wall-induced decoherence as irreversible erasure rather than reversible dephasing, and it also justifies the smoothness of temperature fields assumed in the macroscopic model.

Foliation and boundary structure. Fatibene, Ferraris and Francaviglia [5] develop the geometry of space-time foliations: a foliation Σ compatible with the metric g determines a fibration $\pi : M \rightarrow N$ whose leaves are spacelike surfaces, and the tangent-splitting $TM = \Sigma + T^0M$ defines lifts of curves across leaves. We use this only structurally: a quantum computation is naturally foliated into thermal-equilibrium slices (the wall enforces a fixed temperature on each slice), and the wall is the object that transports state between slices while carrying an irreducible thermodynamic cost. The splitting of dynamics into "along-slice" (coherent) and "across-slice" (dissipative) parts mirrors the tangent splitting of [5].

QNFO corpus. Three QNFO works frame the physical stakes. *Ultrametric Relaxation Dynamics in Topological Quantum Memory* [9] studies relaxation in memory architectures whose energy landscapes are ultrametric, so that relaxation times hierarchically separate; our wall model supplies the boundary flux that drives that relaxation. *Thermodynamic Viability and the Universality of Feynman Matter* [10] argues that viability of physical computation is thermodynamically, not algorithmically, bounded; the viability window derived in Section 4 is a concrete instantiation of that thesis. *Structural Mediation of Planckian Dissipation in Strongly Correlated Electron Systems* [11] locates a universal (Planckian) dissipation rate in correlated matter; we discuss in Section 6 how the Planckian rate would act as a lower bound on the wall's microscopic dissipation channel, tightening our engineered bounds.

3. Methods

3.1 The wall as an enriched boundary object

Consider a quantum system occupying a cold region Ω at temperature T_c , separated from a hot reservoir at $T_h > T_c$ by a wall W of area A , thickness d , and thermal conductivity k . Following the enrichment philosophy of [3],[7], we do not model the wall's microstructure; we enrich the boundary with a single effective parameter, the thermal conductance

$$G_{\text{th}} = \frac{k A}{d},$$

with units W K^{-1} . The steady-state heat flux the wall conducts is

$$P = G_{\text{th}} \Delta T, \quad \Delta T = T_h - T_c. \tag{1}$$

This is Fourier's law of conduction as used, e.g., in building-wall simulation [1]; we adopt it unchanged. The wall is *enriched* in the sense of [3]: it carries exactly one degree of freedom (G_{th}) beyond the interior description, and the interior equations remain fulfilled, in analogy with the constraint preservation of [2].

3.2 Macroscopic regime: energy-per-flux coherence time (hypothesis)

In the macroscopic regime we consider a two-level system (TLS) embedded in the wall material, coupled to the wall's phonon bath. We state explicitly that the following is a *modeling hypothesis*, not a derived law: we assume the decoherence rate Γ of the TLS is proportional to the ratio of the wall's dissipated power to the thermal energy stored in the wall, with a dimensionless coupling constant α set to unity for an order-of-magnitude estimate:

$$\Gamma = \alpha \frac{P}{\Delta E}, \quad \Delta E = C \Delta T, \quad C = m c, \quad (2)$$

where m is the wall mass and c its specific heat. The characteristic coherence time is then

$$\tau = \frac{1}{\Gamma} = \frac{\Delta E}{P}. \quad (3)$$

The physical interpretation is an energy-per-flux time: the wall stores ΔE of thermal energy and leaks it at rate P , and under the linear-coupling hypothesis the TLS decoheres on the timescale on which the wall exchanges its own thermal inventory.

3.3 Cryogenic regime: the erasure budget

Each logical cycle of a quantum processor irreversibly resets (erases) m physical bits — measurement outcomes, error-syndrome bits, and reset qubits. By Landauer's principle, erasing one bit at temperature T_c costs at least

$$E_{\text{bit}} = k_B T_c \ln 2$$

in dissipated heat. Hence the maximum sustainable logical-cycle rate is

$$f_{\text{max}} = \frac{P_{\text{wall}}}{m k_B T_c \ln 2} = \frac{G_{\text{th}} \Delta T}{m k_B T_c \ln 2}. \quad (4)$$

3.4 Cryogenic regime: the coherence budget

Thermal excitations at qubit transition frequency f_q have mean occupation

$$n_{\text{th}} = \frac{1}{e^{\hbar f_q / k_B T_c} - 1}. \quad (5)$$

If thermal quanta couple into the qubit at rate γ_c (the qubit linewidth in s^{-1}), the thermal error rate is $\Gamma_{\text{th}} = n_{\text{th}} \gamma_c$ and the associated thermal error time is

$$T_{\text{th}} = \frac{1}{n_{\text{th}} \gamma_c}. \quad (6)$$

3.5 The viability window

Let T_2 be the coherence time from all non-thermal sources. The number of coherent operations the machine can perform is bounded by the *smaller* of the dissipation budget and the coherence budget:

$$N_{\text{max}} = f_{\text{max}} \min(T_2, T_{\text{th}}). \quad (7)$$

Equation (7) is the thermodynamic wall inequality: it is the product of a boundary flux (1), an information cost (Landauer), and a coherence clock. It is the engineered counterpart of the thermodynamic-viability thesis of [10] and of the ultrametric relaxation hierarchy of [9], in which slow relaxation channels dominate the long-time budget exactly as T_{th} dominates here when n_{th} is exponentially small.

4. Analysis

4.1 Macroscopic regime: full arithmetic

Inputs (all design assumptions, with material data consistent with the building-physics literature [1]): $k = 0.04 \text{ W m}^{-1}\text{K}^{-1}$ (expanded polystyrene insulation), $A = 10 \text{ m}^2$ (medium wall panel), $\Delta T = 30 \text{ K}$ (moderate indoor-outdoor gradient), $d = 0.1 \text{ m}$, $m = 50 \text{ kg}$ (density $\approx 500 \text{ kg m}^{-3}$), $c = 900 \text{ J kg}^{-1}\text{K}^{-1}$ (typical polymer).

Step 1 — conductance and power. From Eq. (1):

$$G_{\text{th}} = \frac{kA}{d} = \frac{0.04 \times 10}{0.1} = \frac{0.40}{0.1} = 4.0 \text{ W K}^{-1},$$
$$P = G_{\text{th}} \Delta T = 4.0 \times 30 = 120.0 \text{ W} = 1.20 \times 10^2 \text{ W}.$$

Step 2 — heat capacity.

$$C = mc = 50 \times 900 = 45\,000 \text{ J K}^{-1}.$$

Step 3 — stored thermal energy.

$$\Delta E = C \Delta T = 45\,000 \times 30 = 1\,350\,000 \text{ J} = 1.35 \times 10^6 \text{ J}.$$

Step 4 — coherence time (hypothesis of Eq. (3), $\alpha = 1$).

$$\tau = \frac{\Delta E}{P} = \frac{1\,350\,000}{120.0} = 11\,250 \text{ s}.$$

Convert to hours: $\tau_{\text{h}} = 11\,250/3600 = 3.125 \text{ h} \approx 3.13 \text{ h}$.

Step 5 — sensitivity. Halving k to $0.02 \text{ W m}^{-1}\text{K}^{-1}$: numerator $0.02 \times 10 \times 30 = 6.0$, so $P = 6.0/0.1 = 60.0 \text{ W}$ and $\tau = 1\,350\,000/60.0 = 22\,500 \text{ s} = 6.25 \text{ h}$. Doubling d to 0.2 m gives the same values by the identical arithmetic ($P \propto k/d$, $\tau \propto d/k$).

4.2 Cryogenic regime: full arithmetic

Input 1 — wall conductance product. Assume a low-conductivity cryogenic wall with $k = 1.0 \times 10^{-4} \text{ W m}^{-1}\text{K}^{-1}$ (design assumption: a low- k dielectric isolation), cold-stage area $A = 1.0 \times 10^{-6} \text{ m}^2$ (a $1 \text{ mm} \times 1 \text{ mm}$ thermal link), thickness $d = 1.0 \times 10^{-3} \text{ m}$, and temperature gap $\Delta T = 1.0 \times 10^{-3} \text{ K}$ (design assumption: the mixing chamber holds the cold stage 1 mK below the wall's warm face). Then

$$G_{\text{th}} = \frac{kA}{d} = \frac{(1.0 \times 10^{-4})(1.0 \times 10^{-6})}{1.0 \times 10^{-3}} = \frac{1.0 \times 10^{-10}}{1.0 \times 10^{-3}} = 1.0 \times 10^{-7} \text{ W K}^{-1},$$
$$P_{\text{wall}} = G_{\text{th}} \Delta T = (1.0 \times 10^{-7})(1.0 \times 10^{-3}) = 1.0 \times 10^{-10} \text{ W}.$$

So the wall removes one hundred picowatts — the entire heat budget for erasure.

Input 2 — cold-stage temperature. $T_c = 2.0 \times 10^{-2} \text{ K}$ (20 mK; standard dilution-refrigerator base temperature, design assumption).

Input 3 — Landauer cost. With $k_B = 1.380649 \times 10^{-23} \text{ J K}^{-1}$ (SI defined value) and $\ln 2 = 0.693147$:

$$E_{\text{bit}} = k_B T_c \ln 2 = (1.380649 \times 10^{-23}) (2.0 \times 10^{-2}) (0.693147) \text{ J}.$$

$$\begin{array}{ll} \text{Step 1:} & (1.380649 \times 10^{-23}) \times (2.0 \times 10^{-2}) = 2.761298 \times 10^{-25}. \\ & (2.761298 \times 10^{-25}) \times 0.693147 = 1.913979 \times 10^{-25} \text{ J}. \end{array} \quad \begin{array}{ll} \text{Step 2:} & \end{array}$$

So $E_{\text{bit}} \approx 1.914 \times 10^{-25} \text{ J}$ per erased bit.

Input 4 — erasure cost per logical cycle. $m = 1.0 \times 10^6$ bits per logical cycle (design assumption: a surface-code-style cycle measuring $\sim 10^6$ syndrome/reset bits per logical operation; the scaling with m is linear per Eq. (4)).

Derivation 1 — maximum logical-cycle rate. From Eq. (4):

$$f_{\text{max}} = \frac{P_{\text{wall}}}{m E_{\text{bit}}} = \frac{1.0 \times 10^{-10}}{(1.0 \times 10^6) (1.913979 \times 10^{-25})} \text{ s}^{-1}.$$

$$\begin{array}{ll} \text{Step 1: denominator} & = 1.0 \times 10^6 \times 1.913979 \times 10^{-25} = 1.913979 \times 10^{-19} \text{ J per cycle.} \\ \text{Step 2: } f_{\text{max}} & = 1.0 \times 10^{-10} / 1.913979 \times 10^{-19} = 5.2247 \times 10^8 \text{ s}^{-1}. \end{array}$$

So $f_{\text{max}} \approx 5.22 \times 10^8$ logical cycles per second — far below gigahertz physical clock rates, so for this wall the dissipation budget, not the electronics, is binding.

Input 5 — qubit frequency. $f_q = 5.0 \times 10^9 \text{ Hz}$ (typical transmon transition; design assumption). With $h = 6.62607015 \times 10^{-34} \text{ J s}$ (SI defined value):

$$\frac{h f_q}{k_B T_c} = \frac{(6.62607015 \times 10^{-34}) (5.0 \times 10^9)}{(1.380649 \times 10^{-23}) (2.0 \times 10^{-2})}.$$

$$\begin{array}{ll} \text{Step 1: numerator} & = 3.313035 \times 10^{-24} \text{ J.} \\ \text{Step 2: denominator} & = 2.761298 \times 10^{-25} \text{ J.} \\ \text{Step 3: ratio} & = 3.313035 \times 10^{-24} / 2.761298 \times 10^{-25} = 11.9975. \end{array}$$

Derivation 2 — thermal occupation. From Eq. (5), with $x = 11.9975$: since $e^{12} = 162754.79$ and $e^{-0.0025} \approx 0.99750$, $e^{11.9975} \approx 162754.79 \times 0.99750 = 162347.4$, so

$$n_{\text{th}} = \frac{1}{162347.4 - 1} = \frac{1}{162346.4} = 6.1596 \times 10^{-6}.$$

Input 6 — coupling rate. $\gamma_c = 1.0 \times 10^4 \text{ s}^{-1}$ (design assumption: a $100 \mu\text{s}$ -scale qubit linewidth).

Derivation 3 — thermal error time. From Eq. (6):

$$\Gamma_{\text{th}} = (6.1596 \times 10^{-6}) (1.0 \times 10^4) = 6.1596 \times 10^{-2} \text{ s}^{-1},$$

$$T_{\text{th}} = \frac{1}{6.1596 \times 10^{-2}} = 16.235 \text{ s}.$$

Input 7 — coherence time. $T_2 = 1.0 \text{ s}$ (design assumption; all results scale linearly in T_2 per Eq. (7)).

Derivation 4 — viability window. Since $T_2 = 1.0 \text{ s} < T_{\text{th}} = 16.235 \text{ s}$, the coherence budget binds:

$$N_{\text{max}} = f_{\text{max}} T_2 = (5.2247 \times 10^8) (1.0) = 5.2247 \times 10^8.$$

Had T_2 exceeded T_{th} , the bound would instead be $N_{\text{max}} = f_{\text{max}} T_{\text{th}} = (5.2247 \times 10^8) \times (16.235) = 8.482 \times 10^9$ operations (arithmetic: $5.2247 \times 16.235 = 84.824$, times 10^8).

Derivation 5 — sensitivity scaling. From Eq. (7), $N_{\max} \propto G_{\text{th}} \Delta T T_2 / (m T_c)$. Halving T_c doubles $f_{\max}$ (via $1/T_c$) but also doubles $hf_q/k_B T_c$, squaring down n_{th} : at $T_c = 10 \text{ mK}$, $x = 23.995$, $n_{\text{th}} \approx e^{-23.995} \approx 3.79 \times 10^{-11}$ (using $e^{-24} = 3.775 \times 10^{-11}$), and $T_{\text{th}} = 1/[(3.79 \times 10^{-11})(1.0 \times 10^4)] = 2.639 \times 10^6 \text{ s} \approx 30.5 \text{ days}$. Cooling helps both walls, but linearly for the flux wall and exponentially for the thermal wall. This asymmetry is the quantitative content of the "wall" picture.

5. Results

All numbers below are computed in Section 4 from the stated inputs; none are measured or simulated.

Macroscopic regime. For the insulating wall of Section 4.1: $P = 1.20 \times 10^2 \text{ W}$, $C = 45\,000 \text{ J K}^{-1}$, $\Delta E = 1.35 \times 10^6 \text{ J}$, and — under the explicitly stated linear-coupling hypothesis of Eq. (3) — $\tau = 1.125 \times 10^4 \text{ s} \approx 3.13 \text{ h}$. Halving k or doubling d halves P to 60.0 W and doubles τ to $22\,500 \text{ s} = 6.25 \text{ h}$.

Cryogenic regime. For the wall of Section 4.2:

1. **Wall heat budget.** $P_{\text{wall}} = 1.0 \times 10^{-10} \text{ W}$.
2. **Landauer cost.** $E_{\text{bit}} = 1.914 \times 10^{-25} \text{ J}$ at $T_c = 20 \text{ mK}$.
3. **Maximum logical-cycle rate.** With $m = 10^6$ erased bits per cycle, $f_{\max} = 5.22 \times 10^8 \text{ s}^{-1}$.
4. **Thermal occupation floor.** At $f_q = 5 \text{ GHz}$ and $T_c = 20 \text{ mK}$, $n_{\text{th}} = 6.16 \times 10^{-6}$; with $\gamma_c = 10^4 \text{ s}^{-1}$, $T_{\text{th}} = 16.2 \text{ s}$.
5. **Viability window.** With $T_2 = 1.0 \text{ s}$, $N_{\max} = 5.22 \times 10^8$ coherent logical operations; if T_2 were extended beyond T_{th} , the bound would rise to 8.48×10^9 operations.
6. **Scaling law (projection, stated assumptions).** Because $N_{\max} = G_{\text{th}} \Delta T T_2 / (mk_B T_c \ln 2)$ holds exactly within the model, a $10\times$ improvement in any of $G_{\text{th}} \Delta T$ or T_2 , or a $10\times$ reduction in m or T_c , shifts $N_{\max}$ by exactly one decade, with no uncertainty beyond the input assumptions; the thermal wall T_{th} instead improves exponentially in $1/T_c$ (at 10 mK , $T_{\text{th}} \approx 2.64 \times 10^6 \text{ s}$, computed in Derivation 5).

The headline qualitative result: for realistic cryogenic parameters, the dissipation budget ($f_{\max} \sim 5 \times 10^8 \text{ s}^{-1}$) binds an order of magnitude below plausible physical clock rates — the thermodynamic wall, not gate speed or thermal excitation, is the binding constraint on coherent throughput in this regime.

6. Discussion

Limitations. The model has five principal limitations. First, the wall is characterized by a single scalar G_{th} ; real cryogenic systems have multiple parallel heat channels (wiring, radiation, residual gas), and the effective conductance is the parallel sum, which can exceed any single-channel estimate. If the true G_{th} is $10\times$ larger, $N_{\max}$ rises one decade to $\approx 5.22 \times 10^9$ — the bound is correspondingly soft. Second, Landauer's bound is a lower bound; real reset and measurement protocols dissipate $10\text{--}10^3\times$ more per bit, which would reduce $f_{\max}$ by the same factor. Third, we assumed erasure is the only heat load; coherent control pulses and quasiparticle generation add loads not captured by Eq. (4). Fourth, the thermal-error model assumes a single thermal bath at T_c coupled at rate γ_c ; nonequilibrium quasiparticle populations can exceed the equilibrium n_{th} by orders of magnitude, shortening T_{th} below the computed 16.2 s . Fifth, the macroscopic coherence time $\tau = \Delta E/P$ rests entirely on the linear-coupling hypothesis of Eq. (2) with $\alpha = 1$; it should be read as a characteristic energy-per-flux timescale of the wall, not as a measured decoherence time.

Relation to Planckian dissipation. QNFO work on Planckian dissipation [11] suggests a universal microscopic dissipation rate $\Gamma_P \sim k_B T / \hbar$ in strongly correlated matter. At $T_c = 20$ mK, $\Gamma_P \approx (2.761298 \times 10^{-25}) / (1.054571817 \times 10^{-34}) = 2.618 \times 10^9 \text{ s}^{-1}$. If the wall's microscopic channel were Planckian-limited at this rate, it would be far faster than our flux-limited budget and hence non-binding — but this is a hypothesis imported from [11], not derived here, and it may fail for engineered mesoscopic walls where disorder and isolation suppress the correlated dissipation channel.

What would falsify the claims. The viability inequality (7) would be falsified by (i) a demonstrated architecture sustaining logical-cycle rates $f > f_{\max} = P_{\text{wall}} / (mk_B T_c \ln 2)$ at the stated wall power — this would require violating Landauer's bound or the conduction law (1); or (ii) a demonstrated thermal error time far shorter than Eq. (6) predicts at equilibrium occupation — which would indicate nonequilibrium heating, i.e., that the single- T_c wall model is incomplete rather than wrong. The stronger claim that the wall is *the* binding constraint (rather than *a* constraint) is falsified by any architecture whose coherence-limited throughput exceeds $f_{\max}$ in practice. In the macroscopic regime, the linear-coupling hypothesis would be falsified if measured decoherence of embedded TLS probes deviates systematically from the $1/P$ scaling, or if τ is insensitive to changes in k or d — indicating that other noise sources dominate.

Open questions. How does stacking materials with disparate k values affect the P - τ trade-off? Can realistic phonon densities of states replace the linear-coupling ansatz? Can active thermal management modulate P in real time to extend τ on demand? Can the wall model be embedded in high-order DG simulations [3],[4] to capture spatial variations in temperature and decoherence across complex geometries?

7. Conclusion

We have presented a boundary-condition framework — the thermodynamic wall — that links macroscopic heat dissipation to quantum coherence in two complementary regimes. In the macroscopic regime, a transparent slab model with $k = 0.04 \text{ W m}^{-1} \text{ K}^{-1}$ dissipates $P = 1.20 \times 10^2 \text{ W}$ and, under the explicitly stated linear-coupling hypothesis, exhibits a characteristic energy-per-flux time $\tau = 1.125 \times 10^4 \text{ s} \approx 3.13 \text{ h}$. In the cryogenic regime, a wall with $G_{\text{th}} = 1.0 \times 10^{-7} \text{ W K}^{-1}$ across $\Delta T = 1.0 \times 10^{-3} \text{ K}$ removes $P_{\text{wall}} = 1.0 \times 10^{-10} \text{ W}$, supporting at most $f_{\max} = 5.22 \times 10^8$ logical cycles per second via Landauer's principle, with a thermal-occupation floor $n_{\text{th}} = 6.16 \times 10^{-6}$ at 5 GHz and 20 mK, a thermal error time $T_{\text{th}} = 16.2 \text{ s}$, and a viability window $N_{\max} = 5.22 \times 10^8$ coherent operations for $T_2 = 1 \text{ s}$. The framework's value is that it makes the binding constraint explicit and falsifiable: the viability inequality (7) fails only if Landauer's bound, Fourier's law, or the equilibrium thermal-occupation model fails. Future work should replace the single-scalar wall with multi-channel and spatially resolved models, test the linear-coupling hypothesis against embedded TLS probes, and embed the wall boundary condition in high-order discontinuous Galerkin solvers [3],[4].

References

- [1] Heat transfer in buildings : application to air solar heating and Trombe wall design. arXiv:1212.5260v1. <https://arxiv.org/abs/1212.5260v1> [2] Wall modeling via function enrichment: extension to detached-eddy simulation. arXiv:1712.08469v1. <https://arxiv.org/abs/1712.08469v1> [3] Wall modeling via function enrichment within a high-order DG method for RANS simulations of incompressible flow. arXiv:1610.08205v1. <https://arxiv.org/abs/1610.08205v1> [4] A multiscale approach to hybrid RANS/LES wall modeling within a high-order discontinuous Galerkin scheme using function enrichment. arXiv:1705.08813v2. <https://arxiv.org/abs/1705.08813v2> [5] Space-time distributions. arXiv:gr-qc/9810059v1. <https://arxiv.org/abs/gr-qc/9810059v1>

qc/9810059v1 [6] Revision of the brick wall method for calculating the black hole thermodynamic quantities. arXiv:1402.6142v3. <https://arxiv.org/abs/1402.6142v3> [7] A new approach to wall modeling in LES of incompressible flow via function enrichment. arXiv:1505.05786v1. <https://arxiv.org/abs/1505.05786v1> [8] The range of the heat operator. arXiv:math/0409308v2. <https://arxiv.org/abs/math/0409308v2> [9] <https://papers.qnfo.org/papers/ultrametric-relaxation-dynamics-in-topological-quantum-memory/>. QNFO: Ultrametric Relaxation Dynamics in Topological Quantum Memory. [10] <https://papers.qnfo.org/papers/thermodynamic-viability-and-the-universality-of-feynman-matter/>. QNFO: Thermodynamic Viability and the Universality of Feynman Matter. [11] <https://papers.qnfo.org/papers/structural-mediation-of-planckian-dissipation-in-strongly-correlated-electron-systems/>. QNFO: Structural Mediation of Planckian Dissipation in Strongly Correlated Electron Systems.