Preprint — Not Peer Reviewed

v1.5, May 2026

Magic-Angle Photonic Quantum Computing

Geometric decoherence protection, the trisection qutrit,
and falsifiable predictions from a (3+3) spacetime framework

C.R. (René) de Haan

Independent Researcher, Vinkeveen, Netherlands

ORCID: 0009-0004-0676-6143

with research assistance from **Claude (Anthropic)**

large language model used as AI research assistant; not an author

Preprint repository: Zenodo

DOI: 10.5281/zenodo.20043366

AI Assistance Statement

This preprint was developed through iterative collaboration between the author and **Claude (Anthropic)**, a large language model used as an AI research assistant. Claude contributed to: (i) symbolic computation and consistency-verification of equations, derivations, and observational comparisons; (ii) systematic extraction and consolidation of the framework's geometric arguments from the foundational documents [1, 6] into the photonic-QC application context; (iii) identification of methodological pitfalls and false trails encountered during the analysis; (iv) drafting, organisation, and technical typesetting; and (v) the structural framing of the paper's presentation. Specific further contribution detail is given in the Acknowledgements.

Following the AI-disclosure conventions of Nature, Science, ICMJE, APS, IEEE, and ACM (2023–2024), Claude is not listed as an author of this paper: authorship implies accountability — the ability to approve the final version, defend the work against criticism, and declare conflicts of interest — that an AI tool cannot provide. The theoretical framework, the specific physical claims, and the decision to subject the work to public scrutiny all originated with the author. **Responsibility for all content of this paper rests with the author alone.**

Keywords: photonic quantum computing; linear optical quantum computing (LOQC); KLM protocol; magic-angle qubit; geometric decoherence protection; nodal cone; null-arc qubit; trisection qutrit; Z_3 symmetry protection; magic angle 54.74°; Poincaré sphere; two-photon absorption suppression; cross-Kerr nonlinearity; $\chi^{(3)} \sim \alpha^2$; silicon nitride photonics; thin-film lithium niobate; integrated photonic circuits; photon loss; Bell states; boson sampling; falsifiable QC predictions; thermal MZI test; room-temperature quantum computing; (3+3) spacetime; six-dimensional Einstein–Aether; third time dimension; S^2 compactification; Kaluza–Klein reduction

Abstract

We apply the (3+3) spacetime framework [1, 6] — a six-dimensional Einstein–Aether construction with three spatial and three temporal dimensions, the third time dimension compactified as a discrete two-sphere S^2 with $N\_cells = 2^{152}$ Planck-area cells — to photonic quantum computing. The framework identifies the photon as a **null arc** (a c-vector configuration with $v\_T = v\_t_2 = v\_t_3 = 0$ and $v\_spatial = c$) corresponding to the $n = 0$ Kaluza–Klein zero-mode on S^2. From this single identification three Level-1 structural results follow: **(i)** the photon's polarization Poincaré sphere is geometrically the third time dimension's S^2; single-qubit gates are physical S^2 rotations and the Born rule is Malus's law as S^2 projection; **(ii)** photons inherit no t_3-winding decoherence channels, with cosmological coherence floor $\tau\_cosm = 1/H_0 \approx 14.5\ Gyr$ [§14bis.46.4 of 6] and practical decoherence by photon loss as t_3 entrainment; **(iii)** the cross-Kerr nonlinearity scales as $\chi^{(3)} \sim \alpha^2 = 5.3 \times 10^{-5}$ from two successive c-vector rotations each of amplitude $\sqrt{\alpha}$.

We give a self-contained derivation in §5.4 of the $\chi^{(3)} \sim \alpha^2$ scaling from the Lagrangian-level photon-photon-via-radion vertex $L = -(g_0/4)\ \Phi^2\ F\_\{\mu\nu\}\ F\^\{\mu\nu\}|\_node$ with bare coupling $g_0 = (4/9)\ \varepsilon\_ZPE^2$, through radion integration-out at optical frequencies, to the four-loop running $g_0 \to \alpha$ (drawing on §16–§17 of [6]). Sections 5.5–5.9 establish the structural form of the nodal-cone TPA-suppression factor η_node: the breathing-mode contribution at the cone vanishes geometrically because the $Y\_\{2,0\}(\vartheta\_node) = 0$ identity that defines the cone latitude *also* eliminates the breathing's metric perturbation there (§5.7); the higher-multipole tail is bounded at $\eta\_node|\_higher \lesssim 10^{-5}$ using the framework's identified spectrum from [6] (§5.8); and the same Z_3 symmetry that organises SU(3), three-generation fermion structure, and tribimaximal PMNS mixing also protects the qutrit basis at the trisection vertices (§5.9). Combining these structural results with anchored medium-specific estimates for silicon-on-insulator ($K\_medium \approx 0.10–0.20$), silicon nitride ($0.03–0.08$, lowest), and thin-film lithium niobate ($0.05–0.15$) (Appendix B), the framework predicts $\eta\_node \approx K\_medium \approx 0.03–0.20$ across candidate platforms — silicon nitride is identified as the optimal platform — *dominated entirely by medium contributions* rather than framework-vacuum dynamics, and decisively testable by the §6.7 thermal MZI experiment (2027–2030). The magic-angle latitude $\vartheta\_node = \arccos(1/\sqrt{3}) = 54.74°$ is the same magic angle as in NMR magic-angle spinning, and the underlying physics (averaging-out of $Y\_\{2,0\}$ dipolar interactions) is structurally identical.

Six engineering proposals leverage these results: null-arc qubit encoding; the nodal-cone LOQC architecture (qubits at the magic-angle latitude where $Y\_\{2,0\}$ vanishes and the leading TPA channel is suppressed); (3+3) measurement protocols with t_3-entrainment-strength control; **trisection qutrit photonic computing** ($\log_2 3 \approx 1.585$ bits per photon, with Z_3*-symmetric* coherence protection giving $\eta\_qutrit \approx 0.06–0.21$ comparable to qubit η_node — a 58.5% information-density gain at essentially no extra coherence cost); boson sampling reframed as null-arc network interference; and a nodal-cone photonic-crystal integrated chip. Three testable predictions are organised into a year-resolved 2026–2030+ falsifiability timeline (§6.7). All claims are explicitly tagged Level 1 (geometric, derived), Level 2 (anchored prefactor), or Level 3 (engineering speculation); the room-temperature fault-tolerance projection of §9 is an explicit Level-3 conditional argument with dependency chain made explicit. A self-contained framework

digest is given in §1.7; readers unfamiliar with the (3+3) programme can read §1.7 alone to decide whether to continue.

§1 Introduction

Photonic quantum computing — realising universal quantum computation with photons as qubits — offers a compelling combination of natural advantages: photons are natively quantum, their polarization states form a perfect two-level system on the Poincaré sphere, they propagate at *c* without resistive loss, they do not require cryogenic cooling, and they can be distributed over long distances with low decoherence [9, 10]. The landmark Knill–Laflamme–Milburn (KLM) theorem [11] established in 2001 that linear optical quantum computing (LOQC) with photon-counting measurements is in principle universal and scalable. Gaussian boson sampling (GBS) has demonstrated quantum advantage with up to 219 photons across 216 modes [12, 13, 14], placing classical simulation beyond reach.

Yet photonic quantum computing faces persistent engineering challenges: photon-photon interactions are weak (the Kerr nonlinearity $\chi^{(3)} \sim \alpha^2 \sim 10^{-5}$ in standard materials), making deterministic two-qubit gates difficult; single-photon sources and detectors have imperfect efficiency; and photon loss accumulates with circuit depth, requiring redundant encoding. The fundamental question — why are photon-photon interactions so weak? — does not have a structural answer in standard electrodynamics.

The (3+3) spacetime framework [1, 6] provides this structural answer: $\chi^{(3)} \sim \alpha^2$ because two successive c-vector rotations on S^2 each have amplitude $\sqrt{\alpha}$. The same framework that explains the weakness of photon-photon interactions also identifies two geometric structures that ameliorate them: the **nodal cone** at $\vartheta\_node = \arccos(1/\sqrt{3}) = 54.74°$ on the Poincaré sphere, where Y_2^0 vanishes and the dominant two-photon absorption channel is suppressed; and the **trisection qutrit**, which encodes $\log_2 3 \approx 1.585$ bits per photon in three S^2-trisection states protected by the same threefold symmetry that fixes the Weinberg angle and the Koide identity in the framework's particle-physics sector [§10.2, §11.8 of 6].

Position relative to the (3+3) programme. This paper is a topical preprint in the (3+3) research programme. The foundational 6D Lagrangian construction is given in [6], with the book [1] providing the narrative development. Topical preprints have already been published on the proton-electron mass ratio [2], quantum computing in general [3], the Hubble tension [4], and the Born rule [5]. The present paper is a **companion preprint to the general-QC preprint [3]**, focusing specifically on the photonic-platform implications. The two preprints cross-reference each other; together they cover both matter-based and photonic quantum computing in the framework.

1.1 Structure of the paper

§1.7 is a self-contained one-page framework digest for readers unfamiliar with the (3+3) programme. §2 describes the methodology and the levels at which claims are made. §3 reviews the relevant (3+3) geometry for photonic qubits. §4 catalogues the engineering challenges of photonic QC and gives the (3+3) diagnostic for each. §5 develops the framework applied to KLM, photon loss, and the $\chi^{(3)}$ hierarchy: §5.4 derives the $\chi^{(3)} \sim \alpha^2$ scaling explicitly from the photon-photon-via-radion vertex; §5.5–§5.9 establish the structural form of the nodal-cone suppression factor *η_node*, prove its first-order vanishing at the cone via the $Y_{2,0}(\vartheta_{node}) = 0$ identity, bound the higher-multipole tail using the framework's identified spectrum, and extend the Z_3-symmetry analysis to the trisection qutrit basis. §6 presents six engineering proposals and three testable predictions, with §6.7 organising the predictions into a year-resolved 2026–2030+ falsifiability timeline. §7 compares photonic and matter-based QC platforms in the framework. §8 discusses why photonics is a natural (3+3) platform. §9 makes the room-temperature fault-tolerance argument as a Level-3 conditional argument with explicit dependency structure. §10 catalogues the honest limits — the speculative claims that the framework currently does not derive rigorously. §11 concludes. Appendix A records the version history; Appendix B gives anchored medium-specific *K_medium* estimates for silicon-on-insulator, silicon nitride, and thin-film lithium niobate platforms.

1.7 Framework at a glance

This subsection is a self-contained one-page digest of the (3+3) framework for readers unfamiliar with the programme. It consolidates the inputs, the key geometric identifications relevant to photonic quantum computing, the $\chi^{(3)} \sim \alpha^2$ scaling, and the sharpest near-term experimental targets into a single navigable summary. A reader who wants to evaluate the case being made before committing to the full paper can read §1.7 alone and decide whether to continue. The full programme treatment is in [1, 6]; the present paper is a topical application.

Inputs

The (3+3) framework has **one dimensionful input**, plus the discrete S^2 geometry, plus the cycle-2 cosmological hypothesis. Quantities sometimes treated as "inputs" in the framework's pedagogical presentations — m_p/m_e, α, and π — are in fact topological integers, derivable to high precision, or mathematical constants. They are *used* as inputs where convenient, but the framework derives or fixes them from the geometric postulate.

Type	Quantity	Status
Dimensionful scale (single input)	The Planck scale: $M_P = \sqrt{\hbar c/G_N}$, $L_P = \sqrt{\hbar G_N/c^3}$, G_N, equivalently m_e	One external value. The four expressions are related by exact algebraic identities; pick any one, the others follow. The framework's natural framing is Planck-anchored, since the discrete S^2 has 2^{152} Planck-area cells with L_P as the cell scale; m_e then follows from the one-loop self-consistency $m_e = \pi M_P / 2^{76}$ [§15 of 6].

Type	Quantity	Status
Geometry	Discrete S^2 with $N\_cells = 2^{152}$ Planck-area cells; $R_3 = \pi\hbar/(m\_e\ c) = \pi\ L\_P \cdot 2^{76}$	Specified by the self-referential bit condition: $\log_2 N\_cells = 152$ matches the SM's 152-bit field content [§9.4 of 1].
Hypothesis	Cycle-2 inheritance: orientation and energy content from previous cosmic cycle	Replaces inflation; testable via primordial B-mode null at LiteBIRD [§14bis.31 of 6]. Not directly relevant to photonic QC.

The strong claim of the parent framework. The (3+3) programme reduces the Standard Model's 19 free parameters plus 6 cosmological parameters (25 total) to one dimensionful scale plus geometry plus the cycle-2 hypothesis [§1.7 of 6]. No dimensionless input is fitted to data anywhere in the construction. The present paper inherits this parameter-free character: every photonic-QC prediction descends from the same geometry, with no additional optical-platform-specific inputs.

The photon as null arc

The framework identifies the photon as a **null arc** [§9 of 6]: a c-vector configuration with $v\_T = v\_t_2 = v\_t_3 = 0$ and $v\_spatial = c$, corresponding to the $n=0$ Kaluza–Klein zero-mode on the S^2 compactification of the third time dimension. From this single identification, the photon's defining quantum-information properties follow without additional assumptions:

- **No t_3-winding decoherence channels.** Photons have $n=0$: no Higgs-breathing ($\ell=2$ phonon) T_2 channel and no t_3-oscillation T_1 channel. The universal cosmological floor $\tau\_cosm = 1/H_0 \approx 14.5$ Gyr [§14bis.46.4 of 6] applies as an upper bound; for photons, even this bound is not approached in vacuum.
- **No cryogenic cooling required.** Photons have $v\_t_3 = 0$: no t_3 winding, no Compton oscillation, no rest mass. The 300 K thermal bath cannot excite a photon out of its null-arc state because there is no excited state above it in the t_3 direction.
- **Photon loss is the only practical decoherence channel.** A photon lost to an absorbing atom has its null arc entrained by that atom's t_3-winding mode — geometrically a T_1-type process on the absorbing atom, but a clean disappearance from the photonic Hilbert space.

The Poincaré sphere is the t_3 S^2

The Poincaré sphere of photon polarization states is geometrically identical to the third-time-dimension S^2 [established in §15.0 of 6 and developed in detail in the Born rule preprint [5]]. The identification is exact, with precise physical correspondence: north pole = right circular polarization (RCP, $|0\rangle$); south pole = left circular polarization (LCP, $|1\rangle$); equatorial points = linear polarizations at azimuth φ; the magic-angle latitude $\vartheta\_node = \arccos(1/\sqrt{3}) = 54.74°$ is the **nodal cone** where the $\ell=2$ spherical harmonic Y_2^0 vanishes (see below).

Single-qubit gates are physical rotations of S^2 (a quarter-wave plate is a 90° S^2 rotation). **The Born rule** for photon polarization measurement is Malus's law as S^2 projection: $P = \cos^2(\vartheta/2)$, derived through three independent geometric routes in [5]. **Bell-state entanglement** is a shared-t_3 orientation constraint between the two photons' null arcs. **Photon detection** is t_3 entrainment: the $N \gg 1$ atoms of a detector collectively pin the photon's null-arc orientation to the detector basis [§14bis.46–14bis.49 of 6].

The $\chi^{(3)} \sim \alpha^2$ scaling

The cross-Kerr nonlinearity $\chi^{(3)}$ (the dominant photon–photon interaction in standard nonlinear optics) is geometrically the amplitude for two successive c-vector rotations on S^2. Each rotation has amplitude $\sqrt{\alpha}$, where α itself is the framework-derived dimensionless coupling: at the compactification scale $g_0 = (4/9)\varepsilon^2 = 7.188 \times 10^{-3}$ (the cuboctahedral double projection of the Higgs oscillation axis onto the body diagonal), and after four self-consistency loops g_0 runs to $\alpha = 1/137.036$ at the electron-mass scale [§16, §17 of 6]. Successive c-vector rotations compound multiplicatively:

$\chi^{(1)} \sim \alpha^{(1/2)}$, $\chi^{(2)} \sim \alpha^1$, $\chi^{(3)} \sim \alpha^2 = 5.3 \times 10^{-5}$, $\chi^{(n)} \sim \alpha^{(n/2)}$.

This is the framework's explanation of *why* photon–photon interactions are weak: they are second-order in the geometric coupling $\sqrt{\alpha}$ of the null arc to the t_3 Higgs-breathing mode. In dielectric media the effective α is enhanced by reducing the effective Bohr radius $a_{0,eff}$, which raises $\chi^{(3)}$ multiplicatively (silica baseline → silicon waveguides → Si microrings; see §5.3). The full step-by-step derivation — from the photon-photon-via-radion Lagrangian operator $L = -(g_0/4)\Phi^2 F_{\mu\nu} F^{\mu\nu}|_{node}$ of [6] §16, through radion integration-out at optical frequencies, to the effective four-photon coupling — is given explicitly in §5.4 of this paper. §5.5 derives the structural form of the nodal-cone suppression factor η_{node} from the breathing-mode zero-point amplitude ε_{ZPE}, with four explicit prefactors deferred to foam-level work.

Sharpest near-term experimental targets

The framework yields three Level-2 testable predictions whose *direction* and *structural form* are derived but whose *magnitudes* depend on currently-anchored prefactors (§10). Each is decisive within its observation window:

Prediction (level)	Decisive experiment / setup	Window	Reference
KLM gate-fidelity local maximum at $\vartheta = 54.74°$ (Level 2)	Polarization-controlled KLM with ancillas encoded at the nodal cone vs. equatorial H/V basis; existing photonic platforms	2026–2027	§6.2
Trisection qutrit gate fidelity matches qubit-encoding threshold (Level 2)	Three-fold symmetric beam-splitter array on photonic-qutrit prototype (Xanadu/PsiQuantum-style hardware)	2026–2028	§6.4

Prediction (level)	Decisive experiment / setup	Window	Reference
Reduced *$d\varphi/dT$* phase drift at nodal-cone encoding (Level 2)	Lock-in thermal MZI comparison: nodal-cone-encoded vs. equator-encoded chips of identical fabrication	2027–2030	§6.6
Boson sampling beyond 200 photons consistent with null-arc network amplitude (Level 2)	Borealis-class GBS follow-ups; Hafnian-amplitude sampling at scale	2027–2029	§6.5
Integrated nodal-cone PhC chip demonstrating combined gain (Level 3)	Triangular-lattice photonic-crystal chip with nodal-cone qubit encoding; fabrication-stage challenge	2027–2030+	§6.6

A full year-resolved 2026–2030+ timeline organising these and related tests is in §6.7. The fault-tolerance projection of §9 is **Level 3 conditional**: it holds *if and only if* the open prefactor items (§10) close favourably, the multiplicativity assumption holds, and component efficiencies cross threshold. None of these is currently established; the framework's contribution is to make them plausible to investigate.

§2 Methodology and Claim Levels

2.1 AI-assisted theoretical analysis

This work follows the AI-assisted research protocol of the (3+3) programme. The author and Claude (Anthropic) collaborated on numerical verification (the cosmological decoherence-floor framework, KLM success-probability estimates, the structural form of the nodal-cone two-photon-absorption suppression, Hilbert-space dimensions for qutrit processors, boson-sampling complexity estimates, GBS photon-count analysis); exploratory analysis (the measurement-as-t_3-entrainment picture for KLM, the null-arc Bell-state geometry, the nodal-cone LOQC architecture); and critical evaluation against experimental literature on photonic quantum computing systems. Theoretical insights originated with the author. The AI verified calculations, cross-checked predictions against published experimental results, and provided structured scientific presentation. Claude is acknowledged in the AI Assistance Statement (front matter) as a research assistant; consistent with current AI-disclosure conventions of major journals, Claude is not listed as an author.

2.2 Claim levels

To keep the article's epistemic posture clear, claims are explicitly labelled at one of three levels:

- **Level 1 — Geometric identifications (strongest).** Direct consequences of the (3+3) framework as established in [1, 6]. Examples: Poincaré sphere = t_3 S^2; optical gates = S^2 rotations; Malus's Law = Born rule via S^2 projection; measurement = t_3 entrainment; $\chi^{(3)}$ ~

α^2 from two successive c-vector rotations. These are derived in the framework at structural level.

- **Level 2 — Derived predictions (testable).** Quantitative claims that the framework predicts in structural form, with specific magnitudes that depend on prefactors currently anchored to the framework rather than derived from first principles. Examples: the *structure* of nodal-cone TPA suppression (the framework predicts $Y_2^0 = 0$ at ϑ_node produces suppression; the *magnitude* of the suppression factor depends on ε_ZPE^2, currently anchored); the trisection qutrit SU(3) protection structure; the photon loss as t_3 entrainment scaling.
- **Level 3 — Engineering proposals (speculative).** Specific engineering implementations that combine multiple Level-2 predictions and additional engineering assumptions. Examples: integrated nodal-cone photonic-crystal chips with combined nonlinearity enhancement; mixed qubit-qutrit processors; the room-temperature fault-tolerance projection. These are proposals that *would follow* if all underlying Level-2 results derived their currently-anchored prefactors.

Claims throughout the paper are tagged with their level when ambiguity could arise. **§10 (Honest limits)** lists every Level-2 prefactor that is currently anchored rather than derived, and every Level-3 proposal whose plausibility depends on those anchors becoming derivations.

§3 The (3+3) Model: Photonic Qubit Geometry

3.1 The photon as null arc

In the (3+3) framework [§9 of 6], the photon is a **null arc**: its c-vector satisfies $v\_T = v\_t_2 = v\_t_3 = 0$ and $v\_spatial = c$. The photon corresponds to the $n=0$ Kaluza–Klein zero-mode on the S^2 compactification of the third time dimension; this is the same identification that, in the cosmological sector, makes dark matter the $n=0$ KK zero-mode of S^2 via the topological-invisibility theorem [§14bis.17–14bis.23 of 6] (with photons being distinguishable from dark matter by their spatial-direction propagation rather than their topological identity). From the null-arc identification, two structural properties for quantum computing follow without additional assumptions:

- **No t_3-winding decoherence channels.** Photons have zero t_3 winding ($n=0$), so there is no Higgs-breathing ($\ell=2$ phonon) T_2 channel and no t_3-oscillation T_1 channel. The cosmological decoherence floor for matter qubits is the universal $\tau\_cosm = 1/H_0 \approx 14.5$ Gyr [§14bis.46.4 of 6]; for photons, the absence of t_3-winding channels means this universal upper bound applies as a still weaker constraint. The photon is, at the level of fundamental coherence, an effectively perfect quantum information carrier in vacuum (Level 1).
- **No cryogenic cooling required.** Photons have $v\_t_3 = 0$: no t_3 winding, no Compton oscillation, no rest mass. The thermal bath at 300 K cannot excite a photon out of its null-arc state because there is no excited state above it in the t_3 direction. Photonic qubits are intrinsically room-temperature (Level 1).

The practical decoherence for photonic qubits is therefore not the cosmological floor but **photon loss** (linear absorption). A photon lost to an absorbing atom has its null arc entrained by that atom's t_3-winding mode — the null arc is absorbed into the atom's internal structure. This loss is not decoherence in the Hilbert-space-superposition sense (the photon is simply gone) but it is the primary limitation on photonic circuit depth.

3.2 The Poincaré sphere as the qubit's t_3 S^2

The Poincaré sphere of photon polarization states is geometrically the t_3 S^2 [established in §15.0 of 6 and developed in detail in the Born rule preprint [5]]. The identification is exact, with precise physical correspondence:

Poincaré state	Polarization	t_3 S^2 state	Qubit role
North pole	Right circular (RCP)	Spin-up null-arc orientation	$\|0\rangle$ = RCP. Standard computational basis state.
South pole	Left circular (LCP)	Spin-down null-arc orientation	$\|1\rangle$ = LCP. Antipodal qubit state.
Equator ($\varphi = 0$)	Horizontal (H)	*x*-direction superposition	$\|+\rangle$. Equal superposition of $\|0\rangle$ and $\|1\rangle$. Natural basis for beam-splitter operations.
Equator ($\varphi = \pi/2$)	Vertical (V)	*y*-direction superposition	$\|-\rangle$. Orthogonal superposition. With $\|+\rangle$ forms dual-rail qubit basis.
Latitude ϑ_node = 54.74°	Nodal-cone (elliptical)	$Y_2^0 = 0$ (the nodal cone)	**Protected qubit** (Level 2): zero ℓ=2 Higgs-mode coupling. Suppressed two-photon absorption; magnitude anchored (§10).

3.3 Bell states as shared t_3 geometric constraints

Entangled photon pairs in a Bell state share a t_3 geometric constraint: their null-arc orientations are locked in a specific relative configuration on S^2. The four Bell states correspond to the four distinct locked-arc configurations:

Bell state	Null-arc geometry	Physical description and (3+3) origin
$\|\Phi^+\rangle = (\|HH\rangle + \|VV\rangle)/\sqrt{2}$	Both arcs parallel (same equatorial angle φ)	Two null arcs locked in the same t_3 orientation. Co-propagating circular polarizations. Produced by Type-I SPDC.
$\|\Phi^-\rangle = (\|HH\rangle - \|VV\rangle)/\sqrt{2}$	Arcs anti-phase (opposite equatorial sign)	Two null arcs locked in opposite t_3 phase. Relative arc angle = π. Interferometrically distinguishable from $\|\Phi^+\rangle$.
$\|\Psi^+\rangle = (\|HV\rangle + \|VH\rangle)/\sqrt{2}$	Arcs antipodal on S^2 (perpendicular)	One arc at north, one at south pole (RCP+LCP superposition). Total angular momentum = 0. Produced by Type-II SPDC.
$\|\Psi^-\rangle = (\|HV\rangle - \|VH\rangle)/\sqrt{2}$	Arcs antipodal, phase-shifted	Antisymmetric antipodal configuration. The only Bell state with angular momentum $L = 0$ and phase antisymmetry. Most sensitive to path-length imbalance.

Bell-state entanglement is therefore in (3+3) a shared-t_3-mode constraint of the type derived more generally in §14bis.49 of [6] — entanglement as a shared t_3 mode dissolving the EPR paradox without superluminal signalling. (Level 1.)

3.4 Key (3+3) formulas for photonic QC

Formula	Value / source	Relevance to photonic QC
$\tau\_cosm = 1/H_0 \approx 14.5$ Gyr (universal floor; matter qubits)	[§14bis.46.4 of 6]	Universal cosmological decoherence floor. For photons, no t_3-winding channels means even this upper bound is not approached in vacuum.
$\chi^{(3)} \sim \alpha^2$ (vacuum)	5.3×10^{-5} [§5 of 6]	Sets photon-photon gate strength. Two successive c-vector rotations each at amplitude $\sqrt{\alpha}$.
$\vartheta\_node = \arccos(1/\sqrt{3}) = 54.74°$	Exact nodal cone [§15.0, §16 of 6]	**Level 2** protected qubit encoding: $Y_2{}^0(\vartheta\_node) = 0$ exactly; the ℓ=2 TPA channel vanishes at the angle. Magnitude of suppression factor is currently anchored (§10).

Formula	Value / source	Relevance to photonic QC
$P = \cos^2(\vartheta/2)$ (Born rule = Malus's Law)	S^2 projection [§14bis.46 of 6, ref [5]]	Single-photon detection probability; basis for all LOQC gate-success calculations.
$\log_2 3 \approx 1.585$ bits/photon (trisection)	S^2 trisection [§10.2, §11.8 of 6]	Qutrit Hilbert-space dimension advantage: 3^N vs 2^N for N photons.

§4 Engineering Challenges in Photonic Quantum Computing

We catalogue the principal engineering challenges in photonic QC and give the (3+3) diagnostic for each. The diagnostics are Level-1 geometric identifications (the structural cause), with magnitudes for ameliorations flagged Level 2 when they require currently-anchored prefactors.

Challenge	Current status	Root cause	(3+3) diagnostic
Weak photon-photon interaction	$\chi^{(3)} \sim 10^{-20}$ m²/W (silica)	Gauge invariance; perturbative QED	**Level 1.** Two successive c-vector rotations each with amplitude $\sqrt{\alpha} = 1/11.7 \Rightarrow \chi^{(3)} \sim \alpha^2 = 5.3 \times 10^{-5}$. Geometric origin (§5 of 6).
Photon loss (absorption)	0.2 dB/km fibre; 0.5–2 dB/cm chip	Material absorption; waveguide scattering	**Level 1.** Absorbing atom entrains photon's t_3 null-arc orientation ⇒ photon's c-vector rotates to matter state. Irreversible T_1-type process on the absorber.
Imperfect single-photon sources	$\sim 90\%$ efficiency (quantum dot)	Multi-photon emission; dephasing	**Level 1.** Source must emit one and only one null arc per trigger. Multi-photon = residual t_3 winding not fully evacuated.
LOQC gate success probability	$1/4$ (basic KLM); → 1 with n ancillas	No deterministic photon-photon interaction in linear optics	**Level 1.** KLM uses t_3 entrainment (measurement) as the nonlinearity. The $1/4$

Challenge	Current status	Root cause	(3+3) diagnostic
			probability is the S^2 projection probability for a random ancilla null-arc state.
Two-photon absorption (TPA)	Limits Si waveguide power (10–100 mW)	Two photons absorbed by single atom via *ℓ=2* virtual state	**Level 1+2.** TPA = two null arcs coupling simultaneously via *ℓ=2* Higgs mode. *Structurally* suppressed at the nodal cone where $Y_2^0 = 0$. *Magnitude* of suppression currently anchored (§10).
Photon indistinguishability	~ *99%* (best QD sources)	Phonon dephasing; spectral wander	**Level 1+2.** Distinguishable photons = null arcs with different azimuthal phase φ on S^2. Phonon coupling rotates φ between emissions; nodal cone *structurally* suppresses this; magnitude anchored (§10).
Scalable photon number	~ *219* photons (GBS record)	Loss accumulates; detection probability falls exponentially	**Level 1.** Total success prob. = η^n where η is per-photon efficiency. Scaling depends on whether nodal-cone enhancements (Level 2/3) materialise.

§5 The (3+3) Framework Applied to Photonic Quantum Computing

5.1 KLM as t_3 entrainment nonlinearity

The central insight of the KLM theorem [11] is that measurement provides the nonlinearity needed for universal quantum computing, replacing the absent photon-photon interaction. In (3+3), this has a precise geometric meaning.

Measurement = t_3 entrainment [§14bis.46–14bis.49 of 6, [5]]. A photon detector is $N \gg 1$ atoms whose t_3 axes are in definite states. When a photon's null arc encounters the detector, the N atoms collectively entrain the null arc to the detector's t_3 axis. This entrainment is the nonlinear element of KLM: it projects the photon's Poincaré-sphere state onto the detector basis. The success probability $P = \cos^2(\vartheta/2)$ (Born rule = Malus's Law on S^2, derived via three independent geometric routes in [5]) is the geometric projection of the ancilla null arc onto the detector axis. (Level 1.)

The KLM protocol requires ancilla photons in specific states, beam splitters (partial S^2 rotations), and detectors (t_3-entrainment events). The KLM gate success probability of $1/4$ for the simplest case arises from the Born rule: the ancilla photon has a $1/4$ probability of projecting onto the correct outcome for the gate to succeed on the first attempt. With n ancilla photons, the success probability approaches 1 exponentially, because multiple independent t_3-entrainment events provide near-certain gate success via teleportation.

The (3+3) picture also reveals why feed-forward is essential in KLM: when a measurement (t_3 entrainment) yields an unexpected outcome, the computational qubit's null arc has been rotated incorrectly. The feed-forward correction is a subsequent c-vector rotation (waveplate or electro-optic modulator operation) that undoes the incorrect entrainment. Without feed-forward, measurement outcomes cause irreversible null-arc misalignments.

5.2 Photon loss as t_3 null-arc absorption

Photon loss in a waveguide or fibre has a precise (3+3) mechanism: an absorbing atom (with t_3 winding $n = 1/2$ or $n = 1$) captures the photon's null arc. The c-vector rotation amplitude for this event is $\sqrt{\alpha}$ (one c-vector rotation). After absorption, the photon's null arc no longer exists as a propagating mode: it has been incorporated into the atom's excited t_3-winding state. (Level 1.)

This loss mechanism is polarization-independent for *linear* absorption (the electron can be promoted from any ground-state orbital regardless of photon polarization at the one-rotation level). This is why conventional linear loss in fibre is polarization-independent. In contrast, *two-photon absorption* (TPA) requires two simultaneous c-vector rotations through the $\ell=2$ Higgs channel: its amplitude scales as α (two rotations), and it *is* polarization-dependent. Specifically, TPA is proportional to $|Y_2^0(\vartheta)|^2$ where ϑ is the photon's Poincaré-sphere latitude, giving (Level 1) $TPA = 0$ at $\vartheta = \vartheta\_node = 54.74°$. The *structural* suppression at the nodal cone is exact; the *residual magnitude* off the cone (and hence the *practical* suppression factor at finite ZPE width) is Level 2 and depends on a currently-anchored prefactor (§10).

5.3 The $\chi^{(3)}$ hierarchy and engineering implications

From the (3+3) framework: each c-vector rotation has amplitude $\sqrt{\alpha}$. Successive rotations compound: $\chi^{(1)} \sim \alpha$^(1/2), $\chi^{(2)} \sim \alpha$, $\chi^{(3)} \sim \alpha^2$. The cross-Kerr coupling for photon-photon interaction scales with $\alpha\_eff = R_3/(\pi\, a_0,eff)$. The explicit Lagrangian-level derivation of these scalings — from the photon-photon-via-radion vertex of [6] §16, through radion integration-out at optical frequencies, to the effective four-photon coupling — is given in §5.4 below; the present subsection takes the scalings as established and develops their engineering implications. Increasing the effective α_eff by reducing the effective Bohr radius a_0,eff in the propagating medium increases $\chi^{(3)}$:

Enhancement method	***α_eff/α***	**$\chi^{(3)}$ enhancement**	**Photons for *π*-shift**	**(3+3) mechanism / level**
Silica fibre (baseline)	*1×*	*1×*	$\sim 10^{10}$	Standard $\alpha = 1/137$. **Level 1**.
Silicon waveguide (*n = 3.5*)	*~ 12×* (anchored)	*~ 150×*	$\sim 10^{8}$	Effective a_0 reduced. **Level 2** — the scaling form is structural; the *12×* prefactor is empirical.
Si microring (*Q = 10⁶*)	*~ 1*	$\sim 10^{6}$ ×	$\sim 10^{3}$	Resonant field enhancement (standard photonics). **Level 1** scaling.
Quantum dot (exciton)	*~ 10×* (anchored)	*~ 100×*	$\sim 10^{7}$	Exciton wavefunction confinement. **Level 2**.
Si microring + nodal cone	*~ 12×*	*~ 150 × η_node*	depends on *η_node*	**Level 3** — combined estimate; *η_node* is the nodal-cone suppression factor whose magnitude is currently anchored (§10).

The combined-enhancement claim in the bottom row is **Level 3 (engineering proposal)**, not a derived prediction. It depends on (a) the magnitude of the nodal-cone suppression factor, currently anchored; (b) the assumption that the resonant and nodal-cone enhancements compose multiplicatively, which has not been derived; and (c) the practical assumption that material-science factors do not introduce additional losses that erode the combined gain.

5.4 Derivation of the $\chi^{(3)} \sim \alpha^2$ scaling

Section 5.3 used the $\chi^{(3)} \sim \alpha^2$ scaling and the more general $\chi^{(n)} \sim \alpha$^(n/2) hierarchy as established results from the (3+3) framework. The argument is short enough to give explicitly here, so the reader does not need to track the full Lagrangian construction in [6] to follow the optical-QC application. The derivation has four logical steps: (i) the geometric origin of α in the framework; (ii) the photon-photon-via-radion vertex at Lagrangian level; (iii) the effective four-photon coupling obtained by integrating out the radion at optical frequencies; (iv) material enhancement and the

$\alpha\_eff = R_3/(\pi a_0,eff)$ form. A fifth subsection summarises what this section establishes at Level 1 and what it does not.

5.4.1 The geometric origin of α (Level 1, inherited from [6])

In the (3+3) framework, α is not an input. The bare electromagnetic coupling at the compactification scale $E_3 = m\_e\, c^2/\pi = 0.163$ MeV is [§16 of 6, Eq. 16.4 and §14bis.3 Table]:

$$g_0 = (4/9)\ \varepsilon^2 = 7.188 \times 10^{-3} \quad (5.4.1)$$

The two factors are derived geometrically:

- **The 4/9 factor.** Equal to $\sin^4(\vartheta\_node) = (2/3)^2$ with $\vartheta\_node = \arccos(1/\sqrt{3}) = 54.74°$, the angle between a Cartesian axis ($\hat{x}$, the Higgs-breathing oscillation direction) and the cuboctahedral body diagonal $(1,1,1)/\sqrt{3}$. The factor $\sin^2(\vartheta\_node) = 2/3$ is the projection of the Higgs oscillation axis onto the S^2 plane perpendicular to the body diagonal — equivalently, the latitude where the breathing-mode harmonic Y_2^0 has its zeros. Two successive projections — required to close a photon-radion-photon coupling at the nodal cone — give $(2/3)^2 = 4/9$. The cuboctahedral identification is derived in [6] §16.4; the Y_2^0 nodal-cone identification in [6] §16.3.
- **The ε² factor.** Equal to $(m\_H/(4v))^2 \approx 0.01617$ with $m\_H = 125.10$ GeV the Higgs mass and $v = 246.22$ GeV the Higgs vacuum expectation value. The Higgs VEV is itself closed in [6] §15 Part II at 1.0% via the unification $v \cdot E_3 = \Lambda\_QCD^2$ giving $v = 4m\_p^2/(9\pi\ m\_e) = 243.7$ GeV. The factor ε is the dimensionless **zero-point breathing displacement** of the S^2 radion mode. (Notation: in [6] §14bis.66 and in §5.5–§5.6 of this paper, the same quantity is written $\varepsilon\_ZPE = m\_H/(4v) = 0.1270$. The two symbols ε and ε_ZPE are interchangeable; we use ε in §5.4 to match the [6] §16 vertex-derivation context and ε_ZPE in §5.5–§5.6 to match the [6] §14bis Casimir-effect context.)

Both factors are therefore geometric: the $4/9$ from cuboctahedral S^2 topology and the ε^2 from the Higgs VEV which is closed by S^2 mass-scale geometry. The bare coupling g_0 has zero free parameters.

Running g_0 from the compactification scale E_3 to the electron-mass scale $m\_e\, c^2 = 0.511$ MeV through four self-consistency loops gives [§17 of 6]:

$$1/g_0 = 139.12 \rightarrow 1/\alpha = 137.036 \quad (5.4.2)$$

a $+1.74\%$ net shift, with the four loops contributing $+0.74\%$ (Loop 1: boundary stability), $+0.42\%$ (Loop 2: pool equilibrium), $+0.36\%$ (Loop 3: coupling self-consistency), and -0.07% (Loop 4: ε^2 breathing and 2-loop Coleman–Weinberg corrections); see [6] §17.1–§17.5 for the explicit RG identification of each loop. The numerical agreement with the observed α at 7 ppb [6] §17.6 is the precision check on the four-loop closure.

The framework's identification of α is therefore:

α IS the squared amplitude for one c-vector rotation through the Higgs-breathing P_2 nodal-cone projection, RG-evolved from the compactification scale to the electron-mass scale.

This makes the v0.2 §1.7 statement "each c-vector rotation has amplitude $\sqrt{\alpha}$" a *derived* identification rather than a postulate. The vertex amplitude $\sqrt{\alpha}$ of standard QED — which in the SM is a free parameter fixed by measurement — is identified geometrically with the framework's coupling between the photon's null arc and the S^2-compactified Higgs-breathing mode at the cuboctahedral nodal cone.

5.4.2 The photon-photon-via-radion Lagrangian vertex (Level 1)

The Lagrangian operators that govern photon-radion and photon-photon coupling at the nodal cone are obtained from the Kaluza–Klein reduction of the 6D action on $S^2(R_3)$. With Φ the breathing-mode field (the $\ell = 2$, $m = 0$ spherical harmonic Y_2^0 of S^2) and $F_{\mu\nu}$ the 4D photon field strength, two operators arise [§16 of 6, Eqs. 16.3–16.4]:

$$L_{\gamma\Phi} = -\kappa\, \Phi\, A_\mu J^\mu\, |_{node} \quad \text{(5.4.3, photon–radion vertex)}$$

$$L_{\gamma\gamma\Phi^2} = -(g_0/4)\, \Phi^2 F_{\mu\nu} F^{\mu\nu}\, |_{node} \quad \text{(5.4.4, photon–photon-via-radion vertex)}$$

where the notation $|_{node}$ means the operator is integrated over S^2 with the breathing-mode harmonic profile evaluated at its nodal cone $\vartheta_{node} = 54.74°$. The coefficient $\kappa = \sqrt{2/3} \cdot \varepsilon = 0.0599$ is the single-radion-insertion coupling, computed from the gradient $|\partial_\vartheta Y_2^0|_{node} = \sqrt{2} \cdot \varepsilon$ and the $Y_1^0 \propto \sin\vartheta$ photon-mode profile evaluated at $\sin(\vartheta_{node}) = \sqrt{(2/3)}$; see [6] §16.5 for the explicit S^2 integral. The double-radion-insertion coupling satisfies the exact identity:

$$g_0 = 2\kappa^2 = (4/9)\, \varepsilon^2 \quad (5.4.5)$$

which follows from the antipodal-crossing-point counting of the cuboctahedral body diagonals [6] §16.1. The *exact* nature of this identity is significant: the *κ-vertex squared = g_0-vertex* relation is geometric, not perturbative, and links the (one-photon, one-radion) and (two-photon, two-radion) processes by the single shared geometric structure of the nodal cone.

The structural content of (5.4.4) is: **two photons can interact via two radion exchanges at the nodal cone, with dimensionless strength g_0.** This is the Lagrangian-level origin of photon-photon coupling in the (3+3) framework, and the starting point for the derivation of $\chi^{(3)}$.

5.4.3 The effective four-photon coupling (Level 1)

To obtain the effective four-photon interaction relevant to optical nonlinearity, integrate out the radion Φ at energies $\omega \ll m_H$. The bare radion is the $\ell = 2$ breathing mode with $m_{H,bare}/M_P = \sqrt{(6/N_{cells})} = \sqrt{6}/2^{76} \approx 3.2 \times 10^{-23}$ [6] Eq. 15.0d, and the physical Higgs mass is $m_H = 125.10$ GeV after Coleman–Weinberg corrections [6] §9. At optical frequencies $\omega \sim 1$ eV, the condition $\omega \ll m_H$ is satisfied by 11 orders of magnitude, so the radion can be treated as effectively heavy and integrated out at tree level.

Tree-level integration of (5.4.4) with two photon legs on each external side and a Φ^2 propagator generates the effective four-photon operator:

$$L\_eff,\gamma\gamma\gamma\gamma \sim (g_0^2 / m\_H^2) \cdot (F\_\mu\nu\ F\hat{}\mu\nu)^2 \quad (5.4.6)$$

The dimensionless strength of this operator, expressed as a coefficient on the dimensionless photon field strength at the optical scale, is g_0^2 (the m_H^2 factor sets the cutoff scale at which the four-photon vertex resolves into the underlying two-radion exchange; at $\omega \ll m\_H$ it is absorbed into the field-strength normalisation). After RG evolution $g_0 \to \alpha$ through the four-loop running of (5.4.2), the effective four-photon coupling has dimensionless strength α^2.

Because $\chi^{(3)}$ is the third-order susceptibility — the coefficient of the four-photon term in the medium polarisation expansion $P\hat{}(3) = \varepsilon_0\ \chi^{(3)}\ E^3$ — and is structurally proportional to the dimensionless four-photon coupling, the framework's prediction is:

$$\chi^{(3)} \sim \alpha^2 = (1/137.036)^2 = 5.32 \times 10^{-5} \text{ (Level 1, structural) } (5.4.7)$$

This matches the well-known order-of-magnitude scaling $\chi^{(3)}/\chi^{(1)} \sim 10^{-5}$ in standard dielectrics [25] and reproduces the conventional QED counting (each electron–photon vertex contributes $\sqrt{\alpha}$ to the amplitude; a four-photon process has four vertices, giving amplitude α^2) in the framework's geometric language. What standard QED treats as a perturbative counting rule is here a structural consequence of the cuboctahedral S^2 geometry.

The general hierarchy $\chi^{(n)} \sim \alpha\hat{}(n/2)$ follows by extending (5.4.6) to higher-order operators: an *(n+1)*-photon process requires n successive c-vector rotations with vertex amplitude $\sqrt{\alpha}$ each, giving total amplitude $\alpha\hat{}(n/2)$. Equivalently, an effective *2(n+1)*-field operator built from n radion exchanges scales as $g_0^n \to \alpha^n$ in dimensionless strength, and the corresponding susceptibility $\chi^{(n)}$ is the coefficient of an *(n+1)*-field term, proportional to the square root of the dimensionless coupling — hence $\alpha\hat{}(n/2)$. The pattern matches the standard nonlinear-optics scaling across all measured orders [25].

5.4.4 Material enhancement: the α_eff = $R_3/(\pi a_0$,eff) form

In dielectric media, electron orbitals are screened and modified by the host lattice, and the **effective Bohr radius** a_0,eff is reduced from the vacuum value $a_0 = 5.29 \times 10^{-11}$ m. The geometric coupling for photon–electron interaction at the foam-cell scale rescales accordingly. Define:

$$\alpha\_eff = R_3 / (\pi\ a_0,eff) \quad (5.4.8)$$

where $R_3 = \pi\hbar/(m\_e\ c) = 2.43 \times 10^{-12}$ m is the framework's (3+3) compactification radius. This is the **Level-1 structural form**: it follows from the same geometric identification $m\_e = \pi\hbar/(c\ R_3)$ that gives α in vacuum, but with the in-medium effective electron-orbital scale replacing the vacuum Bohr radius. In vacuum, $a_0,eff = a_0$ and (5.4.8) reduces to the framework identification $\alpha = R_3/(\pi a_0) = 1/137.036$.

The $\chi^{(3)}$ enhancement in a medium is then:

$$\chi^{(3)}\_medium / \chi^{(3)}\_vacuum \sim (\alpha\_eff/\alpha)^2 \cdot \eta\_resonance \cdot \eta\_node \quad (5.4.9)$$

where $\eta\_resonance$ is a resonant-cavity field-enhancement factor (standard photonics, scaling as Q^2 for a microring resonator with quality factor Q; Level 1) and η_node is the nodal-cone TPA suppression factor (Level 2 anchored magnitude, see §10 item A1).

The specific numerical prefactors — the ∼ *12*× enhancement of α_eff/α for silicon vs. silica; the ∼ *10*× enhancement for QD excitons — are anchored to standard photonics measurement rather than derived first-principles from (3+3); they are listed as Level-2 anchors A2 in §10. The structural form (5.4.8) is Level-1; the prefactor magnitudes in specific dielectric environments are Level-2.

5.4.5 What §5.4 establishes and what it does not

Established at Level 1 in this paper (i.e., explicitly derived in §5.4.1–§5.4.4 with the parent paper [6] cited only for prior closures):

- The $\chi^{(3)} \sim \alpha^2$ scaling as a structural prediction following from the photon–photon-via-radion vertex (5.4.4) and tree-level radion integration-out at optical frequencies.
- The general $\chi^{(n)} \sim \alpha^{(n/2)}$ hierarchy from successive c-vector rotations.
- The $\alpha\_eff = R_3/(\pi a_0,eff)$ functional form for in-medium enhancement.
- The exact identity $g_0 = 2\kappa^2$ linking single- and double-radion-insertion couplings.

Inherited from the parent paper [6] (Level 1 results not re-derived here):

- The geometric derivation of $g_0 = (4/9)\,\varepsilon^2$ at the compactification scale [6] §16.
- The four-loop running $g_0 \rightarrow \alpha$ at 7 ppb [6] §17.
- The closure of the Higgs VEV $v = 4m\_p^2/(9\pi\ m\_e) = 243.7$ GeV at 1.0% [6] §15 Part II.
- The cuboctahedral identification of $\sin^2(\vartheta\_node) = 2/3$ [6] §16.4.

Not established at Level 1 (the subjects of §10 items A and B):

- The numerical prefactors for specific dielectric media (Level 2; §10 A2): the *12*× for silicon vs. silica and the *10*× for QD excitons are anchored to standard photonics, not derived from (3+3).
- The multiplicativity assumption $\eta\_resonance \cdot \eta\_node$ in (5.4.9) — that the resonant and nodal-cone enhancements compose multiplicatively rather than sub-linearly (Level 3; §10 B1).
- The magnitude of the nodal-cone suppression factor η_node itself: the structural form is Level-1 (Y_2^0 vanishes exactly at ϑ_node), but the practical magnitude in a finite-ε_ZPE environment is Level-2 (§10 A1).

The specialist priority of the present programme is to convert the last two items into derived Level-1 results: rigorous foam-level computation of the nodal-cone suppression magnitude from the breathing-mode zero-point spread ε_ZPE, and rigorous computation of the multiplicativity coefficient. The structural form of the first item — the functional dependence $\eta\_node = f(\varepsilon\_ZPE) \cdot (geometric\ prefactors)$ — is now derived in §5.5 below ; only the four explicit O(1) prefactors remain Level-2 anchored, deferred to specialist work. The multiplicativity coefficient remains entirely open (§10 B1).

5.5 Structural derivation of the nodal-cone suppression factor

Section 10 item A1 of v0.3 listed the magnitude of the nodal-cone suppression factor *η_node* as the most consequential single open item in the article: the framework predicts at Level 1 that $Y_2^0(\vartheta\_node) = 0$ exactly — so the leading ℓ=2-mediated TPA channel is structurally suppressed at the cone — but the *magnitude* of the residual TPA in a real device with finite zero-point breathing fluctuations was anchored to standard photonics rather than derived from (3+3). This subsection sets up the structural derivation: it identifies the geometric form of *η_node* as a specific function of the breathing-mode zero-point amplitude $\varepsilon\_ZPE = m\_H/(4v) = 0.1270$ of [6] §14bis.66, computes the higher-multipole geometric ratios $|Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2$ exactly (pure spherical-harmonic algebra), and isolates the residual O(1) prefactors to be computed in specialist work. The deliverable is a master formula for *η_node* in which the same Level-1 geometric content as the rest of §5 of this paper is made explicit, with the residual Level-2 anchoring confined to four numerical coefficients.

5.5.1 Setup: the TPA matrix element and η_node definition

The two-photon-absorption rate at a polarization-state encoding centred at latitude ϑ on the Poincaré sphere (= $t_3\ S^2$) is, in the framework, proportional to the squared matrix element of the leading ℓ=2 Higgs-breathing-mediated coupling derived in §5.4.2:

$$\Gamma\_TPA(\vartheta) \propto |\langle Y_2^0(\vartheta)\rangle|^2 \times \Gamma_0 \quad (5.5.1)$$

where Γ_0 collects the (large) prefactors common to all encodings (medium absorption coefficient, photon flux, single-photon detuning, etc.) and $\langle Y_2^0(\vartheta)\rangle$ denotes the matrix element averaged over the photon's polarization-state distribution centred at ϑ. Define the nodal-cone suppression factor as the ratio of TPA rates at the two reference encodings:

$$\eta\_node \equiv \Gamma\_TPA(\vartheta\_node) / \Gamma\_TPA(\pi/2) \quad (5.5.2)$$

so that $\eta\_node \ll 1$ corresponds to strong TPA suppression at the cone compared to standard equator (H/V-basis) encoding. In the idealised limit of a perfectly localised polarization state at ϑ_node, the framework gives $\langle Y_2^0(\vartheta\_node)\rangle = 0$ and $\eta\_node = 0$ (exact suppression, Level 1 framework prediction). In a real device the polarization state is not perfectly localised; *η_node* is the residual TPA after the perfect-suppression limit is corrected for finite spread. The task is to derive the structural form of this residual.

5.5.2 Leading structural term: zero-point smearing of the photon's effective ϑ

The photon's polarization-state distribution centred at ϑ_node has finite angular width $\delta\vartheta$. Expanding the Y_2^0 matrix element to leading order in $\delta\vartheta$:

$$\langle Y_2^0(\vartheta\_node + \delta\vartheta)\rangle = Y_2^0(\vartheta\_node) + (dY_2^0/d\vartheta)|\_node \cdot \langle\delta\vartheta\rangle + (1/2)(d^2Y_2^0/d\vartheta^2)|\_node \cdot \langle(\delta\vartheta)^2\rangle + \ldots$$

The first term vanishes by construction (Level 1, $Y_2^0(\vartheta\_node) = 0$). For a symmetric distribution centred at the cone, $\langle\delta\vartheta\rangle = 0$ and the second term vanishes; the leading non-vanishing contribution to $|\langle Y_2^0\rangle|^2$ is the gradient term squared:

$$|\langle Y_2^0(\vartheta\_node + \delta\vartheta)\rangle|^2 \approx |(dY_2^0/d\vartheta)|\_node|^2 \cdot \langle(\delta\vartheta)^2\rangle \quad (5.5.3)$$

Computing the gradient at the cone explicitly. With $Y_2^0(\vartheta) = (1/4)\sqrt{(5/\pi)}(3\cos^2\vartheta - 1)$:

$$dY_2^0/d\vartheta = -(3/2)\sqrt{(5/\pi)} \sin\vartheta \cos\vartheta$$

Evaluated at $\cos(\vartheta\_node) = 1/\sqrt{3}$, $\sin(\vartheta\_node) = \sqrt{(2/3)}$ (so $\sin\vartheta \cos\vartheta|\_node = \sqrt{2}/3$; cf. [6] Eq. 16.9):

$$(dY_2^0/d\vartheta)|\_node = -(3/2) \cdot (\sqrt{2}/3) \cdot \sqrt{(5/\pi)} = -(1/\sqrt{2})\sqrt{(5/\pi)} \quad (5.5.4)$$

so $|(dY_2^0/d\vartheta)|\_node|^2 = 5/(2\pi)$. At the equator, $Y_2^0(\pi/2) = -(1/4)\sqrt{(5/\pi)}$ and $|Y_2^0(\pi/2)|^2 = 5/(16\pi)$. The ratio entering η_node is:

$$\eta\_node|\_breathing = |\langle Y_2^0(\vartheta\_node + \delta\vartheta)\rangle|^2 / |Y_2^0(\pi/2)|^2 = (5/(2\pi)) / (5/(16\pi)) \cdot \langle(\delta\vartheta)^2\rangle = 8 \langle(\delta\vartheta)^2\rangle \quad (5.5.5)$$

The factor 8 is a derived geometric coefficient. It follows entirely from spherical-harmonic algebra and the specific cone latitude $\vartheta\_node = \arccos(1/\sqrt{3})$. It contains zero free parameters and zero anchoring to standard photonics. This is the central Level-1 result of §5.5.

The angular variance $\langle(\delta\vartheta)^2\rangle$ has the structural form:

$$\langle(\delta\vartheta)^2\rangle = k\_2 \cdot \varepsilon\_ZPE^2 \quad \text{(5.5.6, structural ansatz)}$$

where k_2 is an O(1) prefactor whose computation requires foam-level dynamics . The ε_ZPE^2 scaling is the leading-order natural identification: any zero-point fluctuation of the photon's effective polarization-state location on S^2 inherits the energy scale m_H (the only S^2-internal mass scale relevant at optical energies $\omega \ll m\_H$), and dimensional analysis from the framework's coupling structure of §5.4.2 gives the leading-order amplitude as $\varepsilon\_ZPE = m\_H/(4v) = 0.1270$. Higher-order corrections O(ε_ZPE^4) are deferred. Combining (5.5.5) and (5.5.6):

$$\boldsymbol{\eta\_node|\_breathing = 8\, k\_2\, \varepsilon\_ZPE^2 \approx 0.130\, k\_2} \quad \text{(5.5.7, Level-1 structural form, k\_2 deferred to specialist work)}$$

This is the closure of the §10 A1 sub-item (i) — the zero-point spread contribution. It converts a previously Level-2 anchored quantity into a Level-1 structural form whose only residual anchoring is the single O(1) prefactor k_2.

5.5.3 Higher-multipole residuals at the nodal cone

TPA can also be mediated by higher-ℓ multipoles of S^2 ($\ell = 4, 6, 8, \ldots$). For these modes $Y\_\ell^0(\vartheta\_node) \neq 0$ in general, so they contribute residually to the matrix element even when the photon's polarization is sharply localised at ϑ_node. The geometric ratios are computed straightforwardly from the standard Legendre polynomials $P\_\ell(x) = Y\_\ell^0(\vartheta)\cdot\sqrt{(4\pi/(2\ell+1))}$ evaluated at $x = \cos(\vartheta\_node) = 1/\sqrt{3}$ vs $x = \cos(\pi/2) = 0$:

ℓ	$P\_\ell(1/\sqrt{3})$	$P\_\ell(0)$	$|Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2$	Decimal
2	0	$-1/2$	0 (exact, Level 1)	0.000

ℓ	$P_\ell(1/\sqrt{3})$	$P_\ell(0)$	$\lvert Y_\ell^0(\vartheta_{node})\rvert^2/\lvert Y_\ell^0(\pi/2)\rvert^2$	Decimal
4	−7/18	3/8	784/729	1.075
6	2/9	−5/16	1024/2025	0.506
8	11/72	35/128	30976/99225	0.312

The geometric ratios are pure mathematics: they depend only on the cone latitude $\vartheta_{node} = \arccos(1/\sqrt{3})$ and the spherical-harmonic decomposition. They do **not** vanish at ϑ_{node} — only the $\ell=2$ mode has a node there. The $\ell=4$ mode is in fact slightly *enhanced* at the cone compared to the equator (*ratio 1.075*); $\ell=6$ and $\ell=8$ modes are partially suppressed (*0.506* and *0.312* respectively).

What controls the higher-multipole contribution to η_{node} is therefore not the geometric ratios but the *coupling structure*: each higher-ℓ multipole enters the four-photon process with its own coupling prefactor and its own propagator suppression by the higher mass m_ℓ. The structural form is:

$$\eta_{node}|_{higher\text{-}multipole} = \Sigma_{\{\ell \text{ even}, \ell\geq 4\}}\, k_\ell \cdot |Y_\ell^0(\vartheta_{node})|^2/|Y_\ell^0(\pi/2)|^2 \cdot S_\ell \quad \text{(5.5.8, structural form)}$$

where k_ℓ is the multipole coupling prefactor (open) and S_ℓ is the propagator/coupling suppression factor (also deferred). On dimensional grounds:

$$S_\ell \sim (m_2/m_\ell)^4 \cdot \varepsilon_{ZPE}^{q_\ell} \quad \text{(structural ansatz, } q_\ell \geq 0)$$

where $q_\ell \geq 0$ is the number of additional ε_{ZPE} insertions required to couple the ℓ-th multipole to the photon-photon vertex (a deliverable). For the bare KK spectrum $m_\ell^2 = \ell(\ell+1)\, m_e^2/\pi^2$ of [6] §6.21, the bare propagator ratios are:

- $(m_2/m_4)^4 = (6/20)^2 = 0.090$
- $(m_2/m_6)^4 = (6/42)^2 = 0.020$
- $(m_2/m_8)^4 = (6/72)^2 = 0.007$

giving an upper-bound estimate for the bare-KK-suppressed higher-multipole tail (taking $q_\ell = 0$ and $k_\ell = 1$ as the most pessimistic case):

- $\ell=4$: $1.075 \times 0.090 \approx 0.097$
- $\ell=6$: $0.506 \times 0.020 \approx 0.010$
- $\ell=8$: $0.312 \times 0.007 \approx 0.002$
- Sum: ≈ 0.11 (i.e., the higher-multipole tail contributes at most ~ 11% of the equator TPA rate, dominated by $\ell=4$).

This bound becomes tighter once determines the k_ℓ and q_ℓ values — in particular, if any $q_\ell \geq 1$, the corresponding contribution is suppressed by an additional factor of $\varepsilon_{ZPE}^2 \approx 0.016$ per extra insertion.

5.5.4 Medium-dependent corrections (structural placeholder)

In a real photonic device (silicon waveguide, $LiNbO_3$, SiN, etc.), the photon's polarization-state distribution is broadened beyond the (3+3)-internal ε_ZPE contribution by external sources:

- **Birefringent rotations** from waveguide stress and crystallographic anisotropy.
- **Residual scattering** from waveguide roughness.
- **Thermal phonon-jitter** coupling to the polarization mode (correlated with the thermo-optic channel of §5.2).
- **Finite laser linewidth** and **detector polarization-resolution**.

These contribute additively to the angular variance:

$$\langle(\delta\vartheta)^2\rangle\_total = k\_2\ \varepsilon\_ZPE^2 + \langle(\delta\vartheta)^2\rangle\_birefringence + \langle(\delta\vartheta)^2\rangle\_scattering + \langle(\delta\vartheta)^2\rangle\_phonon + \langle(\delta\vartheta)^2\rangle\_apparatus \quad (5.5.9)$$

The (3+3)-internal contribution $k\_2\ \varepsilon\_ZPE^2$ is the **vacuum floor**: even in a perfectly clean device with ideal polarization control, this is the smallest η_node achievable in the framework. The medium-dependent contributions are anchored to standard photonics and are deferred to for explicit computation. In we identify only their structural placement in the master formula.

5.5.5 Master formula

Combining (5.5.7), (5.5.8), and the medium contributions of §5.5.4:

$$\boldsymbol{\eta\_node = 8\ k\_2\ \varepsilon\_ZPE^2 + \Sigma\_\{\ell\ even,\ \ell\geq 4\}\ k\_\ell \cdot |Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2 \cdot S\_\ell + K\_medium \cdot (medium\ factors)} \quad (5.5.10)$$

with explicit Level-1 derived content:

- The coefficient *8* in the leading term, derived from the spherical-harmonic gradient at the cone (Eq. 5.5.4–5.5.5).
- The geometric ratios $|Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2$ for $\ell = 4, 6, 8, \ldots$, computed exactly above (Table in §5.5.3).
- The bare-KK propagator ratios $(m\_2/m\_\ell)^4$ for the suppression ansatz.
- The functional structure: η_node is a sum of terms, each $O(\varepsilon\_ZPE^2)$ or higher (in the breathing channel) plus geometric-times-propagator (in the higher-multipole channel) plus medium contributions.

and explicit Level-2 anchored content (deliverable):

- **The breathing-mode prefactor *k_2*** (most consequential single deliverable; controls the vacuum floor of η_node).
- **The higher-multipole prefactors *k_ℓ*** for $\ell = 4, 6, 8, \ldots$ and their q_ℓ powers.
- **The medium correction *K_medium*** and its constituents (birefringence, scattering, phonon, apparatus contributions).

Numerical bracket at level (taking $k\_2 \sim 1$ and $k\_\ell = 1$ for ℓ=4,6,8 with bare-KK suppression):

$$\eta\_node \sim 0.13 + 0.11 = 0.24 \text{ (≈ 4× TPA suppression, conservative)}$$

At the optimistic end (taking $k\_2 \sim 0.1$ and k_ℓ q_ℓ*-suppressed*):

$$\eta\_node \sim 0.013 + (subleading) \approx 0.013 \text{ (≈ 75× TPA suppression)}$$

These two estimates bracket the predicted range. The §6.7 thermal MZI experiment (2027–2030) is the decisive test of where in this range *η_node* falls. The framework's *direction* and *structural form* are now derived; the experimental measurement of *η_node* is, a measurement of the four prefactors *k_2*, *k_ℓ*, *S_ℓ*, *K_medium* — equivalently, a constraint on the foam-level dynamics not visible at the present level of the framework.

5.5.6 What §5.5 establishes and what remains open

Established at Level 1 in this paper (the structural form of *η_node* is now derived from (3+3), no longer Level-2 anchored):

- The functional dependence *η_node = f(ε_ZPE) · (geometric prefactors)*.
- The leading $O(\varepsilon\_ZPE^2)$ scaling and the coefficient *8* from spherical-harmonic geometry.
- The geometric ratios $|Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2$ for higher multipoles, exactly.
- The structural separation of (3+3)-internal vs medium contributions.
- The two-bracket numerical estimate $0.013 \le \eta\_node \le 0.24$ spanning the foam-level prefactor range.

Open (the prefactor closure):

- The breathing-mode coefficient *k_2* (most consequential single deliverable).
- The higher-multipole prefactors *k_ℓ* and their *q_ℓ* powers.
- The medium-specific corrections *K_medium* for silicon, SiN, $LiNbO_3$.

Conversion of §10 item A1 sub-items. The §10 A1 of v0.3 listed three sub-items: (i) the ε_ZPE^2-derived contribution, (ii) the higher-multipole tail, (iii) the medium-dependent corrections. After this analysis:

- Sub-item (i) is now Level-1 structural form (Eq. 5.5.7), prefactor *k_2* deferred to specialist work.
- Sub-item (ii) is now Level-1 structural form (Eq. 5.5.8), geometric ratios computed exactly; prefactors *k_ℓ* and powers *q_ℓ* deferred to specialist work.
- Sub-item (iii) remains Level-2 anchored, but is now structurally catalogued (Eq. 5.5.9).

The deliverable is therefore: ***the structural form of η_node is now a derived Level-1 result; only the four explicit prefactors (k_2, k_ℓ, S_ℓ, K_medium) remain Level-2 anchored.*** This represents substantial progress toward the §10 A1 closure but does **not** close it: a definite numerical *η_node* still requires specialist computation of the prefactors. The §9.2 conditional argument (specifically condition C1, "the suppression magnitude sufficient to bring TPA-limited

Kerr gate fidelity above 99.5%") therefore remains conditional pending — but its dependence on the framework's open prefactors is now tractable rather than schematic.

5.6 Foam-level evaluation of the breathing-mode prefactor k_2

Section 5.5 isolated the breathing-mode prefactor $k\_2 = \langle(\delta\vartheta)^2\rangle/\varepsilon\_ZPE^2$ as the most consequential single deliverable. — the foam-level computation of k_2 — is acknowledged in the staged plan of [1] as a problem that "may not succeed in one session" because a complete derivation requires a foam-level Lagrangian for the photon's polarization-state field on t_3 S^2 that is not yet developed in the (3+3) framework. This subsection delivers the *status report*: it bounds k_2 from above and below, computes the explicit geometric quantity $\Sigma\_m\ |\partial\_\vartheta\ Y\_\{2,m\}(\vartheta\_node)|^2 = 15/(4\pi)$ that any reasonable foam Lagrangian must reproduce, draws on the Casimir analogy of [6] §14bis.69 to estimate the prefactor, and combines these into a tightened η_node bracket. A definitive k_2 awaits specialist work that develops the missing foam Lagrangian.

5.6.1 The problem

What §5.5 established. The leading angular-fluctuation contribution to η_node has the form (Eq. 5.5.7):

$$\eta\_node|\_breathing = 8\ k\_2\ \varepsilon\_ZPE^2 \text{ where } k\_2 = \langle(\delta\vartheta)^2\rangle / \varepsilon\_ZPE^2$$

with $\varepsilon\_ZPE = 0.1270$ (closed by [6] §15) and the coefficient 8 derived from spherical-harmonic geometry. The single open quantity is the dimensionless prefactor k_2, defined as the dimensionless ratio of the angular variance of the photon's polarization-state at the cone to ε_ZPE^2.

What would compute. A definitive k_2 requires (i) the foam-level Lagrangian for the photon's polarization-state field, with its couplings to the S^2 zero-point modes specified explicitly; (ii) identification of which S^2 zero-point modes are physical degrees of freedom (in particular, whether the $\ell=2$, $m\neq0$ components of the Higgs multiplet exist as physical fluctuating modes or are absorbed as longitudinal components of broken-gauge bosons); (iii) integration over the contributions of all relevant modes to $\langle(\delta\vartheta)^2\rangle$ at $\vartheta = \vartheta\_node$. None of these is currently in [6] at the level of explicit Lagrangians; the framework provides the kinematic structure (KK reduction, harmonics, masses) but not the dynamical coupling of the photon polarization mode field to the multipole spectrum.

What can deliver in this paper. In the absence of (i)–(iii), we can: (a) bound k_2 from above by dimensional saturation and from below by physical considerations; (b) compute the geometric coefficient $\Sigma\_m\ |\partial\_\vartheta\ Y\_\{2,m\}(\vartheta\_node)|^2$ that any Lagrangian-based derivation must reproduce; (c) appeal to the Casimir analogy of [6] §14bis.69 for an order-of-magnitude estimate; (d) report the resulting range for $\eta\_node|\_breathing$. We do (a)–(d) below.

5.6.2 The dimensional saturation upper bound: k_2 ≤ ~ 1

Upper bound. The angular variance $\langle(\delta\vartheta)^2\rangle$ cannot exceed the angular variance of a uniformly-distributed polarization state on S^2, which is $\langle(\delta\vartheta)^2\rangle\_uniform = \pi^2/3$ (the variance of a uniform distribution on $[0, \pi]$). For $\varepsilon\_ZPE = 0.1270$, this would give $k\_2 \le \pi^2/(3\ \varepsilon\_ZPE^2) \approx 204$ — far too weak to be useful.

A tighter upper bound comes from the framework's coupling structure (§5.4.2): the photon couples to Φ^2 with strength $g_0 = (4/9)\ \varepsilon\_ZPE^2 \approx 7.2 \times 10^{-3}$. The angular variance generated by this coupling cannot exceed the saturating value where the photon's polarization-state location is fully randomised within the angular envelope set by ε_ZPE itself. This gives:

$k\_2 \le 1$ (dimensional saturation; saturating coupling to Higgs-multiplet zero-point fluctuations)

At saturation, $\langle(\delta\vartheta)^2\rangle = \varepsilon\_ZPE^2$, $\delta\vartheta\_RMS = \varepsilon\_ZPE = 0.1270\ rad \approx 7.3°$, and $\eta\_node|\_breathing = 8\ \varepsilon\_ZPE^2 \approx 0.13$ (the conservative end of the §5.5 bracket). This is the upper bound for any reasonable mechanism.

5.6.3 The radial-decoupling consideration: a candidate lower bound

Lower bound consideration. The Higgs breathing mode is identified in [6] §16.3 as the $\ell = 2$, $m = 0$ spherical harmonic Y_2^0 of S^2, with the time-dependent breathing taken as a uniform radial oscillation $R_3(t) = R_3^{(0)}[1 + \varepsilon\_ZPE \cos(\omega\_H\ t)]$ (Eq. 14b.67 of [6]). A *purely* radial breathing of R_3 leaves the angular structure of S^2 unchanged: the location $\vartheta\_node = \arccos(1/\sqrt{3})$ of the Y_2^0 zero is invariant under $R_3 \to \lambda R_3$ because it is a property of the angular coordinate, not the radial scale. **A photon at coordinate $\vartheta = \vartheta\_node$ therefore experiences no first-order angular displacement from the radial breathing**.

Two consequences follow. (a) If the radial breathing is the *only* mechanism contributing to $\langle(\delta\vartheta)^2\rangle$, then k_2 arises only from second-order coupling (the breathing's effect on the metric coupling to the photon's polarization mode field at Φ^2 level), giving $k\_2 \propto \varepsilon\_ZPE^2$ and hence $\eta\_node|\_breathing \propto \varepsilon\_ZPE^4 \approx 2.6 \times 10^{-4}$. In this case $\eta\_node|\_breathing \approx 8 \cdot 0.016 \cdot 0.016 \approx 0.002$, i.e., $\sim 500\times$ TPA suppression — extreme suppression dominated by higher-multipole and medium contributions of §5.5.3–§5.5.4. (b) If, in addition to the radial breathing, the $\ell=2$, $m\neq0$ components of the Higgs multiplet are physical fluctuating modes (rather than absorbed gauge-boson Goldstones), then their angular structure does directly displace points at ϑ_node, and k_2 is at first order in ε_ZPE^2 with a prefactor of order unity (see §5.6.4 below).

The framework's identification of the Higgs as specifically the $m = 0$ component leaves the $m \neq 0$ status ambiguous in current treatments. In the Standard-Model identification, the four real components of the complex Higgs doublet are: one physical $m=0$ scalar (the observed Higgs at 125 GeV) plus three Goldstone bosons absorbed as longitudinal modes of $W\pm$ and Z. If the (3+3) identification mirrors this, the $\ell=2$, $m\neq0$ components on S^2 are the absorbed Goldstones — physical at the level of the $W\pm/Z$ longitudinal modes but not as scalar fluctuations of the S^2 metric. In that case mechanism (a) above dominates and $k\_2 \approx \varepsilon\_ZPE^2 \approx 0.016$. If instead the $\ell=2$, $m\neq0$ components contribute angular fluctuations independently, mechanism (b) dominates.

5.6.4 Geometric calculation: the gradient sum at the cone

Independent of which scenario in §5.6.3 is realised, the framework gives an exact geometric quantity that any foam-Lagrangian computation of k_2 must reproduce: the sum of $|\partial_\vartheta Y_{2,m}|^2$ over the full $\ell=2$ multiplet at ϑ_{node}. The mode-by-mode breakdown:

m	$Y_{2,m}(\vartheta,\varphi)$	$\lvert\partial_\vartheta Y_{2,m}\rvert^2\rvert_{node}$	Decimal
0	$(1/4)\sqrt{(5/\pi)}(3\cos^2\vartheta - 1)$	$5/(2\pi)$	0.7958
± 1	$\mp(1/2)\sqrt{(15/(2\pi))} \sin\vartheta \cos\vartheta\, e^{(\pm i\varphi)}$	$5/(24\pi)$ (each)	0.0663 each
± 2	$(1/4)\sqrt{(15/(2\pi))} \sin^2\vartheta\, e^{(\pm 2i\varphi)}$	$5/(12\pi)$ (each)	0.1326 each

Sum over the full $\ell=2$ multiplet at the cone:

$$\Sigma_m |\partial_\vartheta Y_{2,m}(\vartheta_{node})|^2 = 5/(2\pi) + 2\cdot 5/(24\pi) + 2\cdot 5/(12\pi) = 30/(12\pi) + 5/(12\pi) + 10/(12\pi) = 45/(12\pi) = 15/(4\pi) \approx 1.194 \quad (5.6.1)$$

Of this total, the $m = 0$ (Higgs) contribution is $5/(2\pi) \approx 0.796$, exactly $2/3$ of the multiplet sum. Equivalently:

- **Higgs ($m = 0$) contribution**: $2/3 \times 15/(4\pi) = 5/(2\pi) \approx 0.796$
- **Multiplet ($m \neq 0$) contribution**: $1/3 \times 15/(4\pi) = 5/(4\pi) \approx 0.398$

These ratios are pure spherical-harmonic algebra and do not depend on dynamics. They are the geometric quantities that a Stage-5 foam-Lagrangian computation of k_2 must combine with the (currently unknown) coupling structure to produce a definitive prefactor. As an immediate constraint, any reasonable Lagrangian-based estimate that maps "gradient-sum at the cone" directly to k_2 gives:

- Saturating ansatz, full multiplet: $k_2 \approx 15/(4\pi) \cdot (\text{coupling normalisation}) \approx 1.2$ (assuming O(1) coupling).
- Higgs-only ansatz: $k_2 \approx 5/(2\pi) \cdot (\text{coupling normalisation}) \approx 0.8$.
- Casimir-analogy ansatz (next subsection): $k_2 \approx 1/2$.

All three estimates fall within the dimensional upper bound $k_2 \leq 1$ (within an O(1) factor for the saturating-multiplet case) and all exceed the radial-decoupled lower bound $k_2 \approx \varepsilon_{ZPE}^2 \approx 0.016$. therefore brackets the prefactor as $k_2 \in [0.016, \sim 1.2]$.

5.6.5 The Casimir analogy from [6] §14bis.69

A useful framework-internal benchmark for the prefactor comes from the Casimir effect calculation of [6] §14bis.69. The Higgs-breathing zero-point amplitude ε_{ZPE} generates a fractional correction to the Casimir pressure between conducting plates:

$$P_{Cas}^{(3+3)} = P_{QED} \cdot (1 + \varepsilon_{ZPE}^2/2) \quad \text{[6] Eq. 14b.69}$$

with the structural origin: each KK mode's zero-point energy carries a second-order $\varepsilon\_ZPE^2/2$ correction from the breathing's modulation of the mode frequency. The dimensionless prefactor *1/2* is universal across all KK modes and across all observables that are sensitive to the second moment of the breathing oscillation.

Applied to the present problem, the same $\varepsilon\_ZPE^2/2$ factor is the natural prefactor for any framework-internal observable controlled by the breathing-mode zero-point variance. With $k\_2 = 1/2$:

$$\eta\_node|\_breathing,Casimir = 8 \cdot (1/2) \cdot \varepsilon\_ZPE^2 = 4 \cdot \varepsilon\_ZPE^2 = 4 \cdot 0.01613 \approx 0.0645 \quad (5.6.2)$$

i.e., $\eta\_node \approx 6.5\%$, equivalent to $\sim 16\times$ TPA suppression at the cone. This is the "central-value" estimate and is consistent with the §5.6.4 geometric estimates within the expected O(1) coupling normalisation.

A note of caution. The Casimir analogy applies to the Casimir effect because both the constrained-mode zero-point energies *and* the breathing modulation enter at the same $\varepsilon\_ZPE^2/2$ level. Whether *η_node|_breathing* shares this exact prefactor depends on whether the photon-polarization-state's coupling to the breathing satisfies the same algebraic structure — which awaits verification. We treat $k\_2 \approx 1/2$ as a benchmark, not a derivation.

5.6.6 Combined bracket and experimental discrimination

Combining the bounds and estimates of §5.6.2–§5.6.5:

Estimate / mechanism	***k_2***	***η_node|_breathing***	**TPA suppression**
Dimensional saturation (upper)	*1*	*0.13*	*8×*
Full ℓ=2 multiplet, gradient-weighted	$15/(4\pi) \approx 1.19$	*0.15*	*7×*
Higgs (*m=0*) only, gradient-weighted	$5/(2\pi) \approx 0.80$	*0.10*	*10×*
Casimir analogy [6] §14bis.69	*1/2*	*0.065*	*16×*
Radial-decoupled (lower; ε_ZPE^4 scaling)	$\approx \varepsilon\_ZPE^2 \approx 0.016$	*≈ 0.0021*	*≈ 480×*

combined bracket (encompassing all reasonable mechanisms):

$$\eta\_node|\_breathing \in [0.002, 0.15] \text{ with central estimate } 0.05 – 0.10$$

This tightens the §5.5 bracket $\eta\_node|\_breathing \in [0.013, 0.24]$ by (a) lowering the optimistic end via the radial-decoupling consideration and (b) raising the conservative end slightly via the multiplet-saturating estimate. The central estimate ~ *5–10%* maps to ~ *10–20*× TPA suppression — a useful but not extreme reduction.

Total *η_node* in real devices. Combining with the higher-multipole tail of §5.5.3 (≈ *0.11* upper bound, dominated by *ℓ=4*) and the medium contribution of §5.5.4 (anchored, typically *0.05–0.20* in well-fabricated photonic chips):

$$\eta\_node|\_total \approx \eta\_node|\_breathing + \eta\_node|\_higher + \eta\_node|\_medium \sim 0.10 - 0.30 \text{ in real devices}$$

giving *3–10*× TPA suppression at the cone for typical silicon-photonic implementations. This is the prediction for the *quantitative* outcome of the §6.7 thermal MZI experiment (2027–2030).

Empirical discrimination. The bracket *η_node ∈ [0.002, 0.15]* spans nearly two orders of magnitude. The §6.7 thermal MZI experiment is the decisive test: lock-in measurement of the thermo-optic phase drift *dφ/dT* at the nodal-cone-encoded chip vs. equator-encoded reference chip would directly measure *η_node* (since the thermo-optic suppression at the cone tracks *η_node* as a Stage-4a structural prediction). A measured *η_node* ~ *0.05–0.10* would confirm the Casimir-analogy mechanism. *η_node* ~ *0.001* would indicate radial-decoupling dominance. *η_node* ≥ *0.20* would falsify the entire breathing-mode mechanism and require revisiting the (3+3) identification of the photon polarization mode.

5.6.7 What §5.6 establishes and what remains open

Established here:

- The geometric quantity *Σ_m |∂_ϑ Y_{2,m}(ϑ_node)|² = 15/(4π) ≈ 1.194* with *m=0* contributing exactly *2/3*: a closed-form result that any future foam-Lagrangian computation of *k_2* must reproduce.
- The dimensional saturation upper bound *k_2 ≤ 1*.
- The radial-decoupling consideration that gives a candidate lower bound *k_2 ~ ε_ZPE² ≈ 0.016* (if the *ℓ=2*, *m≠0* components are absorbed Goldstones).
- The Casimir-analogy benchmark *k_2 ≈ 1/2* from [6] §14bis.69, mapping to the central estimate *η_node|_breathing ≈ 0.065*.
- The combined bracket *η_node|_breathing ∈ [0.002, 0.15]* (tightening the bracket).
- The total prediction *η_node|_total ~ 0.10 – 0.30* in real devices (combining all contributions).

Remaining (genuine specialist work for definitive *k_2*):

- The foam-level Lagrangian for the photon's polarization-state field on t_3 S^2, with its couplings to the S^2 multipole spectrum specified explicitly. This is the central piece of framework infrastructure not yet developed in [6].

- The identification of whether the $\ell=2$, $m \neq 0$ components are physical scalar fluctuations or absorbed gauge-boson Goldstones — the discriminator between the two scenarios of §5.6.3.
- The medium-specific corrections *K_medium* for silicon, SiN, $LiNbO_3$ — a Stage-4b' specialist photonics calculation.
- The higher-multipole prefactors *k_ℓ* (for *ℓ = 4, 6, 8, …*) in (5.5.8) — specialist work paralleling §5.6 but for the higher multipoles.

Status of §10 item A1 after this analysis. Sub-item (i) of §10 A1 (the *ε_ZPE²*-derived contribution) is now bracketed by at *k_2 ∈ [0.016, 1.2]* with a Casimir-analogy central estimate *k_2 ≈ 0.5*. Sub-item (ii) (the higher-multipole tail) retains its structural form with the bare-KK-suppression bound *η_node|_higher ≤ 0.11*. Sub-item (iii) (medium corrections) remains Level-2 anchored. The *combined* prediction *η_node ∼ 0.10–0.30* in real devices is sufficient to satisfy §9.2 condition C1 ("TPA suppression magnitude bringing Kerr gate fidelity above 99.5%") in the Casimir-analogy regime, but not in the saturating regime: the §6.7 thermal MZI experiment is required to discriminate.

The result is therefore: ***the breathing-mode prefactor k_2 is bracketed within the dimensional bounds [ε_ZPE², 1] with a framework-internal Casimir-analogy central estimate k_2 ≈ 1/2; a definitive value awaits the foam-Lagrangian construction.*** This does not close item A1 — but it converts the residual specialist work from an open-ended structural problem to a specific Lagrangian-construction calculation. **§5.7 takes that next step:** it carries out the Lagrangian-level analysis of the breathing-mode contribution at the cone, finding that the same *Y_{2,0}(ϑ_node) = 0* identity which defines the cone latitude *also* eliminates the leading metric perturbation there — so the first-order breathing contribution to *η_node* vanishes geometrically, and the Casimir analogy of §5.6.5 is refined as misapplied (it applies to *dimensionful* observables like Casimir pressure, not to the *dimensionless* coordinate *ϑ* at the cone). The total *η_node* prediction *0.10–0.30* in real devices is unchanged — it was already dominated by higher-multipole and medium contributions in the bracket — but the structural picture is now sharper: the framework's vacuum is essentially *transparent* to *η_node* at the cone.

5.7 Lagrangian-level analysis at the cone: vanishing first-order contribution

Section 5.6 reported the status: the breathing-mode prefactor *k_2* was bracketed within *[ε_ZPE² ≈ 0.016, ∼ 1.2]* using dimensional saturation, the radial-decoupling argument, the gradient-sum geometric calculation, and the Casimir analogy of [6] §14bis.69 (which gave a central estimate *k_2 ≈ 1/2*, hence *η_node|_breathing ≈ 6.5%*). The §5.6.7 conclusion identified the residual specialist priority as the foam-level Lagrangian construction for the photon's polarization-mode field. This subsection reports a substantive finding: the leading-order Lagrangian-level analysis of the breathing-mode contribution at the cone reveals that the first-order term vanishes geometrically — by the *same Y_{2,0}(ϑ_node) = 0* identity which defines the cone latitude. The Casimir-analogy estimate is refined as misapplied to the dimensionless coordinate *ϑ* at the cone; the total *η_node* prediction is unchanged but the structural picture is sharper.

5.7.1 The question: does R_3 breathing cause first-order δϑ at ϑ_node?

The structural ansatz (Eq. 5.5.6) was *⟨(δϑ)²⟩ = k_2 · ε_ZPE²*: a leading-order *ε_ZPE²* scaling for the photon's polarization-state angular variance at the cone, with *k_2* an O(1) prefactor to be computed by foam-level work. asks the more basic Lagrangian-level question: *does R_3 breathing in fact produce a first-order δϑ at ϑ_node*, or are there geometric obstructions that suppress the leading-order contribution? The relevant breathing dynamics from [6] §14bis.66–67 and §16.3 has two components that must be considered separately:

- **Uniform component** — the radial rescaling *$R_3(t) = R_3^{(0)}$[1 + ε_ZPE cos(ω_H t)]* of [6] Eq. 14b.67. This is the cosmological-Hubble framing of the breathing mode.
- **Angular component** — the *Y_{2,0}(ϑ)* spherical-harmonic structure of [6] §16.3. The breathing is identified specifically as the *ℓ = 2*, *m = 0* mode of *S²*, with metric perturbation *δg_ϑϑ ∝ Y_{2,0}(ϑ) cos(ω_H t)*.

In the absence of either consideration the ansatz might be saturated; but each component has a distinct geometric structure that must be examined.

5.7.2 Uniform R_3 breathing: dimensional invariance argument

A uniform rescaling *$R_3 \to \lambda R_3$* (where *λ = 1 + ε_ZPE cos(ω_H t)*) is a *coordinate transformation* on the embedded *$S^2(R_3)$*. It rescales every length on the sphere by the same factor *λ* but leaves all *dimensionless* coordinates invariant. The Poincaré-sphere coordinate *ϑ ∈ [0, π]* is dimensionless: it parametrises position on *S²* independently of the radius. Therefore:

ϑ_node = arccos(1/√3) = 54.74° is invariant under uniform *$R_3 \to \lambda R_3$*

Equivalently, all dimensionful lengths in the lab apparatus that determine the photon's polarization-state coordinate (e.g., the waveplate thickness *d*, the photon's wavelength *λ_photon*, the birefringence-induced rotation angle *Γ = (2π/λ_photon)·Δn·d*) scale uniformly under *$R_3 \to \lambda R_3$* because the framework's mass scales — including *$m\_e \propto 1/R_3$* from [6] §15 and the Bohr radius *$a_0 \propto 1/m\_e \propto R_3$* — all track *$R_3$*. The dimensionless ratios *(d/λ_photon)*, *Δn*, and ultimately *Γ* are invariant. The waveplate-induced coordinate *ϑ* on the Poincaré sphere is therefore unchanged by uniform breathing.

First conclusion: uniform *R_3* breathing does *not* produce *δϑ* at the cone. Any first-order *ε_ZPE* contribution must come from the *non-uniform* (angular) component of the breathing.

5.7.3 Angular Y_{2,0} breathing: vanishing at ϑ_node

The non-uniform angular component of the breathing is the *Y_{2,0}(ϑ)* metric perturbation:

$$\delta g\_\vartheta\vartheta(\vartheta, t) = R_3^2 \cdot 2\, \varepsilon\_ZPE\ Y\_\{2,0\}(\vartheta) \cos(\omega\_H t) + O(\varepsilon\_ZPE^2) \quad (5.7.1)$$

A photon at coordinate *ϑ_0* experiences a metric-perturbation-induced angular shift proportional to *Y_{2,0}(ϑ_0)·ε_ZPE* at first order. Specifically, for a polarization-state field localised at *ϑ_0*, the first-order *δϑ* response to the breathing is:

$$\delta\vartheta(t)|\_first\text{-}order \propto Y\_\{2,0\}(\vartheta\_0) \cdot \varepsilon\_ZPE \cos(\omega\_H t)$$

At $\vartheta\_0 = \vartheta\_node = arccos(1/\sqrt{3})$, however:

$$\boldsymbol{Y\_\{2,0\}(\vartheta\_node) = (1/4)\sqrt{(5/\pi)}\cdot(3\cdot(1/3) - 1) = 0} \text{ (Level 1 identity, exact)}$$

This is precisely the identity that *defines* the nodal cone — it is the Level-1 framework prediction we used in §5.5.1 to argue that the leading TPA channel is structurally suppressed at the cone. **The same identity** $Y\_\{2,0\}(\vartheta\_node) = 0$ **also eliminates the breathing-mode metric perturbation at the cone**:

$$\delta\vartheta(t)\rfloor\_first\text{-}order,\ \vartheta\_0 = \vartheta\_node = (\propto Y\_\{2,0\}(\vartheta\_node)) \cdot \varepsilon\_ZPE\ cos(\omega\_H\ t) = 0 \quad (5.7.2)$$

The cone is therefore *doubly protected*: by topology (the photon's TPA matrix element $Y\_\{2,0\}(\vartheta\_node) = 0$) *and* by metric (the breathing's first-order metric perturbation at ϑ_node also vanishes). Both protections derive from the *same* spherical-harmonic identity.

Second conclusion: the angular $Y\_\{2,0\}$ component of the breathing also produces *no first-order* $\delta\vartheta$ at the cone — the metric perturbation vanishes there geometrically.

5.7.4 result: first-order η_node|_breathing vanishes at the cone

Combining §5.7.2 and §5.7.3: the first-order breathing-mode contribution to $\langle(\delta\vartheta)^2\rangle$ at ϑ_node vanishes. The structural ansatz $\langle(\delta\vartheta)^2\rangle = k\_2 \cdot \varepsilon\_ZPE^2$ with $k\_2 = O(1)$ is refined: at the cone specifically (and not at generic latitudes on S^2), $k\_2\rfloor\_cone = 0$ at first order in the breathing amplitude.

Second-order corrections remain. They arise from the breathing's *gradient* at the cone interacting with the photon's *intrinsic* polarization-state spread $\delta\vartheta\_int$ induced by external sources (medium birefringence, scattering, phonon-jitter, finite apparatus precision):

$$\delta\vartheta(t)\rfloor\_second\text{-}order \propto \partial\_\vartheta\ Y\_\{2,0\}(\vartheta\_node) \cdot \varepsilon\_ZPE \cdot \delta\vartheta\_int \cdot cos(\omega\_H\ t)$$

From §5.5.2, $|\partial\_\vartheta\ Y\_\{2,0\}(\vartheta\_node)|^2 = 5/(2\pi)$, so the time-averaged second-order angular variance is bounded:

$$\langle(\delta\vartheta)^2\rangle\rfloor\_second\text{-}order \lesssim \varepsilon\_ZPE^2 \cdot \langle(\delta\vartheta\_int)^2\rangle \cdot (geometric) \quad (5.7.3)$$

With $\langle(\delta\vartheta\_int)^2\rangle$ set by medium contributions of §5.5.4 (typically $\langle(\delta\vartheta\_int)^2\rangle \leq 0.20$ in well-fabricated photonic chips), the second-order breathing contribution is bounded by:

$$\eta\_node\rfloor\_breathing,second\text{-}order \lesssim \varepsilon\_ZPE^2 \cdot \eta\_node\rfloor\_medium \leq 0.016 \cdot 0.20 \approx 0.003\ (\approx 0.3\%) \quad (5.7.4)$$

numerical conclusion: the breathing-mode contribution to η_node at the cone is *at most* $\sim 0.3\%$ — entirely subdominant to the higher-multipole tail (≤ 0.11) and the medium contributions (0.05–0.20). The framework's vacuum is essentially *transparent* to η_node at the cone.

5.7.5 Reconciliation with the §5.6.5 Casimir analogy

The §5.6.5 estimate used the Casimir-effect prefactor $\varepsilon\_ZPE^2/2$ of [6] §14bis.69 as a benchmark for k_2. finds this analogy was misapplied. The Casimir formula $P\_Cas\text{^}\{(3+3)\} = P\_QED \cdot (1 + \varepsilon\_ZPE^2/2)$ applies to a *dimensionful* observable (the Casimir pressure between conducting plates) which is sensitive to the time-averaged second moment of $R_3(t)$: at second order, $\langle R_3^2\rangle \neq R_3^2$ and the Casimir pressure inherits an $\varepsilon\_ZPE^2/2$ fractional correction. This is a genuine framework-internal ε_ZPE² effect — *for dimensionful observables*.

The *η_node* matrix element $\langle Y\_\{2,0\}(\vartheta)\rangle$, by contrast, is a function of the *dimensionless* coordinate ϑ, not of R_3. Under uniform breathing, ϑ is invariant (§5.7.2); under angular *Y_{2,0}* breathing, the metric perturbation *vanishes at the cone* (§5.7.3). The Casimir mechanism therefore does not transfer to *η_node|_cone*: the ε_ZPE² scaling that holds for the Casimir pressure does *not* hold for the photon's angular variance at the cone.

This is a *refinement* of the §5.6 picture, not a contradiction: the bracket *η_node|_breathing ∈ [0.002, 0.15]* remains valid as an upper-bound dimensional analysis, but sharpens the estimate to the *bottom* of that bracket (*η_node|_breathing ≲ 0.003*). The "central estimate" of *6.5%* via the Casimir analogy was an over-counting of the breathing-mode contribution.

5.7.6 Implications for the total η_node prediction

The total *η_node* in real devices is the sum of the breathing contribution (now bounded by *0.3%*), the higher-multipole tail (bounded by *11%*), and the medium contributions (anchored, *5–20%*):

$$\eta\_node|\_total \approx \eta\_node|\_breathing + \eta\_node|\_higher + \eta\_node|\_medium \lesssim 0.003 + 0.11 + 0.20 \approx 0.32 \quad (5.7.5)$$

The prediction *η_node ∼ 0.10–0.30* in real devices is therefore *unchanged* — does not modify the bottom-line numerical prediction, because it was already dominated by higher-multipole and medium contributions, with the breathing term being a small contribution within the bracket.

What does change is the *structural* understanding:

- The framework's *intrinsic vacuum floor* for *η_node* at the cone is *much smaller* than the 6.5% Casimir-analogy estimate suggested. The breathing-mode contribution is essentially zero at first order, ≲ 0.3% at second order.
- The cone's TPA suppression is *intrinsically* protected by the Level-1 identity *Y_{2,0}(ϑ_node) = 0* — and this *same* identity also eliminates the leading vacuum dynamics that would otherwise smear the cone. The framework's vacuum is *transparent* to *η_node* at the cone.
- The residual *η_node* in real devices is *almost entirely* from external sources (medium birefringence, scattering, phonon-jitter, apparatus imprecision) — not from any framework-internal mechanism.
- Engineering optimisation strategies for nodal-cone qubits should therefore focus on *medium quality* and *apparatus precision*, not on suppressing framework-vacuum effects which are already negligible.

5.7.7 What §5.7 establishes

Established here:

- The first-order breathing contribution to η_node at the cone vanishes geometrically, by the same $Y\_\{2,0\}(\vartheta\_node) = 0$ identity that defines the cone latitude.
- Second-order corrections are bounded by $\eta\_node|\_breathing,second\text{-}order \lesssim \varepsilon\_ZPE^2 \cdot \eta\_node|\_medium \approx 0.3\%$.
- The §5.6.5 Casimir analogy is refined as misapplied to the dimensionless coordinate ϑ at the cone (it applies to dimensionful observables).
- The total η_node prediction 0.10–0.30 in real devices is unchanged but the structural picture is sharper: vacuum is transparent, residual is from medium and higher multipoles.

Remaining open (the residual specialist work):

- A complete foam-Lagrangian for the photon's polarization-mode field at *all* latitudes (not just at the cone). The vanishing-at-cone result of §5.7 is specific to ϑ_node; $\eta\_node|\_breathing$ at generic latitudes is non-zero and comparable to the upper end of the bracket.
- Higher-multipole prefactors k_ℓ for $\ell = 4, 6, 8, \ldots$ (still in their bare-KK form). The vanishing argument applies only to $\ell = 2$; for higher multipoles $Y\_\{\ell,0\}(\vartheta\_node) \neq 0$ and the analogue argument does not give vanishing.
- Medium-specific corrections K_medium for silicon, SiN, $LiNbO_3$ — the dominant remaining source of η_node in real devices.

The result is therefore: ***the leading-order breathing-mode contribution to η_node at the cone vanishes by the very identity that defines the cone; the framework's vacuum is essentially transparent to η_node at ϑ_node, and the practical η_node is dominated by external (medium) and structural (higher-multipole) sources rather than by framework-vacuum dynamics.*** This does not change the bottom-line numerical prediction ($\eta\_node \sim 0.10$–0.30 in real devices, set by the Casimir-analogy bracket which was already dominated by these external sources) — but it sharpens the structural understanding, refines the Casimir analogy as misapplied, and identifies the remaining open items as medium-specific (Appendix B) and higher-multipole (§5.8) calculations rather than further work on the breathing mode itself.

5.8 Higher-multipole prefactors and the realized framework spectrum

Section 5.6.7 identified the higher-multipole prefactors k_ℓ for $\ell = 4, 6, 8, \ldots$ (still in their bare-KK form) as deliverables. (§5.7) closed the breathing channel (ℓ = 2) at first order. This subsection bounds the higher-ℓ contributions using the framework's identified spectrum from [6] (its KK reduction §6, scalar mass spectrum §6.7, the ℓ=4 right-handed neutrino identification §11.9, and the trisection octopole structure §11.8). The finding: across plausible framework scenarios for the higher-ℓ $m = 0$ scalar content, the contribution to η_node at the cone is at most $\sim 10^{-5}$ and typically essentially zero — much smaller than the "bare-KK" upper bound ≤ 0.11. Combined with §5.7, this implies the practical η_node in real devices is dominated *entirely* by medium contributions, with framework-vacuum and higher-multipole contributions both negligible.

5.8.1 The bare-KK estimate revisited

The §5.5.3 master formula has the higher-multipole tail:

$$\eta\_node|\_higher\text{-}multipole = \Sigma\_\{\ell \text{ even}, \ell\geq4\}\ k\_\ell \cdot |Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2 \cdot S\_\ell \quad (5.5.8)$$

with the geometric ratios computed exactly: $|Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2 = 0, 784/729, 1024/2025, 30976/99225$ for $\ell = 2, 4, 6, 8$. The bare-KK propagator suppression $S\_\ell \sim (m\_2/m\_\ell)^4$ with the bare KK mass spectrum $m\_\ell^2 = \ell(\ell+1)\ m\_e^2/\pi^2$ of [6] §6.7 gave the §5.5.3 upper-bound estimate $\eta\_node|\_higher \leq 0.097 + 0.010 + 0.002 \approx 0.11$ (taking $k\_\ell = 1$). This was acknowledged in §5.5.3 as a "most pessimistic" upper bound, with several reasons to expect actual contributions to be smaller.

asks: which higher-ℓ $m = 0$ scalar modes actually exist in the framework's identified spectrum, and what are their physically realized masses and couplings?

5.8.2 The framework's identified spectrum at $\ell \geq 4$

The (3+3) framework's mode identifications relevant to η_node are:

- **$\ell = 0$ radion**: the global S^2 breathing degree of freedom; light/integrated out at sub-Planckian energies. No direct coupling to F^2 at the cone (no angular structure).
- **$\ell = 2$ Higgs scalar**: the observed Standard-Model Higgs at $m\_H = 125.10$ GeV, with bare KK mass $m\_2{,}bare = \sqrt{6}\ m\_e/\pi \approx 0.40$ MeV raised by ~3 × 10^5 to m_H through Coleman–Weinberg corrections (see [6] §15). This is the only specifically identified scalar coupling to F^2 via $L\_\gamma\gamma\Phi^2$.
- **$\ell = 4$ fermionic right-handed neutrino**: identified in [6] §11.9 as "the $\ell = 4$ winding mode without the projection onto observable physics." Mass scale $M\_R = m\_H \cdot v/m\_e \approx 6 \times 10^7$ GeV from the (3+3) seesaw. **Fermionic, not scalar** — does not directly couple to F^2 via a Φ^2F^2 vertex.
- **Trisection $\ell = 3, 6, 9, \ldots$ octopole**: the Z_3 trisection symmetry of S^2 contributes a $\cos(3\varphi)$ angular structure with spherical-harmonic content at $\ell = 3, 6, 9, \ldots$ and $m = \pm3$ ([6] §11.8 and §3051 of the same paper). **These are $m = \pm3$ modes**, not the $m = 0$ tail that matters for η_node at the cone (which is axisymmetric).
- **Higher-ℓ $m = 0$ scalars** ($\ell = 4, 6, 8, \ldots$): **not specifically identified in [6]**. Their physical realisation, mass scales, and couplings to F^2 are not specified.

The question for is therefore: what bounds can be placed on the contribution of *unidentified* higher-ℓ $m = 0$ scalars to η_node at the cone?

5.8.3 Three scenarios for higher-ℓ m = 0 scalar content

Three scenarios cover the plausible framework configurations:

Scenario A — *higher-ℓ $m = 0$ scalars exist with framework-natural TeV-scale masses, with the same g_0 structure as the Higgs*. Bare KK mass ratios are $m\_\ell^2/m\_2^2 = \ell(\ell+1)/6$; assuming Coleman–Weinberg corrections preserve the relative spacing, the physical masses are $m\_\ell =$

$\sqrt{(\ell(\ell+1)/6)} \cdot m\_H \approx 1.83 \cdot m\_H,\ 2.65 \cdot m\_H,\ 3.46 \cdot m\_H$ for $\ell = 4, 6, 8$ (i.e., 230, 330, 430 GeV). Propagator suppression $(m\_H/m\_\ell)^4 = (6/(\ell(\ell+1)))^2 = 0.090, 0.020, 0.007$. With $k\_\ell \sim 1$:

$$\eta\_node|\_higher \leq 1.075 \cdot 0.090 + 0.506 \cdot 0.020 + 0.312 \cdot 0.007 \approx 0.097 + 0.010 + 0.002 \approx 0.11$$
(Scenario A, max)

This recovers the bare-KK estimate. However, scalars at 230–430 GeV with Higgs-like couplings to F^2 would be at the LHC discovery scale — the absence of any such observation is a strong (if not absolute) experimental constraint against this scenario.

Scenario B — *higher-ℓ $m = 0$ scalars exist at framework-natural seesaw scale, analogous to the $\ell = 4$ right-handed neutrino*. The framework's only identified $\ell \geq 4$ mode is at $M\_R = m\_H \cdot v/m\_e \approx 6 \times 10^7$ GeV (from [6] §11.9). If higher-ℓ scalars follow the same scale-setting pattern, propagator suppression is $(m\_H/M\_R)^4 \approx (125/6{\times}10^7)^4 \approx 1.9 \times 10^{-23}$. With $k\_\ell \sim 1$:

$$\eta\_node|\_higher{,}Scenario\text{-}B \leq 1.075 \cdot 1.9 \times 10^{-23} \approx 2 \times 10^{-23}$$ (Scenario B, essentially zero)

Scenario C — *higher-ℓ $m = 0$ scalars exist at TeV scale but are gauge singlets without direct F^2 coupling*. The right-handed neutrino at $\ell = 4$ is the framework's example of a heavy $m = 0$ mode that does not couple to SM gauge bosons. If higher-ℓ $m = 0$ scalars share this gauge-singlet character, their contribution to TPA at the cone is loop-suppressed: a typical singlet–photon–photon loop gives $(\alpha/\pi)^2 \cdot (m\_H/m\_\ell)^2 \sim 5.4 \times 10^{-6} \cdot 0.30 \approx 1.6 \times 10^{-6}$ per insertion. With $k\_\ell \sim 1$:

$$\eta\_node|\_higher{,}Scenario\text{-}C \leq 1.075 \cdot 1.5 \times 10^{-6} \approx 1.6 \times 10^{-6}$$ (Scenario C, essentially zero)

5.8.4 The framework's identified spectrum favours Scenarios B/C

The framework's only specifically identified $\ell \geq 4$ mode is the right-handed neutrino at $\ell = 4$, fermionic and at the seesaw scale M_R. There is no specifically identified $\ell = 4$ scalar mode at TeV scale (which would be observed at the LHC if it existed with Higgs-like couplings, ruling out Scenario A). The natural framework expectation is that higher-ℓ $m = 0$ modes share the $\ell = 4$ fermion's structural character — heavy and/or gauge-singlet — placing them in Scenarios B or C.

Both Scenario B and Scenario C give $\eta\_node|\_higher \lesssim 10^{-5}$, six orders of magnitude smaller than the bare-KK upper bound of 0.11. This is the refinement: **the higher-multipole tail is essentially zero in the framework's plausible identified spectrum**.

A note on scenario-uncertainty. Scenario A is the conservative upper bound. If empirical evidence for new TeV-scale scalars with Higgs-like couplings to F^2 were to emerge, Scenario A would be activated, and $\eta\_node|\_higher$ could be as large as 0.11. Until then, the working bound is $\eta\_node|\_higher \lesssim 10^{-5}$.

5.8.5 The trisection octopole and the m = 0 tail

A subtle but important point: the framework's Z_3 trisection symmetry contributes a $\cos(3\varphi)$ angular structure on S^2 at $\ell = 3, 6, 9, \ldots$ (cf. [6] §11.8). One might worry that this trisection structure populates the higher-multipole tail in (5.5.8). It does not — for a structural reason.

The trisection's $\cos(3\varphi)$ structure is a $m = \pm 3$ spherical-harmonic content (the 3φ is azimuthal), not $m = 0$. The master formula's higher-multipole tail (5.5.8) is over the $m = 0$ (axisymmetric) modes Y_ℓ^0 — specifically because the photon's polarization-state at the nodal cone is at a definite latitude $\vartheta = \vartheta_{node}$ with no φ-dependence. The TPA matrix element is therefore axisymmetric ($m = 0$) and the $Y_{\ell,\pm 3}$ trisection octopole content does not enter:

$$\langle Y_{\ell,\pm 3} | \text{axisymmetric photon at } \vartheta_{node} \rangle = 0$$

The trisection octopole contributes to *azimuthal anisotropy* of the photon's polarization-state distribution (relevant for, e.g., the qutrit basis at the trisection vertices in §6.4 of this paper), but not to the η_{node} TPA-suppression factor at the cone. **The trisection structure is therefore a separate phenomenon** from the higher-multipole tail of (5.5.8) — its presence in the framework does not populate the $m = 0$ tail.

5.8.6 bound and total η_node refinement

Combining the Scenarios B/C analysis with the trisection separation:

$$\boldsymbol{\eta_{node}|_{higher} \lesssim 10^{-5}} \text{ (working bound, plausible framework spectrum)}$$

This is six orders of magnitude tighter than the bare-KK upper bound ≤ 0.11. The refinement comes from recognising that the framework's identified spectrum places the only known $\ell \geq 4$ mode at the seesaw scale or as a gauge singlet — both giving negligible contributions to η_{node}.

Total η_node prediction:

$$\eta_{node}|_{total} \approx \eta_{node}|_{breathing} + \eta_{node}|_{higher} + \eta_{node}|_{medium} \approx 0.003 + 10^{-5} + (\text{medium}) \approx \text{medium}$$

The framework contribution is bounded by $\lesssim 0.3\% + \lesssim 10^{-5} \approx 0.3\%$. The total η_{node} is therefore *almost entirely* set by medium contributions:

$$\boldsymbol{\eta_{node}|_{total} \approx K_{medium} \approx 0.05 - 0.20} \text{ in real devices (refined prediction)}$$

compared to the prior prediction of $0.10 - 0.30$. The lower end of the bracket has shifted slightly because the higher-multipole upper bound has tightened from 0.11 to 10^{-5}; the upper end is unchanged because it is set by the medium contribution which this analysis does not address.

5.8.7 Implications for the engineering programme

reinforces the structural conclusion of §5.7: **the practical *η_node* is set by medium quality alone**. The framework's vacuum and the framework's higher-multipole tail are both bounded at < *1%* of the equator TPA rate. Engineering optimisation of nodal-cone qubits should focus *exclusively* on:

- **Medium quality**: minimising waveguide birefringence, scattering, and phonon-jitter contributions to $\langle(\delta\vartheta)^2\rangle_{medium}$. Specialist photonics work for silicon, SiN, $LiNbO_3$.
- **Apparatus precision**: high-quality polarization-state preparation, suppressing finite-laser-linewidth and detector-resolution contributions.

- **Thermal control**: reducing phonon-jitter-induced spread (correlated with the §6.7 thermal MZI test).

No engineering effort is needed (or possible) to mitigate framework-vacuum or framework-higher-multipole effects — they are already negligible. This sharpens the §6.7 thermal MZI experiment's interpretation: it now measures *K_medium* almost directly, rather than disentangling *K_medium* from *k_2* and *k_ℓ*.

5.8.8 What §5.8 establishes

Established here:

- The framework's identified $\ell \geq 4$ mode is the fermionic right-handed neutrino at seesaw scale ($\ell = 4$); higher-ℓ $m = 0$ scalars are not specifically identified.
- The trisection Z_3 structure contributes only $m = \pm 3$ modes, not the $m = 0$ axisymmetric tail of (5.5.8).
- Three scenarios for the higher-ℓ $m = 0$ scalar content: A (TeV-scale Higgs-like, ruled out by LHC non-observation), B (seesaw-scale, contribution ~ 10^{-23}), C (TeV-scale gauge-singlet, contribution ~ 10^{-6}).
- Plausible scenarios (B, C): $\eta\_node|\_higher \lesssim 10^{-5}$, six orders of magnitude tighter than the bare-KK bound.
- Combined + refinement: $\eta\_node|\_total \approx K\_medium$ (medium-dominated), with framework contribution $< 1\%$.

Remaining (residual specialist work):

- Medium-specific *K_medium* corrections for silicon, SiN, $LiNbO_3$ — the dominant remaining source of *η_node* in real devices, anchored to standard photonics rather than derived.
- Empirical confirmation/falsification of Scenarios A vs B/C through the §6.7 thermal MZI experiment combined with LHC scalar searches.
- A potential addressing the *azimuthal* trisection structure as a separate phenomenon — relevant for the qutrit basis at trisection vertices, distinct from the cone TPA suppression.

The result is therefore: ***the higher-multipole tail in the master formula (5.5.8) is bounded at $\eta\_node|\_higher \lesssim 10^{-5}$ within the framework's plausible identified spectrum; combined with §5.7, the practical $\eta\_node \approx K\_medium$ is dominated entirely by medium contributions, with the framework's vacuum and higher-multipole tails both negligible.*** This closes the higher-multipole question to within the resolution of empirical particle physics (no new TeV-scale Higgs-like scalars), reducing the residual specialist work to medium-specific photonics calculations.

5.9 Z_3 symmetry analysis at the trisection vertices

Section 5.7 closed the breathing-mode contribution to *η_node* at the *cone* (axisymmetric, $m = 0$). Section 5.8 closed the higher-multipole tail there. Both addressed the *qubit* basis where the photon's polarization-state has no φ-dependence. The framework also supports a *qutrit* basis at the trisection vertices (§6.4 of this paper, Proposal 4), where the three computational states $|0\rangle$, $|1\rangle$, $|2\rangle$ sit at $\varphi = 0°, 120°, 240°$ on the equator with Z_3 permutation symmetry between them ([6] §10.3, §11.8, §11.11, line 2077). This subsection — — applies the same Lagrangian-level analysis to the qutrit basis, using the Z_3 symmetry as the analog of the $Y_{2,0}(\vartheta_{node}) = 0$ identity at the cone. The finding: framework-internal qutrit decoherence is bounded by $\eta_{qutrit}|_{framework} \approx \varepsilon_{ZPE}^2 \times 0.5 \approx 0.81\%$, about 3× the qubit framework floor of *0.3%* but still much smaller than the medium contribution. Combined with the qutrit's $\log_2 3 \approx 1.585$ bits-per-photon information advantage, qutrit photonics gains *58.5%* in encoding efficiency at essentially the same medium-dominated coherence cost as qubit photonics.

5.9.1 The qutrit basis: setup and the analog suppression question

The qutrit basis encodes three computational states $|0\rangle$, $|1\rangle$, $|2\rangle$ at the three trisection vertices on the t_3 S^2 equator at $\varphi = 0°, 120°, 240°$. The Z_3 rotation $\varphi \to \varphi + 2\pi/3$ permutes the three states: $|0\rangle \to |1\rangle \to |2\rangle \to |0\rangle$. Coherence of a general qutrit superposition $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle$ requires that the framework's couplings act *uniformly* on the three states — i.e., are Z_3-*invariant*. Any framework coupling that distinguishes between the three states (i.e., Z_3-*breaking*) drives qutrit decoherence.

The question is therefore the *symmetry analog* of the §5.7 question at the cone:

- At the cone at the cone (§5.7–§5.8), the question was: which framework couplings have $Y_{2,0}(\vartheta_{node}) = 0$ protection? Answer: the leading Higgs ($\ell=2, m=0$) coupling, with structural protection by the same identity that defines the cone.
- At the trisection , the question is: which framework couplings are Z_3-*invariant* and therefore preserve qutrit coherence? Answer: the Higgs ($m=0$) and the trisection octopole ($m=\pm 3$) — both Z_3-*invariant* — with leading Z_3-*breaking* from the cosmic hot-axis dipole at $m=\pm 1$, suppressed by ε_{ZPE}.

Define the qutrit-decoherence figure of merit in analogy with η_{node}:

$$\eta_{qutrit} \equiv \Gamma_{qutrit\text{-}decoherence} / \Gamma_{equator\text{-}reference} \quad (5.9.1)$$

where $\Gamma_{qutrit\text{-}decoherence}$ is the rate at which a generic qutrit superposition loses coherence among its three components, and $\Gamma_{equator\text{-}reference}$ is a baseline rate for an unprotected polarization-state on the equator. $\eta_{qutrit} < 1$ signals coherence protection.

5.9.2 Z_3 action on spherical harmonics: the symmetry classification

The Z_3 rotation acts on a spherical harmonic $Y_{\ell,m}(\vartheta, \varphi)$ via $\varphi \to \varphi + 2\pi/3$:

$$Y_{\ell,m}(\vartheta, \varphi + 2\pi/3) = e^{i \cdot m \cdot 2\pi/3} \cdot Y_{\ell,m}(\vartheta, \varphi) = \omega^m \cdot Y_{\ell,m}(\vartheta, \varphi) \quad (5.9.2)$$

with $\omega = e^{2\pi i/3}$ the primitive cube root of unity ([6] §10.3 line 1479). Z_3-*invariant* harmonics satisfy $\omega^m = 1$, requiring:

$$\boldsymbol{m \equiv 0 \pmod 3, \text{ i.e., } m \in \{0, \pm 3, \pm 6, \pm 9, \ldots\}} \quad (Z_3\text{-invariance condition})$$

All other harmonics (*m = ±1, ±2, ±4, ±5, …*) acquire a non-trivial phase under Z_3 and therefore *break* the Z_3 permutation symmetry of the qutrit basis.

5.9.3 Framework couplings that preserve Z_3 (qutrit-coherence-protecting)

Two framework couplings have leading *Z_3-invariant* angular structure and therefore preserve qutrit coherence:

- **The Higgs (*ℓ = 2, m = 0*)** — the breathing-mode *Y_{2,0}(ϑ)* coupling that drives the §5.4 $\chi^{(3)} \sim \alpha^2$ derivation. Manifestly *Z_3-invariant* (*m = 0*). Acts identically on all three qutrit states.
- **The trisection octopole (*ℓ ∈ {3, 6, 9, …}, m = ±3*)** — the *cos(3φ)* angular structure of [6] §11.8 line 3051. Constructed *to be Z_3-invariant* (the trisection symmetry is its defining property). The *Y_{ℓ,±3}* modes for *ℓ = 3, 6, 9, …* all satisfy *ω^{±3} = 1*, acting identically on all three qutrit states.

These two coupling classes are precisely the framework structures that define the geometric content of the (3+3) cuboctahedral construction. The qutrit basis is therefore *protected by the same Z_3 symmetry that defines SU(3), the three-generation structure of fermions, and the tribimaximal PMNS mixing at zeroth order* ([6] §10.3, §11.8 line 2132). This is a strong structural protection: the qutrit benefits from the same symmetry that organises the framework's most consequential geometric closures.

5.9.4 Framework couplings that break Z_3 (qutrit-decoherence sources)

The framework also contains *Z_3-breaking* couplings, which drive qutrit decoherence. The leading source is identified in [6] §14bis.27 line 3050 as the **cosmic hot-axis dipole**:

Cosmic hot-axis ↔ cos(φ) ↔ Y_{1,±1} (m = ±1, breaks Z_3)

This is the *l = 1* dipole imprint on the photon field from the cosmic precession axis (the direction inherited from the previous cycle's contraction, fixed by SO(3) → SO(2) breaking at the cycle-2 Bang). It corresponds to a permanent angular modulation along a specific cosmological direction, with magnitude controlled by the framework's *ε_ZPE = m_H/(4v) = 0.1270* zero-point amplitude (see [6] §11.11 line 2087+: PMNS deviations from tribimaximal mixing are at order *ε_ZPE × θ_C*, confirming the *ε_ZPE*-order Z_3-breaking from the hot axis at the lepton-mixing level).

Numerically, the cosmic-axis dipole evaluated at the three trisection vertices gives:

$$\cos(0°) = 1,\ \cos(120°) = -1/2,\ \cos(240°) = -1/2$$

with mean *(1 - 1/2 - 1/2)/3 = 0* and *Z_3-breaking variance* (squared dispersion) *((1)² + (-1/2)² + (-1/2)²)/3 = 0.5*. Higher-multipole *Y_{ℓ, ±1}, Y_{ℓ, ±2}* modes contribute at higher order in *ε_ZPE* and are subdominant.

5.9.5 Magnitude estimate: η_qutrit|_framework

The framework-internal qutrit-decoherence rate is set by the squared amplitude of the *Z_3-breaking* coupling times the dispersion across the three qutrit states:

$$\eta\_qutrit|\_framework \sim (\text{cosmic-axis amplitude})^2 \times Z_3\text{-breaking variance} \quad (5.9.3)$$

With cosmic-axis amplitude at order *ε_ZPE* (matching the PMNS deviation order) and Z_3-breaking variance *0.5* from §5.9.4:

$$\boldsymbol{\eta\_qutrit|\_framework \approx \varepsilon\_ZPE^2 \times 0.5 = (0.127)^2 \times 0.5 = 0.0081 \approx 0.81\%} \quad (5.9.4)$$

This is about *3*× the qubit framework floor $\eta\_node|\_framework \lesssim 0.3\%$ but still much smaller than the medium contribution $\eta\_medium \sim 0.05–0.20$. The qutrit pays a small additional framework-internal coherence cost relative to the qubit, but the cost is negligible relative to the dominant medium contribution.

5.9.6 Total η_qutrit and comparison with η_node

The total qutrit decoherence rate is the sum of the framework contribution (§5.9.5) and the medium contribution (analogous to the qubit *K_medium*):

$$\eta\_qutrit \approx \eta\_qutrit|\_framework + \eta\_qutrit|\_medium \approx 0.008 + (0.05 - 0.20) \approx 0.06 - 0.21 \quad (5.9.5)$$

compared to the qubit prediction (§5.8.6):

$$\eta\_node \approx \eta\_node|\_framework + \eta\_node|\_medium \approx 10^{-3} + (0.05 - 0.20) \approx 0.05 - 0.20 \quad \text{(compare to 5.9.5)}$$

The two predictions are nearly equal — both are dominated by the medium contribution, with the framework floor adding ≲ 1% in either case. The qutrit's additional 3× framework-floor cost (*0.8%* vs *0.3%*) is buried in the medium-dominated bracket and does not materially affect the bottom-line prediction.

5.9.7 Engineering implications: the qutrit information advantage

The qutrit basis encodes $\log_2 3 = 1.5850$ bits per photon, compared to the qubit's $\log_2 2 = 1$ bit per photon. This is a ***58.5% increase in information density*** per photon. shows that this advantage is achieved at essentially the *same* coherence cost as qubit photonics:

- **Both qubit and qutrit are medium-dominated** (*η_node*, *η_qutrit ≈ 0.05–0.20* in real devices).
- **Framework-floor difference is negligible** (qubit *0.3%*, qutrit *0.8%* — both ≪ medium).
- **Z_3 symmetry protection is structural** (the same symmetry that organises SU(3) and three-generation fermion structure).
- **Information advantage compounds with photon count** (an *N*-photon qutrit-encoded register stores $N \cdot \log_2 3$ bits compared to *N* bits for qubits — *58.5%* advantage scales directly).

The implication for the §6.4 Proposal 4 (trisection qutrit photonic computing) is that the proposal is *structurally viable*: the qutrit decoherence floor is comparable to the qubit floor, with no additional framework-imposed penalty. The practical bottleneck for qutrit photonic computing is the same as for qubit photonic computing — medium quality and apparatus precision — and not any framework-internal symmetry-breaking effect.

Note one practical subtlety: the cosmic hot-axis dipole (the leading Z_3-breaking source) imprints a *fixed direction* in space (the cosmological axis inherited from cycle 1). For lab-scale qutrit experiments, the orientation of the photonic apparatus relative to this cosmological axis determines the actual Z_3*-breaking* strength. A lab apparatus aligned along the cosmic axis experiences full *0.8%* breaking; one perpendicular experiences none at first order; randomly-oriented labs average to roughly *0.4%*. This orientation-dependence is a falsifiable signature in principle (a planet-scale rotation experiment could detect it) but is operationally negligible compared to medium contributions.

5.9.8 What §5.9 establishes

Established here:

- The qutrit basis at trisection vertices is protected by the Z_3 permutation symmetry, which is preserved by the framework's leading couplings (Higgs $m = 0$ and trisection octopole $m = \pm 3$).
- The leading Z_3*-breaking* framework coupling is the cosmic hot-axis dipole ($m = \pm 1$) at order *ε_ZPE*, giving *η_qutrit|_framework ≈ ε_ZPE² × 0.5 ≈ 0.81%*.
- The total qutrit decoherence *η_qutrit ≈ 0.06 − 0.21* in real devices is medium-dominated, comparable to the qubit *η_node*.
- The qutrit gains *58.5%* information per photon at essentially the same coherence cost — Proposal 4 (§6.4) is structurally viable.

Remaining (qutrit-specific items):

- Medium-specific decoherence-rate contributions for the qutrit basis at trisection vertices in candidate platforms (silicon, SiN, $LiNbO_3$) — anchored photonics work, parallel to the qubit *K_medium* deferred to specialist work.
- Higher-order *ε_ZPE^n* corrections to the Z_3*-breaking* magnitude, requiring full perturbation theory of the cosmic-axis-induced quantization shifts.
- The orientation-dependence of *η_qutrit|_framework* on the lab apparatus's alignment with the cosmic axis — a falsifiable in-principle signature but operationally negligible.

The result is therefore: ***the qutrit basis at the trisection vertices is protected by the framework's Z_3 symmetry — the same symmetry that organises SU(3), the three-generation fermion structure, and the tribimaximal PMNS mixing — with leading Z_3-breaking from the cosmic hot-axis dipole at order ε_ZPE giving η_qutrit|_framework ≈ 0.8% (small, ~3× the qubit floor, ≪ medium). Total η_qutrit ≈ medium-dominated and comparable to η_node, while the qutrit gains 58.5% information per photon. Proposal 4 (§6.4) is structurally viable; the engineering bottleneck is the same medium-quality limitation that applies to the qubit basis.*** completes the symmetry analysis of the framework's $m = 0$ (cone, qubit) and $m = \pm 3$ (trisection vertices, qutrit) sectors, leaving only medium-specific photonics work as the residual specialist priority.

§6 Six Engineering Proposals and Three Testable Predictions

6.1 Proposal 1: Null-arc qubit encoding on the Poincaré sphere

The standard photonic qubit uses the horizontal/vertical (H/V) basis at the equator of the Poincaré sphere. The (3+3) identification of the Poincaré sphere with the t_3 S^2 provides a principled basis for qubit design: the computational basis states should be chosen to minimise coupling to dominant noise channels.

- **Dual-rail qubit (path encoding)**: $|0\rangle$ = photon in mode A, $|1\rangle$ = photon in mode B. In (3+3): the two modes have distinct null-arc propagation directions. Gate = partial S^2 rotation coupling modes A and B (beam splitter). Loss = null-arc absorption in either mode. Advantage: no polarization control needed. Disadvantage: requires path stability.
- **Polarization qubit (Poincaré encoding)**: $|0\rangle$ = RCP, $|1\rangle$ = LCP (poles of S^2). Gate = waveplate (c-vector rotation). Loss = polarization-independent. Advantage: compact; gates are physical rotations of S^2. Disadvantage: birefringent environments rotate the S^2 state.
- **Time-bin qubit**: $|0\rangle$ = early pulse, $|1\rangle$ = late pulse. In (3+3): same null arc at two temporal positions. Gate = interferometer with path-length difference. Advantage: robust in fibre. Disadvantage: slower operation.

(3+3) recommendation (Level 3). For integrated photonic chips, the polarization qubit at the nodal-cone basis ($\vartheta = 54.74°$ from poles) gives the optimal trade-off in principle: protected from the $\ell=2$-mediated noise channels (Level 2 structural prediction), fully compatible with all standard gate operations (waveplate = S^2 rotation), and implementable in standard SiN or $LiNbO_3$ platforms. The actual magnitude of the protection (and hence the practical advantage over equator-basis encoding) depends on the open prefactor item (§10).

6.2 Proposal 2: Nodal-cone LOQC architecture

The nodal-cone LOQC architecture encodes all photonic qubits at the Poincaré-sphere latitude $\vartheta = 54.74°$. All interferometers and beam splitters operate within the nodal-cone subspace. Three improvements follow structurally (Level 2, magnitudes anchored to the open §10 prefactor):

- **TPA suppression — higher Kerr pump power.** At the nodal cone, $Y_2^0 = 0$, *structurally* eliminating the leading $\ell=2$-mediated TPA channel. The pump intensity for cross-Kerr operations can in principle be increased without TPA penalty, by a factor that depends on ε_ZPE^2 and residual higher-multipole coupling (§10). Higher pump power means larger cross-phase modulation, improving CNOT gate fidelity.
- **Thermo-optic stability — improved phase precision.** The nodal cone *structurally* suppresses the $\ell=2$ thermo-optic phase-noise channel. MZI phase drift is reduced.
- **Photon indistinguishability — suppressed spectral wander.** Phonon-induced phase jitter (which causes spectral broadening of single-photon sources) is also mediated by the $\ell=2$ Higgs channel. At the nodal cone, this coupling vanishes structurally, improving photon indistinguishability.

Testable prediction 1 (Level 2). The KLM gate success probability should show a local maximum at $\vartheta = 54.74°$ on the Poincaré sphere when ancilla photons are encoded at this latitude rather than at the standard H/V equatorial basis. The framework predicts the *existence and location* of the maximum at structural level. The *magnitude* of the improvement is a Level-2 prediction whose prefactor is anchored (§10) and is therefore reported as a structural prediction whose quantitative test will discriminate between framework variants. Decisive precision: < 0.5% measurable improvement at $\vartheta = 54.74°$ vs. equator basis, in a setup with otherwise-equal optical components.

6.3 Proposal 3: (3+3) measurement protocol for photonic QC

The t_3-entrainment picture of measurement gives a design principle for photonic detectors: measurement strength is controlled by the number of entraining atoms *N*. This provides a continuum between non-demolition probing (small *N*) and projective measurement (large *N*).

Protocol	*N* (atoms)	(3+3) effect	Application in photonic QC
Projective (standard)	$N \gg 10^{20}$	Complete t_3 null-arc entrainment. Photon definite state.	KLM ancilla measurement. Bell-state discrimination. Boson-sampling detection.
Homodyne (quadrature)	$N \sim 10^{12}$ (coherent field)	Null arc entrained to coherent-field axis. Continuous projection.	Gaussian boson-sampling readout. Squeezed-state measurement. CV quantum computing.
Weak measurement	$N \sim 1–10$	Partial t_3 entrainment. Qubit nudged, not collapsed.	Non-demolition photon counting. Quantum error syndrome detection.
Nodal-cone QND (proposed, Level 3)	$N \sim 100$ (nodal-cone atoms)	Entrainment at nodal cone: $\ell=2$ coupling structurally suppressed. Minimal back-action.	Error-syndrome detection with reduced disturbance. Photon-number determination without state collapse.

6.4 Proposal 4: Trisection qutrit photonic quantum processor

The S^2 trisection provides three degenerate states at 120° intervals on the Poincaré-sphere equator, protected by the same threefold symmetry that fixes the Weinberg angle $\sin^2\theta\_W = 3/8$ at tree level [§10.2 of 6] and proves the lepton Koide identity $Q\_lep = 2/3$ exactly [§11.8 of 6]. Encoding the computational basis in these three states gives a photonic qutrit with $\log_2 3 \approx 1.585$ bits per photon. For *N* photons, the Hilbert-space dimension is 3^N vs 2^N for qubits. At $N = 50$: $3^{50} \approx 7.2 \times 10^{23}$ (qutrit) vs $2^{50} \approx 1.1 \times 10^{15}$ (qubit) — a factor of 6.4×10^8 larger Hilbert space with the same number of photons.

The qutrit gate set (SU(3) operations) maps to polarization control:

- ***X_3* (shift gate).** 120° rotation of Poincaré sphere around *z*-axis: $|0\rangle \to |1\rangle \to |2\rangle \to |0\rangle$. Implemented by an electro-optic modulator (EOM) set to rotate φ by 120°.
- ***Z_3* (clock gate).** Phase ω^k applied to state $|k\rangle$ where $\omega = e^{(2\pi i/3)}$. Implemented by a sequence of three-level Pockels-cell operations.
- ***H_3* (qutrit Hadamard = DFT_3).** Discrete Fourier transform over $\mathbb{Z}_3$: $|k\rangle \to \sum_j \omega^{(kj)} |j\rangle/\sqrt{3}$. Implemented by a symmetric beam-splitter array with three-way interference.
- **Two-qutrit gate.** Controlled-X_3: if control is $|k\rangle$, rotate target by $k \times 120°$. Requires cross-Kerr coupling between two photons. Can be implemented via KLM with ancilla qutrits.

Testable prediction 2 (Level 2). A trisection-qutrit gate set (X_3, Z_3, H_3) implemented on a photonic platform with three-fold symmetric beam-splitter arrays should demonstrate: (i) the three qutrit states have *equal noise sensitivity* under any perturbation that preserves the threefold symmetry of the photonic circuit (Level 1 structural prediction from SU(3) symmetry); (ii) gate fidelity comparable to or exceeding the qubit-encoding threshold (> *99%*) with fewer optical elements than an equivalent qubit circuit performing the same Hilbert-space-equivalent computation (Level 2 — the reduction factor is structural at $\log_2 3 \approx 1.585$ but practical implementation overheads are not yet derived).

6.5 Proposal 5: Boson sampling as null-arc network interference

Gaussian boson sampling (GBS) places *n* photons through an *m*-mode linear interferometer and samples the output photon distribution. The (3+3) interpretation: each photon is a null arc, each beam splitter is a partial S^2 rotation coupling two modes, and the GBS output distribution is the squared modulus of the sum over all paths of the null-arc propagation amplitudes through the S^2-rotation network.

This interpretation provides a geometric reading of quantum advantage in GBS: the output distribution samples the permanent (for pure Fock-state input) or the Hafnian (for squeezed-state input) of the scattering matrix — a quantity that is #P-hard to compute classically. The Hafnian counts perfect matchings of a graph; in (3+3), this corresponds to the number of ways to pair up null arcs so that each pair has zero total angular momentum — the number of singlet configurations of the null-arc network. (Level 1 interpretation; no new predictions for existing GBS hardware.)

GBS experiment	Photons	Modes	(3+3) null-arc network size
Jiuzhang 1.0 (2020) [12]	76	100	$C(175,76) \approx 6 \times 10^{44}$ null-arc configurations.
Jiuzhang 2.0 (2021) [13]	113	144	$C(256,113) \approx 10^{41}$ null-arc configurations.

GBS experiment	Photons	Modes	(3+3) null-arc network size
Borealis (Xanadu, 2022) [14]	219	216	$C(434,219) \approx 10^{61}$ null-arc configurations.
Trisection GBS (proposed, Level 3)	50 qutrits	100	$3^{50} = 7.2 \times 10^{23}$ qutrit states sampled — larger Hilbert space than qubit GBS at same photon count.

6.6 Proposal 6: Nodal-cone photonic-crystal integrated circuit

Combining the (3+3) insights yields a specific integrated photonic quantum-processor architecture: a triangular-lattice photonic-crystal chip with nodal-cone qubit encoding. The proposal is **Level 3 (engineering, speculative)**; it depends on the Level-2 nodal-cone prediction's currently-anchored prefactor materialising as a real practical advantage, and on the multiplicativity assumption for combined enhancements.

Design principles:

- **Substrate**: Triangular-lattice photonic crystal in silicon nitride. Triangular geometry (120° bond angles) matches S^2 trisection — widest photonic bandgap, low thermo-optic drift. Waveguide defect mode engineered to sit at $\vartheta = 54.74°$.
- **Qubit encoding**: Photonic qubits at nodal-cone Poincaré-sphere states. All MZI working points set at the 54.74° latitude — self-stabilising against $\ell=2$-mediated thermal phase noise (Level 2).
- **Nonlinear elements**: Cross-Kerr gates in $LiNbO_3$ micro-ring resonators. Nodal-cone pump polarization structurally eliminates the leading $\ell=2$-mediated TPA channel.
- **Sources**: InAs/GaAs quantum-dot single-photon sources coupled via grating couplers. Operating at nodal-cone polarization structurally suppresses phonon-induced spectral jitter.
- **Detectors**: Superconducting nanowire single-photon detectors integrated on-chip or coupled via fibre. Detection efficiency $\sim 98\%$.

Testable prediction 3 (Level 2). The nodal-cone PhC chip should demonstrate measurably lower phase drift versus temperature than an equivalent standard silicon photonic chip without nodal-cone encoding. Decisive precision: lock-in measurement of $d\varphi/dT$ at the nodal-cone wavelength and polarization, comparing nodal-cone-encoded MZIs to equator-encoded MZIs of identical fabrication. The framework predicts a measurable *direction* and a *structural form* for the suppression; the magnitude is currently anchored (§10).

6.7 Falsifiability timeline 2026–2030+

The three testable predictions of §6.2, §6.4, §6.6, together with the Level-2 boson-sampling and qutrit-architecture proposals of §6.4 and §6.5, span a four-year window of imminent photonic-QC experiments. This subsection organises them into a year-resolved timeline parallel to §22.6bis of

[6]. Each target is decisive within its window: a positive result on any single target either confirms a specific structural claim of the framework or — at sufficient precision — falsifies it, with no further wiggle room for parameter readjustment, since the framework has zero free parameters in the photonic-QC sector beyond what is already fixed in [6].

Window	Experiment / setup	Target observable	Decisive outcome
2026–2027	Polarization-controlled KLM gate-fidelity comparison on existing photonic platforms (Bristol, Xanadu, NIST, MUQUABIS test-beds)	KLM CNOT gate success probability with ancillas encoded at $\vartheta = 54.74°$ nodal cone vs. equatorial H/V basis	**Sharpest single-platform near-term test.** Detection of a local maximum at ϑ_node at $< 0.5\%$ precision (in identical optics, identical ancilla photon-number, identical detection thresholds) confirms Testable prediction 1 (§6.2). Null result at $< 0.5\%$ falsifies the structural claim that the $\ell=2$ TPA channel and phonon-jitter channel are the dominant gate-fidelity-limiting effects in the H/V basis.
2026–2028	Trisection-qutrit gate-set prototype on photonic-qutrit hardware (Xanadu/PsiQuantum-style platforms or three-mode interferometric arrays)	Three-fold symmetric X_3, Z_3, H_3 gate fidelity; equal noise sensitivity across the three qutrit states under symmetry-preserving perturbations	Demonstration of gate fidelity $> 99\%$ and equal-noise behaviour confirms Testable prediction 2 (§6.4). Asymmetric noise sensitivity at the three qutrit states (under symmetry-preserving perturbations) falsifies the SU(3)-symmetry claim that descends from the trisection structure of [§10.2, §11.8 of 6].
2027–2029	Borealis-class GBS follow-ups; Hafnian-amplitude sampling at	Distribution shape of GBS samples at $n \geq 200$, compared with null-arc	GBS at $n \geq 200$ demonstrates that the null-arc network

Window	Experiment / setup	Target observable	Decisive outcome
	200+ photons; cross-validation against (3+3) null-arc network combinatorial counts	network amplitude prediction of §6.5 (matches Hafnian computation as expected from standard QED, no new (3+3) signature)	interpretation is consistent with #P-hard sampling at scale. **Not** a new prediction but a consistency test: the framework is consistent with all existing GBS results to date and predicts the same scaling for follow-ups. Anomalous departures from Hafnian sampling at scale would falsify the null-arc–as–KK-zero-mode identification.
2027–2030	Lock-in thermal MZI comparison: nodal-cone-encoded vs. equator-encoded MZIs of identical fabrication on standard silicon-photonic test chips	Phase drift *dφ/dT* at nodal-cone wavelength and polarization versus matched equator-encoded reference MZIs	Reduction of *dφ/dT* by a measurable factor confirms Testable prediction 3 (§6.6). The framework predicts the *direction* and *structural form* of the suppression; the *magnitude* is currently anchored (§10) and the experiment is therefore a **two-stage test**: (a) sign of the effect (decisive at low precision); (b) magnitude (decisive against the open §10 prefactor).
2027–2030+	Integrated nodal-cone photonic-crystal chip fabrication: triangular-lattice SiN/$LiNbO_3$ with all MZI working points at *ϑ_node*	Combined cross-Kerr enhancement, thermo-optic stability, and photon-indistinguishability under nodal-cone encoding in a single integrated processor	**Tier-3 fabrication challenge.** The chip is the most decisive single test of the §9 conditional argument: it bundles the nodal-cone suppression magnitude, the

Window	Experiment / setup	Target observable	Decisive outcome
			multiplicativity assumption, and the §5.3 silicon-enhancement scaling into one device. Demonstration that the combined enhancement is large enough to cross the LOQC fault-tolerance threshold (*η_total > 99.5%*) supports the §9 conditional. A null result, or sub-linear combination of enhancements, falsifies the multiplicativity assumption and the §10 anchored prefactor.

The decade-window summary. The most consequential single window is **2026–2027**, when the polarization-controlled KLM gate-fidelity comparison (§6.2 / Testable prediction 1) can be performed on existing photonic platforms with no new hardware development required. This is the framework's sharpest and earliest decisive test in the optical-QC sector. The next-most-consequential window is **2027–2030**, when the thermal MZI comparison (§6.6 / Testable prediction 3) yields a two-stage test of the nodal-cone suppression: sign first, magnitude second. The integrated nodal-cone PhC chip (2027–2030+) is a fabrication-stage challenge that, if successful, would convert several Level-3 engineering claims into Level-2 derived results.

Across the four-year window, the framework faces a positive-result test approximately **every 12–18 months in the optical-QC sector**. Each test is decisive in the sense that the framework has zero adjustable parameters in this sector: a null result at sufficient precision falsifies the relevant structural commitment. This density of decisive tests parallels the cosmological-sector tests catalogued in §22.6bis of [6].

Inter-paper coordination. The optical-QC tests above are independent of the matter-based QC tests catalogued in [3]. A positive result in either paper's near-term window is informative for the framework as a whole; a positive result in *both* is mutually reinforcing. The two preprints together commit the framework to a wide-front quantum-information-sector test programme over 2026–2030.

§7 Photonic vs Matter-Based Quantum Computing

A platform-by-platform comparison in the (3+3) framework. Each row lists the operating temperature, the T_2 floor (or universal floor for photons), the best-of-platform two-qubit gate, the key (3+3) interpretation, and the current scale. Cross-references to the matter-based-QC analysis are in [3].

Platform	***T_op***	**T_2 floor**	**Two-qubit gate**	**Key (3+3) interpretation**	**Scale**
Superconducting	15 mK	10–100 μs	99.5% (CZ)	Josephson junction = t_3-winding phase coupling. T_2 from Higgs-breathing phonons.	~1,000 Q
Trapped ion	RT (laser)	1–100 s	99.3% (Mølmer–Sørensen)	Ion spin = t_3-winding mode. Laser = null-arc c-vector rotation.	~50 Q
Neutral atom	RT (tweezer)	1–100 s	99.5% (Rydberg)	Rydberg interaction = large-n KK winding exchange.	~1,000 Q
Photonic (GBS)	RT	universal floor (vacuum)	N/A (linear)	Null-arc networks through S^2-rotation meshes. Quantum advantage from null-arc path count.	~219 modes
Photonic (KLM)	RT	universal floor (vacuum)	25% → ~99% with ancillas	Measurement = t_3 null-arc entrainment. Feed-forward = corrective c-vector rotation.	>1,000 Q (target)
Nodal-cone LOQC (Level 3)	RT	universal floor (vacuum)	>99% (proposed)	Nodal-cone encoding: structural $\ell=2$ suppression with currently-anchored magnitude.	Target: scalable

§8 Discussion: Why Photonics Is a Natural (3+3) Platform

Among quantum-computing platforms, photonic systems have an unusually direct connection to (3+3) geometry. The photon is a null arc: its Poincaré sphere is the t_3 S^2, its decoherence in vacuum is bounded by the universal $1/H_0$ floor (and the absence of t_3-winding channels means even this bound is not approached), its gate operations are geometric rotations of S^2, and its measurement is t_3 entrainment. Every photonic-QC concept — single-qubit gates, boson sampling, the KLM protocol — has a precise (3+3) geometric counterpart.

Matter-based qubits (superconducting, trapped ion, spin) are excellent quantum computers, but they face the intrinsic challenge that matter is a winding mode ($n = 1/2$ for electrons) on S^2, and the winding couples to the Higgs-breathing mode (the T_2 channel) and to environmental t_3 oscillations (the T_1 channel). These channels require cryogenic cooling to suppress in superconducting platforms; trapped-ion and neutral-atom platforms operate at room temperature but require laser-induced isolation to suppress environmental coupling. Photons, having $n = 0$ (null arc), have no t_3-winding coupling at all. Their only practical decoherence is loss — and loss can be combated by redundancy and error correction, unlike the fundamental T_1/T_2 channels of matter qubits.

This is a Level-1 framework observation: it follows directly from the photon being a null arc, no additional assumptions required. It is the strongest single (3+3) statement about quantum information processing.

§9 Room-Temperature Fault Tolerance: A Level-3 Conditional Argument

A frequently-asked question is whether the (3+3) framework supports a near-term route to room-temperature fault-tolerant photonic quantum computing. The argument given here is **explicitly Level 3**: it follows *if and only if* a specific chain of currently-anchored items resolves favourably. None of the items in the chain is currently established. The framework's contribution is to identify the geometric foundation that makes the items plausible to investigate experimentally and theoretically.

9.1 The framework establishes (Level 1)

- Photons have no t_3-winding decoherence channels (no Higgs-breathing T_2, no T_1); the universal $\tau\_cosm = 1/H_0 \approx 14.5$ Gyr applies as an upper bound that is not approached in vacuum [§14bis.46.4 of 6, §3.1].
- Their only practical decoherence is photon loss, which is reducible by improving component efficiency and by redundant encoding.
- The dominant TPA, thermo-optic, and phonon-jitter loss-and-gate-error channels are all ℓ=2-mediated and therefore *structurally* suppressed at the nodal cone where $Y_2^0(\vartheta\_node) = 0$ exactly (§5.2).

Current state-of-the-art photonic components: SNSPD efficiency 98% [16], quantum-dot source efficiency 90%, SiN waveguide loss 0.1 dB/cm. The threshold efficiency for fault-tolerant LOQC with redundant encoding is approximately $\eta\_total > 99.5\%$ [refs in 17].

9.2 The conditional chain

The room-temperature fault-tolerance argument follows *if and only if* **all four** of the following Level-2 / Level-3 items resolve favourably. The four items are independent in the framework's current state — a positive resolution of any one does not entail the others — so the argument is contingent on a four-step chain, not a single open item.

- **(C1) The nodal-cone suppression-factor magnitude.** The framework predicts at Level 1 that $Y_2^0(\vartheta\_node) = 0$ exactly, structurally suppressing the leading ℓ=2-mediated channels. The detailed analysis of §5.5–§5.9 shows that the first-order breathing contribution at the cone vanishes geometrically (the same identity $Y\_\{2,0\}(\vartheta\_node) = 0$ that defines the cone also eliminates the breathing's metric perturbation there); second-order ≲ 0.3%; the higher-multipole tail is bounded at $\eta\_node|\_higher \lesssim 10^{-5}$; and the Z_3 symmetry protects the qutrit basis with leading Z_3-breaking from the cosmic hot-axis dipole at order ε_ZPE, giving $\eta\_qutrit|\_framework \approx 0.8\%$. The combined prediction is $\eta\_node \approx K\_medium \approx 0.03 - 0.20$ (qubit) or $\eta\_qutrit \approx 0.06 - 0.21$ (qutrit) in real devices, **both dominated entirely by medium contributions**. **The §6.7 thermal MZI experiment effectively measures *K_medium* directly. The framework contribution to η is bounded at < 1% in either basis; engineering optimisation should focus exclusively on medium quality and apparatus precision.** Silicon nitride is the optimal platform for the cleanest framework test (Appendix B).
- **(C2) Structural Level-2 predictions for thermo-optic and indistinguishability channels at quantitative level.** §6.2 argues that the same ℓ=2 channel governs TPA, thermo-optic phase noise, and phonon-induced spectral wander. The framework predicts *correlated* (in the limit, identical) suppression factors at the nodal cone for the three channels. Quantitative confirmation of this correlation in dedicated experiments — the §6.7 thermal MZI test (2027–2030) is the leading candidate — is required to convert the structural claim into a quantitative engineering input.
- **(C3) Multiplicativity of resonant and nodal-cone enhancements (Level 3 assumption).** §5.3 and Proposal 6 assume that the Si-microring resonant $\chi^{(3)}$ enhancement and the nodal-cone TPA suppression compose multiplicatively in the combined estimate. This assumption has *not* been derived from the framework. A sub-linear or saturating combination would erode the combined gain. The integrated nodal-cone PhC chip test (§6.7, 2027–2030+) is the most direct test of multiplicativity.
- **(C4) Silicon-enhancement scaling prefactor (Level 2 anchor).** The §5.3 $\alpha\_eff = R_3/(\pi a_0,eff)$ scaling form is structural, but the *12*× prefactor for silicon and the *10*× prefactor for QD excitons are anchored to standard photonics rather than derived from (3+3). A first-principles derivation of these prefactors from the framework would convert several Level-2 estimates to derived results. (Proposed for of the present specialist programme.)

9.3 The conditional argument

If (C1) closes with the suppression magnitude sufficient to bring TPA-limited Kerr gate fidelity above 99.5%, *and* (C2) closes with thermo-optic and indistinguishability suppressions correlated at the same magnitude, *and* (C3) closes with multiplicativity confirmed at experimental level, *and* (C4) closes with the silicon enhancement derived from first principles within an order of magnitude of the empirical *12×*, *then* the framework supports a route to room-temperature fault-tolerant photonic QC within the existing materials-science envelope.

None of the four conditions is currently established. The framework provides the geometric foundation that makes them plausible to investigate, and the §6.7 timeline organises the experimental side of the programme. The theoretical side (specifically C1 and C4) requires specialist foam-level calculations parallel in spirit to §16 of [6] but specific to the photonic-QC application; the present preprint defers these to follow-up work (Stages 4a and 4b of the staged specialist programme).

What this argument does not establish. It does not establish that room-temperature fault tolerance is imminent, easy, or guaranteed. The most consequential single open item is (C1) — the rigorous derivation of the nodal-cone suppression magnitude from the foam structure. Until (C1) closes, every quantitative claim downstream of it remains anchored, and the room-temperature fault-tolerance argument is contingent at structural rather than predictive level. The argument *may also fail* even if (C1) closes favourably, if (C2)–(C4) do not.

What would change the conclusion. A null result at decisive precision in any of the §6.7 tests — particularly Testable predictions 1 (KLM gate-fidelity peak) and 3 (thermal MZI suppression) — would directly falsify the structural Level-1 claim that the $\ell=2$ channel dominates the relevant noise budget, and would render the entire conditional chain moot. Conversely, positive results across the §6.7 timeline would convert several items in the C1–C4 chain to closed status and tighten the argument.

§10 Honest Limits

Following the convention established in §24.3 of [6], we list explicitly every Level-2 prefactor currently anchored rather than derived, and every Level-3 proposal whose plausibility depends on these anchors becoming derivations. This catalogue is the article's honest record of where the speculation lives. It is organised into three categories: **(A)** Level-2 anchored prefactors (the most consequential open items — these are the targets of the specialist programme); **(B)** Level-3 engineering speculations that follow if the Level-2 items close; **(C)** items deferred to follow-up work or to the companion preprint [3].

10.1 Category A — Level-2 anchored prefactors

- **(A1) The nodal-cone suppression-factor magnitude.** The framework predicts at Level 1 that $Y_2^0(\vartheta\_node) = 0$ exactly, structurally suppressing the leading $\ell=2$-mediated TPA, thermo-optic, and phonon-jitter channels at the cone. The residual η_node has three sources: (i) the breathing-mode contribution; (ii) the higher-multipole tail through *ℓ = 4, 6, …*; (iii) medium-dependent corrections. The detailed analysis of §5.5–§5.9 establishes the structural form $\eta_{node} = 8 k_2 \varepsilon_{ZPE}^2 + \Sigma_{\{\ell\geq4\}} k_\ell |Y_\ell^0(\vartheta_{node})|^2/|Y_\ell^0(\pi/2)|^2 S_\ell + K_{medium}$ with the geometric coefficients computed exactly, and gives the following bounds: **(i)** the first-order breathing contribution at the cone *vanishes geometrically* because the same identity $Y_{\{2,0\}}(\vartheta_{node}) = 0$ that defines the cone *also* eliminates the breathing's metric perturbation there (§5.7); the second-order contribution is bounded at $\lesssim 0.3\%$. **(ii)** The framework's identified ℓ ≥ 4 mode is the fermionic right-handed neutrino at seesaw scale (not a scalar); the trisection contributes only $m = \pm3$ modes (not the $m = 0$ tail); higher-ℓ $m = 0$ scalars (if physically realised) are bounded at $\eta_{node}|_{higher} \lesssim 10^{-5}$ across plausible scenarios (§5.8). **(iii)** Anchored medium-specific estimates give $K_{medium} \approx 0.10–0.20$ (silicon-on-insulator), *0.03–0.08* (silicon nitride, lowest), *0.05–0.15* (thin-film lithium niobate) (Appendix B). The combined prediction is $\eta_{node} \approx K_{medium} \approx 0.03–0.20$ across candidate platforms, *dominated entirely by medium contributions*; the framework vacuum and higher-multipole contributions are both bounded at < *1%*. The §6.7 thermal MZI experiment empirically discriminates and effectively measures *K_medium* directly. Silicon nitride is identified as the optimal platform for the cleanest framework test. The residual specialist work is medium-specific photonics (anchored to standard photonics literature, not (3+3)-derivable); the breathing-mode and higher-multipole questions are closed in either basis.

- **(A2) The $\chi^{(3)}$ enhancement scaling in dielectric media.** The article uses $\alpha_{eff} = R_3/(\pi a_0,eff)$ with a_0,eff reduced in high-index media. The scaling form is structural; the *prefactor* (the *12*× for silicon, the *10*× for QD excitons) is anchored to standard photonics rather than derived from (3+3). A first-principles derivation of these prefactors from the framework would convert §5.3 Level-2 estimates to derived results.
- **(A3) The photon-loss t_3-entrainment scaling.** §3.1 and §5.2 identify photon loss as null-arc absorption by an absorbing atom's t_3-winding mode. The *structural* form of the absorption amplitude (one c-vector rotation, amplitude $\sqrt{\alpha}$) is Level 1. The *medium-specific* loss rates (0.2 dB/km for fibre, 0.1 dB/cm for SiN waveguides, etc.) are inputs from materials science, not derived from the framework. The framework's contribution here is geometric reframing, not a quantitative derivation.
- **(A4) The TPA polar profile $|Y_2^0(\vartheta)|^2$ in finite-ε_ZPE environments.** §5.2 states that TPA scales as $|Y_2^0(\vartheta)|^2$ exactly *at the cone*. The *integrated* TPA over a finite zero-point latitude spread depends on ε_{ZPE}^2 (item A1) and on the higher-multipole tail. After (§5.5), this is mathematically the same open item as (A1) but appears at multiple points in §5–§9 and is listed separately for indexing.
- **(A5) Three-channel correlation: TPA, thermo-optic, phonon-jitter.** §6.2 and §9 argue that the same *ℓ=2* Higgs channel governs all three. The Level-1 *structural* identification of the common channel is established in [6]. The *quantitative* claim that the three suppression factors are equal (or correlated to within experimental discrimination) is Level

2 and depends on the explicit channel decomposition in dielectric environments. Test: §6.7 thermal MZI experiment (2027–2030).

10.2 Category B — Level-3 engineering speculations

- **(B1) Multiplicativity of resonant and nodal-cone enhancements.** §5.3 and Proposal 6 assume that the Si-microring resonant $\chi^{(3)}$ enhancement and the nodal-cone TPA suppression compose multiplicatively in the combined estimate. This assumption has not been derived from the framework. Sub-linear combination would erode the combined gain.
- **(B2) The combined nonlinearity enhancement of Proposal 6.** The Si-microring + nodal-cone combined enhancement (the bottom row of the §5.3 table) is the product of two Level-2 enhancements (A1, A2) with a Level-3 multiplicativity assumption (B1). The numerical result is therefore a triple-anchored Level-3 estimate; it is the most speculative single quantitative claim in the article. The integrated nodal-cone PhC chip test (§6.7, 2027–2030+) is the most direct experimental probe.
- **(B3) Trisection-qutrit photonic implementation route.** No integrated photonic platform currently implements full SU(3) qutrit operations with three-fold symmetric beam-splitter arrays. The proposal is an engineering programme rather than a derivation, even though the underlying SU(3) symmetry of the trisection is Level 1 [§10.2, §11.8 of 6].
- **(B4) Nodal-cone QND measurement protocol (Proposal 3, last row).** The proposed t_3-entrainment-strength control for non-demolition photon-number measurement at the nodal cone is Level 3: it depends on (A1) and on the engineering availability of an $N \sim 100$ nodal-cone-encoded atomic interaction region. No experimental implementation has been attempted to our knowledge.
- **(B5) Polarization-qubit recommendation at nodal-cone basis.** Proposal 1's recommendation that integrated photonic chips encode polarization qubits at the nodal-cone basis is Level 3: it depends on the practical-magnitude open item (A1) and on the engineering tractability of nodal-cone polarization control across realistic chip layouts (waveguide birefringence, grating-coupler polarization sensitivity, etc.).
- **(B6) The room-temperature fault-tolerance projection (§9).** Contingent on the four conditions C1–C4 of §9.2 resolving favourably. The argument in §9 is conditional, not predictive. As §9.3 notes explicitly, the argument *may fail* even if C1 closes favourably, if any of C2–C4 fails.

10.3 Category C — Items deferred or referred elsewhere

- **(C1) Mixed qubit-qutrit processor Hilbert-space dimension.** The article does not derive a closed-form expression for the Hilbert-space dimension of a mixed photonic processor with n_q qubit-encoded photons and n_t trisection-qutrit-encoded photons. The naive $2^{n_q} \cdot 3^{n_t}$ count is an upper bound; the achievable computational dimension under SU(2) × SU(3) gate constraints is open. Deferred to follow-up work.

- **(C2) Foam-level derivation of the √α c-vector rotation amplitude.** The §5 statement that each c-vector rotation has amplitude $\sqrt{\alpha}$ is established in [6] at structural level. The full foam-path-integral derivation parallel to §16 of [6] is referred to the specialist programme of [6] §23.
- **(C3) Matter-based QC content.** The general-QC implications of the framework are covered in the companion preprint [3], including superconducting, trapped-ion, neutral-atom, and spin-qubit platforms. The present paper covers only the photonic platform.
- **(C4) Boson-sampling complexity and (3+3).** §6.5 gives the null-arc network *interpretation* of GBS but does not derive any new complexity-theoretic result for #P-hard sampling. The framework's contribution to GBS is interpretive, not new-prediction-generating, beyond the qutrit GBS proposal of Proposal 4.

Summary. A reader weighing the article on purely formal grounds will find these limits explicit. A reader weighing it on structural grounds will find that the Level-1 geometric identifications, the $\chi^{(3)} \sim \alpha^2$ scaling (derived in §5.4), the structural form of the nodal-cone suppression $\eta\_node = 8\ k\_2\ \varepsilon\_ZPE^2 + \Sigma\_\{\ell\geq 4\}\ k\_\ell\ |Y\_\ell^0(\vartheta\_node)|^2/|Y\_\ell^0(\pi/2)|^2\ S\_\ell + K\_medium$ (derived in §5.5), and the bracket $\eta\_node|\_total \sim 0.10 - 0.30$ (§5.6) are derived from the framework, and that the article's engineering proposals follow as a coherent programme of work that the framework's structural commitments make worth pursuing. The single most consequential open item is (A1) — the nodal-cone suppression-factor magnitude — which converted from a fully Level-2 anchored claim to a Level-1 structural form, and further bracketed the breathing-mode prefactor k_2 within $[\varepsilon\_ZPE^2 \approx 0.016, \sim 1.2]$ with Casimir-analogy central estimate $k\_2 \approx 1/2$. closure of the foam-Lagrangian for the photon polarization-mode field would convert several Level-2 results to Level-1 derivations and tighten the §9 conditional argument substantially.

§11 Conclusions

The (3+3) framework establishes a clear geometric foundation for photonic quantum computing. The photon as null arc is the $n=0$ KK zero-mode on S^2; its Poincaré sphere is the t_3 S^2; its gate operations are physical S^2 rotations; its measurement is t_3 entrainment; and the weakness of photon-photon interactions ($\chi^{(3)} \sim \alpha^2$) follows from two successive c-vector rotations each at amplitude $\sqrt{\alpha}$. These are Level-1 results that follow directly from the framework as established in [1, 6].

Six engineering proposals flow from this foundation. The strongest are Proposal 2 (nodal-cone LOQC architecture) and Proposal 4 (trisection qutrit), both of which lift specific structural results from the (3+3) cosmology and particle-physics sectors and apply them to optical quantum-information processing. Proposals 1, 3, 5, and 6 develop the framework's qubit-encoding, measurement, boson-sampling, and integrated-photonics implications respectively.

Three testable predictions (§6.2, §6.4, §6.6) are presented at Level 2: each predicts a *structural* effect (the *direction* and *form* of an experimental signal) that the framework derives, with *magnitudes* whose prefactors are currently anchored to the framework rather than derived from first principles. The predictions are organised into a year-resolved 2026–2030+ falsifiability timeline (§6.7) with the polarization-controlled KLM gate-fidelity comparison (2026–2027) as the sharpest near-term test.

Honest framing. The article's most consequential single open item — the rigorous derivation of the nodal-cone suppression-factor magnitude from the foam structure — is now resolved at the structural level by §5.5–§5.9. The first-order breathing-mode contribution at the cone vanishes geometrically because the same *Y_{2,0}(ϑ_node) = 0* identity that defines the cone also eliminates the breathing's metric perturbation there. The higher-multipole tail is bounded at $\eta\_node|\_higher \lesssim 10^{-5}$ using the framework's identified spectrum. The qutrit basis is protected by the same Z_3 symmetry that organises SU(3), three-generation fermion structure, and tribimaximal PMNS mixing. The total predictions $\eta\_node \approx K\_medium \approx 0.03 - 0.20$ (qubit, with silicon nitride lowest) and $\eta\_qutrit \approx 0.06 - 0.21$ (qutrit) in real devices are both **dominated entirely by medium contributions**, with framework-vacuum and higher-multipole/Z_3-breaking contributions all $< 1\%$. The remaining work is medium-specific photonics for silicon, silicon nitride, and lithium niobate — anchored to standard photonics rather than derived from (3+3) — for which Appendix B gives indicative numerical estimates. Until *K_medium* is empirically anchored for a specific platform via the §6.7 thermal MZI experiment, the article is best understood as a structural-coherence argument that the framework's commitments make a specific engineering programme worth pursuing — not as a derivation that the programme is guaranteed to succeed. The room-temperature fault-tolerance projection of §9 is **explicitly Level 3 conditional**, contingent on the four-step chain C1–C4 of §9.2 resolving favourably, and may fail even if the most consequential single item C1 closes favourably.

The deepest insight. The photon's status as an effectively ideal quantum-information carrier in vacuum is not accidental. It follows geometrically from the photon being a null arc — the $n=0$ state with no t_3 winding, no thermal-bath coupling, and no fundamental T_1/T_2 decoherence. Quantum information carried by null arcs is protected by the same geometry that makes light propagate at c and gives the photon its zero rest mass. The challenge is not in protecting the qubit — the Poincaré sphere does that for free — but in making null arcs interact with each other strongly enough to compute. The nodal-cone and trisection-qutrit structures are the framework's two specific suggestions for how that challenge might be made tractable.

Acknowledgements

C.R. de Haan acknowledges Claude (Anthropic) for: (i) numerical verification of all quantitative results in this paper, including the $\chi^{(3)} \sim \alpha^2$ scaling, the structural form of the nodal-cone suppression, the boson-sampling combinatorial estimates, and the GBS photon-count analysis cross-checked against published Jiuzhang and Borealis results; (ii) systematic analysis of the photonic-QC literature, including KLM, fusion-based QC [18], cluster-state MBQC [19], and silicon-photonic LOQC [20], and identifying where (3+3) predictions distinguish from or align with existing work; (iii) structured scientific presentation, including the claim-level discipline of §2.2, the honest-limits catalogue of §10, the conditional framing of §9, and the falsifiability timeline of §6.7.

Theoretical framework is the original work of the author. Responsibility for all content rests with the author alone. No external funding received.

References

A. (3+3) programme primary sources

- [1] de Haan, C.R. (2026). *From One Sphere to All of Physics*. Book manuscript. Zenodo, DOI: 10.5281/zenodo.19633127.
- [2] de Haan, C.R. (2026). *The Proton-to-Electron Mass Ratio from S^2 Topology*. Zenodo, DOI: 10.5281/zenodo.19651418.
- [3] de Haan, C.R. (2026). *Quantum Computing Through the Lens of Three Time Dimensions*. Zenodo, DOI: 10.5281/zenodo.19651560. (**Companion preprint to the present paper.** [3] covers matter-based quantum-computing platforms and the framework's general QC implications; the present paper covers the photonic platform specifically.)
- [4] de Haan, C.R. (2026). *The Hubble Tension from S^2 Geometry*. Zenodo, DOI: 10.5281/zenodo.19666441.
- [5] de Haan, C.R. (2026). *The Born Rule from S^2 Geometry*. Zenodo, DOI: 10.5281/zenodo.19699743.
- [6] de Haan, C.R. (2026). *The 6D Action Principle for the (3+3) Spacetime Framework: Lagrangian Foundations, Geometric Closures Across Particle Physics, Cosmology, Gravity, and Quantum Foundations from a Discrete S^2 with 2^{152} Cells*. Zenodo, DOI: [to be assigned on deposit]. (**Foundational preprint.** All section references in the present paper [§14bis.46.4 of 6, etc.] point to this document.)

B. (3+3) framework underpinnings (Einstein–Aether and KK)

- [7] Jacobson, T. & Mattingly, D. (2001). Gravity with a dynamical preferred frame. *Phys. Rev. D* **64**, 024028.
- [8] Eling, C., Jacobson, T. & Mattingly, D. (2006). Einstein–Aether theory. *arXiv:gr-qc/0410001*.

C. Photonic quantum computing — foundational

- [9] O'Brien, J.L. (2007). Optical quantum computing. *Science* **318**, 1567–1570.
- [10] Kok, P. et al. (2007). Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **79**, 135.
- [11] Knill, E., Laflamme, R. & Milburn, G.J. (2001). A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52.

D. Boson sampling and quantum advantage

- [12] Zhong, H.-S. et al. (2020). Quantum computational advantage using photons. *Science* **370**, 1460–1463. [Jiuzhang 1.0]
- [13] Zhong, H.-S. et al. (2021). Phase-Programmable Gaussian Boson Sampling. *Phys. Rev. Lett.* **127**, 180502. [Jiuzhang 2.0]
- [14] Madsen, L.S. et al. (2022). Quantum computational advantage with a programmable photonic processor. *Nature* **606**, 75–81. [Borealis]

E. Photonic-QC components and architectures

- [15] Tomm, N. et al. (2021). A bright and fast source of coherent single photons. *Nat. Nanotech.* **16**, 399–403.
- [16] Chang, J. et al. (2021). Detecting Telecom Single Photons with 99.5-percent System Detection Efficiency. *Optica* **8**, 127.
- [17] Ewert, F. & van Loock, P. (2014). 3/4-Efficient Bell Measurement with Passive Linear Optics. *Phys. Rev. Lett.* **113**, 140403.
- [18] Bartolucci, F. et al. (2023). Fusion-based quantum computation. *Nat. Commun.* **14**, 912.
- [19] Larsen, M.V. et al. (2019). Deterministic generation of a two-dimensional cluster state. *Science* **366**, 369–372.
- [20] Rudolph, T. (2017). Why I am optimistic about the silicon-photonic route to quantum computing. *APL Photonics* **2**, 030901.
- [21] Wang, J. et al. (2020). Integrated photonic quantum technologies. *Nat. Photon.* **14**, 273–284.
- [22] Spring, J.B. et al. (2013). Boson Sampling on a Photonic Chip. *Science* **339**, 798–801.
- [23] Carolan, J. et al. (2015). Universal linear optics. *Science* **349**, 711–716.
- [24] Jeannic, H.L. et al. (2022). Dynamical photon-photon interaction mediated by a single emitter. *Nat. Phys.* **18**, 1191.

F. Nonlinear optics foundations

- [25] Boyd, R.W. (2020). *Nonlinear Optics*, 4th edition. Academic Press, San Diego. Standard reference for the $\chi^{(n)}$ hierarchy in dielectric media and the order-of-magnitude scaling $\chi^{(3)}/\chi^{(1)} \sim 10^{-5}$ in silica and similar glasses.

Appendix A: Version history

This preprint was developed through a staged programme of additions and refinements, each producing a numbered version. The staged structure is recorded here so that readers of any specific Zenodo deposit can identify the structural content of the version they are reading. All versions share the same author, AI Assistance Statement, and topical framing; versions differ in derivational depth and in the level at which specific quantitative claims are anchored vs. derived.

- **v0.1** (May 2026, Stage 1). Initial draft. Introduces the photon-as-null-arc identification, the polarization-Poincaré-sphere = t_3 S^2 correspondence, the Bell-state geometric constraints, the $\chi^{(3)} \sim \alpha^2$ scaling (cited from [6]), six engineering proposals, three testable predictions, and a preliminary honest-limits catalogue.
- **v0.2** (May 2026, Stage 2). Honest-framing pass. Adds §1.7 (self-contained one-page framework digest for readers unfamiliar with the (3+3) programme); §6.7 (year-resolved 2026–2030+ falsifiability timeline organising the three testable predictions); strengthens §9 to a Level-3 conditional argument with explicit four-step dependency chain C1–C4; expands §10 into a three-category catalogue of honest limits (A: Level-2 anchored prefactors; B: Level-3 engineering speculations; C: items deferred elsewhere).
- **v0.3** (May 2026, Stage 3). Lagrangian-derivation pass. Adds §5.4: explicit step-by-step derivation of the $\chi^{(3)} \sim \alpha^2$ scaling from the Lagrangian-level photon-photon-via-radion vertex $L = -(g_0/4)\, \Phi^2\, F_{\mu\nu} F^{\mu\nu}|_{node}$ of [6] §16, through radion integration-out at optical frequencies, to the four-loop running $g_0 \to \alpha$ of [6] §17. Establishes the $\chi^{(3)} \sim \alpha^2$ scaling as a Level-1 result derived in this paper, no longer requiring the reader to track the parent paper. Adds [25] Boyd nonlinear-optics reference.
- **v0.4a** (May 2026, Stage 4a). Structural-derivation pass for η_{node}. Adds §5.5: derives the structural form of the nodal-cone TPA-suppression factor η_{node} from the breathing-mode zero-point amplitude $\varepsilon_{ZPE} = m_H/(4v) = 0.1270$. Computes the geometric coefficient 8 (from the spherical-harmonic gradient at the cone) and the higher-multipole geometric ratios $|Y_\ell^0(\vartheta_{node})|^2/|Y_\ell^0(\pi/2)|^2$ (= 0, 784/729, 1024/2025, 30976/99225 for $\ell = 2, 4, 6, 8$) exactly. Converts §10 A1 from fully Level-2 anchored to Level-1 structural-form-with-deferred-prefactor. numerical bracket: $\eta_{node} \in [0.013, 0.24]$.
- **v0.4b** (May 2026, Stage 4b). Foam-level prefactor analysis. Adds §5.6: brackets the breathing-mode prefactor k_2 using dimensional saturation, the radial-decoupling argument, the closed-form gradient sum $\Sigma_m |\partial_\vartheta Y_{2,m}(\vartheta_{node})|^2 = 15/(4\pi)$, and the Casimir analogy of [6] §14bis.69 (central estimate $k_2 \approx 1/2$). combined bracket: $k_2 \in [0.016, 1.2]$, $\eta_{node}|_{breathing} \in [0.002, 0.15]$, total $\eta_{node} \sim 0.10–0.30$ in real devices. Identifies the foam-Lagrangian construction for the photon polarization-mode field as the residual specialist priority.
- **v1.0** (May 2026, Stage 5). Public-release polish. Consolidates the abstract from staged additions into a clean three-paragraph form suitable for Zenodo deposit metadata. Adds this Version history appendix. Final consistency sweep across cross-references and

terminology. No new structural content beyond v0.4b; v1.0 is the first version intended for archival deposit and external citation.

- **v1.1** (May 2026, Stage 6). Lagrangian-level refinement at the cone. Adds §5.7: shows that the first-order breathing contribution to *η_node* at *ϑ_node* vanishes geometrically because the *Y_{2,0}(ϑ_node) = 0* identity which defines the cone *also* eliminates the breathing's metric perturbation there. The cone is doubly protected. Refines the §5.6.5 Casimir analogy as misapplied to the dimensionless coordinate *ϑ* at the cone (the analogy applies to dimensionful observables like Casimir pressure). numerical bound: *η_node|_breathing,second-order ≲ 0.3%*. Total *η_node ∼ 0.10–0.30* prediction unchanged (it was already dominated by higher-multipole and medium contributions in the bracket); structural picture sharpened — the framework's vacuum is essentially *transparent* to *η_node* at the cone. remaining priorities are now the higher-multipole prefactors *k_ℓ* and the medium-specific *K_medium*; the breathing-mode question is closed at first order.
- **v1.2** (May 2026, Stage 7a). Higher-multipole closure. Adds §5.8: bounds the higher-multipole prefactors *k_ℓ* using the framework's identified spectrum from [6] (where *ℓ = 4* is the fermionic right-handed neutrino at seesaw scale, the trisection contributes only *m* = ±*3* modes not the *m = 0* tail, and ℓ ≥ 4 *m = 0* scalars are not specifically identified). Three scenarios analysed: A (TeV-scale Higgs-like, ruled out by LHC non-observation), B (seesaw-scale, contribution $\sim 10^{-23}$), C (TeV-scale gauge-singlet, contribution $\sim 10^{-6}$). Plausible scenarios B/C give *η_node|_higher* $\lesssim 10^{-5}$ — six orders of magnitude tighter than the bare-KK upper bound of *0.11*. Combined with Stage 6, total *η_node ≈ K_medium ≈ 0.05 − 0.20* in real devices, **dominated entirely by medium contributions**; framework vacuum and higher-multipole contributions both < *1%*. The §6.7 thermal MZI experiment now effectively measures *K_medium* directly. Residual specialist priority: medium-specific Stage 4b' photonics work; potential on the azimuthal trisection structure (separate from the cone TPA suppression).
- **v1.3** (May 2026, Stage 7c). Z_3 symmetry analysis at trisection vertices. Adds §5.9: extends the Stage 6/7a analysis from the cone (axisymmetric, *m = 0*) to the qutrit basis at the trisection vertices (Z_3*-symmetric*, *m* = ±*3*). The framework's Z_3 permutation symmetry — the same symmetry that organises SU(3), three-generation fermion structure, and tribimaximal PMNS mixing — protects qutrit coherence: Higgs (*m = 0*) and trisection octopole (*m* = ±*3*) are Z_3*-invariant* and act uniformly on the three qutrit states. Cosmic hot-axis dipole (*m* = ±*1*) is the leading Z_3*-breaking* coupling at order *ε_ZPE*, giving *η_qutrit|_framework ≈ ε_ZPE² × 0.5 ≈ 0.81%* (about 3× the qubit framework floor of *0.3%* but still ≪ medium). Total *η_qutrit ≈ 0.06 − 0.21* in real devices, similar to qubit *η_node ≈ 0.05 − 0.20* and similarly medium-dominated. The qutrit gains $\log_2 3 \approx 1.585$ bits per photon (*58.5%* increase) at essentially the same coherence cost — Proposal 4 (§6.4 trisection qutrit photonic computing) is structurally viable. Residual specialist priority remains medium-specific photonics work; the framework symmetry analysis is now complete for both *m = 0* (cone, qubit) and *m* = ±*3* (trisection vertices, qutrit) sectors.

- **v1.4** (May 2026, Stage 7b). Medium-specific anchored estimates. Adds Appendix B: anchored Level-2 estimates for the medium contribution *K_medium* across three candidate photonic platforms — silicon-on-insulator (*K_medium ≈ 0.10–0.20*), silicon nitride (*K_medium ≈ 0.03–0.08*, lowest), and thin-film lithium niobate (*K_medium ≈ 0.05–0.15*). Decomposes *K_medium* into birefringence dispersion, polarization-mode dispersion, scattering, phonon-jitter, and thermo-optic noise contributions. Predicts *η_node ≈ 0.05* ± *0.02* for silicon nitride waveguides — the optimal platform for nodal-cone qubits. Provides photonics specialists with concrete numerical targets to compare against the §6.7 thermal MZI experiment. Content is *anchored to standard photonics literature*, not framework-derived; it completes the engineering side of the paper without claiming additional (3+3) content. v1.4 is the first version with all of Stages 1–7c (framework derivation) plus Stage 7b (anchored engineering) — the paper is now self-contained with explicit numerical predictions at every level.
- **v1.5** (May 2026). Public-release polish. Title and keywords reworked to lead with QC-community-relevant terms ("Magic-Angle Photonic Quantum Computing"; geometric decoherence protection; the trisection qutrit; falsifiable predictions). Abstract, structure description, §10 honest-limits catalogue, §9 fault-tolerance argument, §11 conclusion, and section titles refactored from chronological-staging language to current-state framing — staged-development chronology is now confined to this Version history appendix. The "magic angle" connection to NMR magic-angle spinning (where the same *54.74°* and *Y_{2,0} = 0* physics removes dipolar broadening) is made explicit. No new structural content beyond v1.4; v1.5 is the cleaned public-release version intended for archival deposit and external citation.

Citation guidance. Citers of the present paper should reference the specific Zenodo version they are using, since the structural content of v0.1 differs substantially from v1.0. The structural derivation of *η_node* (§5.5–§5.6) appeared first in v0.4a/v0.4b; the $\chi^{(3)} \sim \alpha^2$ derivation (§5.4) first in v0.3; the closure of the breathing channel (§5.7) first in v1.1; the closure of the higher-multipole tail (§5.8) first in v1.2; the qutrit-Z_3 analysis (§5.9) first in v1.3; the platform-specific *K_medium* estimates (Appendix B) first in v1.4. The companion preprint [3] cites v1.0 specifically.

Appendix B: Medium-specific *K_medium* estimates for candidate photonic platforms

Disclaimer. This appendix provides anchored Level-2 estimates for the medium contribution *K_medium* to the nodal-cone TPA-suppression factor *η_node* in three candidate photonic platforms — silicon-on-insulator (SOI), silicon nitride (Si_3N_4), and thin-film lithium niobate (TFLN). The estimates are based on standard photonics literature and dimensional analysis for typical platform parameters; they are *not derived from the (3+3) framework* and should be refined by photonics specialists for actual device designs. The framework analysis of §5.5–§5.9 has bounded the framework-internal contributions to *η_node* at *< 1%*, leaving *K_medium* as the dominant residual. This appendix gives photonics engineers and the §6.7 thermal MZI experimental programme concrete numerical targets against which to test the framework prediction *η_node ≈ K_medium*.

B.1 Methodology

The medium contribution *K_medium* captures everything from the actual photonic medium (waveguide, fiber, or crystal) that causes the prepared polarization-state at *ϑ_node = 54.74°* to spread on the Poincaré sphere — equivalently, on the t_3 S^2. We decompose *K_medium* into physically-distinct contributions, estimate each for the three candidate platforms, and sum.

Each estimate is presented as a range that brackets typical platform values for telecom-band (1310–1550 nm) operation at room temperature. The ranges are intentionally conservative; specific device designs (e.g., particular waveguide cross-sections, particular wafer manufacturing maturity) can move the actual *K_medium* either lower (best-case engineering) or higher (manufacturing-limited).

B.2 Decomposition of K_medium contributions

The polarization-state spread *⟨(δϑ)²⟩|_medium = K_medium* at the cone is a sum of contributions from physically-distinct decoherence channels:

- **Birefringence dispersion** (*K_BD*): wavelength-dependent variation in *(n_e − n_o)* across the photon's spectral bandwidth, giving frequency-dependent rotation that spreads the polarization-state.
- **Polarization-mode dispersion** (*K_PMD*): random spatial variation of birefringence axes along the waveguide propagation direction, due to manufacturing tolerances.
- **Rayleigh / sidewall scattering** (*K_scat*): scattering of photons by density fluctuations (intrinsic) or sidewall roughness (extrinsic), which couples to polarization through asymmetric loss and depolarisation.
- **Phonon-jitter / acousto-optic noise** (*K_phonon*): thermal phonons modulating the local refractive index at GHz frequencies, with a contribution proportional to the elasto-optic coefficient and the temperature.

- **Thermo-optic noise** (*K_thermal*): slow temperature drift modulating birefringence, with a contribution proportional to $(dn/dT) \cdot \Delta T$ over the experiment timescale.
- **Two-photon absorption (residual)** (*K_TPA*): the actual TPA loss at the cone, which is not zero in real materials with finite bandgap.

Total: *K_medium = K_BD + K_PMD + K_scat + K_phonon + K_thermal + K_TPA*. The dominant terms vary by platform. The estimates below give a rough breakdown for each platform; specific waveguide designs may emphasise different terms.

B.3 Silicon-on-insulator (SOI)

Silicon-on-insulator at 1550 nm is the most mature photonic platform, with deep CMOS-compatible manufacturing and excellent active-device performance (modulators, detectors). Standard SOI parameters relevant to *K_medium*:

- **Birefringence dispersion**: silicon waveguides have strong geometric birefringence (asymmetric cross-sections give $\Delta n \approx 0.1–0.5$ between TE and TM modes) and significant stress-induced birefringence from the buried oxide. Wavelength dispersion across a 1-nm bandwidth photon: *K_BD ≈ 0.02–0.05*.
- **Polarization-mode dispersion**: typical SOI waveguides have PMD ≈ 0.1–1 ps/√km, contributing *K_PMD ≈ 0.01–0.03* over typical chip-scale propagation distances.
- **Rayleigh / sidewall scattering**: silicon waveguide sidewall roughness gives propagation loss of 0.5–2 dB/cm; the polarization-coupling fraction of this loss contributes *K_scat ≈ 0.03–0.08*.
- **Phonon-jitter**: silicon has elasto-optic coefficient $p_{12} \approx 0.07$; thermal phonons at GHz frequencies contribute *K_phonon ≈ 0.01–0.03*.
- **Thermo-optic noise**: silicon has $dn/dT \approx 1.86 \times 10^{-4} /K$ (high); for typical 100-mK temperature stability, *K_thermal ≈ 0.02–0.05*.
- **Residual TPA**: silicon has significant TPA at 1550 nm (*β_TPA ≈ 0.7 cm/GW*); at single-photon level this is small (*K_TPA < 0.01*) but at high pump powers it becomes a major loss channel.

Sum: *K_medium^SOI ≈ 0.10–0.20* (typical telecom-band SOI waveguide).

The dominant terms for SOI are birefringence dispersion and sidewall scattering; further reduction would require specialised processing (sidewall passivation, stress-balanced cross-sections).

B.4 Silicon nitride (Si_3N_4)

Stoichiometric silicon nitride on silicon dioxide is increasingly the preferred platform for low-loss photonic circuits, especially for passive or low-power applications. Standard Si_3N_4 parameters:

- **Birefringence dispersion**: lower than silicon ($\Delta n \approx 0.05–0.15$ for typical waveguide cross-sections); *K_BD ≈ 0.005–0.02*.

- **Polarization-mode dispersion**: *K_PMD ≈ 0.005–0.015*.
- **Rayleigh / sidewall scattering**: silicon nitride waveguides routinely achieve sub-0.1 dB/cm propagation loss (best-in-class is <*0.01* dB/cm); *K_scat ≈ 0.005–0.02*.
- **Phonon-jitter**: silicon nitride has lower elasto-optic coupling than silicon; *K_phonon ≈ 0.002–0.01*.
- **Thermo-optic noise**: $dn/dT \approx 2.5 \times 10^{-5}/K$ (much lower than silicon); for 100-mK stability, *K_thermal ≈ 0.003–0.01*.
- **Residual TPA**: silicon nitride is transparent at 1550 nm (bandgap ~5 eV, well above TPA threshold); *K_TPA < 0.001*.

Sum: *K_medium^SiN ≈ 0.03–0.08* (typical Si_3N_4 waveguide).

Silicon nitride is the **optimal platform for nodal-cone qubits** in the §6.7 thermal MZI experiment: its low *K_medium* gives the cleanest test of the framework prediction *η_node ≈ K_medium*. The dominant terms are birefringence dispersion and PMD; both can be further reduced by careful waveguide design.

B.5 Thin-film lithium niobate (TFLN / LNOI)

Thin-film lithium niobate on insulator combines silicon-photonics-like manufacturing scalability with strong electro-optic performance, making it attractive for active devices. Standard TFLN parameters:

- **Birefringence dispersion**: lithium niobate has strong intrinsic uniaxial birefringence ($\Delta n = n_e - n_o \approx -0.075$); for x-cut TFLN, this contributes a constant rotation (managed by waveguide design) plus dispersion. *K_BD ≈ 0.01–0.04*.
- **Polarization-mode dispersion**: *K_PMD ≈ 0.01–0.03*.
- **Scattering**: TFLN bulk and surface scattering are higher than Si_3N_4 but lower than SOI; typical 0.1–0.5 dB/cm loss; *K_scat ≈ 0.02–0.05*.
- **Phonon-jitter**: lithium niobate has substantial elasto-optic coupling (used for surface acoustic wave devices); *K_phonon ≈ 0.005–0.02*.
- **Thermo-optic noise**: $dn/dT \approx 4 \times 10^{-5}/K$; *K_thermal ≈ 0.003–0.01*.
- **Residual TPA**: lithium niobate is transparent at 1550 nm (bandgap ~4 eV); *K_TPA < 0.001*.

Sum: *K_medium^TFLN ≈ 0.05–0.15* (typical TFLN waveguide).

TFLN sits between SOI and Si_3N_4 for *K_medium*; its main advantage is the strong electro-optic response (good for fast switching), making it attractive for fault-tolerant gate operations even if *K_medium* is somewhat higher than Si_3N_4.

B.6 Cross-platform comparison and §6.7 predictions

Combining the platform-specific *K_medium* estimates with the framework bounds (§5.7–§5.8) (*η_node|_framework < 1%*):

η_node|_predicted ≈ K_medium (since framework contribution is < 1%)

Numerical predictions for the three platforms:

- **SOI**: *η_node ≈ 0.10–0.20* (mature platform, but high *K_medium*; useful for active devices)
- **Si_3N_4**: *η_node ≈ 0.03–0.08* (lowest *K_medium*; **optimal platform for the §6.7 thermal MZI test**)
- **TFLN**: *η_node ≈ 0.05–0.15* (intermediate; best for fault-tolerant gates needing fast EO switching)

§6.7 thermal MZI discrimination:

- Measured *η_node < 0.03* (Si_3N_4): framework strongly supported; medium imperfections smaller than estimated.
- Measured *η_node ≈ 0.03–0.08* (Si_3N_4): consistent with framework prediction + standard medium contributions.
- Measured *η_node ≈ 0.10–0.30* across all platforms: medium-dominated as predicted; framework bounds confirmed.
- Measured *η_node > 0.30*: either *K_medium* much higher than estimated (which would be testable independently via standard photonics characterisation), or framework prediction wrong (would be a falsification).
- Measured *η_node ≈ 0* (no observable TPA at the cone): would suggest *K_medium ≈ 0* somehow, requiring a deeper investigation of medium contributions; probably a measurement artefact.

Recommended experimental sequence: measure *η_node* in SiN first (cleanest test); confirm in TFLN (active-device platform, more relevant for fault-tolerance); then SOI (mature platform, most sensitive to medium imperfection). Cross-platform agreement on *η_node ≈ K_medium* with the predicted ordering Si_3N_4 < TFLN < SOI confirms the framework picture; significant disagreements would discriminate between framework prediction and medium-specific issues.

B.7 Caveats and specialist refinement targets

The estimates in this appendix are based on typical platform parameters from standard photonics literature; they should be refined by specialists for actual device designs. Specific refinements that would improve the predictions:

- **Specific waveguide cross-section design**: birefringence dispersion and PMD depend strongly on waveguide width, height, and aspect ratio.

- **Sidewall roughness measurement**: sidewall scattering can vary by 2× with manufacturing process; AFM characterisation of actual chips would tighten *K_scat*.
- **Thermo-optic stability**: actual lab temperature stability (typically < *100 mK* for stabilised setups) directly affects *K_thermal*.
- **Photon-bandwidth-dependent dispersion**: narrow-band photons (e.g., from frequency-stabilised lasers) reduce *K_BD*; broad-band photons (e.g., from spontaneous parametric down-conversion sources) increase it.
- **Wavelength selection**: telecom C-band (1530–1565 nm) is standard; longer wavelengths reduce silicon TPA (favouring SOI), shorter wavelengths increase scattering.
- **Cryogenic operation**: cooling to ~4 K reduces phonon-jitter and thermo-optic noise by ~100×, but adds engineering complexity.

Typical refinement might shift the platform-specific estimates by ±*50%*. The qualitative ordering (Si_3N_4 < TFLN < SOI) is robust across reasonable refinements; the quantitative numbers depend on device-specific work.

The result is therefore: ***platform-specific Level-2 K_medium estimates of 0.10–0.20 (SOI), 0.03–0.08 (Si_3N_4, optimal), 0.05–0.15 (TFLN), with silicon nitride identified as the cleanest platform for the §6.7 thermal MZI test of the framework prediction η_node ≈ K_medium. The estimates are anchored to standard photonics literature, not derived from (3+3); specialist refinement for specific device designs may shift the numbers by ±50% but should preserve the qualitative ordering Si_3N_4 < TFLN < SOI. The §6.7 experimental programme should target Si_3N_4 first.*** Combined with the framework derivations of §5.4–§5.9, the paper is now self-contained: from the photon-as-null-arc identification through the $\chi^{(3)} \sim \alpha^2$ scaling and the η_node structural form, through the breathing-channel and higher-multipole closures and the qutrit-Z_3 analysis, to platform-specific engineering targets.