V1.0.0

Synthetic Vacuum Photonic Lattices: A Quantitative Feasibility Analysis of Room-Temperature, Gauge-Protected Fibonacci Anyon Computing

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Large-scale fault-tolerant quantum computation currently requires cryogenic operation and an active error-correction layer whose ancilla, syndrome-extraction and magic-state overheads dominate the resource budget of any scalable processor. A conjectured alternative is a room-temperature solid-state photonic lattice—a *synthetic vacuum*—engineered so that an emergent SU(2)-like gauge symmetry hosts Fibonacci anyons as edge excitations, with braiding implementing universal logic protected by the medium rather than by active correction. The specific gap addressed here is that no published analysis quantifies the physical conditions such a substrate must satisfy: whether it can be thermally stable at 300 K, and whether gauge constraints suppress error channels or merely relocate them. We model the lattice as a gapped photonic band structure with a topological edge sector, derive the thermal-stability condition, perform an error-channel accounting by gauge transformation property, and derive a per-hop loss budget for braiding operations. We find that a gap $\Delta \geq 15\,k_B T \approx 0.388\,\mathrm{eV}$ (photon frequency $\approx 93.8\,\mathrm{THz}$, wavelength $\approx 3.20\,\mu\mathrm{m}$) holds the probability of any thermally excited stray anyon below $\approx 6.1 \times 10^{-4}$ for a 10^3-mode device; that gauge invariance reduces a physical error rate p only to $p_{\mathrm{eff}} \approx (2/3)p$, so the raw photonic error rate must reach 1.5×10^{-4} to match the 10^{-4} logical error of active flag-qubit schemes; and that a 10^{-6} logical error over a 100-hop braid requires a per-hop loss probability of 3×10^{-7}. Gauge protection is a rate reducer, not a replacement for correction, and the proposal remains at the level of analytical feasibility.

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

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). Synthetic Vacuum Photonic Lattices: A Quantitative Feasibility Analysis of Room-Temperature, Gauge-Protected Fibonacci Anyon Computing. QNFO Preprint v1.0.0.

1. Introduction

The practical problem is the overhead of fault-tolerant quantum computation. Error-corrected machines demand physical error rates far below what current hardware delivers, and the correction machinery itself—ancilla qubits, repeated syndrome extraction, magic-state distillation—dominates the resource count of any scalable processor. The fault-tolerant framework is well understood theoretically [1], yet the same source notes that experimental efforts had realized arrays of only about ten functional qubits at the time of its writing, a gap between theory and hardware that persists in motivating new substrates. We assume, as a working figure for the overhead of active correction, that scalable code-based architectures require many physical qubits per logical qubit, inflating device size, cycle time and energy budget; the framework itself is well understood theoretically [1], while experimental efforts had realized arrays of only about ten functional qubits at the time of its writing, and uncontrolled error propagation during circuit execution can overwhelm a code [1].

Every proposed route to universal logic carries a specific tax. Transversal gates cannot supply a universal gate set, so non-Clifford resources must be manufactured actively: even the optimized flag-qubit magic-state preparation of [4] uses ten qubits, accepts a state only about 80% of the time, and reaches logical error rates near 10^{-4} at physical error rates near 3×10^{-3}. Flag-based error correction lowers ancilla counts relative to Shor-style schemes, which require as many ancillas as the maximum stabilizer weight [6], but the overhead remains multiplicative in code distance. Who faces this problem: any team designing a scalable quantum processor, and the field's roadmaps, which must budget correction hardware that grows with every reduction in logical error rate. What it blocks today: the correction layer sets the floor on device size, cycle time and energy, and it is the reason scaling arguments terminate at the thermodynamic limits of continuous active surveillance of a quantum state, an argument made explicitly in the non-Archimedean architecture literature [7]. What would change if this paper's central conditional results hold: if a passive, gauge-structured medium could hold topological degrees of freedom stable at room temperature, the correction layer would shrink from an active control stack toward a materials-design problem, and the operating envelope would move from millikelvin cryogenics toward ambient temperature.

The conjecture under examination is that a room-temperature solid-state photonic lattice—a *synthetic vacuum*—can instantiate a non-Abelian, SU(2)-like gauge symmetry such that Fibonacci anyons appear as edge excitations and their braiding implements universal quantum logic intrinsically, with protection arising from the gauge structure of the medium itself. This idea connects to a broader program of making the physics carry the computation: Hamiltonian engineering as computation, formalized as an "Inverted Church-Turing-Deutsch" thesis and tested in silico for Ising-type (Class A) and spectral-synthesis (Class B) engines over problem sizes $N = 3$ to 12 [10]; vectorial topological encoding in which computational results are embodied as Chern-number vectors in moiré superlattices, dissolving the Residue Number System reconstruction bottleneck [8]; and the operationalization of generalized (non-invertible) symmetries in Fractional Chern Insulators into falsifiable experimental protocols, including a halo-radius diagnostic proposed as a proxy for the abstract quantum dimension [9].

This paper does not claim to demonstrate such a device. It performs what a theoretical feasibility paper can do: state the model, derive the quantitative conditions the conjecture must satisfy, and analyze whether gauge constraints genuinely suppress error channels or merely relocate them. The answer developed below is that gauge structure converts some errors into energetically suppressed processes while leaving a residual gauge-respecting channel untouched, so gauge protection is a prefactor on, not a substitute for, error correction.

Contributions.

- A thermal-stability derivation showing that a photonic topological gap of $\Delta \geq 15\, k_B T \approx 0.388$ eV at $T = 300$ K suppresses stray thermally excited anyons below $\approx 6.1 \times 10^{-4}$ per 10^3-mode device, with the corresponding mid-infrared operating frequency ≈ 93.8 THz derived in Section 4.1 and reported in Section 5.
- A fusion-space scaling analysis for Fibonacci anyons showing Hilbert-space growth as φ^n with $\varphi = (1 + \sqrt{5})/2 \approx 1.618$, quantified for $n = 100$ anyons in Section 4.2 and reported in Section 5.
- An error-channel accounting model that splits local errors into gauge-violating (energetically suppressed by the gap) and gauge-respecting (unsuppressed) classes, deriving $p_{\text{eff}} \approx f_{\text{gp}}\, p$ with $f_{\text{gp}} = 2/3$ for a qutrit-based local gauge constraint, and the consequence that the raw photonic error rate must satisfy $p \leq 1.5 \times 10^{-4}$ to reach a 10^{-4} logical error without active correction (Section 4.3, reported in Section 5).
- A per-hop loss budget for anyon braiding: mapping the physical loss probability onto a logical error through the flag-qubit scaling of [4], a 10^{-6} logical error over a braiding path of $L = 100$ hops requires a per-hop photon-loss probability $p_{\text{loss}} \leq 3 \times 10^{-7}$ (Section 4.5, reported in Section 5).
- An identification of the falsifiers and binding engineering constraints for the synthetic-vacuum conjecture, including the specific observation that would refute the thermal-stability and loss-budget claims (Sections 6 and 7).

The analysis may also bear on classical fault-tolerant system synthesis, where fault-tolerant mechanisms are inserted as pattern choices under a fault hypothesis and solved as a game [2], and on reliability modeling for many-core interconnect, where fault models and mitigation strategies are catalogued for 2D bus-based and 3D network-on-chip systems [5]; in both cases the question of whether a structural constraint reduces or relocates faults is the same question this paper asks of gauge symmetry.

2. Background and Related Work

Active fault tolerance and its overhead. The standard framework of fault-tolerant quantum computation rests on quantum error-correcting codes, which must be designed around two quantum-specific obstacles: the no-cloning theorem, which forbids protecting an unknown state by copying it, and the collapse of the state under measurement of an encoded block [1]. The same source emphasizes that uncontrolled error propagation during circuit execution can overwhelm a code, which is why fault-tolerant circuit design, not just code existence, is the operative discipline [1]. Every scheme built in this framework pays an overhead tax at the point where universality is restored. Because transversal gates can never cover a universal set, non-Clifford resources must be prepared fault-tolerantly; the flag-qubit scheme of [4] prepares magic states directly with ten qubits, an acceptance rate around 80%, and logical error rates near 10^{-4} at physical error rates near 3×10^{-3}, an improvement of orders of magnitude over distillation in some error regimes but still a resource-consuming active procedure [4]. On the correction side, flag techniques detect high-weight errors from a few faults using few ancillas, and the distinguishability of the resulting fault set depends on circuit structure, gate ordering, and the choice of stabilizer generators [6]. These works define the baseline against which any passive-protection claim must be measured: the passive scheme must beat a $(2/3)p$-style prefactor protection plus whatever residual correction remains.

Fault tolerance as a boundary problem. In distributed and input/output settings, noise at the first and last circuit layers is unavoidable, and the best achievable fault-tolerant realization of a circuit with quantum input/output reproduces the ideal circuit up to residual noisy channels on the input and output systems, with detrimental effects characterized as adversarial or local stochastic channels and diamond-norm bounds of the form $O(\theta\sqrt{(c\delta)^k})$ for state preparation [3]. This result matters for the synthetic-vacuum conjecture because any photonic device must inject and extract quantum information through its boundary—the very edge where the anyons live—so the boundary-noise results of [3] describe exactly the interface where gauge protection is weakest.

Passive and structural protection proposals. The non-Archimedean quantum architecture program proposes replacing continuous (Archimedean) state spaces with p-adic, ultrametric ones, arguing that the strong triangle inequality natively absorbs low-energy perturbations and that passive geometric fault tolerance was sustained over 10^5 operational cycles in simulation, with the caveat, stated in that work itself, that the exactly 1.0 simulated fidelity is an artifact of Boolean threshold dynamics and physical substrates would yield fidelities of roughly 0.999 [7]. That work also argues that passive geometric protection shifts the energy profile from active microwave pulses to static power, moving thermal dissipation below cryogenic cooling limits [7]. The synthetic-vacuum conjecture is structurally analogous—protection from geometry rather than feedback—but operates in real space with a spectral gap rather than in number space with an ultrametric, and the error accounting of Section 4.3 applies the same skepticism that work applies to itself: simulated threshold behavior is not physical gap behavior.

Computation by physics. The Inverted Church-Turing-Deutsch framework treats computation as the time evolution of engineered Hamiltonians, distinguishing Class A (Ising-based optimization) from Class B (spectral synthesis), and validates scaling claims in silico over $N = 3$ to 12 spin variables with 20 replicates per size, using synthetic relaxation data [10]. The framework explicitly acknowledges thermodynamic constraints on physical realization [10], which is the same constraint this paper quantifies for the photonic case. The morphological-computing framework proposes encoding computational results as vectors of Chern numbers in moiré superlattices, dissolving the serial Chinese-Remainder-

Theorem reconstruction bottleneck of Residue Number System arithmetic, and grounds computation in relaxation of a constrained physical system to a minimum-energy ground state [8]. That grounding is directly relevant: a synthetic vacuum that hosts anyons in its ground-state sector is a computation-by-relaxation device in exactly the sense of [8], and the same source's warning—via its cited critique literature—that claims of physical computation risk unverifiability without a formal process-to-computation map [8] applies with full force to braiding-based logic, where the map is the braid group representation.

Symmetry as an operational resource. The classification of topological phases by generalized symmetries, described by fusion categories rather than groups, provides the language in which non-Abelian anyon content is specified [9]. That work operationalizes these abstract structures into falsifiable protocols for Fractional Chern Insulators in moiré superlattices, proposing the anyon density-wave halo and establishing, through its reported ANOVA simulations with F-statistics in the hundreds, the halo radius as a high-fidelity proxy for the quantum dimension [9]. It also extends the diagnostic to domain walls, predicting structured screening charge along the wall measurable in imaging and transport [9]. For the present conjecture this matters twice: first, the quantum dimension φ of Fibonacci anyons is the quantity whose measurement would confirm the edge sector; second, the halo protocol of [9] is a template for the experimental falsifier proposed in Section 7.

Classical fault-tolerant synthesis and fault modeling. In embedded-systems synthesis, fault tolerance is treated as the algorithmic insertion of predefined fault-tolerant mechanism patterns under an explicit fault hypothesis, translated into a distributed game and solved with incomplete search; the prototype engine reports search results in about 3 seconds and constraint solving within 1 second on small examples, with the authors noting that exponential blowups from variables to states remain unavoidable [2]. The structural lesson for the synthetic-vacuum design flow is that fault tolerance must be specified as a hypothesis plus a mechanism pool before synthesis; a gauge structure is precisely such a mechanism, and Section 4.3 evaluates it as one. On the modeling side, the survey of fault models and fault-tolerance methods for 2D bus-based multi-core and 3D network-on-chip many-core systems frames reliability as requiring explicit fault and failure models before mitigation can be strategized [5]; the present paper follows the same discipline by enumerating error channels by their gauge transformation property before claiming suppression.

3. Methods

This is an analytical feasibility study. No experiment, simulation on hardware, or measurement was performed; all numbers below are derived from stated inputs by explicit arithmetic, or labeled as assumptions.

Model. The synthetic vacuum is modeled as a two-dimensional photonic lattice whose bulk band structure is gapped by Δ, with a topologically ordered edge sector supporting anyonic excitations with fusion rules of the Fibonacci type: two anyon types $\{1, \tau\}$ with $\tau \times \tau = 1 + \tau$ and quantum dimension $d_\tau = \varphi$. The conjectured mechanism is an engineered gauge constraint—locally, a Gauss-law-like projector G_i at each site i enforcing $G_i|\psi\rangle = |\psi\rangle$—such that physical states are gauge-invariant and local gauge-violating operators cost energy at least Δ. Structurally, the gauge constraint can be encoded as a link operator U_{ij} on each lattice bond, with the constraint that the product of U_{ij} around any elementary plaquette equals the identity; this is the form in which the conjecture is stated, and no specific band-structure realization is claimed here.

Assumptions. Each assumption is stated explicitly because no source supplies it:

- (A1) The lattice supports a gapped topological phase with Fibonacci edge anyons; this is the conjecture under test, not a result.
- (A2) The device operates at $T = 300$ K.
- (A3) The device contains $N_{\text{mode}} = 10^3$ photonic modes, each with $g = 2$ internal states contributing to the anyon sector; this is a small-device scale chosen for concreteness.

- (A4) Local physical error processes occur at a per-mode rate p per operation, left symbolic; the gauge structure partitions them into gauge-violating and gauge-respecting classes.
- (A5) The local Hilbert space is qutrit-like, $\dim \mathcal{H}_{\text{local}} = 3$, with a gauge-invariant subspace of dimension 2; this models the minimal gauge constraint that removes one local degree of freedom.
- (A6) A full anyon braiding operation traverses a path of $L = 100$ hopping events; this is a conservative path-length assumption used only in the loss budget of Section 4.5.
- (A7) Photon loss is independent from hop to hop, and the logical-error consequence of a braiding-path loss event can be mapped through the empirical physical-to-logical scaling reported for flag-qubit magic-state preparation in [4].

Reconciliation convention. The feasibility analysis is organized under the convention that the synthetic vacuum must satisfy two independent constraints simultaneously: the thermal-stability condition on the gap (Section 4.1) and the loss-budget condition on braiding paths (Section 4.5). This is the more conservative of the available readings, because satisfying either constraint alone would not establish feasibility.

Derivation procedure. Section 4.1 computes the thermal stability condition from the Boltzmann factor; Section 4.2 computes the Fibonacci fusion-space scaling; Section 4.3 performs the error-channel accounting; Section 4.4 compares against documented active fault-tolerance overheads from [4], [6], and [3]; Section 4.5 derives the per-hop loss budget. All arithmetic is shown step by step.

4. Analysis

4.1 Thermal stability: the required gap

Input numbers. Boltzmann constant $k_B = 8.617333 \times 10^{-5}\ \text{eV/K}$ (physical constant); operating temperature $T = 300\ \text{K}$ (Assumption A2); Planck constant $h = 4.135668 \times 10^{-15}\ \text{eV} \cdot \text{s}$ (physical constant); speed of light $c = 2.99792458 \times 10^{8}\ \text{m/s}$ (physical constant); number of modes $N_{\text{mode}} = 10^3$ and internal states $g = 2$ (Assumption A3).

Step 1: thermal energy.

$$k_B T = (8.617333 \times 10^{-5}) \times 300 = 2.5852 \times 10^{-2}\ \text{eV} = 25.85\ \text{meV}.$$

Step 2: thermal occupation of a gapped mode. A bosonic mode of energy Δ above the band edge has thermal occupation

$$\bar{n}(\Delta) = \frac{1}{e^{\Delta / k_B T} - 1}.$$

At $\Delta = 10\, k_B T$: $e^{10} = 22026.47$, so $\bar{n} = 1/(22026.47 - 1) = 1/22025.47 = 4.540 \times 10^{-5}$.

Step 3: probability of at least one stray thermal anyon. Treating the $N_{\text{mode}} \cdot g = 2000$ modes as independent, the probability that at least one thermally excited anyon exists is, to leading order in $\bar{n}$,

$$P_{\text{thermal}} \approx N_{\text{mode}} \cdot g \cdot \bar{n}(\Delta).$$

At $\Delta = 10\, k_B T$: $P_{\text{thermal}} \approx 2000 \times 4.540 \times 10^{-5} = 9.08 \times 10^{-2}$. This is too high for a computation: roughly one computation in eleven would begin with a stray anyon.

At $\Delta = 15\, k_B T$: $e^{15} = 3.269017 \times 10^{6}$, so $\bar{n} = 1/(3269017 - 1) = 3.059 \times 10^{-7}$, and

$$P_{\text{thermal}} \approx 2000 \times 3.059 \times 10^{-7} = 6.12 \times 10^{-4}.$$

Step 4: the gap in physical units. At $\Delta = 15\,k_B T$:

$$\Delta = 15 \times 2.5852 \times 10^{-2}\ \text{eV} = 0.38778\ \text{eV}.$$

The corresponding photon frequency and free-space wavelength:

$$\nu = \frac{\Delta}{h} = \frac{0.38778}{4.135668 \times 10^{-15}} = 9.376 \times 10^{13}\ \text{Hz} = 93.76\ \text{THz},$$

$$\lambda = \frac{c}{\nu} = \frac{2.99792458 \times 10^{8}}{9.376 \times 10^{13}} = 3.198 \times 10^{-6}\ \text{m} \approx 3.20\ \mu\text{m}.$$

The required operating window is therefore the mid-infrared. Whether mid-infrared photonic platforms can be fabricated at the required tolerance is left as an open engineering question (Section 7); no source consulted supplies a demonstrated device figure.

4.2 Fibonacci fusion-space scaling

Input. Fibonacci fusion $\tau \times \tau = 1 + \tau$ implies the quantum dimension $d_\tau = \varphi$, where φ is the positive root of $\varphi^2 = \varphi + 1$:

$$\varphi = \frac{1 + \sqrt{5}}{2} = \frac{1 + 2.2360680}{2} = 1.6180340.$$

Step 1: fusion-space growth. The dimension of the fusion space of n τ anyons with total charge τ grows as φ^n up to polynomial prefactors. For $n = 100$ anyons:

$$\varphi^{100} = e^{100 \ln \varphi} = e^{100 \times 0.4812118} = e^{48.12118}.$$

$$e^{48.12118} = 7.93 \times 10^{20}.$$

The fusion space of 100 τ anyons therefore has dimension of order 10^{21} — exponentially large, which is the resource that makes braiding universal, and simultaneously the reason classical simulation of the edge sector fails beyond modest anyon numbers, echoing the Hilbert-space barrier stated in [10].

Step 2: braid group coverage. The braid group on n anyons, B_n, is generated by elementary exchanges of neighboring anyons; universality of Fibonacci anyon braiding means the dense image of B_n in the fusion-space unitary group approximates any unitary to accuracy ϵ with a braid word of length polynomial in $\log(1/\epsilon)$. We do not derive the density theorem here; we use it only to set the braid-length scale. A conservative single-qubit-equivalent operation is taken to involve braiding over a path with $L = 100$ elementary hops (Assumption A6), which sets the loss budget of Section 4.5.

4.3 Error-channel accounting: what the gauge constraint suppresses

Input. Local Hilbert space dimension $d = 3$ (Assumption A5); gauge-invariant subspace dimension $d_{\text{inv}} = 2$; per-mode physical error probability p per operation (Assumption A4, symbolic).

Step 1: partition of local operators. A local error operator E acting on a qutrit site is a 3×3 matrix; the operator space has dimension $d^2 = 9$. Operators mapping the gauge-invariant subspace into itself form a subspace of dimension $d_{\text{inv}}^2 = 4$. Operators with at least one matrix element connecting the gauge-invariant subspace to the gauge-violating complement number $d^2 - d_{\text{inv}}^2 = 9 - 4 = 5$.

Step 2: the two error classes.

- *Gauge-violating class* (5 of 9 channels): these create anyon pairs or violate the plaquette constraint; by the model of Section 3 they cost energy at least Δ, so their probability is Boltzmann-suppressed by $e^{-\Delta/k_B T}$. At $\Delta = 15\,k_B T$ this factor is $e^{-15} = 3.06 \times 10^{-7}$, so these channels are effectively closed at room temperature.
- *Gauge-respecting class* (4 of 9 channels): these act within the gauge-invariant subspace — phase errors, dephasing of the encoded qubit, and loss events that remove a photon without violating the constraint. The gauge structure does not suppress them energetically.

Step 3: the effective error rate. Averaging uniformly over the 9 local operator channels, the fraction of channels that survive unsuppressed is

$$f_{\text{gp}} = \frac{d_{\text{inv}}^2}{d^2} = \frac{4}{9}.$$

However, the gauge-respecting channels are weighted by their action on the encoded qubit: of the 4 gauge-invariant operators, one is the identity and two act nontrivially on the encoded two-dimensional subspace (two nontrivial Pauli-like actions). The fraction of *physically harmful* channels that remain is therefore the 2 nontrivial gauge-respecting actions against the 3 nontrivial error actions of an unconstrained qutrit (X-, Y-, and Z-type):

$$f_{\text{gp}} = \frac{2}{3}.$$

We take $f_{\text{gp}} = 2/3$ as the operative convention: it counts the two nontrivial gauge-respecting error actions against the three nontrivial error actions of an unconstrained qutrit (X-, Y-, and Z-type), which is the conservative reading — it credits the gauge constraint with suppressing only the gauge-violating third of the error algebra. The effective logical error probability is therefore

$$p_{\text{eff}} = f_{\text{gp}}\, p = \frac{2}{3} p.$$

Step 4: required raw error rate. To reach a logical error rate of 10^{-4} per operation without active correction, matching the flag-qubit benchmark of [4], the raw photonic error rate must satisfy

$$\frac{2}{3} p \le 10^{-4} \quad \Longrightarrow \quad p \le 1.5 \times 10^{-4}.$$

Step 5: comparison with the active-correction baseline. The flag-qubit scheme of [4] reaches logical error rates near 10^{-4} at physical error rates near 3×10^{-3}. The gauge prefactor $2/3$ improves the physical error budget by only a factor of 1.5; the required raw rate 1.5×10^{-4} is a factor of 20 *below* the physical error rate at which the actively corrected flag scheme already achieves 10^{-4}. Gauge protection alone therefore does not close the gap to active correction: it is a prefactor, not a replacement.

4.4 Overhead comparison with documented active schemes

Input. From [4]: 10 qubits, acceptance $\approx 80\%$, logical error $\approx 10^{-4}$ at physical error $\approx 3 \times 10^{-3}$. From [6]: Shor-style error correction requires as many ancillas as the maximum weight of the stabilizer generators, whereas flag techniques detect high-weight errors from a few faults using few ancillas. From [3]: fault-tolerant realization of circuits with quantum input/output incurs residual noisy channels on input and output systems, with state-preparation error bounded as $O(\theta\sqrt{(c\delta)^k})$.

Step 1: the overhead the synthetic vacuum must beat. The flag scheme's 10-qubit, 80%-acceptance magic-state factory is among the cheapest documented universal-gate resources [4]; any passive scheme must deliver logical error $\le 10^{-4}$ with *zero* active ancilla overhead to be competitive, which requires the raw error condition $p \le 1.5 \times 10^{-4}$ of Section 4.3.

Step 2: the boundary tax. Any photonic device injects and extracts quantum information through its edge — precisely where the anyons live. The input/output results of [3] show that noise in the first and last circuit layers is unavoidable and that the best achievable realization carries residual noisy channels on the input and output systems. For the synthetic vacuum this means the braiding-encoded state must be prepared and read out through the gapless or near-gapless boundary, where the thermal-suppression argument of Section 4.1 applies least. We account for this by requiring the loss budget of Section 4.5 to hold *including* boundary hops; no separate boundary-error number is derivable from the sources, so this is flagged as an open quantity in Section 6.

4.5 Per-hop loss budget for braiding

Input. Braid path length $L = 100$ hops (Assumption A6); target logical error over one braid $P_{\text{logic}} = 10^{-6}$ (chosen as a demanding but standard target, 10^{-2} of the 10^{-4} flag-qubit benchmark); physical-to-logical mapping from [4]: logical error $\approx 10^{-4}$ at physical error $\approx 3 \times 10^{-3}$, i.e., a linear reduction ratio

$$r = \frac{10^{-4}}{3 \times 10^{-3}} = \frac{1}{30}.$$

Step 1: per-hop logical contribution. A photon-loss event on one hop maps to a logical error with the same physical-to-logical ratio r as any local physical error under the flag-qubit scaling (Assumption A7):

$$P_{\text{logic}}^{(\text{hop})} = \frac{p_{\text{loss}}}{30}.$$

Step 2: accumulate over the braid path. With independent hops (Assumption A7), the total logical error over $L = 100$ hops is, to leading order,

$$P_{\text{logic}} \approx L \times \frac{p_{\text{loss}}}{30} = \frac{100\, p_{\text{loss}}}{30} = 3.333\, p_{\text{loss}}.$$

Step 3: solve for the budget.

$$3.333\, p_{\text{loss}} \le 10^{-6} \quad \Longrightarrow \quad p_{\text{loss}} \le 3 \times 10^{-7}.$$

Step 4: physical meaning. A per-hop loss probability of 3×10^{-7} over a hop length ℓ corresponds to a propagation loss coefficient $\alpha = p_{\text{loss}}/\ell$; for a hop length of $\ell = 10\ \mu\text{m}$ (a typical integrated-photonic neighbor spacing, stated as an assumption), this is

$$\alpha = \frac{3 \times 10^{-7}}{10^{-5}\ \text{m}} = 3 \times 10^{-2}\ \text{m}^{-1} = 0.13\ \text{dB/m},$$

using $1\ \text{dB} = \ln(10)/10 \approx 0.2303$ nepers, so $\alpha_{\text{dB}} = 8.686 \times 3 \times 10^{-2} = 0.26\ \text{dB/m}$; we state the conservative 0.13–$0.26\ \text{dB/m}$ range depending on whether loss is counted in power or amplitude. Whether a per-hop loss probability of 3×10^{-7} is achievable in any integrated-photonic platform is not derivable from the sources; the achievable value is left as the binding open engineering question in Section 7.

5. Results

All numbers below are computed in Section 4 from the stated inputs and assumptions; none is measured.

Result	Value	Where derived
Thermal energy at 300 K	$k_B T = 25.85$ meV	Section 4.1, Step 1
Required gap	$\Delta = 15\, k_B T = 0.388$ eV	Section 4.1, Step 4
Operating frequency	$\nu = 93.76$ THz	Section 4.1, Step 4
Operating wavelength	$\lambda \approx 3.20$ μm (mid-infrared)	Section 4.1, Step 4
Stray-anyon probability at $10\, k_B T$ gap	9.08×10^{-2} per 10^3-mode device	Section 4.1, Step 3
Stray-anyon probability at $15\, k_B T$ gap	6.12×10^{-4} per 10^3-mode device	Section 4.1, Step 3
Fusion-space dimension, $n = 100$ anyons	$\varphi^{100} \approx 7.9 \times 10^{20}$	Section 4.2, Step 1
Gauge-protection prefactor	$f_{\text{gp}} = 2/3$, so $p_{\text{eff}} = (2/3)p$	Section 4.3, Step 3
Required raw photonic error rate	$p \leq 1.5 \times 10^{-4}$	Section 4.3, Step 4
Gap to active-correction baseline	factor of 20 below the 3×10^{-3} of [4]	Section 4.3, Step 5
Per-hop loss budget for 10^{-6} logical error over 100 hops	$p_{\text{loss}} \leq 3 \times 10^{-7}$	Section 4.5, Step 3

The headline conclusions are threefold. First, thermal stability at room temperature demands a mid-infrared gap of ≈ 0.388 eV, which reduces stray-anyon contamination to 6.12×10^{-4} per device — feasible as a spectral target, unproven as a topological-phase realization. Second, the gauge constraint suppresses only the gauge-violating third of the local error algebra, so the raw photonic error rate must reach 1.5×10^{-4}, twenty times better than the physical error rate at which actively corrected flag-qubit magic-state preparation already achieves 10^{-4} logical error [4]. Third, braiding over 100 hops at 10^{-6} logical error requires a per-hop loss probability of 3×10^{-7}, an extremely demanding but not physically forbidden target for integrated photonics.

6. Discussion

What the results do and do not establish. The derivations establish necessary conditions, not sufficiency. A gap of 0.388 eV, a raw error rate of 1.5×10^{-4}, and a per-hop loss of 3×10^{-7} are the conditions the synthetic-vacuum conjecture must satisfy; nothing here shows that any photonic lattice, moiré system, or waveguide array can realize a gapped topological phase with Fibonacci edge anyons at those parameters. Assumption A1 — the existence of the phase — remains the conjecture under test and is the largest single risk in the program.

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

- (A1, phase existence) No derivation here supports it; if no photonic platform realizes Fibonacci anyons at any temperature, all downstream numbers are moot. Sensitivity: total.
- (A2, $T = 300$ K) The gap scales linearly with T; at $T = 77$ K the required gap drops to $15 \times 8.617333 \times 10^{-5} \times 77 = 9.95 \times 10^{-2}$ eV, a factor of 3.9 easier. Sensitivity: linear in T.
- (A3, 10^3 modes) The stray-anyon probability scales linearly with $N_{\text{mode}} \cdot g$; a 10^6-mode device at the same gap would see $P_{\text{thermal}} \approx 0.61$, destroying feasibility unless the gap is raised. Sensitivity: linear in mode count.
- (A5, qutrit gauge model) The prefactor $f_{\text{gp}} = 2/3$ follows from a specific minimal gauge constraint; a richer gauge algebra could suppress more channels (smaller f_{gp}) or fewer (larger f_{gp}). Sensitivity: the required raw error rate scales as f_{gp}^{-1}.
- (A6, $L = 100$ hops) The loss budget scales as $1/L$; a 1000-hop braid tightens the per-hop requirement to 3×10^{-8}. Sensitivity: linear in L.
- (A7, independence and flag-scaling mapping) The mapping of loss events to logical errors through the flag-qubit ratio $r = 1/30$ assumes photonic loss behaves like the circuit-level noise of [4]; loss is in reality a leakage process

that may create undetected anyon pairs, which could make the true budget *stricter* than 3×10^{-7}. Sensitivity: potentially order-unity or worse.

Failure modes. Three are prominent. First, thermally excited anyon pairs could be created not in the bulk modes counted in Section 4.1 but at the boundary, where the gap may close; the input/output analysis of [3] shows boundary noise is unavoidable in any fault-tolerant realization, and the synthetic vacuum concentrates its logical degrees of freedom exactly there. Second, photon loss is gauge-respecting: a lost photon may leave behind an unpaired anyon that the gauge constraint does not energetically suppress, converting the loss channel into a deconfinement channel rather than a benign erasure. Third, the uniform averaging over the nine local operator channels that yields $f_{\text{gp}} = 2/3$ may misweight the physical noise; if the physical noise is dominated by gauge-reserving dephasing, the effective prefactor approaches 1 and gauge protection delivers nothing.

What would falsify the claims. The thermal-stability claim is falsified by the observation of thermally generated anyons in a device with a verified gap $\Delta \geq 15\, k_B T$ at 300 K at a rate exceeding 6.12×10^{-4} per 10^3 modes per operational window — such an observation would mean the Boltzmann-suppression model misidentifies the excitation spectrum. The loss-budget claim is falsified by a braiding experiment in which per-hop loss is at or below 3×10^{-7} yet logical error over 100 hops exceeds 10^{-6}, which would show the flag-scaling mapping of Assumption A7 understates the logical consequence of loss. The gauge-protection claim is falsified by measuring $p_{\text{eff}}/p > 2/3$ in a device with a verified gauge constraint, i.e., by finding that the constraint suppresses fewer than a third of the physical error channels.

Open questions. No source supplies: a concrete band-structure realization of the plaquette gauge constraint; the boundary gap profile; the physical loss mechanism spectrum at $3.2\ \mu$m in candidate platforms; or the many-anyon readout scheme. Each is required before the feasibility analysis can be upgraded from necessary conditions to a design.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are research groups in topological photonics, moiré-material engineering, and integrated mid-infrared photonics, for whom the derived targets convert an open-ended conjecture into three checkable design specifications: a band gap of $0.388\ \text{eV}$ at $300\ \text{K}$, a raw error rate of 1.5×10^{-4} per operation, and a per-hop loss of 3×10^{-7}. If the targets are met, the design change is architectural: the fault-tolerance layer of a quantum processor moves from active syndrome-extraction circuits — whose ancilla overhead the flag-qubit work [4] and the capped-color-code work [6] document as the dominant cost — to a passive materials specification. Roadmap planners gain a quantified decision criterion: pursue the synthetic vacuum only if a platform demonstrates the gap and the loss figure; otherwise invest in reducing the overhead of active schemes along the lines of [4] and [6]. The same structural question — whether a built-in constraint reduces or merely relocates faults — recurs in classical fault-tolerant synthesis, where mechanisms are inserted from a predefined pool under an explicit fault hypothesis [2], and in many-core reliability planning, which requires explicit fault models before mitigation can be strategized [5]; the gauge-channel accounting of Section 4.3 is a template for both.

(b) Current readiness. This work is purely analytical and has not been peer reviewed; under the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles identified and concept formulated, with no experimental proof of concept. The underlying conjecture (A1) is itself unproven, so no component of the proposal exceeds TRL 3 even in principle at this stage. No manufacturing readiness is claimed.

(c) The smallest next test. The minimal de-risking calculation is a band-structure computation, not a device build: compute the spectrum of a candidate photonic lattice Hamiltonian (e.g., a moiré or coupled-waveguide design) and check whether a gapped topological sector with Fibonacci fusion rules exists with $\Delta \geq 0.388\ \text{eV}$ at $300\ \text{K}$. The outcome that refutes the thermal-stability claim is a computed gap below $15\, k_B T$ for every candidate Hamiltonian in a

reasonable design family, or an edge-sector gap that closes while the bulk gap opens. The next step beyond that is the halo-style diagnostic of [9]: a study would measure the anyon density-wave halo radius around engineered defects and test whether it tracks the quantum dimension φ, since [9] establishes the halo radius as a high-fidelity proxy for the quantum dimension through its ANOVA simulations with F-statistics in the hundreds. A measured halo radius inconsistent with $d_\tau = \varphi$ would refute the Fibonacci edge-sector claim directly.

(d) Binding engineering constraints. In the order they would bind: (i) *loss*, the 3×10^{-7} per-hop figure is the tightest number derived here and no source supplies a demonstrated mid-infrared waveguide loss at that level; (ii) *materials*, a room-temperature gapped topological photonic phase with non-Abelian edge content is un demonstrated in any source consulted; (iii) *fabrication tolerance*, the plaquette gauge constraint requires the product of link operators around each plaquette to equal identity, implying phase-error tolerances set by the gap, an unquantified open question; (iv) *readout*, no scheme for measuring the fusion state of many anyons without destroying them appears in the sources; (v) *thermal management*, even a passive scheme dissipates static power, and the thermodynamic argument of [7] — that passive geometric protection shifts the energy profile from active pulses to static power below cryogenic cooling limits — applies here only insofar as the gauge constraint is truly static.

(e) Order-of-magnitude projections. Labeled as projections with stated assumptions: if a platform achieves the derived gap and loss figures, a 10^3-mode device performing 100-hop braids at 10^{-6} logical error would operate at $93.76\ \mathrm{THz}$ with a stray-anyon contamination of 6.12×10^{-4} per operational window (from Section 4.1) and a per-braid loss-limited error of 10^{-6} (from Section 4.5). These projections assume A1–A7 hold and that no additional error channels beyond those enumerated in Section 4.3 exist; the uncertainty is dominated by A1 and A7, which could tighten the budgets by orders of magnitude.

(f) Risks and misuse paths. The primary risk is premature claims of room-temperature topological quantum computing based on partial analogies — for instance, reporting a gapped photonic band structure without verifying non-Abelian statistics, which would inflate the perceived readiness of the field. The critique that claims of physical computation risk unverifiability without a formal process-to-computation map [8] applies here: any demonstration must specify the braid-to-logic map explicitly, or the result is unfalsifiable.

8. Conclusion

This paper has converted the synthetic-vacuum conjecture from an open-ended vision into three quantitative necessary conditions. A room-temperature photonic topological medium must exhibit a gap $\Delta \geq 15\, k_B T \approx 0.388\ \mathrm{eV}$ at $300\ \mathrm{K}$, corresponding to $93.76\ \mathrm{THz}$ ($\approx 3.20\ \mu\mathrm{m}$), to hold stray thermally excited anyons below 6.12×10^{-4} per 10^3-mode device. Its gauge constraint suppresses only the gauge-violating fraction of the local error algebra, yielding $p_{\mathrm{eff}} \approx (2/3)p$ and a required raw photonic error rate of 1.5×10^{-4} — a factor of 20 below the physical error rate at which the best documented active magic-state preparation achieves 10^{-4} logical error [4]. Its braiding paths must sustain a per-hop loss probability of 3×10^{-7} to deliver 10^{-6} logical error over 100 hops. The central finding is that gauge protection is a rate reducer, not a replacement for error correction: it multiplies the physical error rate by $2/3$ where an active scheme multiplies the resource count by code distance. The conjecture remains at analytical feasibility, TRL 1–2, and its decisive test is computational — a candidate Hamiltonian with a verified $0.388\ \mathrm{eV}$ topological gap and Fibonacci edge statistics — followed by the halo diagnostic of [9] as the first experimental falsifier. The analysis also offers a transferable method: enumerate error channels by their symmetry transformation property before claiming that a structural constraint protects, a discipline equally applicable to classical fault-tolerant synthesis [2] and many-core reliability modeling [5].

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.3). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 12 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] An introduction to Fault-tolerant Quantum Computing. arXiv:1508.03695v1. https://arxiv.org/abs/1508.03695v1 [2] A Game-theoretic Approach for Synthesizing Fault-Tolerant Embedded Systems. arXiv:1011.0268v1. https://arxiv.org/abs/1011.0268v1 [3] Fault-tolerant quantum input/output. arXiv:2408.05260v3. https://arxiv.org/abs/2408.05260v3 [4] Fault-tolerant magic state preparation with flag qubits. arXiv:1811.00566v2. https://arxiv.org/abs/1811.00566v2 [5] A Survey of fault models and fault tolerance methods for 2D bus-based multi-core systems and TSV based 3D NOC many-core systems. arXiv:2203.07830v1. https://arxiv.org/abs/2203.07830v1 [6] Achieving fault tolerance on capped color codes with few ancillas. arXiv:2106.02649v4. https://arxiv.org/abs/2106.02649v4 [7] https://papers.qnfo.org/papers/computational-simulation-approach-to-non-archimedean-quantum-architectures. QNFO: Computational Simulation Approach to Non-Archimedean Quantum Architectures. [8] https://papers.qnfo.org/papers/functorial-framework-for-morphological-computing. QNFO: Functorial Framework for Morphological Computing. [9] https://papers.qnfo.org/papers/operationalizing-generalized-symmetries. QNFO: Operationalizing Generalized Symmetries. [10] https://papers.qnfo.org/papers/programming-the-vacuum. QNFO: Programming the Vacuum.