

Passive Error Rejection via Cryo-Fluidic Topological Phases: A Feasibility Analysis of the Topological Matter Engine

Passive Error Rejection via Cryo-Fluidic Topological Phases: A Feasibility Analysis of the Topological Matter Engine

Abstract

Active quantum error correction imposes a large and growing overhead of physical qubits, control electronics, and entropy-exporting measurement cycles on every scalable quantum processor. This paper examines a proposed alternative: a "Topological Matter Engine" in which errors are rejected passively, at the hardware level, by engineering a topological phase in a cryogenic fluidic/solid-state stack (diamond, SiN, thin-film lithium niobate, InGaAs) driven by a 3.5 THz intracavity difference-frequency-generation (DFG) field that acts as a synthetic fourth dimension. The specific gap addressed is that no published analysis derives the effective 4D Hamiltonian of such a device, identifies its topological invariant, or quantifies the passive error rejection it would deliver under realistic thermal noise. We derive a 4D two-band lattice Hamiltonian from the DFG drive, show that its gap closings occur at synthetic-dimension masses $m_0 \in \{0, \pm 2, \pm 4\}J$, and infer a second Chern number $C_2 = 1$ on the principal phase. With the gap fixed by the drive, $\Delta = 14.48$ meV, we compute thermal excitation probabilities of 5.8×10^{-19} at 4 K and $\approx 4 \times 10^{-730}$ at 100 mK, and we derive a maximum operating temperature of 36.5 K for thermally quiet operation (defined as excitation probability below 1%). The principal result is a quantitative condition sheet: passive rejection is thermodynamically robust once the gap is open, but the open question shifts entirely to whether DFG-derived synthetic-dimension couplings of order $J/\hbar = 1.75$ THz can be realized and kept coherent in the proposed material stack. This converts an architectural claim into a bounded, testable engineering target.

1. Introduction

Every group building a scalable quantum processor faces the same wall: the number of physical qubits required per logical qubit under active error correction grows with the physical error rate, and every correction cycle consumes energy, control bandwidth, and cryogenic cooling budget while exporting entropy. The control stack, not the qubit, dominates the cost of scaling. A hardware-level mechanism that rejects errors passively — that is, suppresses error propagation through the structure of the Hamiltonian itself, without measurement and feedback — would directly reduce the required qubit count and the thermodynamic load on the cryostat. If such a mechanism is internally consistent and thermally robust, the field gains a concrete device concept with a bounded engineering target; if its required couplings are unreachable, the concept is falsified cheaply, before any hardware is built.

This paper addresses that question for a specific proposed device. The "Topological Matter Engine" (TME) is a cryogenic stack — diamond, silicon nitride (SiN), thin-film lithium niobate (TFLN), and InGaAs — in which an intracavity difference-frequency-generation drive at 3.5 THz is claimed to generate a synthetic fourth dimension, yielding an effective four-dimensional Hamiltonian whose topological phase suppresses errors without any active intervention. The claim, as stated, is unverified: no Hamiltonian has been derived, no invariant identified, no rejection rate computed, and no check performed on whether the material stack can support the required couplings.

The stakes are easiest to see by analogy with rejection as a design pattern in other domains. In clinical machine learning, a model that knows when to abstain can be safer than one that always answers: a diabetic-retinopathy classifier with a confidence-based rejection rule accepted 74.50% of cases at 89.93% accuracy, rejecting 25.50% and calibrating its confidence to an Expected Calibration Error of 0.0217 [2]. Rejection there is a software decision made after inference. The Topological Matter Engine proposes the physical analogue: the device itself refuses to occupy error-prone states because those states lie across a topologically protected gap. The cost profile is the same in kind — completeness is traded for reliability — but the mechanism is moved from post-hoc analytics into the Hamiltonian.

Our approach is analytical throughout. We derive the effective 4D Hamiltonian from the DFG drive, identify the gap-closing structure and the associated topological invariant (a second Chern number in the 4D quantum Hall class), compute the passive rejection rate under thermal noise as a function of temperature, and assess the physical realizability of the required coupling strengths in the proposed stack. We find the mechanism is internally consistent in its thermal and topological structure, but that its viability rests entirely on one unproven engineering assumption, which we state explicitly and propose how to test.

Contributions.

1. A derivation of the effective 4D Hamiltonian generated by a 3.5 THz DFG drive, with the synthetic-dimension coupling J fixed by the drive energy, and the identification of the gap-closing manifold at $m_0 \in \{0, \pm 2, \pm 4\}J$ (Sections 3 and 4.2).

2. An inference of the second Chern number $C_2 = 1$ for the principal topological phase, grounded in the charge conservation of band-touching events (Section 4.3).
3. A computed passive rejection figure: thermal excitation probability $p_{\text{th}} = 5.8 \times 10^{-19}$ at $T = 4$ K and $\approx 4 \times 10^{-730}$ at $T = 100$ mK, with a derived maximum operating temperature of 36.5 K for thermally quiet operation (Section 4.4).
4. A realizability assessment identifying the single load-bearing assumption — synthetic-dimension couplings of order $J/\hbar = 1.75$ THz with coherence sufficient to resolve the gap — and the smallest experiment that would confirm or falsify it (Sections 4.5, 5, and 7).

2. Background and Related Work

The relevant literature falls into four groups: topology of gapless and gapped phases, classification frameworks for topological matter, topological codes and their excitations, and passive (as opposed to active) schemes in adjacent technologies.

Topology of semimetals and node stability. The review of topology beyond band insulators establishes the two facts our construction leans on. First, a Weyl node is stable to small perturbations regardless of symmetry: writing a general two-level Hamiltonian with splitting $\Delta E = \pm \sqrt{\delta E^2 + \psi_1^2 + \psi_2^2}$, annihilating the node requires tuning three real parameters to zero simultaneously, and a three-dimensional Brillouin zone supplies exactly three tunable momenta, so a node, once present, can only move, not disappear [1]. Second, the net charge of Weyl points obeys a conservation law — a single Weyl point cannot disappear on its own — which is the statement that band-touching events carry integer topological charge [1]. We use the first fact to argue that the gap closings of our synthetic-dimension model are structurally isolated and stable, and the second to infer that the second Chern number changes by one unit at each closing. The same review flags, in its outlook, that the physical realization of topological phases, including interacting phases, remains open [1], which is precisely the gap this paper works in.

Free-fermion symmetry-protected topological phases. The pedagogical treatment of free-fermion SPT phases defines a quantum phase of matter as a smoothly connected family of gapped ground states, with the gap required to remain non-zero in the thermodynamic limit, and gives the general free-fermion Hamiltonian $H = \sum_{ij} \left(\mathcal{H}_{ij} c_i^\dagger c_j + \Delta_{ij} c_i^\dagger c_j^\dagger + \Delta_{ij}^* c_j c_i \right)$ with hopping and pairing matrices of dimension $N \times N$ acting on $2^N \times 2^N$ many-body spaces [4]. Our engine is deliberately confined to this non-interacting, charge-conserving sector: the synthetic-dimension lattice model we derive is of the hopping form, and the protection we claim is SPT-type, meaning it persists as long as the gap stays open and the relevant symmetry is respected. This framing also supplies the caution we adopt in Section 6: protection is a property of the gapped ground-state family, and any mechanism that closes the gap destroys it.

Classification frameworks. Two recent works generalize topological classification beyond the standard periodic table. The Weyl C^* -algebra scheme describes topological phases as homotopy classes of sections of fiber bundles of pure states, notes that Kitaev's K -theoretic periodic table is confined to non-interacting systems, and builds the observable algebra from the canonical commutation relations over a phase space $\mathbb{R}^d \times \mathbb{R}^d$ [6]. The states-based approach classifies phases of systems whose observable algebra forms trivial C^* -bundles, via homotopy classes of configurations mapping a quantum parameter space into the space of pure states, and connects crossed-product C^* -algebras to condensed-matter systems [7]. These frameworks matter to us for one specific reason: a synthetic dimension is not a spatial dimension, so the topological invariant of the engine is defined on a parameter space that includes an internal, drive-derived coordinate. The bundle-of-states viewpoint [6, 7] is the natural language for that construction, and we use it in Section 3 to justify treating the DFG phase as a legitimate fourth coordinate rather than a mere modulation.

Topological codes and protected excitations. The study of the topological color code analyzes $(d - 1)$ -dimensional excitations created by transversal application of phase operators and shows they can be viewed as transparent gapped domain walls, with the transformed Hamiltonian $\hat{H} = (\overline{\mathcal{H}}|_R) H (\overline{\mathcal{H}}|_R)^\dagger$ remaining gapped because the transformation is unitary [8]. Two points transfer directly to our proposal. First, protection can be carried by Hamiltonian structure rather than by measurement cycles: the domain-wall construction acts on the Hamiltonian, not the state. Second, the same work warns that certain SPT excitations "do not exist as stable objects since they are superpositions of eigenstate excitations which would decohere immediately" [8] — a caution we apply in Section 6 to the fluidic degrees of freedom of the proposed stack, which are the most likely locus of unprotected, decohering excitations.

Passive schemes in quantum technologies. In quantum key distribution, passive detection setups are "attractive for their simplicity and stability at high protocol repetition rates," and a recent framework bounds the phase error rate for passive setups with detector imperfections and memory effects, showing that protocols can run at higher repetition rates than active schemes because detector memory no longer constrains the design [5]. The structural lesson for us is that removing an active choice (basis selection there, syndrome extraction here) removes a whole class of timing, stability, and overhead constraints — but also that imperfections and memory effects in the passive elements must be bounded before any security or reliability claim stands [5]. Our Section 4.5 performs the analogous bounding for the engine's couplings.

Rejection as a quantified design pattern. The clinical ML result cited above is worth stating precisely because it quantifies the trade our engine makes: accepting only predictions above a 70% confidence threshold yielded 89.93% accuracy at 74.50% coverage with ECE 0.0217, and the authors conclude that reliability "comes at a cost of completeness" [2]. The engine's trade is starker — it rejects error states with probability near unity but occupies a smaller state space — and our Section 4.4 computes the price in thermodynamic terms.

Passive measurement methodology. A methodological parallel comes from network science: passive monitoring of broadband subscribers over three years (March 2014 to March 2017) produced consumption measurements that avoided the unreliable self-reports of survey-based studies and were released as a verifiable anonymized dataset [3]. The analogy is exact in one respect: measuring a system passively, without the perturbation of an active probe, both avoids artifacts and constrains what can be observed. For the engine, the corresponding design principle is that error rejection should require no probe of the qubit at all — which is why we evaluate it purely through the Hamiltonian and the thermal statistics of its excitations.

3. Methods

Device concept. The proposed stack is a heterogeneous cryogenic photonic/solid-state assembly: a diamond platform (for color-center or phononic degrees of freedom), SiN photonics (low-loss waveguiding), TFLN (electro-optic and $\chi^{(2)}$ nonlinear functionality), and InGaAs (single-photon detection and infrared optoelectronics). An intracavity difference-frequency-generation process produces a field at the difference of two optical pump frequencies, specified as 3.5 THz. This THz field is not used as a signal; it is used as a synthetic dimension: the phase of the DFG field plays the role of a fourth lattice coordinate, so that the effective single-particle Hamiltonian lives on a 4D lattice rather than the 3D physical structure.

Modeling the synthetic dimension. We adopt the states-based viewpoint of [6, 7]: the observable algebra of the driven device is treated as a bundle over a quantum parameter space whose base includes the DFG phase ϕ . A periodic drive of frequency $\nu_{\text{DFG}} = 3.5$ THz supplies a compact coordinate $\phi \in [0, 2\pi)$, and hopping amplitudes between physical modes that are mediated by absorption or emission of DFG quanta acquire phases $e^{i\phi}$. Formally, the effective Hamiltonian is

$$H_{\text{eff}}(\mathbf{k}, \phi) = \sum_{\mu=1}^4 2J \Gamma_{\mu} \sin k_{\mu} + m(\mathbf{k}) \Gamma_5,$$

where Γ_{μ} are the five Dirac matrices of the 4D two-band model and $m(\mathbf{k}) = m_0 + \sum_{\mu=1}^4 \cos k_{\mu}$, with m_0 a tunable mass set by the static detuning of the stack. This is the standard 4D generalization of a Qi-Wu-Zhang-type lattice model. We take all four hopping amplitudes equal, as the symmetric design point, and we assume — this is an assumption of the model, not a derived fact — that the DFG drive mediates the fourth-direction hopping with the same strength J as the physical directions.

Fixing the energy scale. The single load-bearing quantitative input is the drive frequency $\nu_{\text{DFG}} = 3.5$ THz, taken from the device concept. We set the synthetic-direction hopping to half the drive quantum energy, $J = \hbar\nu_{\text{DFG}}/2$, so that the bulk gap $\Delta = 2J$ equals one drive quantum. All thermal and rejection figures in Section 4 follow from this single identification.

Topological invariant. The invariant of the 4D two-band model is the second Chern number

$$C_2 = \frac{1}{8\pi^2} \int_{\text{BZ}} \text{tr } \mathcal{F} \wedge \mathcal{F},$$

an integer that changes only when the bulk gap closes. We locate the gap closings analytically (Section 4.2) and infer the integer jumps using the charge conservation of band-touching events established for Weyl-type nodes [1].

Error model. We evaluate passive rejection under thermal noise only: the dominant error mechanism at cryogenic temperature is thermal excitation across the bulk gap Δ . Dephasing and drive-noise mechanisms are assessed qualitatively in Section 6 because no noise spectral densities for the proposed stack are available to us; we state this as a limitation rather than inventing numbers.

Realizability check. We convert the required coupling J into engineering units (frequency, photon energy, temperature) and compare against the operating requirements of the proposed materials, flagging which comparisons can be made from derived quantities and which remain open.

4. Analysis

All input numbers are stated with their source. Derived quantities show every arithmetic step.

4.1 Energy scale of the DFG drive

Input: $\nu_{\text{DFG}} = 3.5 \text{ THz} = 3.5 \times 10^{12} \text{ Hz}$ (device concept). Constants: $h = 6.626 \times 10^{-34} \text{ J}\cdot\text{s}$, $k_B = 1.381 \times 10^{-23} \text{ J/K}$, $c = 2.998 \times 10^8 \text{ m/s}$, $1 \text{ eV} = 1.602 \times 10^{-19} \text{ J}$.

Drive quantum energy:

$$E_\gamma = h\nu_{\text{DFG}} = (6.626 \times 10^{-34})(3.5 \times 10^{12}) = 2.319 \times 10^{-21} \text{ J}.$$

In electronvolts:

$$E_\gamma = \frac{2.319 \times 10^{-21}}{1.602 \times 10^{-19}} = 0.01448 \text{ eV} = 14.48 \text{ meV}.$$

Equivalent temperature:

$$\frac{E_\gamma}{k_B} = \frac{2.319 \times 10^{-21}}{1.381 \times 10^{-23}} = 167.9 \text{ K}.$$

Wavelength of the THz field:

$$\lambda_{\text{DFG}} = \frac{c}{\nu_{\text{DFG}}} = \frac{2.998 \times 10^8}{3.5 \times 10^{12}} = 8.57 \times 10^{-5} \text{ m} = 85.7 \text{ }\mu\text{m}.$$

So the drive quantum is 14.48 meV, thermally equivalent to 167.9 K, at a free-space wavelength of 85.7 μm .

4.2 Gap structure of the 4D model

With $d_\mu = \sin k_\mu$ ($\mu = 1, \dots, 4$) and $d_5 = m(\mathbf{k})$, the effective Hamiltonian can be written as $H_{\text{eff}} = \sum_{\mu=1}^4 2J \Gamma_\mu d_\mu + d_5 \Gamma_5$.

$$d_5 = m(\mathbf{k}) = m_0 + \sum_{\mu=1}^4 \cos k_\mu,$$

the energy eigenvalues of H_{eff} are

$$E_\pm(\mathbf{k}) = \pm \sqrt{\sum_{\mu=1}^4 \sin^2 k_\mu + \left(m_0 + \sum_{\mu=1}^4 \cos k_\mu\right)^2},$$

and the bulk gap closes exactly when both terms vanish simultaneously. The first term vanishes only at the high-symmetry momenta $k_\mu \in \{0, \pi\}$ for all four directions, where $\sin k_\mu = 0$. At those $2^4 = 16$ momenta, $\cos k_\mu = \pm 1$, so the mass term takes the values

$$m(\mathbf{k}) = m_0 + nJ, \quad n \in \{-4, -2, 0, +2, +4\},$$

where n counts the number of momenta set to π (each contributing -1) minus those set to 0 (each contributing $+1$). Defining a dimensionless mass $\tilde{m}_0 = m_0/J$, the mass term can be written as $m(\mathbf{k}) = J(\tilde{m}_0 + \sum_{\mu=1}^4 \cos k_\mu)$. The bulk gap closes when $\tilde{m}_0 \in \{0, \pm 2, \pm 4\}$, equivalently $m_0 \in \{0, \pm 2, \pm 4\}J$.

Between consecutive closings the gap is open and the system is in a gapped phase; this is the discrete set of critical masses quoted in the Abstract and Contribution 1. Each closing is a band-touching event of Weyl type in the 4D parameter space (k_1, k_2, k_3, k_4) : near a closing momentum the Hamiltonian reduces to a 5×5 real antisymmetric form whose vanishing requires all five components to be tuned to zero at once, so the closings are isolated points, structurally stable against small perturbations of the drive or detuning, by the same level-crossing argument that protects Weyl nodes in three dimensions [1].

4.3 The second Chern number of the principal phase

The second Chern number C_2 is constant on each open interval of m_0 between gap closings and jumps by an integer at each closing. The integer carried by a band-touching event is conserved: a single touching event cannot annihilate on its own, which is the momentum-space charge conservation established for Weyl points [1]. We adopt the convention that the trivial phase at $m_0 > 4J$ (all cosines positive, mass large and positive) has $C_2 = 0$, and that each crossing from $m_0 > 4$ down to $m_0 < -4$ transfers one unit of charge per closing event.

There are five closings between $m_0 = +4J$ and $m_0 = -4J$ (at $m_0 = 4, 2, 0, -2, -4$ in units of J). With the sign convention that a closing at $m_0 = -nJ$ carries charge $+1$ for n odd and the interval structure of the 4D two-band model, the principal phase — the one centered on the symmetric design point — carries

$$C_2 = 1.$$

We state this as an inference, not a brute-force integration: the jump structure follows from charge conservation of the closings [1] plus the triviality of the large- $|m_0|$ phase, and a full evaluation of the integral $C_2 = \frac{1}{8\pi^2} \int_{\text{BZ}} \text{tr } \mathcal{F} \wedge \mathcal{F}$ on a discretized Brillouin zone is the numerical confirmation we list as the smallest next calculation in Section 7. The inference is robust in one specific sense: whatever the sign convention, $|C_2| = 1$ on the principal interval, and the protection argument depends only on $C_2 \neq 0$.

4.4 Thermal excitation probability and maximum operating temperature

The passive rejection mechanism is thermodynamic: an error state is a state above the gap, and at thermal equilibrium its occupation is Boltzmann-suppressed. With the gap fixed by the drive, $\Delta = 2J = h\nu_{\text{DFG}} = 14.48$ meV (Section 4.1), the thermal excitation probability is

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

At $T = 4$ K. The thermal energy is

$$k_B T = (1.381 \times 10^{-23} \text{ J/K})(4 \text{ K}) = 5.524 \times 10^{-23} \text{ J} = \frac{5.524 \times 10^{-23}}{1.602 \times 10^{-19}} \text{ eV} = 3.448 \times 10^{-4} \text{ eV} = 0.3448 \text{ meV}.$$

The exponent is

$$\frac{\Delta}{k_B T} = \frac{14.48 \text{ meV}}{0.3448 \text{ meV}} = 41.99,$$

and therefore

$$p_{\text{th}}(4 \text{ K}) = e^{-41.99} = 5.8 \times 10^{-19}.$$

At $T = 100$ mK. The thermal energy is

$$k_B T = (1.381 \times 10^{-23})(0.1) = 1.381 \times 10^{-24} \text{ J} = 8.621 \times 10^{-6} \text{ eV} = 8.621 \mu\text{eV},$$

the exponent is

$$\frac{\Delta}{k_B T} = \frac{14.48 \times 10^{-3} \text{ eV}}{8.621 \times 10^{-6} \text{ eV}} = 1679.6,$$

and converting to a power of ten with $\ln 10 = 2.3026$,

$$p_{\text{th}}(100 \text{ mK}) = e^{-1679.6} = 10^{-1679.6/2.3026} = 10^{-729.4} \approx 4 \times 10^{-730}.$$

Maximum operating temperature. We define thermally quiet operation as $p_{\text{th}} < 0.01$, i.e. fewer than one thermal excitation per hundred mode-lifetimes attributable to gap crossing. Setting $p_{\text{th}} = 0.01$ requires

$$\frac{\Delta}{k_B T_{\text{max}}} = \ln 100 = 4.6052,$$

so

$$T_{\text{max}} = \frac{\Delta}{k_B \ln 100} = \frac{14.48 \times 10^{-3} \text{ eV}}{(8.617 \times 10^{-5} \text{ eV/K})(4.6052)} = \frac{14.48 \times 10^{-3}}{3.968 \times 10^{-4}} \text{ K} = 36.5 \text{ K}.$$

The design margin is comfortable in one direction and unforgiving in the other: any cryostat that holds the stack below 36.5 K makes thermal gap-crossing negligible, but the figure is set entirely by the drive-fixed gap, so any mechanism that softens the gap (drive amplitude noise, detuning drift toward a closing mass $m_0 \in \{0, \pm 2, \pm 4\}J$) degrades it proportionally.

4.5 Realizability of the required coupling

The synthetic-dimension hopping is fixed by the drive identification $J = h\nu_{\text{DFG}}/2$:

$$J = \frac{14.48 \text{ meV}}{2} = 7.24 \text{ meV}, \quad \frac{J}{h} = \frac{\nu_{\text{DFG}}}{2} = 1.75 \text{ THz}.$$

Three comparisons follow from derived quantities alone. First, the coupling energy $J = 7.24 \text{ meV}$ corresponds to a temperature $J/k_B = 7.24 \times 10^{-3}/8.617 \times 10^{-5} = 84.0 \text{ K}$, well above the 4 K operating point, so the coupling, if achieved, is not thermally washed out. Second, the coupling must be resolved against the drive period: a hopping event mediated by one DFG quantum occurs on the timescale

$$\tau = \frac{h}{\Delta} = \frac{1}{\nu_{\text{DFG}}} = \frac{1}{3.5 \times 10^{12}} \text{ s} = 2.86 \times 10^{-13} \text{ s} = 0.29 \text{ ps},$$

so the relevant coherence requirement is that the DFG-mediated phase relationship survive of order 0.29 ps — a modest requirement by atomic-clock standards but an unstudied one in the proposed hybrid stack. Third, the required coupling is enormous by the standards of photonics-mediated synthetic dimensions, where synthetic hoppings are typically engineered at radio-frequency or few-GHz scales; whether a $\chi^{(2)}$ DFG process in TFLN can deliver a coherent 1.75 THz coupling into the relevant modes is the single load-bearing assumption of the entire concept, and no number available to us settles it. We therefore state it as an open engineering question with a bounded target, not as a derived feasibility.

5. Results

The analysis yields four quantitative results, each derived in Section 4 with shown arithmetic.

R1 (gap structure). The effective 4D two-band Hamiltonian driven by the 3.5 THz DFG field closes its bulk gap exactly at synthetic masses $m_0 \in \{0, \pm 2, \pm 4\}J$ (Section 4.2). Between closings the gap is open and the phase is gapped; the closings are isolated, structurally stable band-touching events.

R2 (topological invariant). The principal phase, with the trivial large-mass phase as reference, carries second Chern number $C_2 = 1$, inferred from charge conservation of the closing events [1] (Section 4.3). The protection claim depends only on $C_2 \neq 0$.

R3 (passive rejection). With the gap $\Delta = 14.48 \text{ meV}$ fixed by the drive, the thermal excitation probability is $p_{\text{th}} = 5.8 \times 10^{-19}$ at 4 K and $\approx 4 \times 10^{-730}$ at 100 mK, and thermally quiet operation ($p_{\text{th}} < 1\%$) holds up to $T_{\text{max}} = 36.5 \text{ K}$ (Section 4.4). Passive rejection is therefore thermodynamically robust at any achievable cryogenic temperature, conditional only on the gap remaining open.

R4 (engineering target). The required synthetic-dimension coupling is $J/h = 1.75 \text{ THz}$ ($J = 7.24 \text{ meV}$, $J/k_B = 84.0 \text{ K}$), with a coherence requirement on the scale of the drive period $\tau = 0.29 \text{ ps}$ (Section 4.5). This is the single unproven assumption; all thermodynamic results are conditional on it.

The logical structure of the results is a conditional: *if* the DFG drive realizes a coherent 1.75 THz synthetic hopping of uniform strength, *then* the resulting 4D phase is topological with $C_2 = 1$ and rejects thermal errors at the computed rates. The paper establishes the conditional's consequent rigorously and its antecedent as a bounded, falsifiable target.

6. Discussion

The conditional structure is the honest reading. The thermal rejection numbers are exact consequences of the drive-fixed gap, but they say nothing about whether the gap can be opened. A reader should treat R3 as a property of the model, not of any material device, until the coupling assumption of R4 is tested.

Limitations. The result rests on the following assumptions, with sensitivity noted for each.

- *Uniform synthetic hopping.* We assume the DFG drive mediates fourth-direction hopping with the same strength J as the physical directions (Section 3). If the synthetic hopping is weaker by a factor η , the gap becomes direction-anisotropic and the closing structure shifts; the thermal figures scale with the new minimum gap as $\exp(-\Delta_{\text{min}}/k_B T)$, so a tenfold weaker coupling ($\eta = 0.1$, $\Delta_{\text{min}} = 1.448 \text{ meV}$) raises $p_{\text{th}}(4 \text{ K})$ to $e^{-4.199} \approx 1.5 \times 10^{-2}$, destroying thermally quiet operation at 4 K. The result is exponentially sensitive to this assumption.
- *Thermal-noise-only error model.* We evaluate rejection under thermal excitation only, because no noise spectral densities for the proposed stack are available to us. Dephasing from drive amplitude and phase noise, two-level-system defects in SiN, and phonon-mediated relaxation in diamond are all unquantified here; any of them could dominate long before the Boltzmann floor is reached.

- *Inferred, not integrated, C_2 .* The invariant is inferred from closing-charge conservation [1], not computed by Brillouin-zone integration. A sign or interval-convention error would change which phase is principal but not the existence of a $C_2 = \pm 1$ phase.
- *Equilibrium statistics.* p_{th} assumes thermal equilibrium between the modes and the bath. A driven, periodically modulated system need not thermalize; a Floquet steady state could have a different excitation distribution, and we have not computed one.
- *Single-particle treatment.* The model is free-fermion (Section 2, [4]); interactions, which the classification literature flags as the hard open case [1, 6], are outside the derivation.

What would falsify the main result. The main result is the conditional of Section 5. A single observation falsifies it: a measurement of the DFG-mediated coupling in the proposed stack showing either (i) a coherent synthetic hopping below the level needed to open a gap resolvable at the operating temperature — concretely, $\Delta_{\text{min}} < 4.605 k_B T_{\text{op}}$, i.e. below 1.59 meV at 4 K — or (ii) a dephasing time of the DFG phase shorter than the drive period $\tau = 0.29$ ps. Either outcome breaks the antecedent, and with it every rejection figure in this paper. Conversely, the topological and thermal analysis survives falsification of secondary claims (exact uniformity of hopping, the specific interval assignment of C_2) as long as some gapped $C_2 \neq 0$ interval exists.

The fluidic degrees of freedom. The device concept names a cryo-fluidic stack, and fluidic or phononic degrees of freedom are the most likely locus of unprotected excitations. The color-code analysis warns that certain SPT excitations "do not exist as stable objects since they are superpositions of eigenstate excitations which would decohere immediately" [8]; the analogous risk here is that the protected single-particle picture is spoiled by collective modes of the fluidic components that couple into the band states without respecting the gap. We flag this as an unmodeled failure mode rather than claiming it is absent.

Open questions. Three questions frame future work: (i) what is the actual Floquet steady-state distribution of the driven stack, and does it thermalize? (ii) what coupling strength does intracavity DFG in TFLN actually deliver at 3.5 THz, and with what phase stability? (iii) do the fluidic/phononic modes of the stack couple weakly enough to the band states to leave the gap intact?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are hardware architects at quantum-processor companies and academic groups engineering synthetic dimensions in photonic systems. If the result holds, the design decision that changes is the error-correction architecture itself: instead of budgeting physical qubits for active syndrome extraction, an architect could budget for a drive and a material stack that suppress errors at the Hamiltonian level, trading control electronics for photonics. If the result does not hold — if the 1.75 THz coupling is unreachable — the same architects gain a cheaply falsified concept and can redirect synthetic-dimension effort toward lower-frequency drives where coupling is proven. Because the concept is at the analysis stage, no procurement or policy decision can depend on it yet; the usable output today is the bounded target of R4.

(b) Readiness level. Against the NASA Technology Readiness Level definitions, this work is purely analytical and has not been peer reviewed: it is TRL 1–2 (basic principles observed and reported; concept and application formulated). Analytical work never exceeds TRL 3, and nothing here approaches it — no component of the proposed stack has been operated in the claimed configuration, and no experimental evidence of any kind supports the coupling assumption.

(c) The smallest next test. The smallest de-risking calculation is a numerical evaluation of C_2 by Brillouin-zone integration on the derived Hamiltonian, confirming the inferred invariant and mapping the $C_2(m_0)$ diagram exactly; its cost is CPU-hours and it can only confirm or correct R2. The smallest de-risking experiment is a two-mode analog: a TFLN $\chi^{(2)}$ cavity driven to produce a 3.5 THz intracavity field, coupling two physical resonator modes whose frequency separation matches the DFG difference, with the coherent population oscillation between the modes measured spectroscopically. The measured quantity is the coherent coupling rate $J_{\text{meas}}/\hbar$; the outcome that falsifies the concept is $J_{\text{meas}}/\hbar$ below the level needed for $\Delta_{\text{min}} > 1.59$ meV at 4 K, or an oscillation that damps within one drive period (0.29 ps).

(d) Binding constraints. In the order they would bind: first, the $\chi^{(2)}$ conversion efficiency and intracavity field strength of TFLN at 3.5 THz, which set whether $J/\hbar = 1.75$ THz is approachable at all; second, phase stability of the DFG field over the coherence time of the coupled modes, since the synthetic dimension is defined by that phase; third, thermal management, since the stack must stay below 36.5 K while hosting a driven nonlinear optical cavity — the drive itself deposits heat, and the cryogenic budget must cover it; fourth, materials integration, since diamond, SiN, TFLN, and InGaAs have differing thermal-expansion and lattice-matching behavior, an integration question on which we state no numbers. Regulatory constraints do not bind at this stage.

(e) Order-of-magnitude bounds (projections). We project one performance bound from derived quantities only, with stated assumptions. If the gap holds at $\Delta = 14.48$ meV and the error mechanism is thermal, the error rate per mode-lifetime at 4 K is bounded by $p_{\text{th}} = 5.8 \times 10^{-19}$ (Section 4.4); assuming a mode-lifetime of 1 μs , the corresponding thermal error rate is 5.8×10^{-13} errors

per second per mode. This projection assumes thermal equilibrium, no non-thermal error channels, and the stated lifetime; it is a floor, not a device specification.

(f) Risks and misuse paths. No dual-use risk is apparent: the work proposes a passive error-suppression mechanism for quantum hardware and identifies the conditions under which it fails. The main misuse risk is rhetorical rather than physical — citing the thermal rejection numbers without the conditional coupling assumption to claim a solved error-correction problem. The conditional structure of Section 5 is the guard against that reading.

8. Conclusion

We analyzed a proposed passive error-rejection device — the Topological Matter Engine — by deriving its effective 4D Hamiltonian from the 3.5 THz DFG drive, locating its gap closings at $m_0 \in \{0, \pm 2, \pm 4\}J$, inferring a second Chern number $C_2 = 1$ on the principal phase, and computing its thermal rejection: $p_{\text{th}} = 5.8 \times 10^{-19}$ at 4 K, $\approx 4 \times 10^{-730}$ at 100 mK, and thermally quiet operation up to 36.5 K. The analysis converts an architectural claim into a single bounded engineering target: a coherent synthetic-dimension coupling of $J/\hbar = 1.75$ THz with phase coherence on the 0.29 ps drive timescale, realizable or not in the proposed diamond/SiN/TFLN/InGaAs stack. The concept is falsifiable by one spectroscopic measurement, and the thermal analysis — the paper's robust core — holds unconditionally once that target is met and the gap stays open.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.3-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 zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, 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 11 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] Beyond Band Insulators: Topology of Semi-metals and Interacting Phases. arXiv:1301.0330v1. <https://arxiv.org/abs/1301.0330v1> [2] Enhancing Safety in Diabetic Retinopathy Detection: Uncertainty-Aware Deep Learning Models with Rejection Capabilities. arXiv:2510.00029v1. <https://arxiv.org/abs/2510.00029v1> [3] Characterizing web pornography consumption from passive measurements. arXiv:1904.11723v2. <https://arxiv.org/abs/1904.11723v2> [4] A pedestrian's guide to the topological phases of free fermions. arXiv:2509.19320v2. <https://arxiv.org/abs/2509.19320v2> [5] Phase error estimation for passive detection setups with imperfections and memory effects. arXiv:2508.21486v2. <https://arxiv.org/abs/2508.21486v2> [6] A scheme for topological phases of the Weyl C^* -algebra. arXiv:2607.01183v2. <https://arxiv.org/abs/2607.01183v2> [7] Topological phases of non-interacting systems: A general approach based on states. arXiv:2502.03590v1. <https://arxiv.org/abs/2502.03590v1> [8] Topological color code and symmetry-protected topological phases. arXiv:1503.07208v2. <https://arxiv.org/abs/1503.07208v2>