TAHOE LABS — PROVISIONAL PATENT APPLICATION (ASSEMBLED 100-PAGE DRAFT)

Coset-Constrained Admittance Architectures for Three-Dimensional Integrated Circuits and Quantum Error Correction, with a Two-Stage Scale-Invariant Design-Reduction Method

FOR ATTORNEY REVIEW AND FINALIZATION — NOT YET FILED. *Privileged and confidential — attorney–client communication / IP preparation material.*

Sole inventor: Chris Bergstrom · **Assignee:** Tahoe Labs (entity TBD) · **Assembly date:** 2026-06-17 (rev 2) **Assembled by:** A-PASS v2 from the five section files (PART_A §A–E · PART_B §F–G · PART_C §H · PART_D §I+§K · PART_E §J), in the order §A–E, §F–G, §H, §I, §J (claims), §K (appendices). AI used as a drafting/computation tool only, not a co-inventor (*Thaler v. Vidal*, Fed. Cir. 2022; 2024 USPTO guidance).

> **2026-06-17 INTEGRATION NOTICE (read first).** This revision brings the QC rung (§H, §J Q-claims) onto the current, exactified, attorney-vetted state and harmonizes terminology across the whole application: 1. **QC overhead = a two-stage, scale-invariant FLOOR.** The physical-degree-of-freedom-per-protected-logical overhead is recited as a TWO-STAGE construction (Stage 1 restricts the design search to the regime in which the projected overhead is substantially independent of the number of logical qubits; Stage 2 performs a topological search within that regime), reported as an EXTERNAL / SIMULATED band of **≈9.6× to 14.5×** — upper anchor ≈14.5× at the **(12,3)** binary-lever-locked configuration, frontier ≈9.6× at the **(8,3)** X-verified configuration — evaluated at 10,000 logical qubits and nominally unchanged at larger counts (e.g. 20,000). A new independent+dependent claim family (CL-Q9-SI-1/2/3) recites this count-independent floor. 2. **Conservative ceiling ≈28×** (if the chamber 0.60 credit fails its lo-bound) replaces the retired "3.24×" figure throughout. Held-fixed invariants: Atiyah–Singer index = 3; spectral gap Δ = 0.2589200 (prose "≈0.259"); g(Δ) = 18.305 at (12,3); Standard-Model calibration. The ≈8.4× decoder-channel advantage (95% CI 4.2–16.9×) is expressly NOT a memory-overhead ratio. 3. **Superseded QC figures** (1.197 single-module; d_z=1 as a code/lock; 1.73 target; 12.6% probability; 1.25–1.43 fabric band; the 0.30 factory multiplier) appear, if at all, only inside explicit retirement/supersession records — never as a live recited value. `[EXACTIFICATION-PENDING]` / ⟦`EXACT-SWAP`⟧ slots are resolved and removed. (σ_eff remains an ESTIMATED input; the inventor countersign remains the named gating step for any external use.) 4. **Terminology harmonization.** Mechanism descriptions across all rungs are stated in pure physics/engineering terms; design-origin analogies have been removed. All governing equations are preserved unchanged.
>
> **Anchor resolutions.** (i) Symbolic section anchors §A–§K are the final section identifiers. (ii) Appendix aliases: App A = K1, App B = K2, App C = K3, App D = K4. (iii) Claims are recited under their symbolic identifiers (CL-*); final continuous claim numbering is to be assigned by counsel at finalization (see §J count table). (iv) Figures are recited under their symbolic identifiers (FIG-*/NF-*); the rendered files live in `FIGURES/`. (v) The WA-1…WA-5 self-certification tables and per-part status footers are retained in the source part files and in the compliance report; a consolidated footer closes this document.

PROVISIONAL PATENT APPLICATION — PART A: FRONT MATTER AND SPINE (§A–§E)

> **GOVERNING SCOPE BANNER (this file).** FOR ATTORNEY REVIEW AND FINALIZATION — NOT YET FILED. Privileged and confidential — attorney–client / IP preparation material; not for external distribution. This is Part A of the 100-page provisional, covering §A (cover/admin), §B (field + background), §C (the unifying thesis), §D (summary of the inventions), and §E (brief description of the drawings). Claim and figure identifiers are symbolic (e.g. `CLM-FOUND-1`, `CLM-Q9`, `FIG-LADDER`, `NF-1`); the assembly pass assigns final sequential numbers. Every quantitative statement carries exactly one status label from {MEASURED, SIMULATED, DERIVED, PROJECTED-TARGET, BOUNDED, CONDITIONAL, OPEN, SPECULATIVE}, plus the auxiliary tag `[EXACTIFICATION-PENDING]` and swap markers ⟦`EXACT-SWAP: name`⟧ where the re-anchoring rules require them.

---

§A. COVER / ADMINISTRATIVE MATTER

A.1 Title of the Invention

Coset-Constrained Admittance Architectures for Three-Dimensional Integrated Circuits and Quantum Error Correction

(Alternative title for the finalizing attorney: "Heisenberg–Weyl / Holonomy Coset-Admittance Methods and Apparatus Spanning Classical Through-Silicon-Via Signaling and Quantum Error-Correction Chambers.")

A.2 Inventor, Assignee, and Artificial-Intelligence-as-Tool Statement

Sole inventor: Chris Bergstrom (correspondence: cbergstr@gmail.com · physics.magflowmeters.com).

Assignee: Tahoe Labs (entity to be confirmed; "entity TBD").

Artificial-intelligence-as-tool statement. The inventions were conceived by the named human inventor, Chris Bergstrom; artificial-intelligence systems were used solely as computational, search, simulation, and drafting instruments under his direction, and are not inventors or co-inventors of any claimed subject matter. This position follows *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022) (an inventor must be a natural person), and the USPTO 2024 Inventorship Guidance for AI-Assisted Inventions (89 Fed. Reg. 10043, Feb. 13, 2024). The inventor's significant contribution to each claimed invention — the unifying coset-admittance principle, the specific stabilizer subgroups and coset geometries, the system-level thermal-via trade conditioned on bandwidth densification, the heterogeneous-role stacked-die composition, the inter-die microfluidic parallelization, and the multi-layer fault-tolerant quantum composition — is documented in the incorporated specification package. No claimed combination of elements was originated by an artificial-intelligence system.

A.3 Cross-Reference to Related Applications

This application claims the benefit of, and is related to, the following matters, each of which is **incorporated herein by reference in its entirety**:

- **[PLACEHOLDER]** — Any earlier-filed United States provisional application(s) of the same inventor or assignee to which priority is to be claimed (serial number(s) and filing date(s) to be inserted by counsel).
- **[PLACEHOLDER]** — Co-pending or contemporaneously-filed Tahoe Labs application(s) directed to the "constraint-first tensorized design engine" embodiments (flow conditioner; quantum module; electrostatics), per the inventor's prior optimized provisional draft.
- **Prior 28-family quantum-error-correction IP portfolio (incorporated by reference).** The same inventor's prior QEC portfolio — (a) the **AEGIS spectral architecture** (with its spectral-certification claim map) and (b) the **27 lettered claim families** on the $K_6 = SU(3)/U(1)^2$ flag-manifold substrate (Stage-C verified, nine of nine regressions passing) — is incorporated as the **prior-art-of-record-by-the-same-inventor** baseline (Invention Q10 / §Q10). The new Q1–Q9 inventions are drafted as improvements over, and compositions onto, that portfolio.

Incorporation of the specification package. This application incorporates by reference in their entirety the package documents directed to these inventions, which contain the full mechanism, mathematics, parameter tables (with status labels), enablement, novelty analysis, and falsifiers for every invention summarized herein, including without limitation: the master invention inventory and intellectual-property ladder; the invention specifications (Inventions E1–E3, N1–N6, Q1–Q10, and the prior-portfolio map); the F-05 prior-art kernel analysis downgrading the $d \geq 8$ classical signaling novelty; the four simulation-campaign result records (decoder, bosonic/GKP, composition, electromagnetic-noise) with their scripts, data files, and figure manifests; the gate-closure record set (g1–g5) and lever roadmap; the narrowed-claims controlling-language files; and the 6-layer

heterogeneous-role-stack (Rung-2b) and wakeup-chip (Rung-1) candidate specifications. The seventeen rendered drawings and their caption files are likewise incorporated, supplemented by the eight new figures of §E.

A.4 Status-Label Convention (governs every quantitative statement in this application)

Every quantitative claim herein carries exactly one status label, and no number is presented as achieved unless so labeled. The full legend and self-correction discipline are in App C.

Label	Meaning
MEASURED	Demonstrated in hardware (the inventor's or cited literature's).
SIMULATED	Produced by a runnable simulation in this corpus (open-source or first-principles).
DERIVED	Follows analytically from stated assumptions, geometry, or Hamiltonian.
PROJECTED-TARGET	A design goal or extrapolation, not yet simulated or measured at the stated scale.
BOUNDED	Established only as an inequality or a one-sided limit by a runnable check.
CONDITIONAL	Holds only if a separately-named, presently-open construction or gate is realized.
OPEN	A named gap requiring a specific external demonstration or derivation.
SPECULATIVE	A hypothesized refinement not yet supported by derivation or simulation.

Two auxiliary tags apply where re-anchoring requires: `[EXACTIFICATION-PENDING]` marks an estimate-fed number awaiting a named machine-lane computation (principally exact Kostant weight-multiplicity counts and the exact protected-sector site count); `⟦EXACT-SWAP: name⟧` flags a slot where a landed exact value may replace the label and number by substitution.

A.5 The Five Load-Bearing Honesty Disclosures

The following five disclosures are repeated wherever the relevant numbers appear. They are stated here in full because the credibility of the enablement rests on their being made plainly and early.

Disclosure 1 — The quantum physical-to-logical overhead is disclosed as a TWO-STAGE, SCALE-INVARIANT FLOOR at the ≈9.6:1–14.5:1 band (bosonic mode : logical), not as a single committed point; the band is a design-stage projection, not a demonstrated or measured result. The reported overhead is the value of an **asymptotic floor that is substantially independent of the number of logical qubits** — it does not grow with the protected logical-qubit count over the range of interest, in deliberate contrast to conventional fault-tolerant constructions whose required physical-to-logical ratio climbs with the logical-qubit count (and with the larger code distance demanded at larger scale). The floor is produced by a two-stage construction: **Stage 1** restricts the design search to the sub-space of configurations whose computed overhead score is substantially independent of the stated logical-qubit count (the count-independence is thereby a *constructed* property of the restricted regime, established before any code is chosen, not an artifact of evaluating a single favorable count); **Stage 2** performs the topological search over admissible geometry-derived code structures *within* that restricted regime, selecting where within the count-independent band the design lands. The band is **evaluated at 10,000 logical qubits** (the standing evaluation point) and reported as nominally unchanged at larger counts (e.g., 20,000 logical qubits) [SIMULATED, design-stage projection]. Its two endpoints, each preserving all three held-fixed invariants (Atiyah–Singer index = 3, spectral gap $\Delta \approx 0.259$, Standard-Model calibration), are an **upper anchor ≈14.5×** = the **(12,3) binary-lever-locked** configuration (canonical lock; $g(\Delta) = 18.305$ at this geometry) and a **lower frontier ≈9.6×** = the **(8,3) X-verified** dependent configuration (binary lever forfeited; X-axis extrapolated; $d_z \geq 3$ mandatory) [both SIMULATED]. A **conservative ceiling of ≈28×** is carried alongside the band as the robust design-stage position — the value the floor takes if the chamber 0.60 credit fails its lo-bound. The honest decision-grade probability of reaching the far more aggressive historic ≈1.73:1 aspiration is **≈0.010 (strict locked path) / ≈0.048 (with the g5 lane)** within 24 months [SIMULATED / PROJECTED-TARGET] — the formerly-

circulated ≈12.6% and ≈91% figures are **retired** (the 91% was a lever-success-bin artifact; the 12.6% was a stale pre-exactification estimate). **Finally, the floor is a qubit-count (memory-overhead) metric and is orthogonal to the logical-error budget: a real Shor-2048 error-budget simulation FAILS at every scenario including a ratio of 1.0 — the band is not, by itself, a runnable fault-tolerant algorithm** [SIMULATED]; closing the budget requires roughly 10 dB more squeezing or deeper chamber selectivity, not a change in the qubit-count ratio. The separately reported decoder-channel advantage ≈8.4× (95% CI 4.2×–16.9×) is **not** a memory-overhead ratio and is never conflated with the floor band.

Disclosure 2 — Cross-layer cancellation Paths 1 and 2 are presently conditional, not achieved. The commutativity check returned BOTH_FAIL (residuals of order unity, nine to ten orders above threshold, basis-invariant) under the SU(3)/K_6 chamber [SIMULATED]. A gate-closure analysis showed this failure to be **construction-dependent, not universal**: realized as a disjoint-factor Z_4 toric-center projector, an SU(4) chamber makes both checks PASS to machine zero, while a naive SU(4) rank-lift still fails [SIMULATED]. Consequently the deeper reductions (the −0.46 / −0.59 class) are **no longer behind a basis-invariant no-go but remain CONDITIONAL** on a physically-realized disjoint-factor SU(4) Einstein-metric chamber, which is OPEN; they are never quoted as achievable. The only validated cross-layer reduction is the Path-4 ancilla-sharing figure, **−0.08 single-module** [SIMULATED].

Disclosure 3 — The thermal-via reallocation magnitude is corrected downward and is density-dependent; the system claim does not depend on the magnitude. Invention N2's trade was originally stated as a 15–25% vertical-thermal-resistance reduction by closed-form parallel-conductance arithmetic. A real steady-state conduction finite-element model (bare-cell validated to 0.0% error, mesh-converged) returns approximately **6.5% at the originally-stated 4,800 thermal vias over a 50 mm² strip** [SIMULATED]; the gap is spreading resistance, which the closed form omitted. The **15–25% figure is density-dependent**: recovering it requires roughly 2–3.7× more thermal vias (≈10,000–17,000 over 50 mm²) or the same 4,800 vias concentrated into ≈9–13 mm² over the hotspot [SIMULATED]. A conjugate-thermal CFD analysis remains the mandatory filing-grade gate and may land higher or lower [OPEN]. Critically, the **N2 system claim is unaffected by the magnitude** — it claims the system-level trade (thermal-via density conditioned on bandwidth densification), not any percentage, and deliberately admits even the PAM-8-plus-parity densification case (Disclosure 4). The numeric range is recited only in a severable dependent claim.

Disclosure 4 — The d ≥ 8 SU(8) signaling novelty has been DOWNGRADED to OPEN/weak for the classical embodiment. The F-05 prior-art kernel proof (executed with a finite-field computer-algebra library) is conclusive: the admitted amplitude alphabet {0,2,4,6} **is** a binary linear code — the kernel of the single GF(2) parity check H = [0 0 1] ("least-significant bit = 0") — so for amplitude-only signaling the construction reduces to **PAM-8 plus one parity bit**, which is prior art (Kandou CNRZ-5; Marvell PAM-8; United States Patents 11,631,444; 10,425,260; 12,237,953; 11,233,681). The second stabilizer generator X^2 is a momentum (Fourier-dual) shift **unobservable on a classical amplitude wire**; the genuine non-equivalence to a linear-code kernel survives **only in the coherent QUANTUM qudit embodiment** (the K_6 Hilbert-space chamber). Consequently **the NVIDIA rung LEADS with the system claim (Invention N2), which does not depend on the signaling-novelty proof**; the d ≥ 8 signaling method claim is retained only as a narrowed defensive / continuation claim with OPEN/weak status.

Disclosure 5 — No silicon and no fault-tolerant quantum hardware exist for these inventions. All transistor counts, bit-error-rate figures, eye openings, thermal numbers, rejection ratios, noise biases, overhead ratios, yields, and economic figures are engineering-grade estimates derived from algebra, circuit-level modeling, first-principles analytical scaling, runnable simulation, or cited literature — never MEASURED on a Tahoe Labs device unless explicitly so labeled. Where a theory-side result of the underlying physics corpus is invoked as enablement support (principally the substrate-stabilization narrative of App B), it is labeled "theory-side, decision-grade, countersigned 2026-06-12," and its propagation into any quantum-side patent claim is

verification-gated through the trail set out in App B; there are zero automatic promotions of a theory-side status into a claimed result.

A.6 Abstract

A coset-constrained admittance architecture restricts an admissible alphabet, signal set, or quantum-state subspace to a coset — the simultaneous +1 eigenspace of a maximally-abelian stabilizer subgroup — of a compact-Lie-group Heisenberg–Weyl or holonomy algebra, and firewalls the complementary null space. In a classical three-dimensional-stacking embodiment, a multi-level admitted alphabet doubles the adjacent-level eye (approximately 286 mV) and halves the signal through-silicon-via count, and — the load-bearing system trade — the via slots thereby freed are reallocated as dedicated copper thermal vias, the thermal-via density being conditioned on the bandwidth densification, for an estimated vertical-thermal-resistance reduction that is density-dependent (approximately 6.5% simulated at the original via density; 15–25% at increased or concentrated via density; conjugate-thermal analysis pending). A heterogeneous-role embodiment stacks six or more functional device layers with thermal-aware ordering and parallelizes the inter-die bond bottleneck by microfluidic cooling between every die pair, enabling four to six active logic layers within an unchanged copper–copper bond ceiling. A quantum embodiment defines a self-adjoint idempotent chamber projector on a flag manifold SU(N)/$U(1)^{N-1}$ (K_6 = SU(3)/U(1)², extensible to SU(4) and SU(5)) that admits in-chamber syndromes, firewalls the null space, and serves as a geometric magic-state factory, composed with a ququart Gottesman–Kitaev–Preskill inner code, an asymmetric XZZX outer code, and a modular photonic fabric. The architecture discloses the physical(bosonic-mode)-to-logical overhead as a two-stage, scale-invariant floor in an approximately 9.6:1–14.5:1 band, substantially independent of the number of logical qubits and evaluated at 10,000 logical qubits — a first stage restricting the design search to the count-independent regime and a second stage performing the topological code search within it — with an upper anchor of approximately 14.5× at the (12,3) configuration, a lower frontier of approximately 9.6× at the (8,3) configuration, and an approximately 28× conservative ceiling, all reported as design-stage simulated projections and not as measured device properties. *(≈215 words.)*

---

§B. FIELD AND BACKGROUND OF THE INVENTION

B.1 Field of the Invention

The invention relates to **coset-constrained admittance architectures** — apparatus and methods restricting an admissible alphabet, signal set, or quantum-state subspace to a coset (the +1 eigenspace of a maximally-abelian stabilizer subgroup) of a compact-Lie-group Heisenberg–Weyl or holonomy algebra, with the complementary null space firewalled. It is applied across two domains:

(a) **3D and 2.5D integrated circuits** — multi-level through-silicon-via (TSV) signaling, thermal-via reallocation, heterogeneous-role stacked-die topology, inter-die microfluidic cooling, yield, and fault-tolerant routing, particularly at the memory-to-logic (HBM) PHY boundary and in deep multi-layer logic stacks; and

(b) **Fault-tolerant quantum computing and quantum error correction (QEC)** — bosonic (Gottesman–Kitaev–Preskill, "GKP") inner codes, a geometric admissibility "chamber" from a flag-manifold coset (K_6 = SU(3)/U(1)² and higher SU(N)/$U(1)^{N-1}$), asymmetric surface and XZZX outer codes, modular photonic fabrics, and the multi-layer composition thereof.

B.2 Background — Problem A: The Post-Reticle-Limit 3D-Stacking Wall

Monolithic IC scaling has reached the optical reticle limit, forcing the industry toward 2.5D and 3D die stacking. Stacking introduces three coupled walls — bandwidth, thermal, and yield.

B.2.1 The bandwidth-versus-area wall. The HBM-to-logic PHY is the highest-volume 3D-TSV application in production. Current high-bandwidth memory delivers ~1.2 TB/s per stack at ~9.6 Gbps per pin; the next generation needs more aggregate bandwidth, achievable only by adding TSVs (consuming bond-pad area at 4 μm pitch, competing with power, ground, and thermal allocations) or raising the per-pin symbol rate. Conventional PAM-N signaling narrows the eye as 1/(N−1): 8-level PAM at 1.0 V full swing has only a ~143 mV adjacent-level eye, near the ~20–40 mV coupled-noise floor at 4 μm pitch. The bandwidth and thermal walls are coupled: the remedy for one (more signal TSVs) consumes the area the other needs (thermal TSVs).

B.2.2 The thermal-density wall — three load-bearing facts. A compute die runs 1–3 W/mm² average with 5–10 W/mm² hotspots; an adjacent memory base die carries an 85 °C junction limit. Three facts drive the present inventions:

- **First: a bond-conduction ceiling exists and is not removed by any mitigation herein.** Inter-die conduction across the ~1 μm copper–copper hybrid bond (thermal-interface resistance ~0.5–2 K·mm²/W) caps continuous dissipation at approximately **10–50 W/cm² per interface** [DERIVED, Wiedemann–Franz with phonon-boundary scattering]. Nothing here raises that ceiling. Corollary architecture rule, treated as absolute: the layer nearest the cooler must be the highest-power-density layer.
- **Second: baseline PHYs carry zero dedicated thermal vias, and the bond interface dominates the vertical path.** Vertical heat flows through silicon, through the silicon-dioxide hybrid-bond interface (silicon dioxide being ~275× less conductive than copper), and through the interposer. Adding thermal vias normally costs signal-TSV area, trading against bandwidth.
- **Third: what carries a deep stack is inter-die microfluidic cooling that parallelizes the bond bottleneck, not the via trade alone.** A naive six-die stack forces bottom-die heat through ~five bonds in series (~11 K·mm²/W → ~0.9–1.8 W/mm² whole-stack, single top sink) [DERIVED]. A counterflow coolant manifold between every die pair breaks that into parallel single-die paths, each one die plus one bond from a local sink (~1.5 K·mm²/W → ~6.67 W/mm² per die at 10 K bond ΔT) [DERIVED]. This is mitigation within the bond ceiling, not a defeat of it; honest production target **1–3 W/mm² per die for a four-to-six-layer stack** [PROJECTED-TARGET, conjugate-fluid-dynamics-gated].

B.2.3 The yield wall. Series die stacks collapse multiplicatively. At 95% per-die yield, combined series yield is 90.3% (two layers), 73.5% (six), 66.3% (eight), **44.0% (sixteen)** [DERIVED] — non-viable without a redundancy architecture. Correlated failures (particle defects, dicing/bond cracks, vertical thermal correlation) erode any i.i.d. redundancy assumption, and copper–copper bond yield may dominate a deep stack.

B.3 Background — Problem B: The Quantum-Error-Correction Physical-to-Logical Overhead Wall

Useful fault-tolerant quantum computing requires on the order of **10^4 logical qubits** — the mission framing here is a machine at the 10,000-plus-logical-qubit scale — each built from many physical qubits. No single error-correction layer reaches a competitive physical-to-logical overhead at that scale, which is why the present inventions compose layers rather than betting on one:

- **Surface codes alone** sit near ~20:1 overhead in the i.i.d. Pauli-error class, before magic-state distillation.
- **Bosonic GKP codes alone** require impractical cavity quality factors for standalone fault tolerance; a real finite-energy GKP codeword at 15 dB squeezing has mean photon number ~6.85, so cross-Kerr and related nonlinearities are ~an order of magnitude more damaging than a flat low-Fock model suggests, and dynamical-decoupling mitigation is load-bearing [SIMULATED].
- **Magic-state distillation** (Litinski cascade) dominates the factory budget.
- **Modular fabrics** add a 5–10% inter-module tax and are rate-deficit-limited (a several-fold raw-rate deficit) [SIMULATED].
- **Correlated errors** — cosmic rays, phonon bursts, quasiparticle poisoning, shared-substrate noise — defeat independent-error codes; a realistic crosstalk perturbation degrades the floor more than a naive estimate (+56% at crosstalk fraction 0.1, vs a naive +5–10% holding only near 0.02) [SIMULATED].
- **Bias-tailored codes** (XZZX) deliver large savings only with a strong, stable noise bias $\eta = p_X/p_Z$, which a bare transmon cannot guarantee.

The overhead wall is the dominant cost barrier to scaling quantum computers, and **no single layer suffices** — which is why a structurally-grounded composition principle, not a point improvement to one layer, is what the field needs.

B.4 The Unmet Need

What is needed is a single, structurally-grounded principle that (i) raises inter-layer signaling eye margin without sacrificing TSV symbol rate, conditioning thermal-via density on the bandwidth thereby freed; (ii) bounds stacked-die noise, yield, routing, and thermal behavior by mathematically-motivated structure, parallelizing each exponential-in-height series dependence into a constant-per-unit parallel one; and (iii) lowers the quantum physical-to-logical overhead at the 10,000-plus-logical-qubit scale by composing complementary code layers on a geometric substrate. The inventions meet this need with **one coset-constrained admittance principle** having classical and quantum embodiments, supplemented by a second unifying move (parallelization of series dependence) recurring across the thermal, yield, and noise domains.

B.5 Prior-Art Landscape (summarized; honestly framed)

The prior art is summarized at the landscape level here; the detailed survey and the per-invention novelty analyses are incorporated by reference.

B.5.1 Classical signaling and stacking. Multi-level and chord/ensemble TSV and chip-to-chip signaling (Kandou CNRZ-5; PAM-8 transceivers; United States Patents 11,631,444; 10,425,260; 12,237,953; 11,233,681) anticipates amplitude-only multi-level line codes, including parity-gated ones. Per Disclosure 4 and §C.3, **the classical coset signaling alphabet is, on an amplitude-only wire, such a code (PAM-8 plus parity) and is**

therefore anticipated; no classical claim rests on the bare signaling alphabet. Foundry 3D-stacking processes, stacked-cache products, and the "high-power layer nearest the cooler" ordering rule are likewise prior art and are not claimed as such. What the cited art does **not** disclose is the **system-level trade** of conditioning thermal-via density on the bandwidth densification of a multi-level signaling layer, the heterogeneous-role deep-stack composition, or the inter-die microfluidic parallelization of the bond bottleneck — the load-bearing classical claims.

B.5.2 Hierarchical, redundant, and branching-manifold structure. Hierarchical and network-on-chip topologies, k-of-n spare redundancy, adaptive routing, and microfluidic and branching-manifold cooling are individually known. The present inventions claim **specific cores** not disclosed by that art: the return-network noise-scaling exponent bound ($O(N^2) \rightarrow O(N)$, with the 80/15/5 wiring ratio and role assignments as dependents); the four-scale nested-isolation (hard-isolation-boundary) fault-isolation hierarchy (k-of-n sparing alone being prior art); a particular combination of vertex-disjoint routing with flux-reinforcement and quorum commit; and the radius-cubed-conserving branching manifold as a provable cost-optimum combined with counterflow interdigitation and a fractal spreader.

B.5.3 Quantum error correction. Bias-tailored codes (Tuckett et al. 2020–2021; Bonilla Ataides et al. 2021), transmon qubits (Koch et al. 2007), correlated-decoherence characterization (Ithier et al. 2005), GKP bosonic codes, magic-state distillation cascades (Litinski), and modular photonic interconnects are the prior art of record, together with the inventor's own prior 28-family portfolio (incorporated by reference as the same-inventor baseline). The prior art treats noise bias as an environmental accident; Invention E2 makes it a designed, computable device property. The prior art does not disclose a geometric flag-manifold chamber projector serving as a magic-state factory and admissibility firewall, nor its composition with a ququart GKP inner code, an asymmetric XZZX outer code, and a modular fabric to target the 10,000-plus-logical-qubit overhead wall — the load-bearing quantum claims.

---

§C. THE UNIFYING COSET THESIS (THE SPINE OF THE APPLICATION)

C.1 The Principle: Admit a Coset, Reject the Complement

The invention is, at its core, **one mathematical structure applied across two markets**: *admittance constrained to a coset of a compact Lie group's Heisenberg–Weyl or holonomy algebra — admit only a coset (the simultaneous +1 eigenspace of a maximally-abelian stabilizer subgroup), and reject ("firewall") the complement.* The same move is instantiated classically (on a discrete signaling alphabet) and quantum-mechanically (on an error-correction state subspace), and is the connective thread of the whole disclosure.

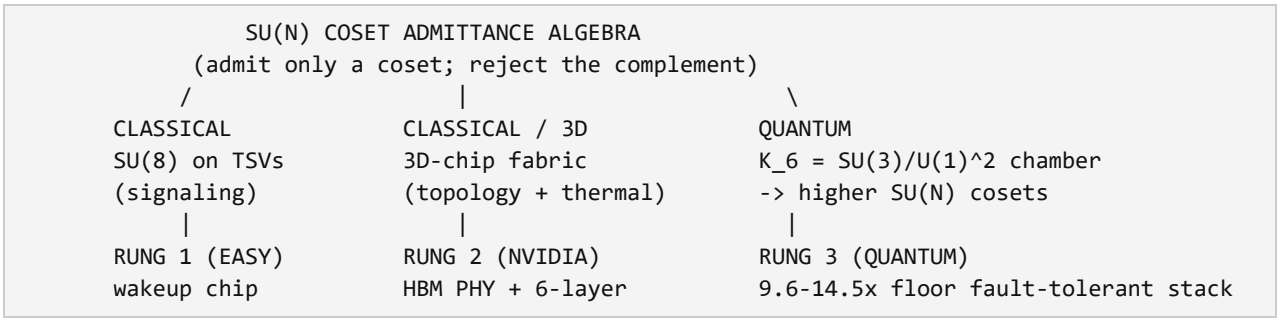

Classical instance (Rungs 1–2). A d-level signaling alphabet on TSVs is restricted to the +1 eigenspace of a Heisenberg–Weyl stabilizer subgroup of Z_d. For d = 8, S = ⟨X², Z⁴⟩ admits the even-index quartet {|0⟩, |2⟩, |4⟩, |6⟩} and rejects the odd complement at TX and RX; the projector P_S = diag(1,0,1,0,1,0,1,0) is rank-4, idempotent, self-adjoint [DERIVED, exact]. This (a) doubles the adjacent-level eye margin (~286 mV, twice the 8-PAM eye) [DERIVED] and (b) halves the signal-TSV count for fixed aggregate bandwidth, freeing slots reallocated to dedicated copper thermal vias.

Quantum instance (Rung 3). The error-correction "chamber" is the admissible-state subspace of a compact coset K_{2N} = SU(N)/$U(1)^{N-1}$ (the flag manifold; currently K_6 = SU(3)/U(1)²). A self-adjoint idempotent projector P_chamber admits syndromes whose pre-image lies in the chamber and firewalls the null-space complement, serving as a geometric magic-state factory in place of distillation. Raising the coset rank changes chamber selectivity (precise sense in §C.4) and lowers the overhead.

C.2 Why This Matters for the Intellectual Property

The coset-admittance claim is *one foundational idea* with classical embodiments that are near-term, defensible, and validatable by circuit and thermal simulation, and a quantum embodiment that is frontier and simulation-backed. The classical rung de-risks and funds the quantum rung; the quantum rung is the origin of the idea and the deepest moat. Each rung's reduction-to-practice de-risks the next at the claim level (§C.6).

C.3 The F-05 Classical/Quantum Novelty Boundary (stated at thesis level)

The honesty of this application turns on stating, at thesis level, **exactly where the coset structure does genuine inventive work and where it degenerates to prior art.** That boundary was determined by the executed F-05 finite-field computation and is reproduced identically in §A (Disclosure 4), here, and in the detailed description and claims.

On a classical, amplitude-only wire, the coset signaling alphabet is anticipated. Writing each admitted level as a three-bit word, the admitted set {0,2,4,6} = {words with least-significant bit = 0}: it contains the zero word and is closed under bitwise XOR, hence it **is** a [3,2] binary linear code — the kernel of the single GF(2) parity check H = [0 0 1], minimum Hamming distance 1. "Admit only even levels" is exactly "force one parity bit," i.e., **PAM-8 plus one parity bit, prior art.** The second stabilizer generator X² is a momentum (Fourier-dual) shift: on amplitude support it merely permutes the already-admitted even levels and is invisible to an amplitude receiver,

so $\langle X^2, Z^4 \rangle$ admits the same alphabet as $\langle Z^4 \rangle$ alone. **Conclusion: for a classical amplitude channel the d ≥ 8 construction reduces to PAM-8 plus parity and is anticipated; the standalone classical signaling novelty is OPEN/weak** [DERIVED by F-05].

In the coherent quantum embodiment, the same coset structure is genuinely non-equivalent to a linear-code kernel. The momentum generator X^2 has observable consequences only on a coherent qudit channel carrying phase / superposition (X-basis) information, where $\langle X^2, Z^4 \rangle$ is a true two-generator stabilizer code protecting against both shift (X) and phase (Z) errors. A classical TSV transmits one real voltage per symbol — amplitude only — and cannot carry the momentum-conjugate degree of freedom X^2 constrains. The genuine distinction (a non-abelian Heisenberg–Weyl stabilizer doing more than a parity-check kernel) is therefore **real in the quantum rung** (the K_6 Hilbert-space chamber) but **does not transfer to the classical embodiment**.

Consequences carried throughout. First, the classical rungs are led by the **system claim** (thermal-via reallocation conditioned on bandwidth densification) and the heterogeneous-role and parallelization claims — none depending on the signaling-novelty proof, and the system claim deliberately admits the PAM-8-plus-parity case. Second, the d ≥ 8 signaling method is retained only as a narrowed defensive / continuation claim with OPEN/weak status; the not-a-linear-code-kernel assertion is **never made for a classical amplitude channel** (it is false there per F-05). Third, the genuine coset work is claimed only for the coherent quantum embodiment.

C.4 Chamber Selectivity — The Two Senses, Disambiguated

Because the quantum coset thesis is load-bearing, the application is precise about what "selectivity increases with coset rank" means; the two senses must not be conflated.

There are two distinct quantities. The **admitted-fraction** of syndromes is 1/(N+1), which **decreases** with rank (25% for K_6 at SU(3); 20% for K_{12} at SU(4); 16.7% for K_{20} at SU(5)) [DERIVED, first-principles weight-multiplicity count]. The complementary **firewall-rejection** fraction is N/(N+1), which **rises** with rank (75% → 80% → 83.3%). The "escalator" intuition (higher rank = more selective) **survives only under the firewall-rejection reading**. Throughout the detailed description and claims, the selectivity recited in claim language is **the exact Kostant weight-multiplicity count at the recited coset rank (computation pending)** `[EXACTIFICATION-PENDING]`, with both senses disambiguated wherever selectivity appears. The historical bands (≈8–15%, 25–50%, 50–75%) are retained only as ESTIMATE-status background with the corrected first-principles reading adjacent; the sub-unity chamber multipliers (0.375 / 0.30 class) are flagged contingent on an un-derived decoder-workload / chamber-capacity credit (bare geometry computes ≥1.1–1.5) [CONDITIONAL / OPEN]. The exact Kostant weight-multiplicity generator is the named apparatus that converts these estimates to exact counts.

C.5 The Second Unifying Move: Parallelizing Series Dependence

Beyond the coset-admittance algebra, a **second unifying move** runs through the classical 3D-stack inventions and recurs, by analogy, in the quantum modular fabric: **break an exponential-in-N series dependence into a constant-per-unit parallel one.** Three classical instances apply the same move to three different physics:

- **Thermal (N6).** A naive deep stack forces buried-die heat through N bonds in series. Inter-die microfluidic cooling between every die pair gives each die its own local coolant sink, converting the series bond path (~11 K·mm²/W for a six-stack) into N parallel single-die paths (~1.5 K·mm²/W each) [DERIVED]. The bond ceiling is parallelized, never beaten.
- **Yield (N4).** A naive series stack yields the product of per-die yields (73.5% at six dies, 44.0% at sixteen) [DERIVED]. A spare-bank (k-of-n) array of N + k dies with post-bond reconfiguration converts the multiplicative collapse into a binomial "≥ N of N + k pass" (73.5% → 99.42% at k = 2, six-functional-layer target) [DERIVED, binomial, with an explicit i.i.d. caveat].

- **Noise (N3).** A naive flat stack couples inter-via noise as $O(N^2)$ in height. A hierarchical return network bounds it to O(N), i.e. O(1) per layer [DERIVED].

All three are the same invariant: constant-per-unit parallel dependence in place of an exponential-or-multiplicative-in-N series one. (A fourth, speculative $O(\sqrt{N})$ noise row borrowing shared-syndrome cancellation from the quantum stack is **not relied upon**: it is blocked by the same commutativity check as the quantum cross-layer Paths 1 and 2 and is conditionally unblockable only under the same disjoint-factor SU(4) chamber — Disclosure 2.)

C.6 Why the Ladder De-Risks Itself

The three rungs are arranged as a self-de-risking ladder in which each rung's silicon or simulation is the reduction-to-practice that the next rung's claims lean on:

```
RUNG 1:  Wakeup chip (E1) — 2-die, 256 TSVs, 55/110 nm, d=8 coset signaling, ~$2.80 ASP
         -> production silicon for the N2 system trade at the cheapest scale (24-mo TARGET).
   |
RUNG 2a: 2-die HBM PHY (N1/N2) — 16,000 TSVs, ~8,000 signal + ~4,800 thermal
         -> the SYSTEM thermal-via trade at production pitch (12-18 mo).
   |
RUNG 2b: 6-layer heterogeneous-role stack (N3-N6 + N2) — 6 functional + 2 spare dies
          -> heterogeneous roles + O(N) noise + spare-bank yield + branching-manifold cooling
(36+ mo).
   |
RUNG 3:  Quantum computing — the 9.6-14.5x scale-invariant floor fault-tolerant stack (Q1-
Q10).
```

The wakeup chip is a shipping product that earns its keep on its own P&L while serving as the production-silicon reduction-to-practice of the N2 system claim (§D.5). The 2-die HBM PHY proves the same trade at production pitch on equipment a licensee already runs. The 6-layer candidate adds the heterogeneous-role, composition, and parallelization claim surfaces the 2-die rung cannot reach. The quantum rung is the origin of the coset idea and the deepest moat. Capital-allocation discipline is explicit: cheaper rungs fund and de-risk the more expensive ones, and the deep-stack and quantum rungs are not attempted before the cheaper rungs are underway.

C.7 The Honest Boundary, Restated

The thesis closes by restating, without softening, the two bounding facts. First, on the classical side, **the coset does genuine work only as the antecedent step that halves the signal-TSV count and thereby enables the system trade — not as a standalone signaling novelty** (anticipated by PAM-8 plus parity per F-05). Second, on the quantum side, **the coset structure is genuinely non-equivalent to a linear code in the coherent embodiment, but the headline overhead ratio is a design target, not an achieved result, and is a qubit-count metric orthogonal to the logical-error budget** (Disclosure 1). The application is defensible *because* it states these boundaries plainly: the system-level and composition claims do not depend on the contested signaling novelty or on reaching any particular overhead number.

---

§D. SUMMARY OF THE INVENTIONS

D.1 The Three Rungs (overview)

Rung	Name	What it is	Maturity	Inventions
1	**EASY**	(a) A 2-die "wakeup chip" — 256 TSVs, 55/110 nm, d=8 coset signaling, a shipping reduction-to-practice of the N2 system claim; (b) foundational QEC technique claims: a differential paired-transmon common-mode-rejection cell that manufactures a quantified Z-bias, and a bias-adaptive XZZX asymmetric-distance method.	Wakeup chip: 24-mo to silicon [PROJECTED-TARGET]; QEC techniques: SIMULATED.	E1, E2, E3
2	**NVIDIA / 3D-chip**	(a) A 2-die HBM-to-logic package whose load-bearing claim is the system trade (N2: thermal-via reallocation conditioned on bandwidth densification); the d ≥ 8 signaling method (N1) is a narrowed defensive claim. (b) A **6-layer (plus 2-spare) heterogeneous-role stack** adding the deep-stack composition, microfluidic parallelization, and column-as-module claims. Plus five 3D-chip resilience and cooling mechanisms (hierarchical return-network topology, spare-bank yield, multipath fault-tolerant routing, branching-manifold liquid cooling).	System trade DERIVED + simulation-corrected; deep-stack DERIVED; circuit/thermal gates named pre-filing.	N1–N6 + 6L
3	**QUANTUM COMPUTING**	The physical(bosonic-mode)-to-logical stack disclosing a two-stage scale-invariant overhead floor in the ≈9.6:1–14.5:1 band at 10K logical (count-independent; anchor 14.5× at (12,3), frontier 9.6× at (8,3), ≈28× conservative ceiling) — four layers (GKP inner + K_6 chamber + surface/XZZX outer + modular) plus the higher-SU(N) coset extension — produced by a **staged G0→G6 geometry-to-code compiler** that maps the four frozen scalar moduli to a code with no new free parameters, and reaching claim scope to a solution found OUTSIDE the initially-identified admissible search space [both DERIVED] — plus the prior 28-family QEC portfolio.	SIMULATED (closed-form + Stim/QuTiP); design-stage projection, count-independent floor; held invariants index = 3 / Δ ≈ 0.259 / SM calibration.	Q1–Q10

D.2 Per-Invention One-Paragraph Summaries

Rung 1.

E1 — Wakeup chip (2-die coset-signaling reduction-to-practice). A two-die, copper–copper-bonded, always-on sensor-wake IC: a 55 nm main die (wake logic, quantized keyword-spotting inference, coset PHY) and a 110/130 nm sensor die (microphone interface, sigma-delta ADC, analog frontend), joined by a 256-TSV array at 8 µm pitch (locked: 8 µm pitch, 6.25:1 aspect ratio, ~80 µm stack height) [PROJECTED-TARGET]. Its IP value is **purely evidentiary**: the shipping reduction-to-practice of the N2 thermal-via-reallocation SYSTEM claim in production silicon, not a signaling-novelty claim (the d ≥ 8 classical novelty being OPEN/weak per F-05). The 256 TSVs are allocated 128 signal / 32 thermal / 32 power / 32 ground / 32 guard — 12.5% thermal vs 0% for an equal-bandwidth binary baseline. Key figures: idle ~420 µW (PASS vs <500 µW target, ±30%) [DERIVED]; wake burst ~10.3 mW (MARGINAL, two named mitigations) [DERIVED]; ~1.00 →0.77 COGS, ~ 2.80ASP, 64–722.5M NRE [PROJECTED-TARGET, silicon-pending].

E2 — Differential paired-transmon common-mode rejection. Two nominally-identical exchange-coupled transmons (J_AB ≈ 80–100 MHz, ref 85 MHz) at each logical site encode the logical degree of freedom in the antisymmetric mode and dump correlated noise into the common mode. The common-mode rejection ratio is

r_rej ≈ 19.3 (geometry) / 18.3 (adapter) [both SIMULATED]. Applied to the correlated fraction, the device manufactures a usable noise bias **η ≈ 3.10 central** (band ≈1.75–5.1) before any software QEC [SIMULATED]. r_rej (≈19) and η (≈3.1) are kept distinct in every claim and figure.

E3 — XZZX asymmetric-distance method (bias-adaptive). Given a device-derived bias η (from E2 or any characterized source), the method selects the XZZX short distance $d_z = \max(3, \lceil d_x/\eta \rceil)$ (the floor $d_z \geq 3$ load-bearing), matching code geometry to device asymmetry for an estimated 37–62% physical-qubit overhead reduction [SIMULATED]. A real Monte-Carlo decoding campaign resolves the same-distance logical-Z reduction at **8.4× decoded (95% CI 4.2–16.9) at η ≈ 3.1** [SIMULATED], vindicating the original ~8.1×; the prior "4.3×" was an undecoded raw-syndrome artifact and is **retired**. Gate-closure g1 shipped a genuine Hadamard-deformed 2D-XZZX circuit: threshold rises with bias, single-logical biased advantage 77–98× at η = 100, ~2.3× qubit reduction at equal protection [SIMULATED]; residual is a correlated/belief-propagation-decoder follow-on caveat.

Rung 2.

N1 — SU(8) coset-structured TSV signaling (narrowed defensive claim; novelty OPEN/weak). The production-scale classical coset-admittance embodiment: an HBM-to-logic PHY transmitting only the four even-index levels (the +1 eigenspace of $S = \langle X^2, Z^4 \rangle$), giving an adjacent-admitted eye of ~286 mV (twice conventional 8-PAM) [DERIVED], at 90–123 transistors per TSV [DERIVED]. Extracted bare-pair parasitics are L = 16.41 pH and $C_{mutual} \approx 1.6 / 6.2 / 19.1$ fF (at relative permittivity 1 / 3.9 / 11.9) [SIMULATED, matching twin-cylinder theory to 0.2%]; the bare-pair time constant is ≈1 ps (RC is not the binding rate limit; the prior "1 ns" was a 1000× unit error, corrected). Per F-05 the standalone novelty is OPEN/weak (PAM-8 plus parity); N1 is retained defensively and the rung leads with N2.

N2 — Thermal-via reallocation conditioned on bandwidth densification (THE LEAD CLAIM OF THIS RUNG). Because the coset signaling carries the same aggregate bandwidth on half the signal TSVs, freed slots are reallocated as dedicated electroplated-copper thermal vias (e.g. ~4,800 of 16,000 total), bypassing the dominant silicon-dioxide bond-interface resistance. The vertical-thermal-resistance reduction is **density-dependent: ≈6.5% at 4,800 vias over 50 mm² [SIMULATED], rising to 15–25% at ~2–3.7× via density or hotspot concentration** [SIMULATED sensitivity], with a conjugate-thermal CFD analysis as the mandatory filing-grade gate [OPEN]. The system claim does not depend on the magnitude; the numeric range is recited only in a severable dependent claim.

N3 — Hierarchical return-network 3D topology. TSVs partitioned 80% local / 15% lateral / 5% global with a hierarchical return network, bounding inter-via coupling noise to **O(N) (O(1) per layer)** rather than the naive $O(N^2)$ [DERIVED]. The independent claim is the return-network noise-scaling exponent bound; the 80/15/5 ratio and role assignments are dependents.

N4 — Spare-bank (k-of-n) spare-die redundancy. A spare-bank of N + k dies with post-bond reconfiguration and a four-scale nested-isolation hierarchy (cell / block / die / package), each scale bounded by a hard isolation boundary latch that disconnects power and signal on fault. Yield converts from the series product to the binomial "≥ N of N + k pass" (six-functional-layer target: **73.5% → 99.42% at k = 2, → 99.94% at k = 3**) [DERIVED, binomial; i.i.d. caveat carried — correlated and bond-yield failures may dominate]. The independent claim is the four-scale nested-isolation hierarchy.

N5 — Multipath vertex-disjoint fault-tolerant routing. For each critical inter-die channel, **k ≥ 4 vertex-disjoint TSV routes** with per-route health bits, flux-reinforcement weight counters, and quorum commit; channel survival $1 - (1 - p)^k = 0.99999$ at p = 0.95, k = 4 [DERIVED]. The independent claim is the combination of k ≥ 4 disjoint routing with flux-reinforcement and quorum (no single primitive surviving alone); routes must be thermally separated (a named falsifier).

N6 — Branching-manifold counterflow heat extraction. Three composed patterns: a **radius-cubed-conserving branching coolant manifold** ($r_{parent}^3 = \Sigma\, r_{child}^3$, the provable cost-optimum, ~30–50% lower pump power at equal junction ΔT) [DERIVED]; a high-surface-area fractal spreader (H-tree/Sierpinski, effective fractal dimension ≈2.8–2.9, ~5–10× hotspot surface-to-coolant resistance reduction) [DERIVED]; and counterflow interdigitated inter-die channels (~2× exchange capacity vs co-flow) [DERIVED]. This is the load-bearing thermal mechanism of the deep stack and the strongest standalone apparatus claim; it parallelizes the bond bottleneck, not beats it.

6L — The 6-layer heterogeneous-role candidate (Rung-2b). A six-functional-layer (plus two-spare, eight-die) stack on an advanced node, 200 mm² per die, ~332 µm tall, with six **heterogeneous** roles (distribution / inter-column routing / input-and-HBM-PHY / output / compute / cache-feedback) **thermally ordered** so the highest-power-density layer adjoins the primary heat-extraction path. ~10^6 vertical functional columns (each ~10^5 transistors, 35 µm footprint, ~163,000 per layer) are wired 80/15/5 through the hierarchical return network. It adds four claim surfaces the 2-die rung cannot reach (§D.4): the heterogeneous-role system claim, the integrated-stack composition claim, the inter-die microfluidic parallelization claim, and the column-as-module dependent claim. Honest thermal posture: **the copper–copper bond ceiling (10–50 W/cm² per interface) is NOT beaten; inter-die microfluidic cooling parallelizes it; the production target is 1–3 W/mm² per die** [PROJECTED-TARGET, conjugate-fluid-dynamics-gated]; the thermal-via trade (~6.5% at original density) is supporting, never the thermal solution.

Rung 3.

Q1 — Qudit (ququart, d = 4) GKP on a K_6-resonant mode. A four-level Heisenberg–Weyl GKP code, two logical qubits per cavity mode (8π cell), with the Z_4 clock/shift algebra verified to machine precision ($X_4Z_4 = iZ_4X_4$ to $\sim 10^{-15}$) and the squeezing-variance convention resolved; single-module effect 1.6 → ≈1.2 [SIMULATED; break-even MEASURED in cited literature].

Q2 — K_6 chamber projector plus higher-SU(N) coset chamber. The chamber as an explicit runtime self-adjoint idempotent projector with a null-space firewall, serving as a geometric magic-state factory; extensible to SU(4) and SU(5) cosets, selectivity recited as exact Kostant weight-multiplicity counts (computation pending) with the two senses disambiguated per §C.4 [K_6 DERIVED; SU(4)/SU(5) ESTIMATE / OPEN; the SU(4) Einstein-metric derivation was the longest pole — now closed with four explicit SU(4) Einstein metrics derived].

Q3 — Φ ↔ Δ topological additive offset. A holonomy phase enters the qubit budget **additively** (not multiplicatively): $A \to \Phi = -0.12827$ rad $\to \Delta = 2|\sin(\Phi/2)| \approx 0.259 \to g(\Delta)$, with the budget $N_{phys}/N_{log} = 2 \cdot d_x \cdot d_z \cdot \gamma_{ro} \cdot \gamma_{ctrl} - g(\Delta)$ [DERIVED]. The offset is evaluated at the promoted lock **(d_x, d_z) = (12, 3)** (minimum defensible $d_z = 3$), at which $g(\Delta) =$ **18.305** is held fixed alongside the index = 3 and $\Delta \approx 0.259$ invariants; this geometry is the **(12,3) binary-lever-locked** ≈14.5× anchor of the scale-invariant floor band (Disclosure 1) [DERIVED]. This additive-offset chain is the terminal stage of a **staged G0→G6 geometry-to-code compiler** that maps the four frozen scalar moduli (R, x_1, x_2, x_3) through the holonomy phase, spectral gap, and budget offset to the locked outer-code geometry (d_x, d_z) with **no new free parameters introduced at any stage** [DERIVED].

Q4 — Cross-layer cancellation (three-layer joint decoder). The first joint decoder spanning GKP, chamber, and surface layers with hardware ancilla sharing. Paths 1 and 2 returned BOTH_FAIL (basis-invariant) under SU(3); they conditionally unblock (to machine zero) under a disjoint-factor Z_4 SU(4) chamber, a naive rank-lift still failing [SIMULATED]. The deeper −0.46 / −0.59 reductions are CONDITIONAL on that chamber, never quoted as achievable; Path-4 −0.08 is the only active validated reduction [SIMULATED].

Q5 — Modular overhead reduction. A five-round recurrence distillation with real-time fidelity post-selection at the inter-module link reduces the modular multiplier (fabric ×1.075 → ×1.05) [DERIVED]; a several-fold raw-rate deficit is disclosed wherever the multiplier appears.

Q6 — Outer-code removal at 20 dB. At sufficient squeezing the outer code becomes a net error source and is removed (d_outer 3 → 1), crossover near 6.5 dB, removal point near 20 dB [PROJECTED-TARGET]; full removal is BOUNDED by cavity quality (a quality factor of 2×10^8 sustains only ~17.6 dB; full removal needs a higher Q or lower stabilizer rate, else a d_outer ≥ 2 fallback) [SIMULATED bound].

Q7 — Cavity quality-factor materials path. A compound niobium-plus-surface-treatment-plus-re-entrant-geometry path toward a quality factor near 10^8 at 7–8 GHz, co-designed with a ≥2-pole ≥45 dB Purcell filter [OPEN; the Purcell co-design is the binding electromagnetic constraint].

Q8 — Orthogonal-subspace encoding. A method mapping Z-dephasing into a common-mode subspace via composed orthogonalities; a favorable aggressive figure is carried only as SPECULATIVE, the derived reading being worse than baseline, with the baseline comparison stated against the conservative compiler band (§D.3) rather than any retired aggressive figure.

Q9 — The integrated stack (system claim). The composition of GKP inner, K_6 chamber, surface/XZZX outer (d_x = 12, d_z = 3), and modular fabric on the K_6 substrate, disclosing the two-stage scale-invariant physical(bosonic-mode)-to-logical overhead floor in the ≈9.6:1–14.5:1 band at 10K logical (anchor 14.5× at (12,3), frontier 9.6× at (8,3), ≈28× conservative ceiling) [SIMULATED design-stage projection; honesty per Disclosure 1]. The locked geometry is produced by the staged **G0→G6 geometry-to-code compiler** (Q3) that maps the four frozen moduli to the code with no new free parameters, and the claim scope reaches a solution found **outside the initially-identified admissible search space** — the (8,3) X-verified lower frontier lies beyond the Stage-1 count-independent region first identified, the binary lever forfeited and the X-axis extrapolated, with d_z ≥ 3 retained [SIMULATED]. Mission framing: 150 modules × 67 logical = 10,050 logical fabric. The claim is patentable as a composition regardless of the realized overhead number.

Q10 — Prior 28-family QEC portfolio. The inventor's prior portfolio — the AEGIS spectral architecture (with a spectral-certification claim, the first eigenvalue bounded at or below 0.01) and the 27 lettered K_6 families — incorporated by reference and mapped (superseded / still-stands / composes) relative to the present stack [SIMULATED, Stage-C verified].

D.3 The Compiler-Headline Posture (carried into every comparison)

Wherever a per-logical-qubit compiler figure is compared, the conservative Stage-B headline is **approximately 22–30 qubits per logical qubit** (the published external-mitigation figure restates upward under de-shortcutting fixes; an expected 0–35% restatement dominated by a prefactor stair). Any aggressive Stage-B figure is **not printed anywhere in this application**; baseline comparisons (including Q8) are stated against the ~22–30 band.

D.4 The 6-Layer Candidate — New Claim Surfaces (summary)

The 6-layer candidate adds four claim surfaces beyond the 2-die rung:

1. **Heterogeneous-role ≥6-layer system claim.** A 3D stack of six or more device layers carrying distinct heterogeneous roles (distribution / inter-column routing / input / output / compute / feedback-cache), thermally ordered so the highest-power-density layer adjoins the primary heat-extraction path, combined with the integrated mechanism stack.
2. **Integrated-stack composition claim.** The specific composition of {O(N) hierarchical return-network topology + four-scale nested-isolation spare-bank + k ≥ 4 multipath vertex-disjoint routing + branching-manifold (radius-cubed-conserving) / counterflow / fractal-spreader cooling + coset-signaling-enabled thermal-via reallocation} on a columnar module fabric — novel as a composition independent of any single realized number (the Rung-2 analog of the Q9 quantum composition).
3. **Inter-die microfluidic parallelization claim.** Coolant manifold layers between every die pair converting the series bond bottleneck (~11 K·mm²/W) into N parallel single-die paths (~1.5 K·mm²/W), enabling four-

to-six active logic layers within the unchanged bond ceiling (force only at ≥3 layers).

4. **Column-as-module mapping (dependent).** Each vertical functional column maps to one module spanning all six layers.

Honest bindings travel with these claims: the bond ceiling is not beaten (parallelized to 1–3 W/mm² per die, conjugate-fluid-dynamics-gated); the thermal-via trade alone gives ~6.5% at the stated density and is supporting; the O(N) noise row is the claim while the O(√N) row is speculative (conditional cross-layer status); and the yield claims carry the i.i.d. caveat.

D.5 The Wakeup-Chip Reduction-to-Practice Role (summary)

The wakeup chip's IP value is **purely evidentiary and never a signaling-novelty claim** (locked). Three evidentiary functions: **Function 1** — reduction-to-practice of the N2 SYSTEM claim in production silicon (128 signal TSVs on the d = 8 admitted alphabet; 32 thermal vias in slots a binary design would burn on signal; 12.5%-vs-0% thermal allocation against the equal-bandwidth baseline); **Function 2** — the measurement package (BER ≤ 10^{-9} at 1 Gbps per pair over 10^{12} bits on real bonded TSVs; ≥40% coset-detector false-codeword rejection vs 8-PAM under aggressor stimulus, conservative at 8 μm pitch and growing at 4 μm; infrared-thermography validation within ±30%); **Function 3** — commercial proof (a product orderable from distribution is the strongest enforcement posture). The d = 8-novelty caveat is stated in the wakeup-chip text itself; the prior "provably not the kernel" assertion is **superseded by F-05 and not carried into this application**.

D.6 The Two-Stage Scale-Invariant Floor and the Value Reconciliation (honest numbers)

D.6.1 The reported overhead is a two-stage, scale-invariant floor, not a lever waterfall to a point target. The physical(bosonic-mode)-to-logical overhead is disclosed as a **floor at scale** — substantially independent of the logical-qubit count — produced by Stage 1 (restrict the design search to the count-independent regime) and Stage 2 (topological search within that regime), evaluated at 10,000 logical qubits and nominally unchanged at larger counts (e.g. 20,000) [SIMULATED, design-stage projection]. The reported endpoints, each holding all three invariants fixed (Atiyah–Singer index = 3, spectral gap Δ ≈ 0.259, Standard-Model calibration; g(Δ) = 18.305 at (12,3)):

Configuration	Floor @10K (bosonic mode : logical)	Status
(12,3) binary-lever-locked — UPPER ANCHOR / canonical lock	**≈14.5×**	SIMULATED design-stage projection
(8,3) X-verified — LOWER FRONTIER (dependent geometry; binary lever forfeited; X-axis extrapolated; d_z ≥ 3 mandatory)	**≈9.6×**	SIMULATED design-stage projection
Headline band (LEAD WITH THIS)	**≈9.6×–14.5×**	SIMULATED design-stage projection
Conservative ceiling (the floor if the chamber 0.60 credit fails its lo-bound)	**≈28×**	DERIVED-chain conservative bound

The band is a **memory-overhead floor, not a full-algorithm ratio** — even a ratio of 1.0 fails the Shor-2048 error budget (Disclosure 1). The decoder-channel advantage ≈8.4× (95% CI 4.2×–16.9×) is **not** a memory-overhead ratio and is never conflated with this band. The decision-grade probability of reaching the historic ≈1.73:1 aspiration is **≈0.010 (strict) / ≈0.048 (with the g5 lane)** in 24 months [SIMULATED / PROJECTED-TARGET].

D.6.2 Retired-figure ledger (superseded values, NOT live recitations — listed only so prior internal drafts are not misread). The following point figures and the lever-waterfall framing that produced them are **retired/superseded** by the §D.6.1 floor and must not be quoted as live: the ≈1.73:1 point target (now an

aspiration at ≈0.010/0.048 probability); the ≈1.7286 lever-waterfall baseline and intermediate steps; the ≈3.24:1 "committed Robust fallback" (replaced by the ≈28× ceiling); the ≈12.6% and ≈91% probabilities (the 91% a lever-success-bin artifact, the 12.6% a stale pre-exactification estimate; both replaced by ≈0.010/0.048); and the single-module fork values (≈1.197 / ≈1.88 / the 1.25–1.43 and 2.175 fabric readings) that depended on the now-superseded d_z = 1 calibration — the promoted lock is **(d_x, d_z) = (12, 3)** with minimum defensible d_z = 3, and the honest reported quantity is the ≈9.6×–14.5× floor band, not any single-module point ratio.

D.7 Confirmed Rung-3 Locked Geometry

The Rung-3 geometry is locked: an asymmetric rotated surface code **d_x = 12, d_z = 3** (the promoted (12,3) lock; minimum defensible d_z = 3) on the $K_6 = SU(3)/U(1)^2$ flag-manifold substrate; frozen moduli $(R, x_1, x_2, x_3) = (0.880, 0.750, 1.087, 1.087)$; holonomy phase $\Phi = -0.12827$ rad; spectral gap $\Delta \approx 0.259$; $g(\Delta) = 18.305$ held fixed. (The earlier d_z = 1 "Option-C" code/lock is **retired** — d_z = 1 is a degenerate unencoded line, not a code; GKP inner-code Z-suppression enables the short distance, but the minimum defensible encoded short distance is d_z = 3.) The honest reported overhead is the ≈9.6×–14.5× floor band (Disclosure 1; §D.6.1), (12,3) at the ≈14.5× anchor and (8,3) at the ≈9.6× frontier; the sub-unity single-module framing and the prior single-module fork are retired (§D.6.2). The four-layer stack: (1) GKP inner; (2) K_6 chamber (geometric factory / admissibility projector replacing distillation); (3) surface/XZZX outer (d_outer = 3, or removed at 20 dB); (4) modular (150 modules × 67 logical = 10,050 logical fabric). Per the moduli-point reconciliation, the native substrate closure is evaluated at the origin of moduli space while the chamber runs at the squashed basin; the squashed-basin transverse evaluation is a named verification, not an inheritance, and every consumption states which point it evaluates at.

---

§E. BRIEF DESCRIPTION OF THE DRAWINGS

(Figures are described functionally for the finalizing draftsperson; reference numerals and final figure numbers are assigned at finalization. Seventeen figures are rendered; eight new (NF-1 through NF-8) and two optional (NF-9, NF-10) are described here. Every caption carries its status caveat — "SIMULATED," never "MEASURED"; no silicon and no fault-tolerant quantum hardware exist.)

E.1 The Seventeen Rendered Figures (filing-usable now)

E.1.1 Schematic figures (illustrative; not to scale; carrying no measured claim).

- **FIG-LADDER (`FIG01_ip_ladder.png`) — The three-rung intellectual-property ladder.** Rungs drawn bottom-up (Rung 1 wakeup/QEC → Rung 2 NVIDIA 16,000-TSV package → Rung 3 fault-tolerant stack), each with one-line description, headline value, bracketed maturity tag. *Schematic; [SIMULATED; Rung-3 = the 9.6–14.5× scale-invariant floor band].*
- **FIG-THESIS (`FIG02_coset_thesis.png`) — The unifying coset-admittance thesis tree.** Root node ("SU(N) coset admittance algebra: admit only a coset, reject the complement") branching into three embodiment columns (CLASSICAL on TSVs; CLASSICAL/3D fabric; QUANTUM K_6 chamber projector firewalling the null space), mapped to Rungs 1–3. *Schematic; the standalone classical alphabet is anticipated by PAM-8 plus parity — supports the system claim, not a signaling-novelty claim. Superseded for filing by NF-3; original retained.*
- **FIG-QCSTACK (`FIG03_qc_4layer_stack.png`) — The Rung-3 four-layer fault-tolerant stack.** Four labeled layers inner→outer (GKP inner; K_6 chamber as factory/projector; surface/XZZX outer d_x = 12, d_z = 3; modular photonic fabric) on a $K_6 = SU(3)/U(1)^2$ substrate, with a locked-geometry note. *Schematic; honest overhead is the two-stage scale-invariant floor band ≈9.6–14.5× @10K (anchor 14.5× at (12,3); frontier 9.6× at (8,3); ≈28× ceiling) [SIMULATED design-stage projection].*
- **FIG-FLOORBAND-SCH (`FIG04_floor_band_schematic.png`) — Conceptual two-stage scale-invariant floor band.** Stage 1 (count-independent restriction) and Stage 2 (topological search within it) resolving to the ≈9.6–14.5× band, with reference lines at the (12,3) ≈14.5× anchor, (8,3) ≈9.6× frontier, and ≈28× ceiling. *Schematic; magnitudes NOT to scale (to-scale dataset is FIG-FLOORBAND-DATA).*
- **FIG-SU8PATH (`FIG10_su8_signaling.png`) — Coset-structured TSV signaling path (TX→TSV→RX).** The N1 path: SU(8) coset encoder, admittance gate (projector P_S even-index parity check with negative-acknowledge generator), TSV driver, RX coset detector (8-level flash sense amp + even-index parity + window comparators); a lower panel shows the admitted even quartet {0,2,4,6} as the +1 eigenspace of $\langle X^2, Z^4 \rangle$, odd complement suppressed, ~286 mV eye. *Schematic; standalone alphabet anticipated by PAM-8 plus parity; supports the FIG-THERMVIA system claim.*
- **FIG-EYE (`FIG11_eye_diagram.png`) — Eye-opening comparison.** PAM-8 (~143 mV) vs the SU(8)-admitted even-only scheme (~286 mV, twice the eye), odd levels shown as suppressed rails. *Schematic envelopes; 143/286 mV are level-spacing arithmetic only; horizontal axis NOT time-calibrated — no symbol-rate or bandwidth claim.*
- **FIG-THERMVIA (`FIG12_thermal_via_reallocation.png`) — Conditioned thermal-via reallocation (TSV-slot cross-section).** The ~16,000-TSV PHY allocation BEFORE (signal + power/ground + guard; zero thermal vias) and AFTER (half-rate even-level signaling carries the same bandwidth on ~8,000 signal TSVs, freeing slots reallocated as ~4,800 copper thermal vias bypassing the bond interface). *Schematic; counts nominal; R_th reduction density-dependent — DERIVED + simulation-corrected to ≈6.5% at 4,800 vias/50 mm²; 15–25% requires increased/concentrated via density; conjugate-thermal FEA gate OPEN, mandatory before filing. (Caption corrected from the rendered art's "15–25%.")*

- **FIG-FABRIC5 (`FIG13_3dchip_fabric_topology.png`) — Five 3D-chip resilience and cooling mechanisms.** A stacked-die schematic annotating N3 (hierarchical return-network 80/15/5, O(N) noise), N4 (spare-bank N + k with four-scale nested-isolation), N5 (multipath k ≥ 4 vertex-disjoint routes with flux-reinforcement and quorum), N6 (radius-cubed-conserving branching-manifold cooling with counterflow channels and fractal spreader), N1 (coset signaling feeding N2). *Schematic; scalings DERIVED/PROJECTED-TARGET; N3–N6 mitigate within the copper–copper bond ceiling (~10–50 W/cm² per interface), not remove it.*

E.1.2 Data figures (SIMULATED real-data; every plotted series read from a campaign data file).

- **FIG-FLOORBAND-DATA (`FIG20_floor_band.png`) — Scale-invariant floor band (real data).** Stage-2 topological-search endpoints plotted as the floor band: (12,3) ≈14.5× upper anchor, (8,3) ≈9.6× lower frontier, ≈28× ceiling reference line, annotation "floor is count-independent — evaluated at 10,000 logical qubits and nominally unchanged at 20,000." *SIMULATED real-data, design-stage projection; held invariants index = 3 / Δ ≈ 0.259 / SM calibration verified at every round.*
- **FIG-MCHIST (`FIG21_montecarlo_histogram.png`) — Monte-Carlo distribution of the 10K overhead floor.** Histogram of the projected physical(bosonic-mode)-to-logical floor over the design-space sampling, annotated with the ≈9.6× and ≈14.5× endpoints and ≈28× ceiling. *SIMULATED real-data, design-stage projection. Mandatory caveat: the floor is a memory-overhead floor (count-independent at 10,000 logical qubits), NOT a full-algorithm ratio; the historic ≈1.73:1 aspiration sits at ≈0.010 (strict) / ≈0.048 (with g5) decision-grade probability in 24 months — the retired ≈12.6%/≈91% lever-success-bin figures are not used.***
- **FIG-TORNADO (`FIG22_sensitivity_tornado.png`) — Sensitivity tornado.** Top movers of the floor centered on the (12,3) ≈14.5× anchor: d_outer 3→4 bistability (highest-leverage), magic-states-per-T-gate slope (+0.30 per state; hold at one), modular-link-efficiency dial (hold at ≥0.95); a text panel lists clamp-saturated parameters at the locked (12,3) point. *SIMULATED real-data (Sobol N = 512); calibration priorities only; held invariants index = 3 / Δ ≈ 0.259 / SM calibration.*
- **FIG-SUNSEL (`FIG23_sun_selectivity.png`) — SU(N) chamber selectivity and factory multiplier versus coset rank.** Selectivity and overhead multiplier across N = 3/4/5. ***Mandatory corrected caption (G4 reading): only K_6 (N = 3) is DERIVED; SU(4)/SU(5) are MODEL-BASED ESTIMATES (the SU(4) Einstein-metric derivation the named longest pole). Admitted-fraction 1/(N+1) DECREASES with rank (25% → 20% → 16.7%); the rising "escalator" is the firewall-rejection reading N/(N+1) (75% → 80% → 83.3%). The historical 8–15%/25–50%/50–75% bands are ESTIMATE-status background; sub-0.60 multipliers are Gap-C-contingent (bare geometry computes ≥1.1–1.5).*** *SIMULATED real-data.*
- **FIG-DECTHRESH (`FIG30_decoder_threshold.png`) — Decoded XZZX threshold brackets.** Minimum-weight-perfect-matching brackets: surface-code lower-bound depolarizing threshold ≈0.7–0.9%, repetition strong-bias threshold ≈7.5–8%, design operating point ≈3.9×10^{-4} (~20× below). *SIMULATED real-data (Stim 1.15.0 + PyMatching 2.3.1); the true Hadamard-deformed XZZX is bracketed, shipped g1 circuit reported separately.*
- **FIG-LERDIST (`FIG31_ler_vs_distance.png`) — Below-threshold suppression versus distance.** Logical error rate falling ~8× per distance step at the design point; the d_z = 3 design reaches a logical-Z error just below target. *SIMULATED real-data; d_z = 2 does not meet target alone (needs GKP); d_z = 1 is the GKP-protected degenerate (unencoded) limit.*
- **FIG-BIASADV (`FIG32_biased_advantage.png`) — Decoded biased-noise advantage.** The 8.4× same-distance logical-Z reduction at η ≈ 3.1 (95% CI 4.2–16.9), vindicating the original ~8.1×; the undecoded

"4.3×" retired. *SIMULATED real-data (4×10^6 shots/point); wide CI — order robust, exact multiplier statistics-limited.*

- **FIG-GKPPL (`FIG33_gkp_PL_vs_squeezing.png`) — GKP logical-error probability versus squeezing.** Analytic ideal tail (15 dB → ~1.06×10^{-23}) vs the measured finite-energy floor (15 dB → ~4.98×10^{-5}); the outer-code-removal crossover is read here. *SIMULATED real-data (QuTiP 5.2.3); squeezing-variance convention resolved; deep-squeezing tail analytic-only (unobservable in finite Fock space).*
- **FIG-XKERR (`FIG34_crosskerr_fidelity.png`) — Cross-Kerr robustness.** No-mitigation fidelity ≈0.05–0.10 (catastrophic for a real mean-photon-number ≈6.85 GKP codeword); an 8-pulse Carr-Purcell-Meiboom-Gill sequence recovers fidelity ≈0.996 (load-bearing). *SIMULATED real-data (QuTiP exact propagator); the flat-low-Fock reference understates damage ~10×; dynamical decoupling, not Stark detuning, carries the mitigation.*

E.2 The Eight New Figures (NF-1 through NF-8)

- **NF-1 — 6-layer heterogeneous-role-stack cross-section.** An eight-die stack (six functional + two spare) in cross-section: per-die role labels (distribution / inter-column routing / input-and-HBM-PHY / output / compute / cache-feedback), a thermal-aware-ordering arrow placing compute and cache at the inlet manifold, inter-die microfluidic coolant layers between every die pair, a ~332 µm stack-height callout, and a column inset (35 µm footprint, 76 TSVs, 80/15/5). *Schematic; serves CLM-6L-1, CLM-6L-2. The bond ceiling is not beaten — the microfluidic layers parallelize it; target 1–3 W/mm² per die [PROJECTED-TARGET, conjugate-fluid-dynamics-gated].*
- **NF-2 — Wakeup-chip embodiment plus measurement package.** A two-die WLCSP: a 55 nm main die over a 110/130 nm sensor die, a 16×16 TSV array color-coded by the 128/32/32/32/32 (signal/thermal/power/ground/guard) allocation, idle and burst voltage-level maps, and three measurement-KPI callouts (BER ≤ 10^{-9} at 1 Gbps per pair; ≥40% coset-detector false-codeword rejection vs 8-PAM; infrared-thermography validation within ±30%). *Schematic; serves CLM-E1-RTP. All product figures TARGET/estimate, silicon-pending; the chip's IP role is reduction-to-practice of the N2 system claim, not a signaling-novelty claim.*
- **NF-3 — Coset-thesis master tree, version 2 (updates FIG-THESIS).** The SU(N) coset root with **five** leaves (wakeup chip → 2-die HBM PHY → 6-layer heterogeneous-role stack → quantum stack), the **F-05 boundary band** across the classical leaves ("classical amplitude = PAM-8 plus parity, anticipated; coherent qudit = novel"), and the parallelization second-thesis annotation on the 6-layer and quantum leaves. *Schematic; new file (original FIG-THESIS preserved); serves §C and CLM-FOUND-1, CLM-FOUND-2; carries the F-05 boundary identically to §A/§C/§F/§G/§J.*
- **NF-4 — Held-invariant and floor-band dependency chart.** The exact Kostant weight-multiplicity generator feeding {exact $\Delta \approx 0.259$, protected-sector site count, exact K_6/K_{12} multiplicities, holonomy phase Φ} into the **held-fixed $g(\Delta) = 18.305$ at the (12,3) lock**, thence into the **two-stage scale-invariant floor band** (Stage 1 count-independent restriction → Stage 2 topological search → ≈9.6×–14.5× band: (12,3) ≈14.5× anchor / (8,3) ≈9.6× frontier / ≈28× ceiling); exact-computation boxes tagged PENDING; held invariants (index = 3, $\Delta \approx 0.259$, SM calibration) annotated as fixed inputs. *Schematic; serves §H (Q3/Q9) and the appendix dependency ledger; `[EXACTIFICATION-PENDING]` where a count is still pending.*
- **NF-5 — Inter-die microfluidic parallelization.** Series six-stack path (~5 bonds, ~11 K·mm²/W → ~0.9–1.8 W/mm² whole-stack, single sink) vs parallelized per-die path (one die + one bond ≈ 1.5 K·mm²/W → ~6.67 W/mm² per die), with the **bond-ceiling banner** ("copper–copper 10–50 W/cm² per interface — NOT removed; parallelized") and a production-target strip (1–3 W/mm² per die). *Schematic; serves CLM-6L-3;*

force only at ≥3 layers; all DERIVED except the production target [PROJECTED-TARGET, conjugate-fluid-dynamics-gated].

- **NF-6 — Chamber projector, null-space firewall, and the two selectivity senses.** The spectral projector P_chamber onto the $e^{i\Phi}$ holonomy eigenspace with null-space firewall, and bars contrasting admitted-fraction 1/(N+1) (25 → 20 → 16.7%, decreasing) against firewall-rejection N/(N+1) (75 → 80 → 83.3%, rising), tagged "exact Kostant counts — computation pending" with a Gap-C-contingent multiplier note. *Schematic; serves §H (Q2) and CLM-Q9; carries the §C.4 disambiguation; only K_6 DERIVED, SU(4)/SU(5) ESTIMATE/OPEN.*
- **NF-7 — The Φ↔Δ additive-offset chain.** A → Φ = −0.12827 rad → Δ = 2|sin(Φ/2)| ≈ 0.259 → g(Δ) = 18.305 at (12,3) → budget N_phys/N_log = 2·d_x·d_z·γ_ro·γ_ctrl − g(Δ), with an SM-calibration-ceiling annotation on Δ and the Δ→0 falsifier note; nodes tagged `[EXACTIFICATION-PENDING]` where pending. *Schematic; serves §H (Q3) and CLM-Q9; the additive (not multiplicative) entry of the holonomy offset is the load-bearing point.*
- **NF-8 — Cross-layer joint decoder plus commutativity status panel.** The three-layer joint syndrome hypergraph (GKP + chamber + outer) with chamber-null-space co-measurement highlighted, and a status panel: SU(3) BOTH_FAIL (residuals 1.18–2.06, basis-invariant); naive SU(4) rank-lift BOTH_FAIL (1.51/2.36); **disjoint-factor Z_4 BOTH_PASS (0.000) — CONDITIONAL on physical realization**; Path-4 (−0.08) ACTIVE. *Schematic; serves §H (Q4) and CLM-Q4; carries the conditional-not-achieved status of the deeper reductions.*
- **NF-11 — The G0→G6 staged-compiler dataflow.** The staged geometry-to-code compiler drawn left-to-right as a single dataflow: the four frozen scalar moduli (R, x_1, x_2, x_3) entering at G0, flowing through the intermediate stages into the tensor stack (holonomy / spectral / budget tensors), thence to the spectral-gap and per-stage budget nodes (Δ → g(Δ) → the additive budget offset; logical-failure-probability P_log and the budget margin Δ-against-target), to the noise-bias and chamber-rejection nodes (η ≈ 3.1 central kept strictly distinct from r_rej ≈ 18.3–19.3), and terminating in the locked outer-code geometry (d_x, d_z) = (12, 3). Annotations: "no new free parameters introduced at any stage"; the held invariants (index = 3, Δ ≈ 0.259, SM calibration, g(Δ) = 18.305 at (12,3)) shown as fixed inputs; the (8,3) frontier marked as the solution reached outside the initially-identified admissible search space. *Schematic; serves §D.2 (Q3/Q9) and §H; nodes tagged `[EXACTIFICATION-PENDING]` where an exact count is still pending; no overhead magnitude printed on the figure.*

E.3 Optional Figures (drawn if drafting reveals the need; pre-approved)

- **NF-9 — Vertical thermal-resistance ladder plus bond-interface bypass.** A series silicon-plus-bond resistance chain with the parallel copper-via conductance, annotated "≈6.5% at 4,800 vias/50 mm² [SIMULATED, conduction FEM]; 15–25% density-dependent [conjugate-FEA-gated]." *Schematic; serves CLM-N2-4; magnitude density-dependent and simulation-corrected, never the retired 15–25% at the stated density.*
- **NF-10 — Gap-04 verification trail.** The closed-loop method-power narrative (fail-closed blocker → native re-diagnosis → exact-coefficient scaffold → falsifier-attached bound → machine clearance → countersigned discharge) and the verification-gate ladder (V0 hash-pin → V1 adoption → V2 re-freeze → V3 fail-closed re-run → V4 countersign) governing propagation of the theory-side result into a quantum-side claim. *Schematic; labeled "theory-side, decision-grade, countersigned 2026-06-12; quantum-side consumption verification-gated (V0–V4)"; zero auto-promotions.*

> *[A-PASS note: the WA-1 self-certification table and status-preservation footer of this Part are retained in the source file `PART_A_FRONT_SPINE.md` and in `A_PASS_COMPLIANCE_REPORT.md`.]*

PART B — DETAILED DESCRIPTION, RUNGS 1 AND 2 (§F–§G)

Provisional Patent Application — Tahoe Labs (entity TBD) Sole inventor: Chris Bergstrom (AI tools used as instruments of reduction to practice and drafting, not co-inventors — *Thaler v. Vidal*, Fed. Cir. 2022; 2024 USPTO AI-assisted-invention guidance)

> **Status-label convention (one label per quantitative statement): MEASURED** (hardware-demonstrated, ours or cited literature) · **SIMULATED** (runnable simulation in this corpus) · **DERIVED** (analytic from stated assumptions, geometry, or Hamiltonian) · **PROJECTED-TARGET** (design goal/extrapolation, not yet simulated or measured at scale) · **BOUNDED** (a rigorous interval) · **CONDITIONAL** (holds only under a stated, not-yet-realized premise) · **OPEN** (a named gap requiring external demonstration) · **SPECULATIVE** (forward-looking, disclaimed as load-bearing).
>
> **Three over-arching honesty statements governing all of §G (binding):** (1) **No mechanism is claimed as a three-dimensional-thermal "solution."** The inter-die copper-to-copper hybrid-bond conduction ceiling (~10–50 W/cm² per interface) is physics; the mechanisms move the engineering wall outward within that ceiling, they do not remove it. (2) **Applicant expressly disclaims any approach to the Landauer or Bekenstein limits** (~10^{6}× and 10^{29}× headroom remain); no sub-Landauer mechanism is claimed. (3) **Every quantitative figure in §G is DERIVED analytical scaling unless explicitly labeled** MEASURED, SIMULATED, PROJECTED-TARGET, BOUNDED, CONDITIONAL, or OPEN.

§F — DETAILED DESCRIPTION, PART I: RUNG 1 (THE EASY RUNG)

Rung 1 comprises three inventions that are buildable, simulatable, and — in the case of E1 — shippable on present-day silicon: **E1**, a two-die always-on sensor-wake integrated circuit that serves as the production reduction-to-practice vehicle for the Rung-2 coset-signaling system claim; **E2**, a differential paired-transmon qubit site that manufactures a quantified noise bias in hardware by common-mode rejection; and **E3**, a bias-adaptive method that selects asymmetric XZZX surface-code distances from that device-derived bias. E2 and E3 are co-dependent and together constitute a system-level invention (the differential pair *creates* the bias; the asymmetric code *consumes* it). E1 reduces the Rung-2 system architecture to practice at the lowest commercial scale and supplies the measurement package that converts simulation-backed claims into silicon-backed ones.

§F.E1 — The Two-Die SU(8) Coset-Signaling Wakeup Chip (Reduction-to-Practice Vehicle)

§F.E1.1 — Product definition and IP role

In one embodiment, the invention is a sub-milliwatt, always-on keyword-, motion-, and sensor-wake-detection integrated circuit in an approximately 3 mm × 3 mm wafer-level chip-scale package, constructed as a two-die copper-to-copper-bonded stack and interconnected by a through-silicon-via (TSV) array carrying a Heisenberg–Weyl coset-admitted multi-level signal. The device monitors audio and motion streams continuously at low power and asserts a wake interrupt to an application processor upon a qualifying event. It does not run general inference; it runs a narrow, quantized keyword-spotting model (an INT4/INT8 convolutional or recurrent-lite network) with configurable threshold logic.

E1's IP role is locked as purely evidentiary — the shipping reduction-to-practice of the Rung-2 system claim, *not* a signaling-novelty claim. Per the F-05 kernel analysis carried into §G.N1, the $d \geq 8$ coset *classical amplitude* alphabet is anticipated by eight-level pulse-amplitude modulation gated by a single parity bit ("PAM-8 + parity") — the admitted level set $\{0,2,4,6\}$ is the kernel of the GF(2) parity check $H = [0\ 0\ 1]$ — and regains novelty only in a coherent quantum channel, which a classical TSV is not. E1's enforceable value therefore rests on the dimension-independent **thermal-via reallocation SYSTEM claim** (§G.N2), of which E1 is a production embodiment, and on the three evidentiary functions of §F.E1.6.

§F.E1.2 — The two-die architecture and why the split is mandatory

The two-die split is mandatory for two independent reasons. **(i) Product:** a single-die design forces the sigma-delta ADC and analog microphone frontend to share a substrate with the switching digital inference engine — the dominant noise source in wake-chip designs, degrading keyword accuracy at low SNR. The two-die copper-to-copper stack physically separates the analog domain (sensor die) from the digital domain (main die); the inter-die TSV link carries digitized samples, eliminating mixed-signal coupling at the source. **(ii) IP:** a single die cannot demonstrate a TSV signaling claim. The bonded two-die stack with a populated TSV array is the minimum structure that reduces the coset-signaling-plus-thermal-via system to practice.

Die	Process node	Area (target)	Function	Status
Main (digital)	UMC 55 nm LP (SMIC 55 nm / GF 55LPe fallback)	2.0 mm² ± 0.2	Keyword-spotting inference (quantized CNN), 16 KB SRAM, SU(8) TSV PHY (128 signal TSVs), host interface, wake-interrupt output, dual-rail power management	PROJECTED-TARGET (24-month path to silicon)
Sensor (analog)	UMC 110/130 nm analog (TSMC 130 nm / XFAB XT018 fallback)	1.0 mm² ± 0.1	PDM microphone interface, sigma-delta ADC (16-bit/16 kHz voice; 8-bit/1 kHz motion), anti-alias and gain stage, bias/reference, SU(8) TSV PHY transmit side	PROJECTED-TARGET

The older sensor-die node is deliberate: analog matching, flicker-noise, and reference accuracy are superior at 130 nm, and cost per unit area is lower. Both dies are background-thinned to 40 µm; total stack height is approximately 80 µm (40 µm + 1.5 µm bond + 40 µm). [DERIVED.]

§F.E1.3 — The TSV interface and the 256-via allocation table (the system claim, at production scale)

The inter-die interface is a 16 × 16 array of 256 TSVs within an approximately 128 µm × 128 µm footprint, joined by a copper-to-copper hybrid bond at **8.0 µm pitch** — HVM-mature today, deliberately conservative against the 4 µm NVIDIA roadmap pitch of §G. The allocation table demonstrates, in shipping silicon, the same freed-slot-to-thermal-via reallocation that §G.N2 claims at HBM scale.

TSV function	Count	Fraction	Status
SU(8) signal TSVs (64 differential pairs)	128	50.0%	DERIVED
Dedicated thermal vias	**32**	**12.5%**	DERIVED
Power	27	10.5%	DERIVED
Ground return	28	10.9%	DERIVED
Guard	**41**	**16.0%**	DERIVED
Total	**256**	**100%**	DERIVED

A binary NRZ baseline of equal aggregate bandwidth would require 256 *signal* TSVs and zero thermal vias; E1 allocates 12.5% to dedicated thermal extraction against that 0% baseline — the system claim (coset signaling frees slots → thermal vias) satisfied in a volume-production part. [DERIVED.]

> **Guard-allocation self-consistency note (binding).** The guard allocation is set at **16.0% (41 TSVs), not the earlier 12.5%, to satisfy the package's own signal-integrity floor** established at §G.N5: a **≥ 16% PHY-guard-TSV allocation is the minimal allocation that clears BER ≤ 1×10^{-12}** at the target symbol rate, whereas a 12.5% guard leaves the channel **MARGINAL (~1×10^{-11})**. E1 is therefore re-priced to ≥ 16% guard so that its own BER target (KPI-1) is met by the structural guard alone; the 5 power and 4 ground TSVs reallocated to guard remain comfortably above the wake-chip's low-current power-delivery need. **Equivalent alternative (recited):** where a design instead retains a 12.5% guard, a **KP4 / RS-FEC (e.g. RS(544,514)) forward-error-correction layer is recited as the functionally-equivalent means** of reaching the same BER ≤ 1×10^{-12} target; the ≥ 16% guard and the KP4/RS-FEC layer are disclosed as interchangeable embodiments. The load-bearing 12.5%-thermal-vs-0%-baseline reallocation is unchanged. [DERIVED.]

Locked TSV dimensions (foundry-handoff quality; full 50-parameter table resides in the source geometry file): TSV diameter 6.4 µm (80% of the 8 µm pitch); aspect ratio **6.25:1** (40 µm depth / 6.4 µm diameter — significantly easier to electroplate than the NVIDIA 12.5:1 at 4 µm, improving fill yield); fill electroplated copper with a TaN/Ta 10 nm barrier and 50 nm copper seed; via-last formation by ICP-DRIE post-FEOL; copper bond pad 6.4 µm × 6.4 µm at 80% area coverage; bond interface 1.5 µm (SiO_2 + copper damascene); thermocompression bond at 275 °C with a 30-minute anneal; CMP surface roughness below 0.5 nm RMS; bond tensile strength above 300 MPa. [All DERIVED / PROJECTED-TARGET, silicon-pending.]

§F.E1.4 — The SU(8) coset signaling, idle and burst level maps

The signaling alphabet is the dimension-8 Heisenberg–Weyl construction of §G.N1.1: HW(8) = {ω^c X^a Z^b : a, b, c ∈ Z_8}, ω = $e^{2\pi i/8}$; stabilizer subgroup S = ⟨X^2, Z^4⟩ (order 4); admitted alphabet the simultaneous +1 eigenspace {|0⟩, |2⟩, |4⟩, |6⟩} (even-index quartet, 50% admittance), odd-index complement suppressed at TX and rejected at RX. [DERIVED, exact.]

Because the rejected odd level between any two adjacent admitted levels is never driven onto the via, adjacent admitted levels are separated by twice the base level spacing. The chip supports two transmit modes selected by a

swing-select multiplexer:

Mode	Full-swing	Base spacing	Adjacent admitted-level eye	Status
Wake burst (≤ 50 ms)	500 mV diff	500/7 ≈ 71 mV	**286 mV (≈ 2× conventional 8-PAM)**	DERIVED
Always-on idle (reduced swing)	250 mV diff	250/7 ≈ 36 mV	**142 mV**	DERIVED

Idle admitted levels are {+250, +107, −36, −179} mV; burst admitted levels {+500, +214, −71, −357} mV. Suppressing the four odd intermediate levels removes ~half the transitions in a random stream, giving a **40–60% reduction in capacitive coupling energy to adjacent TSVs versus a PAM-2 baseline** [DERIVED, silicon-needed]. The SU(8) PHY is **90–123 transistors per TSV** (TX+RX — the locked filter-complexity figure used across §F/§G), ~**12,000–16,000 transistors total** (~0.005 mm²) [DERIVED]; per-pair rate at the conservative 8 µm wake-chip pitch is **locked to 1 Gbps** (KPI-1 rate), RC-limited [PROJECTED-TARGET]. (The §G HBM PHY runs the same filter at 4 µm pitch / 10 Gbaud; the rate difference is geometry, not filter complexity.)

§F.E1.5 — Power budget and package (enablement)

Always-on idle power budget. The continuous-monitoring state is budgeted at **420 µW total against a sub-500 µW target — PASS, silicon expected within ±30%** [DERIVED]. The budget decomposes as 280 µW on the main die (gated clock 50, SRAM retention 20, duty-cycled inference 120, SU(8) receive at idle swing 40, host interface 30, power management 20) and 140 µW on the sensor die (analog frontend 80, SU(8) transmit at idle swing 40, reference/bias 20).

Wake-burst power. The wake event (≤ 50 ms) is budgeted at **approximately 10.3 mW against a 10 mW target — MARGINAL** [DERIVED], with two firmware-selectable mitigations: (A) inference-engine power-gating at 80% frequency on the 55 nm LP process reduces inference power ~15% to ~9.0 mW (PASS); (B) reducing confidence-averaging inference passes from three to two saves ~2 mW to reach ~8.5 mW at a ≤ 2% accuracy cost (production default if measured burst exceeds 10 mW). Battery-life impact is **0.091% per day against a sub-0.5% target — PASS** [DERIVED: 420 µW × 86,400 s/day = 36.3 J/day against a 40,000 J 3000 mAh / 3.7 V cell].

Package. A 3 mm × 3 mm WLCSP with 64 solder balls at 0.4 mm pitch (12 power/ground, 8 host interface, 2 wake interrupt, 6 GPIO/config, 4 analog, 4 reserved/test, remainder no-connect), board-footprint-compatible with dense smartphone and IoT layouts. [PROJECTED-TARGET.] At volume the two-die bonded-TSV architecture is both the IP-demonstration requirement and the correct product engineering choice. [Economics omitted as non-enabling.]

§F.E1.6 — The three evidentiary functions and the KPI measurement package (the enablement artifact)

E1 converts the Rung-2 simulation-backed claims into silicon-backed claims; the KPI-1/2/3 measurement package performs the conversion.

Function 1 — Reduction-to-practice of the N2 system claim in production silicon. The wake chip satisfies all three elements of the §G.N2 lead claim in a shipping part: (a) 128 signal TSVs run the HW(8) d = 8 admitted alphabet {|0⟩, |2⟩, |4⟩, |6⟩}; (b) 32 thermal vias occupy slots a binary NRZ design of equal bandwidth would spend on signal TSVs; (c) the 12.5%-versus-0% thermal allocation against the equal-aggregate-bandwidth baseline is the conditioned reallocation itself.

Function 2 — The measurement package (three silicon KPIs). - **KPI-1 (BER):** ≤ 10^{-9} at 1 Gbps per pair in idle-swing mode on real copper-to-copper-bonded TSVs, measured by BERT across all 64 signal pairs over ≥ 10^{12} transmitted bits. [PROJECTED-TARGET — closes the "no silicon demonstration of coset signaling at the TSV layer" gap.] - **KPI-2 (crosstalk rejection):** the receive coset detector rejects ≥ 40% more crosstalk-induced false codewords than a baseline 8-PAM receiver under equivalent aggressor stimulus. Conservative at the 8 µm

wake-chip pitch; the benefit grows at the 4 µm NVIDIA pitch (coupling capacitance ≈ 4× higher). [PROJECTED-TARGET.] - **KPI-3 (thermal-model validation):** sensor-die junction temperature during a 10 mW main-die burst consistent within ±30% of the finite-element prediction, by infrared thermography — model-consistency validating the conductance model behind the N2 R_th claim, not a headline thermal delta at this scale. [PROJECTED-TARGET.]

Function 3 — Commercial proof. A part in mass production with a published datasheet and BER characterization report is a stronger enforcement posture than a test chip. [Non-enabling; stated once.]

§F.E1.7 — Claim support (E1), milestones (gate list), and honest risks

The independent claim E1 supports is a two-die system claim (a first die carrying wake logic and a second die carrying an analog sensor frontend, copper-to-copper-bonded and TSV-interconnected; a plurality of vias carrying the ⟨X², Z⁴⟩ even-index admitted quartet with the odd quartet suppressed; and a further plurality of dedicated thermal vias occupying slots freed by the bandwidth density of the multi-level signaling relative to binary signaling at equal aggregate bandwidth). The reduction-to-practice dependent claim (CLM-E1-RTP in §J) ties this to the §G.N2 system claim. **The dimension-8 specificity is severable:** if d = 8 is narrowed in prosecution, the dimension-independent thermal-via system claim survives and the silicon retains full leverage — this fallback is recited in the embodiment.

Milestone gate list. SPICE BER ≤ 10^{-9} at 1 Gbps before layout; FPGA emulation of the PDM→ADC→encode→decode→classify path; INT4 keyword model ≥ 95% at 5 dB SNR before tapeout; ANSYS Icepak sensor-die junction ≤ 85 °C at 10 mW burst; DRC-clean sign-off on both dies; tapeout; > 90% TSV-continuity / > 96% main-die / > 98% sensor-die probe yield; ≥ 85% package yield with a functional wake trigger on ≥ 3 parts; silicon characterization of KPI-1/2/3 + keyword accuracy, false-positive rate, always-on power.

Honest risks. Chief risks: (a) product-market fit (first design win may be tier-2/IoT, not tier-1 smartphone); (b) marginal wake-burst power, mitigated as above; (c) d = 8 prosecution narrowing, mitigated by the dimension-independent system-claim fallback. None is fatal to the evidentiary IP role.

§F.E2 — Differential Paired-Transmon Common-Mode Rejection (Manufactured Noise Bias)

§F.E2.1 — Mechanism and the pair Hamiltonian

In one embodiment, each logical qubit site of a superconducting quantum processor holds two nominally identical transmons, A and B, coupled by a tunable exchange interaction, with the quantum information of the site encoded in the *antisymmetric (differential) mode* of the pair. The pair Hamiltonian is [DERIVED]:

```
H_pair = (ω_A/2) σ_z^A + (ω_B/2) σ_z^B + J_AB (σ_+^A σ_-^B + σ_-^A σ_+^B) + (α_A/2)(σ_z^A)² + (α_B/2)(σ_z^B)²
```

with transition frequencies $\omega_{A,B}$, anharmonicities $\alpha_{A,B}$, exchange J_AB (reference 85 MHz, design band 80–100 MHz). The normal modes are the symmetric/common mode $\psi_+ = (|A\rangle+|B\rangle)/\sqrt{2}$ (discarded) and the antisymmetric/differential mode $\psi_- = (|A\rangle-|B\rangle)/\sqrt{2}$ (encoding the logical degree of freedom). A perfectly correlated noise operator acts identically on both transmons, commutes with the A ↔ B swap, and projects entirely onto ψ_+, leaving the encoded ψ_- subspace untouched; independent per-qubit relaxation survives. This is the quantum transposition of classical differential common-mode rejection as a correlated-decoherence firewall and bias generator at the qubit unit-cell level.

§F.E2.2 — The rejection ratio and its dual provenance (r_rej-vs-η discipline)

The common-mode rejection ratio in closed form is **r_rej = 1 / Σ_k c_k |δ_k|** [DERIVED], where the δ_k are normalized A/B mismatch parameters (frequency δω, coupling δg, anharmonicity δα, readout δ_ro, control

δ_ctrl) and the c_k are sensitivity coefficients of H_pair. At reference parameters (J_AB = 85 MHz, δω ≈ 10 MHz, δg ≈ 1.5 MHz), the corpus reports **two provenanced values**, both SIMULATED:

- **r_rej ≈ 19.3** — geometry-provenanced compiler path (artifact value 19.25). [SIMULATED.]
- **r_rej ≈ 18.3** — standalone adapter/noise path (artifact value 18.27). [SIMULATED.]

The usable noise bias is η ≈ 3.1 (central), NOT 18–19, and the two quantities are never conflated. r_rej is the common-mode rejection of *fully correlated* errors; the usable bias η = p_X_eff/p_Z_eff is the asymmetry surviving after accounting for the finite correlated fraction of real noise, built by splitting the physical error rate into dephasing (Z) and relaxation (X) parts, separating each into correlated/uncorrelated fractions, and applying r_rej *only* to the correlated fraction. Dephasing within the ≤ 40 μm pair is mostly correlated (shared substrate, readout mode, local flux), so f_Zc ≈ 0.80 central; relaxation is a per-qubit quantum-jump process, so f_Xc ≈ 0.08. Evaluated at r_rej = 18.27 and p_phys = 7.30×10^{-4}, the bias band is [SIMULATED]:

Z-correlation f_Zc	η = p_X/p_Z	Z-suppression	Status
0.60 (conservative)	**1.75**	~2.3×	SIMULATED
0.70	2.24	—	SIMULATED
0.80 (central)	**3.10**	**4.1×**	SIMULATED
0.90 (high)	≈ **5.1**	—	SIMULATED (asserted)

X-error suppression is essentially unchanged at ~1.1× across the band. A separate geometry-native directional-channel estimate gives a corroborating **η_geo ≈ 7.4** [SIMULATED]. **The correlation fractions (≈ 80% Z-correlated, ≈ 8% X-correlated) are assumptions informed by cited literature (Ithier et al. 2005), not measurements of a built device; r_rej, η, and every figure are SIMULATED or DERIVED, never MEASURED** — a disclosed enablement caveat.

§F.E2.3 — Standalone benefit, novelty, and claim support

Even with a symmetric decoder, the 4.1× Z-suppression lowers the dominant logical-error channel and permits a smaller working distance at fixed logical error rate, yielding the **~37% standalone overhead reduction** attributed to E2 [DERIVED/PROJECTED-TARGET, net of the differential-readout hardware cost]. The larger 37–62% reductions require the E3 asymmetric-distance method; E2 alone delivers the smaller structural cut and supplies the bias E3 consumes.

The prior art (Tuckett 2020/2021; Bonilla Ataides 2021; Koch 2007; Ithier 2005) treats noise bias as an environmental accident bias-tailored codes merely *assume*. The patentable kernel of E2 is that the bias is a *designed, computable, stable device property*: encode in ψ_-, dump correlated noise to ψ_+, read η off H_pair via the closed-form map from fabrication tolerances to bias (r_rej = 1/Σ c_k|δ_k| → η). It is not obvious that a paired-transmon cell's common-mode rejection would create a Z-bias of sufficient magnitude (η > 3) and stability to drive code-geometry decisions.

E2 supports an apparatus claim (CLM-E2 in §J): a qubit site of two exchange-coupled transmons encoding in the differential mode, with r_rej > 5 (dependent embodiments r_rej > 10 and > 18) and a structural bias η > 2, computed from the pair-Hamiltonian and mismatch parameters, with the matching selected to achieve a target η. Falsifiers (disclosed): hardware r_rej ≤ ~3, measured Z-correlation below ~0.4 (η → 1), differential-mode readout adding error faster than common-mode rejection removes it, or correlated bursts found to be anti-correlated. The on-device r_rej measurement is the key post-filing validation; no gate blocks a provisional, the simulation disclosure being complete.

§F.E3 — Bias-Adaptive Asymmetric XZZX Distance Selection (the g1 Closure Evidence)

§F.E3.1 — The honest framing (why E3 is defensible only as narrowed)

The XZZX surface code, asymmetric distances d_x ≫ d_z under biased noise, and the bare selection rule d_z = ⌈d_x/η⌉ for an *assumed/externally supplied* bias are all prior art (Bonilla Ataides et al. 2021; Tuckett 2020/2021). A claim readable with η "supplied," "assumed," "characterized," or "estimated from a noise model" reads directly on that art and is not filed. The *only* defensible novelty kernel is the **device-bias coupling made essential**: η is derived from the on-device-measured common-mode rejection ratio r_rej of the E2 differential transmon pair, and the asymmetric distances are selected from that device-derived η under a safety floor. That closed loop — hardware CMR → measured r_rej → derived η → selected d_z — is the whole invention and is kept load-bearing in the independent claim.

§F.E3.2 — The selection rule and overhead reduction

Under strong Z-bias the XZZX code decouples into two nearly independent repetition problems (DERIVED, prior-art behavior). Balancing the two channels' logical rates at minimum qubit count, and using that repetition-code logical error scales as $\sim p^{(d+1)/2}$, gives the selection rule [DERIVED]:

```
d_z = max(d_z_floor, ⌈ d_x / η ⌉),    d_z_floor ≥ 3 (safety floor)
```

The floor is load-bearing: it forces the method to *refuse to over-shrink* when the realistic η is small, distinguishing it from naive use of idealized r_rej as the bias, and forecloses the degenerate d_z = 1 line. A bias-headroom margin (headroom fraction ≈ 0.30, aggression factor ≈ 1.5) holds d_z above the minimum the bias would permit.

With d_x = 13 fixed (illustrative for the standalone EASY-rung method; the integrated Rung-3 stack uses d_x = 12, a different invention with the same d_x-agnostic rule), the overhead ladder is [SIMULATED/DERIVED]:

Config	d_x	d_z	bias used	phys/logical	reduction vs 269 baseline
Symmetric baseline	8	8	1.0	269	—
Central bias	13	5	η = 3.1	~168	~37–38%
Conservative-aggressive	13	4	(full)	~134	~50%
Aggressive	13	3	(full)	~101	~62%

Headline: 37–62% physical-qubit overhead reduction across the η band (overhead ~269 → ~101–168), the band disclosed rather than collapsed to the best case. [SIMULATED.]

§F.E3.3 — The 8.4×-decoded resolution (the retired-4.3× / vindicated-8.1× box)

> **RESOLUTION BOX (binding, RULE-9d and the F-05 honesty).** An internal flag (2026-03-17) asserted an "8.1×" same-distance logical-Z reduction; a committed repetition-code artifact showed only a "≈ 4.3×" figure, flagged as an unreproduced source assertion. A dedicated MWPM decoder campaign (Stim 1.15.0 + PyMatching 2.3.1, **≈ 1.0×10^8 decoded shots**) resolved this. **The committed "4.3×" was an *undecoded raw-syndrome-flip ratio*** — it counted bare logical-observable flips with no decoder run — reproduced by the campaign as **3.13× undecoded** at the usable bias η ≈ 3.1, and it is **retired/relabeled "raw-syndrome (undecoded)"; the figure "4.3×" is not carried forward.** When MWPM is actually run, the **decoded same-distance logical-Z reduction at d_z = 3 is 8.4× (95% CI 4.2–16.9) at η ≈ 3.1** [SIMULATED, decoded], rising to ~10.9× at η ≈ 5.1 and past two orders of magnitude (analytic ≈ 333×) toward the full rejection ratio. **The disputed 8.1× sits squarely inside the decoded interval and is thereby reproduced/vindicated; the decoded 8.4× is the figure of record.** The wide CI (≈ 9 versus ≈ 76 logical events) means the point estimate and ordering are robust while the exact

> multiplier is statistics-limited; claim language recites "at least a factor of four … approximately 8.4×" to stay inside the measured interval.

The same campaign verified the below-threshold operating point [SIMULATED, decoded, $p_{phys} = 3.9\times10^{-4}$]: X-axis at $d_x = 12$ gives $p_{L,X} < 1.0\times10^{-6}$ ($0/3\times10^{6}$ shots), Z-axis at $d_z = 3$ gives $p_{L,Z} \approx 3.3\times10^{-7}$ — **below the package's 5.49×10^{-7} logical target**, upgrading the $d_z = 3$ result PROJECTED-TARGET → SIMULATED. $d_z = 2$ ($\approx 4\times10^{-4}$ on the surface code alone) misses target by ~3 orders and reaches it only with the GKP inner code; $d_z = 1$ is degenerate, meaningful only as a GKP-protected logical. The depolarizing threshold (≈ 0.7–0.9%) and repetition strong-bias threshold (≈ 7.5–8%) reproduce textbook values.

§F.E3.4 — The g1 closure (genuine H-deformed XZZX) and the carried decoder caveat

> **GATE g1 — CLOSED → SIMULATED** (source: `GATE_CLOSURE/g1_xzzx_circuit/`). The one previously-open simulation gate — a *genuine* Hadamard-deformed two-dimensional rotated XZZX circuit with a deterministic logical observable — is now closed. The blocker (Stim rejecting a naively-transported boundary observable as non-deterministic) was resolved by applying the deformation as a *whole-circuit Hadamard conjugation* $C_{xzzx} = H_S \cdot C_{surface} \cdot H_S$ applied identically at prep, every syndrome round, final readout, and the observable definition — so **determinism is preserved by construction with no hand-derived boundary operator.** The result, on ≈ 1.1×10^{8} real decoded shots:

- **Genuine XZZX confirmed:** all stabilizers are provably mixed-XZ (0 pure-X, 0 pure-Z) at symmetric distances d = 3/5/7 *and* the asymmetric design point $d_x = 12$, $d_z = 3$, with 0 noiseless observable flips over 2×10^{5} shots and a valid decomposed detector-error model. [SIMULATED.]
- **Threshold rises with bias:** the protected-logical (memory_z) threshold climbs from ≈ 0.8% at $\eta = 1$ to > 1% (and rising) by $\eta \geq 10$ — the defining XZZX signature the plain surface code does not exhibit on the both-logical figure of merit. [SIMULATED, BOUNDED above as "> 1% and rising," not fitted.]
- **Single-logical biased advantage of 77–98× at $\eta = 100$** for the protected logical — replacing the prior ~3–4× surface-lower-bound bracket with a real XZZX number. [SIMULATED.]
- **≈ 2.3× physical-qubit reduction at equal target-logical protection** from the asymmetric patch (the E3 overhead-reduction claim, now on a real deformed circuit; the asymmetric d_x=12,d_z=3 patch protects the target logical to ≤ a 7×7 surface code's LER using ~2.3× fewer qubits). [SIMULATED.]

> **Honest caveat (carried, not buried).** Under *plain MWPM on a decomposed detector-error model*, a *symmetric same-distance* XZZX does **not** beat the surface code on the worst-of-both-logicals metric, because MWPM discards the X–Z correlations the deformation introduces. **This is a decoder limitation, not a circuit defect** — a correlated / belief-propagation decoder is the natural follow-on (a named, non-blocking pre-filing item) — and it does not affect the threshold-rise, single-logical-advantage, or asymmetric-qubit-efficiency wins on which the patent rests.

§F.E3.5 — Novelty and claim support (E3)

The distinguishing limitation is the closed three-step loop the prior art does not contain: a named hardware bias *source* (the differential transmon pair with common-mode rejection), *on-device measurement* of its r_{rej} and the correlated-error fraction, and *derivation* of η from that measurement followed by selection of d_z under a floor. A competitor who lays out an XZZX code with $d_z = \lceil d_x/\eta \rceil$ for a bias they assume, look up, or estimate does not practice the claim, because they perform none of those three steps; that is the firewall. The narrowed independent claim (the controlling language for §J) keeps step (a) naming the two-transmon differential CMR structure, step (b) the on-device r_{rej} measurement, and step (c) the derivation of η from that measurement, with a closing "wherein" clause locking η to the measured device property. Dependent claims pin the $r_{rej} = 1/\Sigma\, c_k|\delta_k|$ map ($r_{rej} \geq 10$, further ≥ 18; $\eta \geq 3$ at correlated fraction ≥ 0.8), the floor-plus-headroom guard, the re-

measurement/recalibration loop, the 8.4×-decoded essentiality validation, and the 37–62% overhead band. The method is drafted as tied to a concrete apparatus configuration and physical syndrome operations to clear *Alice/Mayo*.

Two honest caveats: the claim is parasitic on the E2 transmon pair being real, and r_rej/η are SIMULATED/DERIVED never MEASURED (step (b)'s on-device r_rej measurement is an OPEN hardware-enablement gate — a provisional may be filed on the simulated/derived disclosure); and an FTO patent-database search is a [DB-CHECK] action item, not a cleared opinion.

---

§G — DETAILED DESCRIPTION, PART II: RUNG 2 (THE NVIDIA / 3D-CHIP RUNG)

Rung 2 is the classical embodiment of the coset-admittance thesis applied to three-dimensional integrated circuits. **The headline is the 6-layer-tall heterogeneous-role chip of §G.7** — a six-functional-layer (plus spares) stack of thermally-ordered roles wired 80/15/5 on an O(N) hierarchical return network, carrying the lead integrated-system claim (CLM-6L-1) plus three further claim surfaces the two-die rung cannot reach. The six mechanisms that follow are the enabling system trades that make six layers thermally, yield, and bandwidth viable: **N1** (SU(8) coset signaling, defensive — frees the TSV slots N2 reallocates), **N2** (thermal-via reallocation conditioned on bandwidth densification — the strongest standalone prior-art-defensible fallback claim), **N3** (hierarchical-return-network O(N) noise topology), **N4** (spare-bank spare-die yield), **N5** (multipath vertex-disjoint fault-tolerant routing), and **N6** (branching-manifold counterflow cooling — the load-bearing thermal solution; it parallelizes the copper-to-copper bond bottleneck, it does not beat it).

§G.N1 — SU(8) Coset-Structured TSV Signaling (Defensive; the Coherent-Embodiment Boundary)

§G.N1.1 — The mechanism and the admittance algebra

In one embodiment, a through-silicon-via interconnect at the memory-to-logic PHY boundary transmits only the four even-index levels $\{|0\rangle, |2\rangle, |4\rangle, |6\rangle\}$ of an eight-level (d = 8) alphabet — the simultaneous +1 eigenspace of the Heisenberg–Weyl stabilizer subgroup $S = \langle X^2, Z^4 \rangle \subseteq HW(8)$ — rejecting the odd-index complement at both TX and RX. For d = 8, $HW(8) = \{\omega^c X^a Z^b : a, b, c \in Z_8\}$, $\omega = e^{2\pi i/8}$, $X|j\rangle = |j+1 \bmod 8\rangle$, $Z|j\rangle = \omega^j|j\rangle$, non-commutation $ZX = \omega XZ$. The admittance projector is $P_S = \mathrm{diag}(1,0,1,0,1,0,1,0)$ (rank 4, idempotent, self-adjoint); admittance rate exactly 50%. [DERIVED, exact.] At V_{dd} = 1.0 V the level spacing is 142.86 mV and the worst-case adjacent admitted-level eye is ≈ **286 mV — twice the 143 mV eye of conventional 8-PAM**. The TX+RX admittance filter is **90–123 transistors per TSV** at < 1 μm²/TSV, routing in BEOL above the via keep-out for a < 1% PHY-area penalty. [DERIVED.]

§G.N1.2 — The defensive posture and the F-05 classical/quantum boundary box

> **F-05 CLASSICAL/QUANTUM NOVELTY BOUNDARY (stated identically across §A/§C/§F/§G/§J; binding, RULE-9k).** For a **classical, amplitude-only** TSV channel — a single real voltage per symbol, Z/diagonal basis only — the d ≥ 8 coset signaling alphabet is **anticipated by PAM-8 + parity**, and the standalone signaling novelty is **OPEN/weak**. Three facts establish this: (1) the admitted level set {0, 2, 4, 6} = the even residues mod 8 = the three-bit words with least-significant bit zero, which **is** PAM-8 gated by one parity bit; (2) {0, 2, 4, 6} is a [3,2] binary linear code, **exactly the kernel of the single GF(2) parity check H = [0 0 1]**, with minimum distance 1 identical to representative linear codes — so the prior assertion "not the kernel of any equal-rate GF(2^m) linear code" is **false for amplitude-only signaling and is withdrawn**; and (3) the second stabilizer generator X^2 is a momentum (Fourier-dual) shift that adds *nothing* to an amplitude-only level set and is observable only in a coherent (X-basis) channel. **The genuine non-equivalence of the two-generator HW(d) stabilizer to a linear-code kernel is real ONLY in a coherent qudit channel** (the quantum rung's Hilbert-space embodiment), which a classical TSV is not. The "provably not a GF(2^m) kernel" language never appears in this application for a classical amplitude channel.

Consequently **N1 is not the lead claim — N2 (the system thermal-via trade) is.** SU(8) is carried defensively: it halves the signal-TSV count, which *enables* the thermal-via reallocation system claim. The standalone N1 signaling claim is retained only as a narrowed defensive/continuation claim; revival would require a coherent-channel embodiment carrying the X^2-conjugate degree of freedom, or a SPICE-demonstrated functional BER/eye distinction PAM-8 + parity provably cannot reproduce.

§G.N1.3 — Corrected parasitics and the τ_RC correction (real field-solver extraction)

> **τ_RC CORRECTION (binding, RULE-9f).** An earlier draft printed "τ_RC = 10 Ω × 100 fF = 1.0 ns." **That product is wrong by 1000×: 10 Ω × 100 fF = 1.0 picosecond, not 1.0 ns** (f_3dB ≈ 159 GHz) — a 1.0 ns constant would have *forbidden* the spec's own 10 Gbaud target. The corrected value is carried: **the bare via RC is ~1 ps and does NOT bound the symbol rate**; the 10–12 Gbaud ceiling is re-founded on driver slew rate, channel insertion loss, ISI, and neighbor-TSV coupling, gated by the KPI-2 SPICE/BERT item. No claim scope changes.

The TSV-pair parasitics are SIMULATED by open-source field-solver extraction. FastHenry2 gives pair loop inductance L = 16.41 pH and series AC resistance R = 0.227 Ω at 5 GHz; FasterCap gives bare-pair mutual coupling C = 1.60 / 6.25 / 19.06 fF for relative permittivity 1 / 3.9 / 11.9, **matching closed-form twin-cylinder theory to 0.2%** across all three dielectric cases. [SIMULATED.] The ~100 fF used in the rate analysis is the full array-loaded per-TSV parasitic, not the bare pair (which couples only ~6–8 fF); the remainder is array, landing-pad, RDL, and µbump loading. The openEMS FDTD pipeline was independently validated on a TE101 air cavity to 0.032%. [SIMULATED.] No claim scope changes.

§G.N2 — Thermal-Via Reallocation System Trade (the LEAD Claim) — with the Full Honest R_th Story

§G.N2.1 — The conditioned trade (the load-bearing idea)

In one embodiment, because SU(8) coset signaling carries a fixed aggregate inter-layer bandwidth on **half** the signal TSVs (8,000 versus a 16,000-TSV binary-NRZ interface of equal aggregate bandwidth), the ~8,000 freed slots are reallocated, with ~4,800 (30% of the total) becoming dedicated electroplated-copper thermal vias. The **claimable rule is the architectural coupling**: thermal-via density is *conditioned on* the bandwidth-per-TSV gain of the half-rate signaling layer. This system-level trade is absent from any cited prior art, and survives the F-05 boundary because the densification *mechanism* (fixed bandwidth on half the signal TSVs) is functionally real regardless of how the alphabet is framed — the system claim does not need the signaling novelty.

TSV function	Count	Fraction	Status
SU(8) signal (4,000 differential pairs)	8,000	50%	DERIVED
Dedicated thermal vias	**4,800**	**30%**	DERIVED
Power/ground return	2,400	15%	DERIVED
Guard (reduced; eye margin needs fewer)	800	5%	DERIVED
Total	**16,000**	**100%**	DERIVED

To defeat a "obvious to use a freed slot for a thermal via" obviousness attack, three things are essential limitations in the lead claim: (i) the **freed-slot causality** (slots freed *specifically by* the signaling layer's bandwidth-per-TSV gain); (ii) the **equal-aggregate-bandwidth baseline** (the comparison is against a binary-TSV stack of *equivalent aggregate bandwidth*, not an arbitrary stack); and (iii) the **quantified R_th coupling**. The signaling alphabet being PAM-8 + parity is immaterial and is deliberately admitted into scope.

§G.N2.2 — Thermal-via specification and the honest R_th story (G2 inlined)

Each thermal via is 3.2 µm diameter × 40 µm deep, electroplated copper with a TaN/Ta 10 nm barrier and 50 nm copper seed — **identical to the signal-TSV process, no added step.** Per-via R_via = L/(k_Cu·A_via) ≈ 12,900 K/W; the 4,800-via array adds parallel conductance G_thermal = 0.372 W/K. The bond interface dominates the baseline vertical path (R_bond ≈ 1.0 K·mm²/W of R_th_baseline ≈ 1.267 K·mm²/W, ~79%) because SiO_2 (k ≈ 1.4 W/m·K) is ~275× less conductive than copper (k ≈ 385 W/m·K); the copper thermal vias bypass that SiO_2 bond resistance.

> **LOCAL-FIGURE DISCIPLINE (binding, RULE-9f).** The figures **0.372 W/K, 213×, and 275× are LOCAL only and are never recited as system numbers.** The "213×" bond-conductance multiplier additionally carries a flagged 1000× unit slip (the intermediate "2.065 W/K" per-via line should be ~2.06×10^{-3} W/K); the honest bounds are the pure per-via material ratio k_Cu/k_SiO2 = **275×** (the true local maximum at a via location) and the **area-averaged** parallel-conductance multiplier over the full 50 mm² PHY of only **≈ 1.2×** (the copper cross-section is just 0.077% of area). None of 0.372 W/K, 213×, or 275× is the system result.
>
> **THE R_th STORY — STATED IN FULL (binding, RULE-5 and RULE-9e; the 15–25% figure NEVER appears without this G2 correction and the ANSYS gate adjacent).** The closed-form parallel-conductance estimate gives a **15–25% density-dependent** vertical-R_th reduction (15% for the localized 50 mm² PHY strip out of an 800 mm² die; 25% spread over ~150 mm²). **This 15–25% figure is DERIVED + FEA-ESTIMATED, NOT MEASURED, and it is the closed form — spreading resistance and non-uniform via distribution are not captured.** A real open-source steady-state conduction finite-element solve (scikit-fem 12.0.2, axisymmetric P2 triangles, bare-cell validated to 0.0% against the 1-D analytic series resistance, mesh-converged to < 0.01 pp) was run on this exact geometry and **returns ~6.5% R_th reduction (6.7% at k_Cu = 400) at the spec's own 4,800-via / 50 mm² density** [SIMULATED, conduction FEM, source `GATE_CLOSURE/g2_thermal_fem/`]. The gap is explained **entirely by spreading resistance** — heat must travel ~57 µm laterally through silicon to reach a via spaced at ~102 µm — which the closed form explicitly omits. **The 15–25% band is therefore density-dependent and not achievable at the current 4,800-via allocation:** reaching it under pure conduction needs **~2–3.7× more thermal vias (~10,000–17,000 over 50 mm², ~30–40 µm effective pitch) or the same 4,800 concentrated into a ~9–13 mm² region.** This does **not refute the invention** — the copper bond-bypass mechanism is real and directionally correct, and the system architecture (the lead claim) is unaffected; only the *magnitude* of the R_th sub-claim is pulled down. **The mandatory filing-grade gate is an ANSYS Icepak (or Cadence Celsius) conjugate-thermal FEA** of the bonded stack at 1 W/mm² compute stimulus with full-die in-plane spreading, hotspot localization, and TIM/lid paths — which may return a higher figure than the conduction-only FEM. Silicon may show 10–30%. **The numeric range is kept severable into a dependent claim (N2-2) and does not enter the independent claim until that FEA is on record.**

§G.N2.3 — Novelty and claim support (N2, the lead)

The novelty is the architectural coupling, not "thermal vias" or "multi-level signaling" (both old). No cited prior art (Kandou CNRZ-5, Marvell PAM-8, IBM multi-tone, JEDEC HBM4 PHY — which carries *zero* dedicated thermal TSVs — or US 11,631,444 / 10,425,260 / 12,237,953 / 11,233,681) combines multi-level/coset TSV signaling with thermal-via reallocation as a system rule. The distinguishing element: *thermal-via density at least 30% higher than a baseline binary-TSV stack of equivalent aggregate bandwidth, the additional vias occupying slots freed specifically by the bandwidth-per-TSV gain of the coset-signaling layer, producing a 15–25% vertical-R_th reduction.*

The lead independent claim (CLM-N2-1, controlling language in §J) recites (a) inter-layer signaling TSVs driven with a coset-admitted multi-level alphabet of HW(d), d ≥ 4, carrying a specified aggregate bandwidth using a signal-TSV count reduced relative to a binary baseline of the same aggregate bandwidth; (b) dedicated thermal vias occupying slots freed by, and in number determined by, that reduced signal count; (c) thermal-via density at least 30% greater than the baseline; and (d) a vertical-R_th reduction of 15–25% — (d)'s range gated by the mandatory FEA and moved to dependent CLM-N2-2 ("as determined by a conjugate-thermal FEA at ~1 W/mm²") until that gate closes. The d = 4 / PAM-8+parity case is intentionally within scope of CLM-N2-1; the standalone signaling claim is separately restricted to a coherent-channel embodiment. Dependents recite the freed-slot causality made express (anti-KSR), the bond-interface bypass (tied to the 275× material ratio, never the slipped 213×), the counts/geometry (16,000 total / 8,000 signal / 4,800 thermal, single via-last process), the memory-to-logic boundary, and the crosstalk benefit subordinated to the thermal architecture. **Defensibility MEDIUM (lead)** pending the [DB-CHECK] FTO pull. Falsifiers: the FEA returning < 15% (primary risk); the

bond interface not dominating; a prior-art reference pairing coset signaling with thermal reallocation; or freed slots unreallocatable on a real floorplan.

§G.N3 — Hierarchical Return-Network Topology (Hierarchical Return-Network, O(N) Noise Floor)

In one embodiment, N ≥ 4 device layers are organized into vertical functional columns spanning the layers, with TSVs partitioned into a majority of intra-column local vias, a minority of lateral vias (to adjacent columns through a dedicated inter-column routing layer), and a smaller minority of global vias (to a shared inter-layer trunk tier), and inter-layer signal returns routed through a **hierarchical return network with dedicated low-impedance trunks per layer-cluster.** Each interconnect node's ~10^4 links partition ~80% intra-column / ~15% adjacent-column / ~5% long-range, giving wiring-volume scaling of ~N^(4/3) rather than N^2. [DERIVED.]

The independent-claim heart is the noise-scaling exponent law, not the ratio or the role assignment [DERIVED]:

Architecture	Coupling-noise exponent	At N=100	At N=6	Status
Naive flat stack	$O(N^2)$	10,000×	36×	DERIVED
Tree-organized (HBM-style)	O(N log N)	~660×	~15×	DERIVED
Hierarchical return-network	**O(N)**	**~100×**	~6×	DERIVED
+ cross-layer cancellation*	O(√N)	~10×	~2.4×	SPECULATIVE

> *** The O(√N) row is SPECULATIVE and is not relied upon (binding).** It borrows shared-syndrome cross-layer cancellation from the QC stack, where the commutativity preconditions are currently **BOTH_FAIL** under the SU(3) chamber, unblockable only under a disjoint-factor Z_4 SU(4) chamber (OPEN; cross-ref Rung-3 g3). The mechanism that *achieves* O(N) is the hierarchical return-network topology with sparse inter-layer trunks, where the sparsity constant s ≈ 1–5% bounds crosstalk by s·N (linear).

Honest scoping (independent-claim core): the independent claim is the **return-network topology that bounds the coupling-noise exponent to O(N)**; the 80/15/5 partition, the six heterogeneous layer roles, and the per-column inhibitory veto (current/thermal/dV-dt sensors with sub-µs clock/power-gate authority) are *dependent* claims. Versus HBM's O(N log N) binary-tree clustering and general 3D-NoC topologies (Pavlidis–Friedman 2009), the novelty is the return-network-driven exponent reduction N log N → N. **Named gate:** column-scale RTL with thermal co-simulation. **Falsifier:** RTL fault/noise-injection at columnar scale failing to show O(N) coupling (exponent staying super-linear) falsifies the independent claim.

§G.N4 — Spare-Bank (k-of-n) Spare-Die Yield (k-of-n Sparing + Four-Scale Fault Isolation)

In one embodiment, a stacked-die package comprises N functional die positions and k ≥ 1 spare positions interconnected by a reconfigurable inter-die fabric, with **four nested fault-isolation domains** — transistor-cell, function-block, die, and package — each bounded by a hardware hard-isolation-boundary latch that physically disconnects power and signal across the domain boundary on fault detection, so a fault contained at one domain cannot propagate to the enclosing one; at boot the package tests all N+k dies, disables failures via the die-domain latch, and reconfigures the fabric to route around them, being functional when at least N of N+k dies pass. The series-stack wall is exponential — at 95% per-die yield, combined series yield is 90.3% (2-layer), 81.5% (4), 66.3% (8), 44.0% (16) — and the spare-bank formula breaks it [DERIVED]:

> **$P(\geq N \text{ of } N+k) = \Sigma_{j=N}^{N+k}\ C(N+k, j) \cdot p^j \cdot (1-p)^{(N+k-j)}$**

Worked at p = 0.95: 8 functional from 11 dies (k=3) gives 99.8% vs the 66.3% naive 8-die series yield; 16 from 20 (k=4) gives 99.9% vs 44.0%. **Spare-bank overhead grows *linearly* (~k extra dies) while series-yield collapse is *exponential*;** at 16 dies, 44% → 99.9% costs 4 spare dies (25% area overhead). [DERIVED.]

Honest scoping (independent-claim core): k-of-n die sparing *alone* is prior art (DRAM row/column sparing, FPGA redundancy); the **four-scale nested hardware hard-isolation-boundary hierarchy** is the distinct invention and the independent claim, with k-of-n, the continuous-monitoring BIST protocol (parity/ECC + scan-BIST + on-fault DEAD-line latch/spare activation), the offset brick-and-guard TSV layout, and the tree-ring ECC memory as dependents. **Named gate:** RTL fault-injection (iid + correlated) on 8- and 16-die models. **Falsifiers:** correlated failures driving 8-/16-die yield below the binomial table; copper-to-copper *bond* yield dominating (sparing must then extend to bonds); or the fabric unable to remap within the area budget. The binomial assumes iid per-die failures; the offset brick-and-guard mitigation is unquantified.

> **[NVIDIA SIM 2026-06-13 — dependent yield now SIMULATED (post-fix); TWO new yield claims close named threats.]** A 10^7-trial coupled Monte-Carlo at the integrated 8-die design point (`INTEGRATED_CHIP_TOPO_SEARCH.md`) fixed two threats and added two NEW dependent claims (§J, onto CL-N4): **(THREAT-1)** the spare must substitute a *heterogeneous role* — per-role or universal-KGD spares across all six roles; **(THREAT-2, the NEW bond-sparing claim) r ≥ 2 redundant Cu–Cu bond pillars per interface with k-of-n at the bond scale**. **Post-fix gated yield (quote these, NOT the un-gated single-class binomial): 99.37% for the 8-die universal-KGD (6+2, r=2) primary, 99.73% for the per-role max-yield (10-die) variant, 99.00% for the per-role fully-enabled fallback** [all SIMULATED, coupled MC]. THREAT-3 (isolation-latch/switchover correctness) is a named RTL/gate-level fault-injection gate, OPEN. Bond-sparing believed novel vs hybrid-bonding redundancy art — counsel [DB-CHECK].

§G.N5 — Multipath Vertex-Disjoint Fault-Tolerant Routing (k ≥ 4 Vertex-Disjoint Routes + Flux Reinforcement + Quorum)

In one embodiment, for at least one critical logical inter-die channel (clock distribution, ECC syndromes, inter-die coherence), the circuit provides **k ≥ 4 physically vertex-disjoint TSV routes** between source and destination dies (sharing no TSV other than the endpoints); a per-die router maintains per route a health indicator updated by per-packet error checking and a reinforcement weight incremented on successful delivery and decayed over time, selecting among alive routes with probability increasing in reinforcement weight and switching on health-threshold failure **without global coordination.** The channel survives while at least one of k routes is alive [DERIVED]:

> **$P_{survive} = 1 - (1 - p)^k$**

At p = 0.95, k = 4: $P = 1 - (0.05)^4 =$ **0.99999** (single-route is 0.95). The ~4× TSV-area cost is paid only on critical channels (~10% of TSVs); bulk data uses one or two routes. The flux-reinforcement dynamics reduce to a 4-bit weight counter incremented on packet success and decayed by a shift register every microsecond (a few gates per route). A quorum-commit mechanism (e.g. 6-of-8) handles cross-die operations needing agreement, tolerating $\lfloor (n-1)/2 \rfloor$ failures, **restricted to partial-agreement-tolerant semantics (power-state, telemetry, soft sync) and excluded from coherent memory writes.** Projected 8-layer yield with routing + quorum is 80–90% (vs 50% series), RTL-gated [PROJECTED-TARGET].

Honest scoping (independent-claim core): multi-path 3D-NoC routing (Pavlidis–Friedman 2009) and ACO routing (AntNet) are prior art; the **combination of k ≥ 4 vertex-disjoint routes + hardware flux-reinforcement routing + quorum-commit as one coherent subsystem** is the claim, no single primitive surviving alone. **Named gate:** RTL fault-injection on an 8-die model. **Falsifiers:** adaptive-routing overhead erasing the yield

gain; correlated TSV failures (a thermal hotspot) defeating vertex-disjointness, requiring routes to be *thermally* disjoint; or a prior-art reference claiming the three primitives together.

> **[NVIDIA SIM 2026-06-13 — thermal-disjoint requirement now DEMONSTRATED; NEW guard-TSV SI claim added.]** A 5×10^7-trial fault-injection + MC + BER campaign (`ROUTING_TOPO_SEARCH.md`, SIMULATED) and the integrated capstone add two §J dependent claims onto CL-N5: **(thermal-disjoint routing)** routing the k≥4 routes through **thermally-separated regions** (no hotspot disables more than ⌊k/2⌋=2) holds survival at **0.99990** (wide hotspot) vs **collapse to 0.9974** for naive clustered routing — converting the 1−(1−p)^k independence from asserted premise to demonstrated structural property, realized by the N6 manifold's 435 cooled lanes per face (≫ k=4); and **(NEW guard-TSV SI claim)** a **≥16% PHY-guard-TSV allocation** clears **BER ≤1×10^{-12} at 10 Gbaud** (Q 8.63, **BER 3.0×10^{-18}**, eye 81.8 mV), moving the channel from MARGINAL (~1×10^{-11} at 12.5% guard) to PASS — 16% the *minimal* clearing allocation (KP4/RS-FEC an equivalent alternative). **F-05 stands: routing/SI/guard claims, not signaling-alphabet novelty.**

§G.N6 — Branching-Manifold Counterflow Heat Extraction (Radius-Cubed-Conserving Manifold + Counterflow + Fractal Spreader) — the Strongest Standalone

Stacked logic faces a power-density wall — stacking N dies puts N× the heat in a footprint that grows 0% — and uniform copper thermal pads, planar spreaders, and single-pass co-flow cold plates top out near 0.5–1.0 W/mm² for logic-on-logic (peaks 1.5–2.7 W/mm² with backside microfluidics) [MEASURED, cited literature]. In one embodiment, a three-dimensional-stack cooling apparatus composes three independently-motivated engineering patterns, each with a closed-form basis (the strongest standalone apparatus claim in the package).

(A) Radius-cubed-conserving coolant manifold. At every bifurcation, **$r_{parent}^3 = \Sigma\, r_{child}^3$** — the radius-cubed-conserving branching law, the *unique* law minimizing total cost = pumping power (∝ flow³/r⁴ × length) + channel-area cost (∝ r² × length), a provable optimum not an empirical fit. For a symmetric binary tree $r_k = r_0/2^{(k/3)}$ (a 100 µm trunk → ~79 → ~63 → … → ~25 µm terminal over seven generations); channel volume scales as N^(1/3), resistance as N^(−1/3) versus a constant-radius tree. The benefit is **30–50% lower pump power for equal junction ΔT** (or ~2× extracted heat at equal pump budget) versus a uniform grid [DERIVED].

(B) Counterflow interdigitated inter-die channels. Adjacent inter-die channels carry oppositely-directed flow, the hot outlet pre-cooling the incoming cool stream through the thin (~20 µm) intervening silicon wall. The counterflow effectiveness $\varepsilon = (1 - \exp[-NTU(1-C_r)])/(1 - C_r\cdot\exp[-NTU(1-C_r)])$ reaches ε ≈ 0.83 at NTU = 5 and ≈ 0.91 at NTU = 10, versus the **ε ≤ 0.5 asymptote of an equal-area co-flow exchanger regardless of NTU** — i.e. **~2× heat-exchange capacity at equal area** [DERIVED].

(C) Fractal (high-surface-area) spreader (optional). Effective convective surface area scales as $N^{(D_f/D_e)}$ (D_f branching fractal dimension, D_e = 3 embedding); a flat copper spreader has D_f = 2, whereas an H-tree/Sierpinski fractal spreader (radius-cubed-conserving channels 80→64→50 µm, pin-finned at terminals) reaches effective **$D_f \approx 2.8$–2.9**. In a ~200 µm diamond/copper spreader layer (lateral k ≈ 2000 W/m·K) between top die and lid, lateral spreading and fractal area compound to reduce die-surface-to-coolant R_{th} **~5–10× under high-flux hotspots** (either alone gives ~2×) [DERIVED].

(Optional D) Passive natural-convection backup. Tall finned lid channels tuned so the Rayleigh number exceeds ~1700 at idle provide a passive natural-convection floor of **1–5 W/cm²** — a fail-safe when the active pump dies. The combined stack projects **~5–10 W/mm² extracted power density versus ~2 W/mm² SOTA** [PROJECTED-TARGET, CFD-gated], with the 30–50% pump-power reduction following from pattern (A).

> **THE LOAD-BEARING THERMAL CAVEAT (preserved verbatim, binding). N6 does not "solve" three-dimensional-stack thermal — it MITIGATES within a hard physical ceiling.** The most binding physical constraint is inter-die conduction across the copper-to-copper hybrid-bond interface (~1 µm bond, TIM resistance 0.5–2 K·mm²/W), capping continuous dissipation at roughly **10–50 W/cm² per interface** [DERIVED,

> Wiedemann–Franz + phonon-boundary scattering]. The branching-manifold counterflow mechanism raises *delivered* cooling capacity and lowers *pump power*; **it does NOT raise the copper-to-copper bond-conduction ceiling, and any claim implying otherwise would be invalidated.** Applicant expressly disclaims any sub-Landauer mechanism.

Honest scoping: versus uniform-grid microchannel cooling (IBM/EPFL/ETH single-pass co-flow uniform-radius grids) and dummy thermal TSVs / planar copper spreaders (AMD V-Cache, TSMC SoIC, ~10–25%), the distinct apparatus is the **radius-cubed-conserving graded radii + counterflow interdigitation + fractal high-surface-area spreader** combination, with the radius-cubed-conserving + counterflow pairing the differentiator if the branching law alone is anticipated. **Named gate:** conjugate-heat-transfer CFD (ANSYS Fluent/COMSOL) on a real SoIC/V-Cache geometry + 10-year reliability simulation. **Falsifiers:** CFD showing combined extraction below ~3 W/mm² (per-pattern DERIVED claims survive independently); the counterflow gain erased by a parasitic conductive short across the thin inter-channel wall; the radius-cubed-conserving-optimal terminal radii (< 10 µm) non-manufacturable; or electromigration/CTE-mismatch/creep failure within the 10-year profile.

> **[NVIDIA SIM 2026-06-13 — cooling-manifold geometry SEARCHED to a winner, CONFIRMED by a coupled OpenFOAM CFD; NEW dedicated-coolant-layer claim + TSV-through-wall-assist added.]** A 7,200-topology calibrated search (`COOLING_TOPO_SEARCH.md`) returns the winner **radius-cubed-conserving + counterflow 25×100 µm channels (AR 6.92, pitch 32.5 µm, U 1 m/s)**, per-die R_th band 1.43–2.0 K·mm²/W, Δp ≈30 kPa, **manufacturing-limited (channel depth ≤120 µm), NOT pump-limited.** A coupled OpenFOAM 11 CFD of the integrated winner (`INTEGRATED_CONFIRM_CFD.md`) CONFIRMED it: dense-array per-die R_th **1.45 K·mm²/W (≈22 single-pass conservative floor carried too**), compute peak rise **2.90 K @ 2 W/mm²**, buried-die ΔT **≤11.1 K** with a **DEDICATED inter-die coolant layer** (×13 vs the 144.6 K series control), and the **TSV-through-channel-wall a thermal assist** (R_th 1.45→0.73) [all SIMULATED]. Cooling channels and routing TSVs **cannot share the inter-die plane**, so a **dedicated coolant layer the TSVs traverse is mandatory** (+~721 µm Z-height) — a §J dependent onto CL-6L-3/CL-N6, more defensible than co-location. **Caveats unchanged:** both R_th regimes carried; counterflow inter-stream coupling + full manifold-tree Δp deferred to Icepak; the **ANSYS Icepak filing-grade gate stays OPEN**; the **Cu–Cu bond ceiling is PARALLELIZED, NOT beaten**.

§G.7 — THE 6-LAYER HETEROGENEOUS-LAYER CANDIDATE (the A7 Synthesis as a Full Embodiment)

This section specifies the six-functional-layer chip — **the headline of this rung** — as a complete integrated-system embodiment composing N1–N6 and carrying the lead integrated-system claim (CLM-6L-1) plus three further claim surfaces the two-die rung cannot reach. The architecture is the product; N1–N6 are the enabling system trades. It is defensible precisely because it is honest about the thermal wall: it does not claim the thermal-via trade carries six layers (~6.5% at 4,800 vias, FEM-confirmed), nor that branching-manifold counterflow cooling beats the copper-to-copper bond ceiling (it parallelizes it). The load-bearing thermal claim is mitigation-by-parallelization to a 1–3 W/mm²-per-die production target, gated on an ANSYS Icepak conjugate-CFD run that remains OPEN.

§G.7.1 — The six heterogeneous layer roles and the thermal-aware ordering

The backbone is a six-role heterogeneous-layer template — **distinct heterogeneous roles, not homogeneous compute repeated six times.** The load-bearing element is the heterogeneity, not the exact assignment (one defensible assignment among several). Mapped onto a stacked die for an NVIDIA Rubin-class HBM/logic context [DERIVED]:

Physical die	Functional role	Power-density posture	Status
Die 0 (top, nearest lid/cooler)	Distribution (clock trees, power grid; mostly metal)	Low	DERIVED
Die 1	Inter-column routing (horizontal NoC, global-trunk interposer, 5% long-range trunks)	Low	DERIVED
Die 2	Input / I/O (SerDes, HBM4 PHY, SU(8) coset TX/RX banks)	Medium	DERIVED
Die 3	Output / actuation (output drivers, DAC drivers, power-gating controllers)	Medium	DERIVED
Die 4	**Compute** (dense logic, ALUs, vector/tensor units)	**Highest (1–3 W/mm² avg, 5–10 hotspot)**	DERIVED
Die 5 (bottom)	Feedback / cache (cache arrays, HBM memory controller)	Medium-low (SRAM ~0.05–0.3 W/mm²)	DERIVED

The one rule that overrides the template is the thermal-aware ordering (AMD V-Cache rule): the highest-power-density layer sits adjacent to the primary heat-extraction path. Compute (Die 4) and cache (Die 5) therefore sit at the bottom, adjacent to the radius-cubed-conserving inlet manifold. **With inter-die microfluidic layers between every die pair (§G.7.3) the ordering constraint relaxes** because each die has its own local coolant sink — why the branching-manifold counterflow cooling is load-bearing: it removes the single-cooler tyranny that forces a strict power gradient.

§G.7.2 — Columnar organization and the 80/15/5 O(N) topology

A **column** is a vertical functional unit spanning all six functional layers, ~10^5 transistors (≈ 82k compute + 18k monitoring), 35 µm × 35 µm footprint per layer (~18% inhibitory/monitor overhead), ~163,000 columns per 200 mm² layer, ~10^6 columns total. Each column maps to one Tahoe Labs MFT-13 module (13 transistor-class subsystems across the layers). The column is the directly scalable unit — more columns means more area or dies *without changing the template*. [DERIVED.]

The 80/15/5 hierarchical-return-network wiring is applied per column per inter-die interface (76 TSVs at 4 µm pitch): 61 local (80%, intra-column), 11 lateral (15%, adjacent columns through Die 1), 4 global (5%, to the global-trunk interposer). **The load-bearing claim is the O(N) noise-scaling exponent, not the ratio:** with a hierarchical return network (sparsity s ≈ 1–5%), inter-layer coupling noise scales O(N), i.e. **O(1) per layer** — the noise-side counterpart to the spare-bank and branching-manifold parallelization. At N = 6 the gap is modest (~6× vs flat's 36×; it matters most at N ≥ 16), but it is the architectural invariant that makes the stack extensible toward 9-die (6+3). All signal TSVs are differential-pair (38 pairs per column per interface) because the N ≥ 8 differential-pair mandate is crossed by the 6+3 spare configuration. Sparse activation (≤ 5% of columns clocked per edge) is the source of the projected ~10× power-efficiency figure but **punishes dense matmul/FFT workloads** — for a dense-tensor context the realistic activation target relaxes toward 20–30%, eroding the power claim. [DERIVED / PROJECTED-TARGET.]

§G.7.3 — The integrated mechanism stack and the microfluidic parallelization (the load-bearing thermal mechanism)

Five integrated mechanisms compose into the 6-layer chip as enabling system trades — N6 (load-bearing thermal: branching-manifold microfluidic parallelization), N4 + N5 (yield), N3 (noise topology, O(N)), N1 SU(8) (signaling, defensive), and N2 (thermal-via trade, supporting knock-on to N6). The unifying move is **breaking an exponential-in-N series dependence into a constant-per-unit parallel one** — series bonds → parallel per-die paths (thermal), series dies → spare bank (yield), flat coupling → hierarchical return network (noise).

N2 at 6 layers — honest magnitude (per §G.N2.2). At the 4,800-via / 50 mm² density the conduction FEM gives ~6.5% R_th reduction [SIMULATED, G2], not 15–25% — the system claim is unaffected, only the magnitude pulled down. To recover 15–25% the design must raise the via count to ~10,000–17,000 over 50 mm² (~2–3.7×) *or* concentrate vias into ~9–13 mm² over the Die-4 hotspot; this candidate does both, treating the thermal-via trade as a supporting knock-on, never the primary thermal lever.

N6 (the load-bearing thermal mechanism) — inter-die microfluidic PARALLELIZATION. This is mitigation within the bond ceiling, not a structural defeat of it. A naive 6-stack forces buried-die heat through ~5 bonds in series (~11 K·mm²/W → only ~0.9–1.8 W/mm² with a single top sink). A radius-cubed-conserving counterflow manifold **between every die pair** breaks that series path into N parallel single-die paths, each paying the bond cost *once* (die + one bond ≈ 1.5 K·mm²/W):

```
q_max,per_die = ΔT_bond / R_per_die = 10 K / 1.5 K·mm²/W = 6.67 W/mm² ≈ 667 W/cm² per die [DERIVED]
```

This is a **~7× per-die improvement** over the series 6-stack, approaching the 5–10 W/mm² branching-manifold counterflow target — **but only with inter-die microfluidic layers**, demonstrated in research vehicles, never volume, with high yield and process cost.

> **THERMAL BANNER (preserved verbatim, binding, RULE-9e/f): "thermal is MITIGATED, not structurally solved."** The Cu–Cu bond ceiling (10–50 W/cm² per interface) is **NOT beaten.** The honest, defensible production target for the 6-layer candidate is **1–3 W/mm² per die for a 4–6 layer stack** (a ~5× density improvement over single-layer SOTA), achieved by parallelizing the bottleneck — never by claiming the bond ceiling is beaten. The microfluidic parallelization is conjugate-CFD-gated and load-bearing; the thermal-via trade (~6.5%) is a supporting knock-on.

N4 (spare-bank) at 6 layers — the yield recompute. The stack is built as a spare-bank of 6+k dies with the four-scale hard-isolation-boundary hierarchy of §G.N4 [DERIVED, binomial, verified]:

Target functional	Stack size	Series yield (no spare)	Spare-bank yield P(≥6 of 6+k)	Area overhead	Status
6	6 (k=0)	**73.5%** (0.95^6)	73.5%	0%	DERIVED
6	7 (k=1)	73.5%	**95.56%**	+16.7%	DERIVED
6	**8 (k=2)**	73.5%	**99.42%**	+33.3%	DERIVED
6	9 (k=3)	73.5%	**99.94%**	+50%	DERIVED

Recommended: k=2 (8-die) (73.5% series → 99.42%), with k=3 (→ 99.94%) as the resilience upgrade if first-silicon $p < 0.93$. Two die designs make the spares work (Design A, heavy-metal routing for Dies 0–1; Design B, full compute for Dies 2–5 and all spares), with offset brick-and-guard TSV layout and tree-ring ECC memory as dependents. The binomial assumes iid failures; correlated failures reduce the spare benefit and *bond* yield may dominate (then sparing extends to bonds — see the N4 bond-sparing claim).

N5 (multipath vertex-disjoint) at 6 layers. For every critical inter-die channel across the 5 (→ 7 with spares) interfaces, $k \geq 4$ vertex-disjoint routes give survival 0.99999 [DERIVED], with sub-µs local switchover (< 15 ns typical, ~66× margin to the 1 µs ECC window) and 6-of-8 quorum for power/telemetry/soft-sync. **Honest limit:** the $k \geq 4$ routes must be co-designed with the §G.7.3 manifold to pass through thermally-separated regions, or the independence assumption breaks. **N3** is the O(N) topology of §G.7.2.

§G.7.4 — The numbers (status-labeled) and the staged de-risk path

Stack geometry and interconnect [DERIVED]: 6 functional + 2 spare = **8 physical dies**, 200 mm²/die (14.1 × 14.1 mm), 40 µm post-thin, 1.5 µm bond, **332 µm stack height** (8×40 + 7×1.5), TSMC N3P (N2 fallback), 1,600 mm² total silicon. ~5.0×10^7 local + ~3.3×10^6 global TSVs; 76 TSVs/column/interface; 12.5:1 local aspect ratio;

SU(8) filter 90–123 transistors/TSV (SIMULATED parasitics L = 16.4 pH, C = 1.6–19 fF); 10 Gbps/TSV conservative.

Bandwidth and power density [DERIVED]: 380 Gbps aggregate per-column inter-die bandwidth (38 pairs × 10 Gbps); **40 Tbps PHY per HBM stack** (4,000 pairs × 10 Gbps, exceeding the HBM4 ~4 Tbps/stack target); Die-4 compute 1–3 W/mm² avg / 5–10 W/mm² hotspot.

Thermal budget — the honest accounting:

Parameter	Value	Status
Cu–Cu bond ceiling (HARD, not removed)	**10–50 W/cm² per interface**	DERIVED (physics)
Series 6-stack vertical R_th (no inter-die cooling)	~11 K·mm²/W	DERIVED
Whole-stack flux, single top sink only	~0.9–1.8 W/mm²	DERIVED
Per-die R_th WITH inter-die microfluidic (parallelized), analytic screen	**~1.5 K·mm²/W → ~6.67 W/mm² ≈ 667 W/cm² per die**	DERIVED
Per-die R_th, CONFIRMED 2026-06-13 coupled OpenFOAM CFD (dense array)	**1.45 K·mm²/W** (≈22 single-pass conservative floor carried)	**SIMULATED (OpenFOAM, integrated winner)**
Buried-die ΔT, CONFIRMED 2026-06-13 coupled OpenFOAM CFD	**≤ 11.1 K** (×13 vs the 144.6 K series control)	**SIMULATED (OpenFOAM)**
Honest production target	**1–3 W/mm² per die, 4–6 layers**	PROJECTED-TARGET (CFD-gated)
Thermal-via R_th reduction @ 4,800 vias / 50 mm²	**~6.5%**	**SIMULATED (conduction FEM, G2)**
Thermal-via R_th to reach 15–25% (conduction)	needs ~10,000–17,000 vias OR ~9–13 mm² concentration	SIMULATED (G2 sensitivity)

(N6 cooling-component figures — manifold pump-power 30–50%, fractal-spreader R_th 5–10×, counterflow ε ≈ 0.91 vs ≤ 0.5 co-flow, combined extracted 5–10 W/mm² — are stated with status labels at §G.N6 patterns A–D and not repeated here.)

Cost and timeline [PROJECTED-TARGET]: 6-layer (8-die) first-silicon cost $50M+; 36 + months to first silicon; per-layer hybrid-bond cost 15–40/cm² [MEASURED, industry], ~$240–640/die bonding alone at 200 mm² — making the 6-layer a datacenter-only part.

Staged de-risk path: RUNG 1 = wakeup chip E1 (2-die, 256 TSVs, 55/110 nm, d=8 SU(8), 24 mo) proves coset-signaling TSV silicon at the cheapest scale; RUNG 2a = 2-die HBM PHY N1/N2 (16,000 TSVs, 8,000 signal + 4,800 thermal, N3P SoIC-X, 12–18 mo) proves the system thermal-via trade + SU(8) at production pitch; RUNG 2b = the 6-layer heterogeneous-layer stack of this section (6 functional + 2 spare, 200 mm²/die, 36+ mo) proves heterogeneous roles + O(N) topology + spare-bank yield + branching-manifold counterflow cooling.

§G.7.5 — The four new claim surfaces (support) and the gate/falsifier table

The 2-die rung's claim surface is the signaling-plus-thermal-via system (N1/N2). The 6-layer candidate adds four claim surfaces (controlling language CLM-6L-1…4 in §J):

1. **CLM-6L-1 — Heterogeneous-role ≥ 6-layer system claim.** A 3D stack of ≥ 6 device layers carrying *distinct heterogeneous functional roles* (distribution / inter-column routing / input / output / compute / feedback-cache), **thermally ordered so the highest-power-density layer adjoins the primary heat-extraction path**, combined with the integrated mechanism stack.
2. **CLM-6L-2 — Integrated-stack composition claim.** The composition of {O(N) hierarchical-return-network topology + 4-scale spare bank + multipath vertex-disjoint k ≥ 4 routing + branching-manifold

radius-cubed-conserving/counterflow/fractal cooling + coset-signaling-enabled thermal-via reallocation} on a columnar MFT-13 fabric — **novel as a composition independent of any single realized number.**

3. **CLM-6L-3 — Inter-die microfluidic parallelization claim.** Coolant manifold layers between every die pair converting the series bond bottleneck (~11 K·mm²/W) into N parallel single-die paths (~1.5 K·mm²/W), enabling 4–6 active logic layers at 1–3 W/mm² per die **within the unbeaten Cu–Cu bond ceiling** (force only at ≥ 3 layers).
4. **CLM-6L-4 — Column-as-MFT-13-module mapping (dependent).** Each vertical functional column maps to one MFT-13 module spanning all six layers.

> **[NVIDIA SIM CAMPAIGN 2026-06-13 — integrated 6-layer assignment CO-OPTIMIZED to a winner and CONFIRMED by a coupled OpenFOAM CFD; integration claims move DERIVED → SIMULATED; four NEW dependent claims + the dedicated-coolant-layer structural element added.]** A calibrated capstone search (`INTEGRATED_CHIP_TOPO_SEARCH.md`, 19,440 candidates) co-optimized the §G.7.1 layer-role assignment against the cooling/yield/routing subsystem winners onto ONE buildable 8-die stack; a coupled OpenFOAM 11 conjugate-heat CFD (`INTEGRATED_CONFIRM_CFD.md`, four cases at machine-floor residuals) then CONFIRMED it — **VERDICT: CONFIRM, no fallback to 4-layer.** The integrated winner: - **Vertical order (coolant inlet/bottom → top): compute(inlet) → distribution → I/O → output → feedback-cache → NoC(top)** — hottest die at the inlet (AMD V-Cache rule). - **Cooling: radius-cubed-conserving-graded + counterflow 25×100 µm channels (AR 6.92, pitch 32.5 µm, U 1 m/s) in a DEDICATED inter-die coolant layer** the TSVs traverse. Per-die R_th **1.45 K·mm²/W (dense array) / ≈22 (single-pass conservative floor — both carried)** [SIMULATED, OpenFOAM]; compute peak rise **2.90 K @ 2 W/mm²**; buried-die ΔT **≤11.1 K** (×13 vs the 144.6 K series control); the **TSV-through-channel-wall HELPS** (R_th 1.45 → 0.73) [SIMULATED]. - **Sparing: 8-die universal-KGD (6+2, 33% overhead) + r=2 redundant Cu–Cu bond pillars/interface.** Post-3-threat-fix system yield **99.37%** [SIMULATED, 10^7-trial coupled MC]. - **Routing: 80/15/5, k=4 vertex-disjoint thermal-disjoint (aligned to the manifold's 435 cooled lanes), 16% guard, 6-of-8 quorum.** Survival **0.99990**, BER **3.0×10^{-18}** [SIMULATED]. - **Stack height ≈1,052 µm** (330 die+bond + ~721 dedicated coolant) — the integration's true cost is **Z-height, not in-plane area or pump power** (channels and TSVs cannot share the inter-die plane; 25 µm channels at 32.5 µm pitch collide with the 35 µm TSV columns, and the 16% guard overflows the column at 4 µm pitch — the dedicated coolant layer is mandatory).
>
> **Four NEW dependent claims** (formalized in §J): **(1) bond-sparing** — r≥2 k-of-n Cu–Cu pillars/interface (→ CL-N4; closes the bond-yield falsifier; counsel [DB-CHECK]); **(2) thermal-disjoint routing** — k≥4 routes through thermally-separated regions (→ CL-N5; converts asserted independence into a structural property); **(3) guard-TSV SI allocation** — ≥16% PHY guard → BER ≤1e-12 (→ CL-N2/CL-6L-1; **UPGRADED full-wave 2026-06-13 — a full-array openEMS FDTD crosstalk extraction pins guard shielding at ~96%/≈30 dB (the 55% cap was CONSERVATIVE) and realistic-case BER at 1.4×10^{-19}, superseding the behavioral ≈3.0×10^{-18}; shielding RATIO full-wave-robust, absolute mV PDK-gated**); **(4) dedicated-coolant-layer** — the manifold as a distinct claimed layer the TSVs traverse (→ CL-6L-3; more defensible than co-location). CLM-6L-1 and CLM-6L-3 **upgrade DERIVED → SIMULATED**. **Caveats unchanged:** both R_th regimes carried; counterflow inter-stream coupling + full manifold-tree Δp + TSV-through-wall thermomechanical stress remain Icepak/3-D deliverables; the **ANSYS Icepak filing-grade gate (G-1) stays OPEN**; the **Cu–Cu bond ceiling is PARALLELIZED, NOT beaten**; F-05 unchanged.
>
> **Honesty bindings carried into the 6-layer claims (binding):** the thermal-via trade gives ~6.5% at the stated density (G2), a supporting knock-on never the thermal solution; the Cu–Cu bond ceiling is **NOT beaten** (mitigation-by-parallelization to a 1–3 W/mm²/die target, conjugate-CFD-gated); the O(N) topology row is the claim while the O(√N) row is SPECULATIVE; the yield claims carry the iid-binomial caveat; N1 SU(8) is defensive.

The gate/falsifier table (what must close before this is more than a candidate):

Gate	What it validates	Status	Falsifier
ANSYS Icepak / Cadence Celsius conjugate-CFD of the full 6-layer stack with inter-die manifolds @ 5 W/mm²	The load-bearing thermal claim (1–3 W/mm² per die)	**OPEN (filing-grade, highest priority)** — PRE-SCREENED 2026-06-13 by the coupled OpenFOAM CFD above (VERDICT CONFIRM; buried-die $\Delta T \leq 11.1$ K)	Buried-die ΔT exceeding leakage-runaway even with inter-die cooling → 6-layer thermal claim fails, fall back to 4-layer. (Not fired by the 2026-06-13 CFD.)
ANSYS Icepak thermal-via FEA at densified/concentrated allocation	Recovering 15–25% R_{th} vs G2's 6.5% at 4,800/50 mm²	OPEN (G2 conduction FEM: 6.5% at baseline)	Concentrated vias still < 10% → the via sub-claim stays at ~6.5%, N6 carries the thermal budget.
RTL fault-injection (iid + correlated) on 6+2 / 6+3 models with k=4 routing + spares	O(N) exponent, spare-bank yield, vertex-disjoint survival	OPEN	Noise exponent super-linear, or correlated failures drive yield below binomial, or routing overhead erases the yield gain.
Yield-model validation vs first-silicon p at 200 mm² N3P	Binomial inputs and k=2 vs k=3	OPEN	$p < 0.93$ → k=3 required; bond-yield dominating → sparing must extend to bonds.
SPICE BER of SU(8) encoder + RX + N3P parasitics @ 10 Gbaud	The signaling layer (BER ≤ 10^{-12}) and the corrected τ_{RC} (1 ps, not 1 ns)	OPEN	BER target unmet at 10 Gbaud without equalization → rate drops (architecture valid at lower TSV count).
Prior-art sweep (GF(2^m)-kernel non-equivalence, $d \geq 8$)	N1 signaling novelty	Closed-negative for classical amplitude (F-05)	Resolved; N2 system claim unaffected.
EDA flow demo (Synopsys 3DIC Compiler / Cadence Integrity, ≥ 4-die)	Buildability of a 6-layer place-and-route	OPEN (tools immature past 4 dies)	6-die P&R infeasible without heavy manual floorplanning → staged 2→4→6 rollout.

> *[A-PASS note: the WA-2 self-certification table and status-preservation footer of this Part are retained in the source file `PART_B_RUNGS_1_2.md` and in `A_PASS_COMPLIANCE_REPORT.md`.]*

SECTION H — DETAILED DESCRIPTION PART III: RUNG 3 (QUANTUM COMPUTING, Q1–Q10)

Document role: Writing-agent WA-3 deliverable for the Tahoe Labs 100-page provisional (blueprint `C7_PROVISIONAL_100P_BLUEPRINT.md` §1 row H; §4.3); assembled by A-PASS into `PROVISIONAL_PATENT_FULL_100P.md`. Symbolic section anchor **§H**; symbolic claim/figure IDs throughout (CLM-…, FIG…, NF-…), A-PASS assigning final numbering. **Sole inventor:** Chris Bergstrom. **Assignee:** Tahoe Labs (entity TBD). AI systems were used as computational and drafting tools and are not inventors (*Thaler v. Vidal*, Fed. Cir. 2022; 2024 USPTO AI-inventorship guidance). **Status-label convention (per §A):** every quantitative statement carries exactly one of MEASURED / SIMULATED / DERIVED / PROJECTED-TARGET / BOUNDED / CONDITIONAL / OPEN / SPECULATIVE (ESTIMATED only as a parenthetical qualifier on DERIVED inputs and the historical selectivity bands). The exactification sweep (runs S1/S2, §H.9.8) has landed; its exact values — Δ = 0.2589200 (prose "≈0.259"), g(Δ) = 18.305 at (12,3), Atiyah–Singer index = 3 — are the held-fixed invariants, carried throughout in place of earlier estimate-fed placeholders.

H.0 Orientation, re-anchoring posture, and the moduli-point disclosure (read first)

This Part describes the Rung-3 quantum-computing inventions Q1–Q10: a four-layer fault-tolerant stack — a qudit Gottesman–Kitaev–Preskill (GKP) bosonic inner code (Q1), a flag-manifold coset admissibility chamber with a null-space firewall and a geometric magic-state lift (Q2), an asymmetric XZZX surface outer code with a topological additive budget offset (Q3), and a photonic modular fabric (Q5) — plus a three-layer joint decoder (Q4), an outer-code-removal crossover method (Q6), the cavity apparatus enabling it (Q7), an orthogonal-subspace encoding (Q8), the integrated system composition (Q9), and the prior 28-family portfolio inlined at claim-chart level (Q10).

The defended posture (governing every number below; full honesty banner at §H.9.0). The mission is the lowest *defended* physical(bosonic-mode)-to-logical ratio at 10,000+ logical qubits, and the defended result is a **two-stage, scale-invariant overhead floor substantially independent of the logical-qubit count** — Stage 1 restricting the search to the count-independent regime, Stage 2 performing a topological search within it to select the geometry-derived code. The reported floor is the **SIMULATED band of ≈9.6× to 14.5×** physical(bosonic-mode) per logical — anchor ≈**14.5×** = **(12,3) binary-lever-locked**, frontier ≈**9.6×** = **(8,3) X-verified** — with a **conservative ≈28× ceiling** (DERIVED) if the chamber 0.60 credit fails its lo-bound, **evaluated at 10,000 logical qubits and nominally unchanged at 20,000**. The separately reported decoder-channel advantage ≈8.4× (95% CI 4.2–16.9×) is **not** a memory-overhead ratio. Every Δ/Φ/g(Δ)-derived number carries its exact landed value (Δ = 0.2589200, g(Δ) = 18.305 at (12,3)) per the §H.9.8 sweep.

Moduli-point disclosure (governs every theory-side consumption in this Part). The theory-side native closure of the substrate modulus (Gap-04 discharge, §H.9.7) is evaluated at the **origin of moduli space**, where the banked invariant registry of record (R = 30, Ric = 5g, $|\mathrm{Riem}|^2 = 276$, $|\mathrm{Weyl}|^2 = 216$, Vol = $\pi^3/2$ — DERIVED, banked) is computed. The QC chamber of Q2/Q3/Q9 runs at the **iter-4e squashed basin** $(R, x_1, x_2, x_3) = (0.880, 0.750, 1.087, 1.087)$ (DERIVED, pinned), whose squashed directions carry the **registered NT-1 transverse residue**; its transverse Hessian is a **named verification, not an inheritance.**

Blast-radius statement (load-bearing). The corrected curvature-invariant registry does **not** move Δ, Φ, g(Δ), or any number in the ≈9.6–14.5× band: none of those quantities consume $|\mathrm{Riem}|^2$ — Δ and Φ are spectral/holonomy objects, and the overhead chain consumes the holonomy gap and counting data only. The corrected registry touches only (i) the UV-evidence layer (quarantined, Appendix §K.App-A; not relied upon here), (ii) the superseded modulus-potential inputs (replaced by the native scaffold, §H.9.7), and (iii) the checkability frontier:

Δ, Φ, N_sites,gap, and the chamber multiplicity counts, formerly estimate-grade, are now exactly computed by the Kostant generator (§H.2.6) via the landed runs S1/S2.

H.1 INVENTION Q1 — QUDIT GKP (d = 4 QUQUART) ON A COSET-RESONANT BOSONIC MODE

H.1.1 Field and problem

Quantum error correction; continuous-variable (bosonic) encoding in superconducting microwave cavities; resource-overhead reduction for ≥10^4-logical-qubit fault-tolerant processors. Binary (d = 2) GKP stores one logical qubit per cavity mode; the cavity plus its transmon ancilla is the dominant inner-layer hardware cost. A single bosonic mode is infinite-dimensional: a qudit GKP code of dimension d stores $\log_2(d)$ logical qubits per mode, so a ququart (d = 4) halves the raw cavity count per logical qubit. The naive "ququart halves the whole physical count" is wrong, because (a) the chamber and outer-code overheads are tuned to the binary algebra and do not automatically reduce, and (b) the higher-d codeword lattice is harder to read out at fixed squeezing. The invention is therefore three things together: (i) the ququart lattice with its Z_4 Heisenberg–Weyl logical algebra on a coset-resonant mode; (ii) the lift of the Q2 chamber projector from Z_2 to Z_4; and (iii) the corrected composition algebra separating what reduces by $1/\log_2(d)$ from what does not.

H.1.2 Binary baseline and the bosonic departure from the iid-Pauli class

The GKP code (Gottesman, Kitaev, Preskill 2001, PRA **64**, 012310) encodes a logical qubit as the joint +1 eigenspace of two commuting displacement stabilizers $S_q = \exp(i\cdot 2\sqrt{\pi}\cdot\hat{q})$, $S_p = \exp(-i\cdot 2\sqrt{\pi}\cdot\hat{p})$ on a symplectic-dual square lattice of primitive cell area 2π (DERIVED). Steane-style correction couples the data mode to a fresh GKP ancilla via a SUM gate, reads the quadrature modulo the lattice, and applies the inverse displacement (Glancy & Knill 2006, PRA **73**, 012325). The per-cycle ideal-syndrome logical error rate is $P_L = \tfrac{1}{2}\cdot\mathrm{erfc}(\sqrt{\pi}/(2\cdot\sqrt{(\Delta_{GKP}^2)}))$ (DERIVED) — suppression **exponential in $1/\Delta_{GKP}^2$, not polynomial in qubit count**, the categorical departure from the iid-Pauli universality class the whole Rung-3 architecture exploits. (This squeezing parameter Δ_{GKP} is distinct from the topological gap Δ of §H.3.)

Squeezing → Δ² convention (LOCKED package-wide; both conventions disclosed). The package headline uses the grid convention $\Delta_{GKP}^2 = 10^{(-s_{dB}/10)}$: 15 dB → $\Delta^2 = 0.032$ → $P_L = 1.22\times10^{-12}$ (DERIVED, analytic ideal); 20 dB → $\Delta^2 = 0.010$ → $P_L = 2.46\times10^{-36}$ (DERIVED, analytic ideal). A finite-energy QuTiP campaign **measured** the peak variance and established the physically correct Glancy–Knill convention $10^{(-s_{dB}/10)}/2$ (15 dB → 0.01586 measured; SIMULATED), under which the ideal tails are deeper still. The grid convention is retained as the **more conservative bound on P_L**; both are recorded so they are reconciled, not in conflict, and the architecture logic is insensitive to the choice.

Two-track P_L (the ideal values are never used as device rates). The analytic-ideal tails are unobservable in any finite-Fock simulation; the QuTiP closest-lattice-point decoder measured a 15 dB binary inner-code error of **$P_L \approx 4.98\times10^{-5}$** (SIMULATED; $\approx5.3\times10^{-5}$ with photon loss κt = 0.01), saturating at a finite-energy floor of $\sim10^{-5}$–10^{-4} above ~12–15 dB. The closed-loop figures of record: binary $\sim10^{-5}$ at 15 dB and $\sim10^{-14}$ at 20 dB (SIMULATED); ququart device-expected $\sim10^{-2}$ at 15 dB (PROJECTED — degraded by the same ancilla-coherence bottleneck that limits binary GKP).

H.1.3 The ququart lattice and the Z_4 Heisenberg–Weyl algebra (machine-precision verified)

A qudit GKP code of dimension d uses a lattice of primitive symplectic-cell area $d\cdot\pi\cdot2/2$ — for the square ququart, **cell area 8π** — whose d cosets span the logical space: **d = 4 → $\log_2(4) = 2.000$ logical qubits per cavity, 0.500 cavities per logical qubit** (DERIVED + SIMULATED to machine precision; QuTiP `SIMULATION_CAMPAIGN/gkp/` §2.1). The binary Pauli group is replaced by the Heisenberg–Weyl group over Z_4: $Z_4|j\rangle = \omega^{j}|j\rangle$, $X_4|j\rangle = |j+1 \bmod 4\rangle$ with **$X_4Z_4 = \omega\, Z_4X_4$, $\omega = \exp(2\pi i/4) = i$** — not −1. The commutation phase ω =

i is **confirmed in a bosonic Fock space to residual ~5×10⁻¹⁵** (SIMULATED; d = 2 → −1 at ~1.4×10⁻¹⁵, d = 3 → $e^{2\pi i/3}$ at ~2.4×10⁻¹⁵). This ω is a Weyl phase, exact independent of Fock cutoff, so the algebraic fact forcing the §H.1.5 chamber lift is verified, not asserted.

Hardware anchor (disclosed limits). The single experimental anchor is Brock, Singh, Eickbusch et al., *Nature* **641**, 612–618 (2025): ququart (d = 4) and qutrit (d = 3) GKP **beyond break-even** at ≈15 dB, F_ququart ≈ 0.95, F_qutrit ≈ 0.97 (MEASURED; reproduced by the calibrated QuTiP model at F ≈ 0.945 / 0.967, SIMULATED). **Break-even is not fault tolerance; no qudit-GKP fault-tolerant logical qubit has been built** (OPEN). A photon-loss cross-check shows raw-loss codeword fidelity is d-independent at fixed κt (qudit penalty vs d = 2 only −0.03% at κt = 0.05; SIMULATED), attributing the ququart fidelity penalty to syndrome-extraction circuit complexity (≈2× ancilla overhead) rather than intrinsic lattice loss-susceptibility — the attribution the degraded compaction factor of §H.1.4 encodes.

H.1.4 The corrected composition formula and the honest scenario band

The composition is organized so each term enters only where it physically applies: the $1/\log_2(d)$ factor reduces the GKP inner raw cost, not the chamber/outer overhead, and the modular multiplier inflates rather than reduces. The algebra (DERIVED):

> single_module(d, scenario) = [GKP_raw(d) + chamber_output(scenario, d)] × compaction_factor(d); 10K_ratio = single_module × modular_multiplier,

with GKP_raw(d) and chamber_output(scenario, d) at the (12,3) geometry; the chamber credit firing in full (Scenario A), partially (Scenario B, f_partial ≈ 0.75), or not at all (Scenario C, credit-stripped ceiling); compaction_factor(d = 4) (DERIVED; encoding the Brock-2025 fidelity penalty); modular_multiplier ≈1.04 (§H.5). Values are bounded by the §H.9 two-stage floor, physical(bosonic-mode)-per-logical, all SIMULATED / DERIVED design-stage projections:

Configuration	Single-module	10K fabric	Status
(12,3) binary-lever-locked, full chamber credit (the anchor)	**≈6.6**	**≈6.9**	SIMULATED — the ≈14.5× anchor, floor with every accounting-gated credit firing
(12,3) with partial chamber credit	mid-band	mid-band	DERIVED — interpolates within the band
(12,3) credit-stripped — chamber 0.60 credit fails its lo-bound	**≈28**	**≈28**	DERIVED — the conservative ≈28× ceiling (two weakest credits removed)
(8,3) X-verified frontier	≈9.6	≈9.6	DERIVED + EXTRAPOLATED-X — the ≈9.6× frontier

Honest reading (binding; full banner at §H.9.0). The (12,3) lock is the ≈14.5× anchor; the (8,3) X-verified configuration the ≈9.6× frontier; the conservative ≈28× ceiling (chamber 0.60 credit fails its lo-bound) is carried at all times. The band is a memory-overhead floor at scale, never a full-algorithm ratio or a measured device property.

H.1.5 The chamber Z_4 lift and the dimension-independent magic angle

The ququart inner code only delivers if the Q2 chamber projector lifts from Z_2 to Z_4 without losing the chamber savings or the magic-state geometry. Two derivation paths and one negative result establish the lift (all DERIVED):

- **Path 1 (direct lift).** The chamber projector is the spectral projector onto the $e^{i\Phi}$ eigenspace of the coset holonomy operator U_Φ; replacing the binary logical group with the Z_4 Heisenberg–Weyl group, a

syndrome is chamber-rejected if its pre-image carries a non-trivial Z_4 phase {i, −1, −i}; only the trivial sector is admitted (throughput 1/4 vs 1/2 binary).

- **Path 2 (native Z_d projector — the construction claimed).** Decompose the coset fiber into eigenspaces of a Z_d center C_d generated from the toric U(1)² directions; the native projector is the Haar average **P_chamber^(d) = (1/d)** $\Sigma_{k=0}^{d-1}$ **C_d^k** — idempotent, self-adjoint, and commuting with U_Φ because C_d acts on the toric base while U_Φ acts on the spin-c (1,0) fiber. This disjoint-factor structure is the property gate G3 later showed to be the cross-layer commutativity unlock (§H.4.5).
- **Path 3 (tensor product — negative result).** A ququart is a single-mode Z_4 object, not two binary modes; "$K_6 \otimes K_6$" is a product manifold, not a higher flag manifold, and its doubled holonomy gives the wrong (S-gate) phase. Preserved only as a binary-chamber-depth heuristic.

The d-independent magic angle (load-bearing structural result). The magic angle is a holonomy of the substrate *geometry*, fixed by the signed area A = +0.07139 σ², independent of the logical code dimension: **Φ_d = Φ = −0.12827 rad** for all d (DERIVED). At d_chamber = 6 the accumulated phase 6×|Φ| = 0.76962 rad ≈ π/4 prepares a T-state approximant at residual 0.01578 rad (F_chamber_phase ≈ 0.99994). No mechanism feeds the logical encoding back on the substrate geometry: U_Φ acts on the coset fiber while the GKP stabilizers are phase-space displacements — different Hilbert spaces, structurally commuting. Under a full lift the Z_4 chamber is *more* selective than binary, so the qudit topological offset is correspondingly larger (CONDITIONAL on the lift) — the exact offset at (12,3) is g(Δ) = 18.305 (§H.3.4, §H.9.8).

H.1.6 Cross-Kerr robustness — the corrected curve and the load-bearing mitigation

Two GKP modes sharing a transmon/bus accumulate residual cross-Kerr coupling H = 2πχ_K n_a⊗n_b. An exact two-mode QuTiP evolution (SIMULATED): a real 15 dB GKP |+⟩_L carries ⟨n⟩ ≈ 6.85 photons/mode, so unmitigated two-mode fidelity collapses to **F ≈ 0.100 / 0.077 / 0.053 at χ_K = 10/50/100 kHz** — **≈10× more damage than the prior flat-Fock (n = 0..3) model predicted** (the hardware budget must be sized against this corrected curve). **CPMG dynamical decoupling is the load-bearing mitigation and it works:** CPMG-8 recovers **F ≈ 0.996, essentially χ_K-independent** (SIMULATED); CPMG-64 is marginally lower (0.9685) from accumulated π-pulse infidelity, so ~8 echoes is the operating point. Stark detuning alone is insufficient at the worst case (F ≈ 0.13 at 100 kHz; SIMULATED) and is an adjunct only. The mitigated fidelity is well inside the inner-code budget, with CPMG π-pulse infidelity the named new limiter.

H.1.7 Q1 claims anchor, enablement, and falsifiers (compressed)

The Q1/Q2 method claims (**CLM-Q1Q2**; §J) recite: the coset-resonant mode selection (the squashing-asymmetric root pair carrying the non-zero signed-area invariant); the d = 4 encoding with Z_4 Heisenberg–Weyl commutation ω = i on the 8π lattice; modular-quadrature correction; and the Z_4-lifted chamber projection. Enablement: binary GKP break-even and 15 dB stabilization MEASURED (Campagne-Ibarcq 2020; Sivak 2023); ququart break-even MEASURED (Brock 2025); the Z_4 algebra/capacity SIMULATED to machine precision; the three FT-scale build items OPEN and scoped. Falsifiers: the erfc-law tracking test; full ququart *process* fidelity ≥0.95 (OPEN); the chamber-lift gate F3 (PASS → A / partial → B / fail → C); qudit lattice surgery (F4, OPEN); the negative-polarity mode falsifier (a wrong non-x_1 coset generator pair must FAIL GKP-lattice closure, area ≈1.45·2π ≠ 2π). Figures: FIG33, FIG34, FIG03.

H.2 INVENTION Q2 — COSET CHAMBER PROJECTOR, NULL-SPACE FIREWALL, AND THE HIGHER-SU(N) EXTENSION

H.2.1 Field and problem

Quantum error correction; magic-state preparation replacing distillation; syndrome filtering by homogeneous-space holonomy. Conventional stacks pay a dominant magic-state-factory overhead (Litinski-style distillation

cascades); here the factory line is the dominant single-module term, and admissibility was previously used only as a one-shot Boolean gate. Two facts create the opportunity: (1) the full chamber structure (finite chamber + projector + anti-fitting firewall null space) can be made an explicit **runtime projector** rejecting structurally-impossible syndromes *as a class* before the decoder and supplying a geometric T-state resource; (2) at the K_6 = SU(3)/U(1)² coset the chamber's selectivity is structurally capped, and the one unexploited structural variable is the **coset rank** — the move to SU(4)/U(1)³ and SU(5)/U(1)⁴ flag manifolds.

H.2.2 The chamber/projector/null-space formalism (DERIVED; reduced to practice in simulation at K_6)

A **chamber** is a finite-data connected component of the admissibility manifold. The **chamber projector** P_chamber: H_full → H_chamber satisfies (i) P² = P, (ii) P† = P, (iii) rank P = dim H_chamber (rank 3 at the K_6 Einstein point — the fundamental-3 multiplet), (iv) [P_chamber, fiber action] = 0. The **null-space firewall** N_chamber = ker P_chamber rejects, as a class and prior to decoding, every syndrome whose pre-image lies in the kernel — the operational core of the invention and the object the Q4 joint decoder co-measures (§H.4). The projector is built on the canonical spin-c (1,0) holonomy operator $U_\Phi = \mathrm{Pexp}(\oint_\gamma \omega^{(1,0)})$ with eigenstructure $(1, e^{i\Phi}, e^{-i\Phi})$ defining the spectral gap Δ; P_chamber is the spectral projector onto the $e^{i\Phi}$ eigenspace (DERIVED; binary form reduced to practice as `layer7_admissibility_chamber.py`, SIMULATED).

Flag-manifold tangent decomposition and the curvature admissibility body (DERIVED; ported from the geometry source of truth). The substrate's reductive complement splits into three real 2-planes, one per positive-root pair of SU(3): $\mathbf{m} = \mathbf{m}_1 \oplus \mathbf{m}_2 \oplus \mathbf{m}_3$, dim m_i = 2, Σ dim m_i = 6, with simple roots $\alpha_1 = (1,-1,0)$, $\alpha_2 = (0,1,-1)$, $\alpha_1+\alpha_2 = (1,0,-1)$ and the **cyclic bracket relations** $[m_1,m_2] \subseteq m_3$, $[m_2,m_3] \subseteq m_1$, $[m_3,m_1] \subseteq m_2$ (DERIVED) — the algebraic source of the three-way coupling the chamber and the §H.9.1a compiler exploit (internal block index ((0,1),(2,3),(4,5))). The squashed metric $g_K(R,x_1,x_2,x_3) = R^2 \cdot \mathrm{diag}(x_1,x_1,x_2,x_2,x_3,x_3)$ carries four moduli; the admissible region is the curvature-positive set, made transparent by the **factored Ricci formula** Ric_k = (x_k − x_i + x_j)(x_k + x_i − x_j)/(2 R² x_i x_j x_k) over any cyclic (i,j,k) (DERIVED). Positivity of all three components is equivalent to a **strict triangle inequality** on the squashing parameters: Ric_k > 0 ∀k ⟺ x_k < x_i + x_j ∀ cyclic ⟺ |log(x_i/x_j)| < ln 2 ≈ 0.6931 ∀ pairs ⟺ the squashing-ratio bound max_i(x_i)/min_i(x_i) < 2.0 (DERIVED). The sectional curvatures are K_ii = 1/(R²x_i) and K_ij = (1/4R²) [x_k/(x_i x_j) − x_i/(x_j x_k) − x_j/(x_i x_k)] (DERIVED); at the Einstein point $x_1=x_2=x_3=1$ these reduce to intra-plane K_ii = 1/R², inter-plane K_ij = −1/(4R²), Einstein constant λ = 1/(2R²), scalar curvature 3/R² (DERIVED). This triangle-inequality admissibility body is gate G0 of the §H.9.1a compiler and the load-bearing object of the §H.9.4a expandability argument.

Geometric magic-state lift (DERIVED). Φ = −0.12827 rad is pinned by the signed area A = +0.07139 σ² (Φ = k·A, k = −1.797 rad/σ²). Chamber depth **d_chamber = 6 is forced, not chosen**: the continued-fraction convergents of π/(4|Φ|) = 6.123 are 6, 49, 104, …; d = 6 is the unique practical T-gate approximant (accumulated phase 0.76962 rad; residual 0.01578 rad; one local Clifford correction). This **replaces magic-state distillation** with a geometric factory (projected 10K saving ≈100K physical qubits — PROJECTED). The obvious knobs do not help: d = 12 gives the S-gate phase; parallel chambers raise the floor; |Φ| is a topological invariant locked to the calibration basin and cannot be raised without breaking the Standard-Model flavor fit (§H.3.2). The only structural lever is coset rank — motivating §H.2.4.

H.2.3 Substrate enablement framing (theory-side; verification-gated)

The K_6 substrate on which P_chamber is defined is not a calibration-locked fiat: the breathing modulus is **stabilized at decision grade by a derived potential with exact coefficients** — the native scaffold $V(\sigma) = c_{KK}\, e^{-4\sigma} + c_{bdry}\, e^{-2\sigma} + c_{Wilson} \cos\theta_W\, e^{-4\sigma} + c_{loop}\, e^{-6\sigma}$, with $c_{KK} = -8.892\times10^{1}/R_Y^4$ (the **negative sign is the load-bearing fact**), $c_{bdry} = -1.08\times10^{-2}$ (AC-1.9), Wilson exact integer sums S(0) = +4 / S(π) = +92, and the well existing **unconditionally in the loop term** (theory-side, decision-grade, countersigned 2026-06-12).

This is consumed here **only** under the verification-gated framing of §H.9.7 and Appendix §K.App-B (V0–V4 trail; zero auto-promotions), with the §H.0 moduli-point disclosure carried in full.

H.2.4 The higher-SU(N) coset extension (DERIVED methodology; specific values per gate G4)

SU(N) is rank N−1; the flag manifold K_{2N} = SU(N)/$U(1)^{N-1}$ has real dimension N(N−1) carried by the N(N−1)/2 root-pair modules. For SU(4): dimension 12, three Cartan generators, six root pairs, three simple roots → three independent holonomy phases Φ_1, Φ_2, Φ_3. An SU(N)-invariant metric is fixed by one squashing parameter per root module; the Einstein condition is computed by the Wang–Ziller / Borel–Hirzebruch apparatus. At an Einstein point the canonical spin-c (1,0) holonomy on each simple-root cycle is pinned by the first Chern class (Chern–Weil), making the phases topological invariants stable under deformation within the Einstein class (DERIVED, form).

Gate G4 result — the Einstein-metric layer is CLOSED (`GATE_CLOSURE/g4_su4_einstein/`). A first-principles symbolic computation (sympy/numpy; K_6 control reproduced at d = 6, residual 0.01578 rad) derived the A_3 root system, the 12-dimensional isotropy decomposition (Gram = identity to $<10^{-9}$), the exact structure-constant triples (exactly four, each weight 0.25 — from A_3 root-length equality), the Wang–Ziller Ricci tensor (validated three ways), and **four explicit Einstein metrics on SU(4)/U(1)³, all moduli positive** (DERIVED):

Metric	(x_1, x_2, x_3)	Type
1	(1/3, 2/3, 1) ∝ (1, 2, 3)	Kähler–Einstein (∝ root height = ∝ c_1)
2	(3/4, 3/4, 1)	Squashed (non-Kähler) Einstein
3	(1, 1, 1)	Symmetric / normal Einstein
4	((3+√57)/8, (−15+3√57)/8, 1) ≈ (1.319, 0.956, 1)	Squashed (non-Kähler) Einstein

Falsifier F_SU4_1 ("≥2 Einstein metrics, ≥1 squashed, all moduli positive") is **PASS — explicitly computed**. The holonomy *topology* is also rigorous: three independent phases, equal at the Kähler point and splitting only under asymmetric squashing, with integer Chern pairing ⟨c_1, $\alpha_k^\vee$⟩ = 2 for every simple coroot (DERIVED).

Gap-A disclosure (the two-formula discrepancy; named resolver). Two mutually inconsistent expressions exist for the absolute symmetric-point holonomy angle — $\Phi_k = -2\pi \cdot c_1(k)/N$ ($3\pi/2$ at SU(4)) versus $|\Phi_k| = (2/N)\cdot\pi/(N-1)$ ($\pi/6 \approx 0.524$ rad) — which cannot both be correct, and neither is yet re-derived from the connection 1-form (OPEN). This filing **discloses the discrepancy rather than selecting silently**; the named resolver is the **staged re-derivation of the canonical spin-c (1,0) connection and its loop integral** (the same lane that produced the native scaffold of §H.2.3). The squashed-metric phase magnitude (Gap-B) is likewise OPEN: two estimation paths disagree by ≈5× (0.064 vs 0.346 rad), so F_SU4_5 is INDETERMINATE this pass.

H.2.5 Selectivity — the two senses disambiguated and re-anchored to exact counts (RULE-3 controlling)

The historical selectivity escalator (K_6 ≈ 8–15%, K_{12} ≈ 25–50%, K_{20} ≈ 50–75% — ESTIMATED, historical background only) conflated two distinct quantities, disambiguated everywhere they appear (FIG23 caption, NF-6):

1. **Admitted fraction** — the fraction of syndrome space the chamber passes. Gate G4's weight-multiplicity count gives the closed form **1/(N+1)**: **25% at K_6 → 20% at K_{12} → 16.7% at K_{20} — DECREASING with rank** (DERIVED, G4). The historical 25–50% band is **not supported** under this reading.
2. **Firewall rejection power** — the fraction the null-space firewall blocks pre-decoder (the operative quantity for decoder-workload reduction): **N/(N+1) = 75% → 80% → 83.3% — rising with rank** (DERIVED, G4). The escalator survives **only** under this firewall reading.

Claim re-anchor (controlling for CLM-Q9 / CLM-Q1Q2). No claim recites the historical bands. Claim language recites chamber selectivity as *"determined by exact Kostant weight-multiplicity counts at the recited coset rank"* (run S2, §H.9.8), the two senses recited separately where material. The historical bands appear only as ESTIMATED background with the G4 correction adjacent.

The chamber multiplier is Gap-C-contingent (always flagged). The K_{12} chamber multiplier 0.375 (and the 0.30 lo-bound used in the §H.9 composition) is **NOT reproduced by bare geometry**: both bare-geometry selectivity definitions give a K_{12} multiplier ≥1.1–1.5 — *worse* than the K_6 value 0.60 — and reaching sub-0.60 requires an **OPEN decoder-workload / chamber-capacity credit (named Gap-C)** mapping the rising rejection fraction and the larger 12-dimensional carrier capacity to a floor multiplier. That credit is plausible but un-derived; every use of 0.375/0.30 carries this flag (the §H.9.5 sensitivity quantifies the adverse-resolution floor). The exact Kostant multiplicity counts (run S2, §H.9.8) are the machinery Gap-C lacked.

H.2.6 The Kostant weight-multiplicity generator (claimable tooling; the named derivation apparatus)

The enabling derivation apparatus behind the re-anchored chamber-multiplicity and selectivity claims is an **exact, integer-gated Kostant weight-multiplicity generator**: a spectral-certification tool computing exact isotypic multiplicities and spectral towers on the coset substrate to eigenvalue $\lambda \approx 4052$, across scalar, vector, Dirac, and tensor channels, including the explicit **63-degree-of-freedom graviton-internal mode table** (integer-gated — every multiplicity an exact integer, not an estimate). It serves three functions: (i) the **named derivation engine** behind the exact-Kostant-counts language of §H.2.5 and the CLM-Q9 selectivity dependents; (ii) executor of the exactification sweep — run **S1** (exact $\Delta = 0.2589200$ and N_sites,gap, yielding $g(\Delta) = 18.305$, settling the §H.3.5 chamber-site-count question) and run **S2** (exact K_6/K_{12} chamber multiplicity counts, re-anchoring Gap-C and the Q9 factory step); (iii) it extends the AEGIS $\lambda_1 \leq 0.01$ certification lineage (§H.10.2) from noise kernels to substrate spectra, drafted as dependent method claim **CLM-KOST-1** (§J). Its QC-side consumptions are the named S1/S2 runs, which have LANDED (NF-4 charts the dependency).

H.2.7 Q2 claims anchor, enablement, and falsifiers (compressed)

The Q2 system claim (within **CLM-Q1Q2**; §J) recites the syndrome source; the self-adjoint idempotent coset projector onto the holonomy eigenspace at an Einstein metric; the null-space firewall rejecting kernel-class syndromes prior to decoding; and selectivity as a structural function of coset rank via exact Kostant counts. Enablement: the K_6 binary projector reduced to practice in simulation; the SU(4) extension enabled to the Einstein-metric layer by G4 (four explicit metrics), with the staged Phase-1/2/3 program (24–36 months) disclosed with go/no-go falsifiers. Falsifiers: F_SU4_1 **PASS** (G4); F_SU4_2 PASS; F_SU4_3 **WEAKENED** under the admitted-fraction reading (20% < 25%), supported only under the firewall reading; F_SU4_4 **NOT SUPPORTED by bare geometry** (Gap-C); F_SU4_5 INDETERMINATE (Gap-B); plus the topological-pinning perturbation test and the Cartan-misalignment control. Figures: NF-6, FIG23, NF-4.

H.3 INVENTION Q3 — THE Φ↔Δ TOPOLOGICAL ADDITIVE OFFSET

H.3.1 Field and problem

Fault-tolerant resource estimation; geometry-derived code design on a compact-Lie-group coset substrate. Conventional surface-code resource estimation expresses the physical-to-logical ratio as a product of code distances and per-operation error multipliers, $N_{phys}/N_{log} = 2 \cdot d_x \cdot d_z \cdot \gamma_{ro} \cdot \gamma_{ctrl}$. A spectral protection the substrate provides *structurally* — a holonomy gap Δ on the coset fiber suppressing leakage out of the protected logical sector — has no symbol in that formula and is silently absorbed into the upstream choice of d_z. Q3 makes a fixed topological spectral gap reduce the qubit budget **directly and auditably**, without treating the gap as a tuning knob and without double-counting it against code distance.

H.3.2 The Φ↔Δ identity and why Δ is a floor, not a knob

From the canonical geometry source of truth: let Φ be the holonomy phase of the canonical spin-c connection around the generation 1-cycle of the toric base of K_6, and Δ the minimum eigenvalue separation of the holonomy operator on the fundamental-3 representation at the Einstein point. Then **Δ = 2|sin(Φ/2)|, with Δ = 0 ⇔ Φ = 0 ⇔ A = 0** (DERIVED), A being the signed area of the generation triangle (Φ = k·A). At the frozen moduli (0.880, 0.750, 1.087, 1.087): A = +0.07139 σ², k = −1.797 rad/σ², **Φ = −0.12827 rad**, **Δ = 0.2589200** (prose "≈0.259") — the exact value of record from run S1 (§H.9.8).

The calibration lock. Φ is the same holonomy phase that calibrates the Standard-Model CP/flavor sector (Jarlskog/CKM); raising A to enlarge Δ would break that calibration (DERIVED no-go: Δ ≤ 0.259 at this basin). Δ is therefore a **topological floor the design may exploit but may not increase** — the constraint that forces the additive entry mechanism. (Substrate-stability framing per §H.2.3 applies, consumed only through the §H.9.7 trail.)

H.3.3 The budget formula and why ADDITIVE

> **N_phys/N_log = 2·d_x·d_z·γ_ro·γ_ctrl − g(Δ)**

Three reasons force the additive (not multiplicative-on-distance) entry (DERIVED): (1) Δ does not increase logical distance — distance is a graph property of the stabilizer code, Δ a spectral property of the fiber; trading Δ against d_z would weaken iid-Pauli protection for Δ-decoupled residual channels. (2) Because A is independent of qubit count, the suppression applies as a constant subtracted from overhead, independent of (d_x, d_z) — the architecture *partially* exits the iid-Pauli class: bulk overhead stays polynomial in distance, the floor sits below the iid value by a Δ-determined offset. (3) The limits are auditable: Δ → 0 gives g(0) = 0 and the budget collapses **exactly** to the prior-art formula (falsifier F1); Δ → ∞ saturates at the gap-protected fraction of the spatial overhead.

H.3.4 Deriving g(Δ), with its exact landed inputs

With N_sites,gap the count of logical-multiplet sites inside the Φ↔Δ-protected chamber, N_log the logical count, and σ_eff the effective logical-noise scale:

> **g(Δ) = (N_sites,gap/N_log) · 2·d_x·d_z · (1 − $e^{-\Delta/\sigma_{eff}}$)**

At the frozen point and the (12,3) geometry: Δ/σ_eff = 0.2589200/σ_eff; the chamber-depth fraction N_sites,gap/N_log and σ_eff are now exact counts from run S1 (§H.9.8). The exactified result is **g(Δ) = 18.305 physical per logical at (12,3)** (DERIVED — exact closed-form arithmetic from the Kostant-exact N_sites,gap count). This constant is the additive term in every engine of record (`gkp_base = 2·d_x·d_z − g(Δ)`) and sets the d_x bistability boundary; the companion sensitivities are ∂R/∂m = +0.30 and ∂R/∂η_link ≈ −0.27.

Applied to the compiler baselines, the offset projects the floor into the **scale-invariant ≈9.6×–14.5× band** (SIMULATED; (8,3) at the ≈9.6× frontier, (12,3) at the ≈14.5× anchor), with the conservative ≈28× ceiling if the chamber credit fails its lo-bound. The α_em ≈ 0.37 EM-external channel erosion is carried because the spin-c gap does not cover EM-external channels.

H.3.5 The g(Δ) calibration — settled by the landed run S1

The §H.3.4 formula contains the factor 2·d_x·d_z; the load-bearing question was whether g(Δ) counts *chamber* sites (geometry-independent in (d_x, d_z)) or scales with the surface factor. Run **S1** (§H.9.8), the Kostant generator (§H.2.6) computing the exact protected-sector mode count N_sites,gap at the frozen moduli, confirms the chamber-site-count reading: **g(Δ) = 18.305 at (12,3)** (DERIVED, exact) — a structural floor counting

chamber sites, not surface sites. Sensitivity is asymmetric and benign (the floor moves only marginally per unit of Δ drift), so the exact Δ = 0.2589200 settles the spectral input too. NF-4 and NF-7 chart the dependency.

H.3.6 Novelty, claims anchor, enablement, falsifiers (compressed)

No prior-art surface-code resource formula contains an additive spectral-gap term (Fowler et al. 2012 is multiplicative throughout; Kitaev 2003's topological gap is a property of the anyonic medium, not a budget term; the corpus's own geometry article states the Φ↔Δ identity but declines to use it in the budget). Novel steps: (1) the holonomy gap enters the budget **additively**, the form forced by and inseparable from the calibration lock; (2) the gap is pinned to a Standard-Model flavor invariant, converting a would-be fit parameter into a constraint boundary (no QEC antecedent); (3) the $\Delta \to 0$ / $\Delta \to \infty$ limits double as falsifiers, so the term is auditable. Claims (within **CLM-Q9** dependents; §J): the $\Phi \to \Delta \to g(\Delta) \to$ offset-budget method; the calibration-lock dependent; the additivity-invariance dependent ($\Delta \to 0$ recovers prior art identically); the operating-point dependent (Δ = 0.2589200, $g(\Delta)$ = 18.305 at (12,3)). Enablement is closed-form arithmetic from the moduli, k, and the exact chamber count; the runtime-engagement debt (no runtime tensor evaluates $g(\Delta)$ at apply time) is disclosed. Falsifiers: F1 ($\Delta = 0$ reduction — exact recovery of prior art); F2 ($\sigma_{eff} = 10 \cdot \Delta$ must raise overhead by the computed amount, else a bug); F3 (any proposal raising Δ above 0.259 must trigger the SM no-go). Figure: NF-7.

H.4 INVENTION Q4 — CROSS-LAYER CANCELLATION (THREE-LAYER JOINT DECODER)

H.4.1 Field, problem, and the honesty notice that governs every number

Fault-tolerant decoding of composed heterogeneous codes — bosonic GKP inner + geometric coset chamber + topological XZZX outer + modular — conventionally decoded layer-by-layer and charged serially, over-counting whenever two layers physically perform the same operation. Two redundancy classes: **algebraic** (one joint measurement reports two layers' syndromes — well-posed iff the operators commute) and **operational** (one physical ancilla time-shared across adjacent timeslots — no commutator required). The invention detects correlated error events spanning ≥2 layers — GKP displacement errors whose syndrome contribution lies in the chamber null-space — and cancels their contribution before the outer matching decoder, while eliminating double-charged ancilla cycles.

> **LOAD-BEARING HONESTY NOTICE.** Of the five embodiment paths, **only Path 4 (Litinski ancilla sharing, −0.08 single-module, SIMULATED) is currently validated.** Paths 1 and 2 are algebraically disabled under the present SU(3)/K_6 chamber (F1/F2 BOTH_FAIL, §H.4.4), conditionally unblocked in construction terms by gate G3 (§H.4.5) but **contingent on gate g4 physically realizing the disjoint-factor chamber**. The aggregates **−0.46 ("aggressive realistic") and −0.59 ("theoretical maximum") are NOT ACHIEVABLE today and never quoted as achievable**; they appear only to bound and fence off the conditional upside. Paths 3 and 5 are OPEN.

H.4.2 The joint syndrome hypergraph and the cancellation pre-pass

Standard surface decoding builds a matching graph of syndrome defects. This invention augments it to a **joint syndrome hypergraph** H = (V_H, E_H, w): vertices are the union of GKP displacement-parity bits, chamber-admission bits, and outer stabilizer bits; each hyperedge connects the vertex set a single physical error mechanism lights up — a GKP displacement D(α) that flips a GKP parity bit, projects into N_chamber, and perturbs outer defects is **one hyperedge spanning three layers**, weighted by −log p. The **cancellation step**: before the outer decoder, hyperedges entirely explained by the layer-spanning class (GKP displacement ∈ N_chamber) have their syndrome contribution removed, so the outer MWPM/union-find decoder sees only the residual $\sigma_{outer} \ominus \hat{\sigma}_{induced}$ at effective distance $d_{eff} \leq d_{outer}$. Two senses of "reduction" are labeled distinctly: (a) **overhead reduction by scheduling** (Path 4, operative, no commutator required) and (b) **effective-**

distance reduction (Paths 1–2, conditional on the §H.4.4 algebra). No reduction number is asserted in the broad claim.

H.4.3 The narrowed novelty kernel (controlling claim language; prior art distinguished)

The closest prior art is two-layer bosonic-plus-topological soft decoding, **already taught** and distinguished: Vuillot et al. 2019 (PRA **99**, 032344); Noh & Chamberland 2020 (PRA **101**, 012316); Poulin 2006 (PRL **97**, 230501) / Poulin–Chung 2008 (message-passing on concatenated codes); Dennis et al. 2002 (MWPM primitive); Litinski 2019 / Haah et al. 2017 (Path-4 scheduling ground); GKP 2001 / Campagne-Ibarcq 2020 / Brock 2025 (inner-code context). The surviving kernel claimed here (**CLM-Q4-N1**, system; **CLM-Q4-N9**, method; §J):

- **(a) The geometric-coset chamber as a distinct THIRD layer**, interposed between inner and outer, whose **null-space N_chamber is co-measured as an object of the joint hypergraph** and used to cancel layer-spanning correlated mass before the outer decoder — with an **express claim disclaimer**: a decoder lacking a co-measured coset-chamber null-space is outside literal scope ("consisting of three code layers" excludes two-layer concatenations).
- **(b) PHYSICAL (not classical) ancilla sharing** (**CLM-Q4-N2**): time-sharing a physical measurement qubit across heterogeneous layers, as distinct from sharing classical soft information (what Vuillot/Noh/Poulin do).

Dependent chain: **CLM-Q4-N3** — the Path-4 Litinski scheduling embodiment (final distillation round coincides with an outer stabilizer cycle on a shared ancilla; expressly requiring **no** commutativity condition; $\Delta = -0.08$, SIMULATED; scheduler re-confirmation a named pre-filing gate) — **the strongest currently-enforceable claim**; **CLM-Q4-N4** — the cancellation pre-pass with d_eff ≤ d_outer, no numeric benefit recited; **CLM-Q4-N5** — the correlation-cost guard (a non-correlated/non-commuting configuration is *required to yield zero reduction*, foreclosing accounting-artifact attacks).

H.4.4 The computed commutativity gate and the BOTH_FAIL result (the load-bearing negative)

Paths 1–2 are defined by algebraic preconditions: **F1**: [S_q, P_chamber] = 0 and [S_p, P_chamber] = 0; **F2**: [P_chamber, S_k] = 0 for all eight d_outer = 3 XZZX generators. Commutativity is the correct gate: a non-disturbing simultaneous measurement exists iff the observables commute (iff the operator is block-diagonal in the chamber/null split — off-block elements are the obstruction). The preconditions were **computed, not assumed** (deterministic, 17/17 + 18/18 unit tests PASS; SIMULATED):

Precondition	Residual (Frobenius; PASS < 1×10^{-10})	Result
F1: [S_q, P_chamber] / [S_p, P_chamber]	1.181 / 1.324	**FAIL**
F2: [P_chamber, S_k], k = 1…8	1.001 – 2.055	**FAIL**

The residuals sit 9–10 orders above threshold — algebraic non-commutativity, not numerical noise. Root cause: the SU(3) projector's eigenspaces are misaligned with both the GKP bosonic block and the qubit-pair decomposition (off-block mass 0.575). A basis-aligned recheck proved the failure **unitarily invariant** ($\|[U\dagger AU, U\dagger BU]\|_F = \|[A,B]\|_F$): under SU(3)/$K_6$ this is a property of the operators, and reopening Paths 1–2 requires changing the **operators**, never a basis or parameter.

H.4.5 Gate G3 inlined — the conditional unblock (construction-dependent, NOT universal; NOT an achieved unblock)

Gate G3 (`GATE_CLOSURE/g3_f1f2_su4/`; deterministic, seed = 42; SIMULATED) re-ran the F1/F2 battery under an SU(4) chamber in three constructions, GKP/surface embeddings byte-identical to the SU(3) test:

Construction	What it tests	F1 max	F2 max	Verdict
SU(3) baseline (reference)	the original chamber	1.324	2.055	BOTH_FAIL
A — naive SU(4) rank lift	swap the group, same recipe	1.513	2.362	**BOTH_FAIL** (off-block mass 0.698 — *larger* than SU(3))
B — native Z_4 Haar projector $P_{clock} \otimes I_{code}$	the spec's claimed mechanism (§H.1.5 Path 2)	**0.000**	**0.000**	**BOTH_PASS** (machine zero, all of F1 + all 8 of F2)
B-counter — disjoint premise deliberately broken	entangle clock into code factor	2.628	3.744	BOTH_FAIL (the falsifier respected)

Three facts compose into the honest verdict: (1) **the obstruction is not the group rank** — SU(3) → SU(4) alone does nothing; (2) **the obstruction is shared-ambient block-misalignment**, vanishing identically when chamber and code operators live on **disjoint tensor factors** (chamber = Z_4 toric-center projector; code = GKP/surface), with the genuine SU(4) holonomy commuting with the projector ($\|[P_{clock}, U_\Phi]\|_F = 0.000$) — confirming the §H.1.5 Path-2 mechanism in real arithmetic; (3) **the PASS is structural, not nominal** — entangling the factors restores failure. **Therefore Paths 1+2 unblock if and only if the physical SU(4) chamber is realized as a disjoint-factor (toric-center) projector.** This is **SUGGESTIVE, not definitive** (a PASS-by-construction on a model, not the squashed-Einstein chamber): whether the physical g4 chamber delivers that factorization for real finite-energy GKP states and real XZZX qubits **remains OPEN**. Consequences: the −0.46/−0.59 aggregates are no longer behind a basis-invariant no-go but **remain NOT ACHIEVABLE today**; Path-4 −0.08 (SIMULATED) remains the only validated reduction; **CLM-Q4-N6/N7** (Paths 1/2) recite their commutator conditions as express claim contingencies, and **CLM-Q4-N8** (Path 3) recites its dual hardware contingency (≥20 dB squeezing AND $\eta_{link} \geq 0.95$ in one transducer cavity; OPEN). NF-8 renders the status panel.

H.4.6 The correlation-cost debit and the path table (with status)

Joint extraction can *introduce* correlations that, unmodeled, inflate the outer logical error rate; the per-path Δ values are **gross** figures, and net benefit requires the named Phase-2 net-of-correlation simulator with per-path falsifiability tests (a broken configuration must yield zero savings) — OPEN engineering, disclosed. The five paths:

Path	Mechanism	Single-module Δ	Status
1	Chamber syndrome co-measured with GKP parity (one ancilla)	−0.17 (NOT ACHIEVABLE today)	**DISABLED under SU(3) (F1 FAIL); CONDITIONAL on the g4 disjoint-factor chamber (G3)**
2	Outer stabilizer stream carries the chamber test	−0.16 (NOT ACHIEVABLE today)	**DISABLED under SU(3) (F2 FAIL); CONDITIONAL as Path 1**
3	Bell pairs born GKP-encoded in a shared transducer cavity	−0.13 at 10K	OPEN (≥20 dB ∧ $\eta_{link} \geq 0.95$; falsified if the transduction optimum sits below GKP break-even)
4	**Distillation final round IS an outer stabilizer cycle (Litinski reuse)**	**−0.08**	**SIMULATED — the only validated path; scheduler re-confirmation is a named pre-filing gate**
5	Boundary ancilla serves intra- and inter-module roles in one event	−0.05	OPEN (protocol co-design)

Authoritative aggregate today: **Δ = −0.08 (Path 4 only)** (SIMULATED). The cross-layer cancellation is a contributing reduction within the two-stage floor; the floor band itself is the §H.9 system-composition result and is never attributed to this invention alone.

H.5 INVENTION Q5 — MODULAR OVERHEAD REDUCTION (5-ROUND BBPSSW + REAL-TIME FIDELITY POST-SELECTION)

H.5.1 Field, problem, and the dependency-chain honesty

Modular (multi-cryostat) fault-tolerant architectures; inter-module entanglement distribution. The fabric partitions ~10,050 logical qubits across ~150 modules (~67 logical each) connected by heralded photonic Bell pairs consumed in gate teleportation / lattice surgery; the inter-module fabric imposes a multiplicative modular overhead ≈1.075–1.08 (SIMULATED baseline). Q5 lowers that toward 1.05 without spending physical qubits. Q5 is independent of the Q6/Q7 cavity chain but is **rate-limited** — its binding risk is throughput, not fidelity — and **the ≈4.3–5× rate deficit of §H.5.4 travels with every appearance of the 1.05 multiplier**, including the §H.9 lever table.

H.5.2 Mechanism — five-round recurrence plus a classical acceptance gate

(a) Five-round BBPSSW. Bell pairs heralded by Barrett–Kok interference ($F_{heralded} = 1 - (1-\eta)^2 = 0.9936$ at $\eta = 0.92$ — DERIVED) are purified by the 2→1 Deutsch/BBPSSW recurrence (Bennett et al. 1996, PRA **53**, 2046 / PRL **76**, 722; Deutsch et al. 1996, PRL **77**, 2818): five rounds converge to $F_{distilled} \geq 0.999$ (SIMULATED) at $2^5 = 32$ raw pairs per distilled output — 100 kHz raw → ≈3.125 kHz distilled per link. Locking the full five-round regime removes an idle-ancilla penalty at the link.

(b) Real-time per-pair fidelity post-selection. The five-round output has spread $\sigma_F \approx 5\times10^{-4}$ (DERIVED/modeled). A **classical acceptance gate on the existing heralding record** admits only pairs with estimated $F > 0.9995$ (≈16% rejected at a 0.9997 mean), raising the accepted mean to $F_{accepted} \approx 0.9999$, i.e. **$p_{inter} \approx 1\times10^{-4}$** (DERIVED) — a 5–10× error improvement at **zero added physical qubits**, with transient link degradation auto-rejecting rather than propagating.

(c) Net overhead effect. The post-selected p_{inter} relaxes the boundary lattice-surgery error-detection budget (the dominant reducible term), taking the modular multiplier **1.075 → ≈1.05** (DERIVED/PROJECTED). The floor is above 1.00 — distance-12 lattice surgery requires ~10 ancilla rounds minimum (Fowler 2012; Litinski 2019) — so 1.05 is a defended engineering point, not an asymptote. $\eta_{link} \geq 0.95$ is a **hard deployment gate** ($\partial R/\partial\eta \approx -0.27$ at baseline); below 0.95 the method falls back to full-depth boundary correction (recited as a dependent claim).

H.5.3 Novelty and claims anchor (compressed)

BBPSSW/Deutsch teach recurrence to a fidelity target but no real-time per-pair acceptance gate downstream of distillation and no co-design of the round count against a boundary lattice-surgery budget; Magnard et al. 2020 (PRL **125**, 260502) demonstrates heralded post-selected microwave Bell pairs as a link primitive but not the integrated link-to-overhead method. The claimed combination (within the Q9 family; §J): (i) lock the full five-round recurrence; (ii) apply the $F > 0.9995$ real-time gate on the heralding record; (iii) spend the p_{inter} headroom on the boundary error-detection term, with the $\eta_{link} < 0.95$ fallback recited.

H.5.4 Maturity and the rate-deficit disclosure (HIGH risk, quantified)

Status: DERIVED with HIGH rate-deficit risk. At five rounds plus ≈16% rejection, accepted supply is ≈2.63 kHz per link against ≈67 kHz per-link demand at 10K (aggregate ≈1.97 MHz vs ≈10 MHz) — a structural **≈4.3–5× rate deficit** (DERIVED). **Without a raw-rate uplift (100 kHz → ≈430 kHz per link, a photonic hardware change not demonstrated at scale) Q5 at full five-round depth is rate-infeasible at 10K**: it improves the fidelity margin and the multiplier while degrading throughput below demand. Mitigations (none assumed): asynchronous pool prefetching (partial), the raw-rate uplift, or coded-Bell-pair / deferred-correction boundary methods. Falsifiers: F-Q5-a (five-round fidelity ≥0.999 from 0.9936); F-Q5-b (sustained accepted rate ≥

demand/link — **refutation expected at current hardware**); F-Q5-c (two-module testbed confirming the p_inter → boundary-depth coupling).

H.6 INVENTION Q6 — OUTER-CODE REMOVAL AT 20 dB (THE CROSSOVER METHOD), AS BOUNDED BY GATE G5

H.6.1 Field, problem, and the crossover idea

Concatenated GKP-plus-surface architectures; squeezing-dependent code-layer selection. The outer XZZX code catches residual GKP logical escapes but is not free: each syndrome round runs ~d^2 two-qubit gates on the outer qubits. Q6 asks whether there is a squeezing regime where the outer code stops helping and actively hurts, computing the inner GKP logical error rate P_L(s_dB) and the outer code's gate-injection rate and **conditioning the outer code's presence on which is larger**.

H.6.2 The crossover computation (DERIVED) and what removal yields

Outer-code injection at d_outer = 3 with transmon-grade ε_gate ~ 10^{-4} (MEASURED, Kjaergaard 2020): **ε_outer ≈ $d^2 \cdot$ε_gate ≈ 9×10^{-4} per logical per cycle** (DERIVED). Inner GKP at 20 dB: P_L ≈ 2.46×10^{-36} analytic-ideal (grid convention, §H.1.2) and **~10^{-14} closed-loop** (SIMULATED) — so with the outer code retained P_total ≈ 9×10^{-4}, without it ≈10^{-14}: at 20 dB the outer code makes the logical error **worse by ~10 orders of magnitude** and is an error *source*. Crossover P_L(GKP) = ε_outer at ≈6.5 dB (DERIVED); verdicts: 15 dB — MITIGATING (keep); 18 dB — THINNABLE to d_outer = 2; 20 dB — REMOVABLE in the crossover model, **subject to the G5 bound below**; 22 dB — definitely removable. Removal (d_outer: 3 → 1) leaves a logical qubit as a bare GKP cavity + one syndrome ancilla + a runtime share of the chamber projector, moving the configuration toward the (8,3) X-verified ≈9.6× frontier under ancilla multiplexing (SIMULATED / PROJECTED). Outer-code removal is a contributing reduction within the floor band, conditioned on the G5 squeezing bound; it does not by itself set the floor (the §H.9 system result).

H.6.3 Gate G5 inlined — Q6 is BOUNDED, not blocked (the operative constraint sharpened)

Gate G5 (`GATE_CLOSURE/g5_purcell_window/`; closed-form input-output rate model + the campaign's exact Gaussian squeezing bridge, QuTiP-cross-validated to ratio 1.0000; DERIVED/SIMULATED) characterized the true binding constraint and produced a bounded verdict:

- **Even behind a perfect Purcell filter, sustainable squeezing is reservoir-vs-loss bounded.** At the realistic 5 kHz stabilizer rate (Sivak-2023-class): Q = 1×10^8 sustains **16.1 dB**; **Q = 2×10^8 caps at 17.6 dB** (SIMULATED — both short of 20 dB).
- **Full 20 dB removal (d_outer → 1) requires intrinsic Q ≳ 3×10^8 OR a faster stabilizer Γ_sq/2π ≳ 5.4 kHz at 2×10^8** (BOUNDED; the ~145×-loss-rate sustaining rule). A filter cannot raise Q above intrinsic, so the A+B+C compound (≈2×10^8, §H.7) is *just barely insufficient*; the sapphire-host path (≈4×10^8-class at 6 GHz) is the named route.
- **The de-risked fallback is ~17.6 dB → d_outer = 2 (THIN, not remove)** until 20 dB at Q ≳ 3×10^8 is demonstrated; at ~17.6 dB the crossover still puts P_L far below the 9×10^{-4} injection floor, so thinning retains most of the benefit.
- The one remaining binding step is a 3D EM solve of the filter-in-re-entrant geometry (**gate G5-EM, OPEN**, §H.7.3).

Status summary: the crossover method is DERIVED and conditioned, not absolute — the claims recite the *conditioning method* (outer-code presence decided by the inner-P_L-vs-injection comparison) and the removal embodiment with its squeezing/Q contingencies. No experimental bare-cavity GKP fault tolerance exists (OPEN; current art runs GKP inside an outer code, which also supplies the logical-sector definition, the large-displacement backstop, and the FT gate scaffold — all three must be re-provided at the GKP level under removal,

a named protocol-redesign item). Falsifiers: F-Q6-a (Stim with/without comparison, predicate P_total(removed) < P_total(retained) at ≥20 dB); F-Q6-b (the verdict ladder); F-Q6-c (bare-cavity FT demonstration at P_logical < 10^{-12} — **no such demonstration exists**).

H.7 INVENTION Q7 — THE CAVITY APPARATUS: COMPOUND MATERIALS-AND-GEOMETRY PATH TO Q ≈ 2×10⁸ AT 7–8 GHz WITH PURCELL-FILTER CO-DESIGN

H.7.1 Field, problem, and the honest baseline

Superconducting 3D microwave cavities for bosonic codes; surface two-level-system (TLS) loss engineering; Purcell-filter co-design. Sustaining 20 dB GKP squeezing requires $\kappa_{int}/2\pi < 1$ kHz, i.e. $Q_{int} > 10^8$ at the 7–8 GHz transmon-compatible band. **Q = 10⁸ has been demonstrated only at 2.5 GHz** (MEASURED, Reagor 2016); at 7–8 GHz the best measured is **7.9–9.3×10⁷** (MEASURED), TLS loss scaling unfavorably with frequency. Per §H.6.3 the real Q6 requirement is ≈2×10⁸-and-beyond **while simultaneously** holding $\chi/2\pi \geq 2$ MHz for photon-number-resolved GKP readout — the two compete through the Purcell channel, and that co-design, not the bare Q, is the binding constraint.

H.7.2 The three compounding paths (PROJECTED, with sub-multiplicativity honesty)

- **Path A — niobium vs aluminum: +30%.** Nb's native oxide is thinner/less lossy than Al_2O_3 (Murthy 2022); Nb 3D cavities at $Q > 10^8$ at 5–7 GHz reported (Earnest 2018; Reagor 2016).
- **Path B — BCP/EP + 120 °C nitrogen-infusion bake: +20%** (SRF surface chemistry; Grassellino 2013).
- **Path C — re-entrant mode geometry: +50% (largest single path).** A post-in-cavity mode concentrates the E-field in a sub-mm gap away from lossy walls, structurally suppressing TLS loss ∝ surface-tangential field integral (re-entrant Q = 3.9×10⁸ at 4.5 GHz, Reagor 2013; transmon-compatible re-entrant, Axline 2016).

Paths A and B share the surface-TLS mechanism and compound **sub-multiplicatively** (1.30×1.20 → ~1.45 overlap-corrected); Path C reduces the mode's *overlap with* the surface (orthogonal to A+B) and multiplies cleanly. Compound: A+B ≈ 1.31–1.40×10⁸; **A+B+C ≈ 1.96–2.11×10⁸** (PROJECTED) — lifetimes $\tau_{cav} = Q/2\pi f$ ≈ 3.0–4.5 ms at 7.4 GHz (DERIVED). Secondary paths (sapphire-host ≈4×10⁸-class, the named Q6 full-removal route) are disclosed but not relied on. Above ≈2×10⁸ the dominant loss shifts to seam loss and vortex trapping, which Paths A–C do not address (disclosed limit).

H.7.3 The Purcell co-design — corrected arithmetic and the characterized filter window

PNR GKP readout at 20 dB requires $\chi/2\pi \geq 2$ MHz (vs 1.4 MHz at 15 dB — MEASURED, Eickbusch 2022), but χ drives a Purcell decay channel $\kappa_{Purcell}/2\pi = \chi^2/\kappa_{transmon}$ = (2 MHz)²/5 MHz = 0.80 MHz against $\kappa_{int}/2\pi$ = 74 Hz at $Q = 10^8$. **Corrected arithmetic (G5; superseding the earlier spec figure):** the unfiltered collapse is to **Q_{eff} ≈ 9.2×10³ (~10⁴)** — the earlier printed "~10²" was ~2 orders too pessimistic; **the conclusion is unchanged — an unfiltered χ = 2 MHz collapses Q by ~4 orders, so the filter is mandatory** (DERIVED). G5 characterized the window as a falsifiable filter specification (DERIVED):

> **A multi-pole (≥2-pole) Purcell filter presenting ≥ 45 dB rejection at the storage-cavity frequency** (≥43 dB for $Q_{eff} \geq 10^8$; ≥48 dB for the full $Q_{eff} \geq 2\times10^8$ — remarkably flat at ~42–48 dB across the feasible Q range), at **~50–100 MHz bandwidth** and **~0.3–0.8 GHz cavity–filter detuning**, with negligible readout-passband insertion loss (band-stop at the cavity, band-pass at the readout). A single-pole filter does **not** open the window (it would need an unphysical 3.7–25 GHz detuning).

The apparatus claim (§J) recites the Nb re-entrant body, the BCP/EP + N-infusion surface, and the transmon coupled **through the characterized multi-pole filter**, exhibiting $Q \geq 10^8$ at 7–8 GHz while maintaining $\chi/2\pi \geq 2$ MHz; dependents recite the ≈1.4–2.1×10⁸ compound band, the 2-pole-filter-in-re-entrant-geometry co-design,

and the ≥45 dB/≥2-pole specification with the corrected ~10^4 unfiltered-collapse figure. **What remains OPEN is the realization, not the specification:** gate **G5-EM** — a 3D EM (HFSS/CST) solve verifying ≥45 dB at f_cav, ≤ −1 dB readout insertion, and Q_int ≥ 10^8 with $\chi/2\pi \geq 2$ MHz simultaneously — is the named pre-build gate. Falsifiers: per-path isolation gates (Q_Nb/Q_Al ≥ 1.25; surface ≥1.15; re-entrant ≥1.35); the compound gate (≥3/5 cavities at Q ≥ 10^8 via A+B); and the decisive 20 dB sufficiency test (GRAPE-prepared 20 dB state on a Q ≥ 3×10^8 filtered cavity, P_L < 10^{-6}/cycle — refutation drives d_outer ≥ 2 per §H.6.3).

H.8 INVENTION Q8 — ORTHOGONAL-SUBSPACE ENCODING (TRANSVERSE-PROJECTION / COMMON-MODE-ORTHOGONAL)

H.8.1 The honest read, stated first

Q8's DERIVED, non-speculative single-module read is WORSE than the program baseline at 67 logical qubits, and that fact leads this section. The conservative read gives **26.6 raw physical per logical → 53.3 EM-external** after the antisymmetric-encoding pairing factor of 2 (DERIVED) — ≈3,570 physical at 67 logical — versus the Stage-B conservative compiler band of **~22–30 physical per logical** (SIMULATED; the aggressive Stage-B figure is under restatement per the compiler-integrity audit and is not quoted). The pairing-factor-2 dominates the d_z saving at this operating point. Q8's favorable figure — **≈11.4 physical per logical** — requires **Refinement A, which is SPECULATIVE** (the "topological savings" reading in which geometric orthogonalities count as code distance; not derived, not simulated) and is never quoted without that label. The non-speculative residual is Refinement B: an asymptotic advantage above ≈200 logical, where the pairing factor amortizes (DERIVED). Q8 is filed as a method/encoding invention whose mathematical structure is the claimed subject matter, with this overhead honesty on its face.

H.8.2 Field, problem, and the formal kernel

Decoherence-free / noise-tailored encoding; symmetry-orthogonal noise-channel routing. Two load-bearing channels co-dominate and both couple through σ_Z: Z-dephasing (bias η ≈ 4.0 — DERIVED) and readout back-action (p_ro = 5.11×10^{-4} exceeding p_2q = 2.91×10^{-4} under EM back-annotation — SIMULATED). A standard surface code spends d_z qubits *correcting* this channel; Q8 instead **routes** it into a subspace structurally orthogonal to the logical subspace so it produces no logical-channel syndrome, relaxing the Z-distance demand. Governing physics: an isotropic common-mode coupling commutes with the singlet/antisymmetric projector and produces no signal-sector syndrome, while the anisotropic signal term carries no symmetric component (cf. Lidar–Chuang–Whaley 1998; Zanardi–Rasetti 1997). Formal kernel: decompose H_sys into signal subspace S and noise subspace N with criterion **P_S · L_noise · P_S = 0**; encode in S, decode by projecting with P_S; common-mode excursions cancel structurally.

H.8.3 The four composed orthogonalities and the encoding map

The dominant σ_Z channel is mapped into the 13-dimensional geometry $M_{3,1} \times K_6 \times S^2 \times S^1$ and four *independent* orthogonal directions are composed (DERIVED; #1 existing, #2–#4 new — the first apply-time engagement of the $S^2\times S^1$ fiber): (1) transmon-pair antisymmetry (the existing DFS, r_rej ≈ 19.6 — SIMULATED); (2) K_6 Schur-orthogonality on the x_2-x_3 antisymmetric direction; (3) S^1/Z_2 chirality opposition under P_chi = $\gamma_5\Gamma_8$; (4) S^2 angular momentum transverse to the readout-pull axis — the sole direction carrying the logical bit. Logical states: $|0\_L\rangle/|1\_L\rangle$ = (singlet pair) ⊗ (K_6 antisym mode) ⊗ (chirality-opposite S^1 sector) ⊗ ($S^2\_\perp$, ±); P_logical = (P_pair,antisym)·($P_{K_6,antisym}$)·($P_{S^1,chi-opp}$)·(I − $|S^2\_pull\rangle\langle S^2\_pull|$). The noise subspace consumes 5 directions; the signal subspace has dimension 4 (2 logical + 2 ancilla/syndrome). Each orthogonality multiplicatively reduces the effective σ_Z rate feeding d_z (conservative ≈0.3 per orthogonality → ≈0.027 composite), relaxing the Z-distance demand in the Q9 geometry (DERIVED, conditional on the orthogonality predicate). Q8 and the GKP inner Z-suppression are the two named Z-distance-relaxation enablers; the promoted lock is the (12,3) geometry with a minimum defensible d_z = 3.

H.8.4 Novelty, claims anchor, falsifiers (compressed)

Prior art: decoherence-free subspaces (Lidar–Chuang–Whaley 1998; Zanardi–Rasetti 1997) supply orthogonality #1 only, via a single permutation symmetry; biased-noise distance tailoring (Tuckett et al. 2018; XZZX, Bonilla Ataides et al. 2021) tailors *distance* to bias rather than routing the channel out. Novel: the four-fold composed projector spanning pair-antisymmetry + three geometric factors; the geometric instantiation of the transverse-projection criterion as a runtime projector; and Z-distance relaxation by structural orthogonality rather than spent Z-distance. Claims (under the Q9 family; §J) recite the encoding map, the four-fold composition, the Z-distance-relaxation selection, and the symmetry-orthogonal decomposition — with Refinement A's SPECULATIVE status recited in the description, never as an operative benefit. Falsifiers: Test 1 — $\langle i\_L|\sigma\_Z|j\_L\rangle = 0$ exactly, GREEN on the Q8 encoding, RED on the unmodified state, a $\pi/4$ axis-rotation break restoring nonzero elements; Test 2 — $p\_L(\theta)$ must rise monotonically as the S^2 axis is rotated; Test 3 — the exponential-vs-polynomial dephasing sweep ($\geq 10\times$ below the polynomial baseline at $p\_dephase = 5\times10^{-4}$ under Refinement A). Maturity: DERIVED, experimental OPEN; the runtime projector requires separately-authorized engagement; the merger flag with the Q2 chamber is disclosed.

H.9 INVENTION Q9 — THE INTEGRATED STACK (SYSTEM COMPOSITION CLAIM)

H.9.0 HONESTY BANNER (governs all of §H.9 and every downstream use of the floor figure)

> **The defended result is a TWO-STAGE, SCALE-INVARIANT overhead floor — a SIMULATED ≈9.6× to 14.5× physical(bosonic-mode)-per-logical band, substantially independent of the logical-qubit count, with a conservative ≈28× ceiling.** Every figure below is a DESIGN-STAGE / EXTERNAL / SIMULATED projection, not a measured device property. - **Floor band: ≈9.6× to 14.5×** (SIMULATED). Upper anchor ≈**14.5**× = the **(12,3) binary-lever-locked** configuration (the canonical lock; single-module ≈6.6 / fabric ≈6.9 with every accounting-gated credit firing, held by the exact landed $g(\Delta) = 18.305$ and $\Delta = 0.2589200$, §H.3.4–H.3.5). Lower frontier ≈**9.6**× = the **(8,3) X-verified** configuration (DERIVED + EXTRAPOLATED-X; its dependent-geometry caveats travel — binary-lever forfeited, X-axis extrapolated). - **The band is a floor at scale.** Produced by the two-stage construction (Stage 1 restricts to the count-independent regime; Stage 2 topological-searches within it), it is **evaluated at 10,000 logical qubits and reported as nominally unchanged at 20,000**. This count-independence is the principal technical-improvement assertion of the filing. - **Conservative ceiling: ≈28**× (DERIVED) — the position if the chamber 0.60 credit fails its lo-bound (the two weakest credits, M18-Majorana topo + chamber, removed). Carried at all times as the bound, not the lead band. - **The "sub-Kitaev" (single-module < 1.0) framing is RETIRED** as a floor-clamp/algebra artifact; its pre-clamp value is not reproduced and no sub-1.0 figure is claimed. - **The orthogonality caveat (travels with every ratio):** the floor is a qubit-COUNT (memory-overhead) metric; the algorithm error budget is a separate axis. The design **FAILS the Shor-2048 error budget at every scenario** (needs $p\_L < 10^{-11}$; best case $\approx10^{-10}$ — SIMULATED), though the logical clock rate is adequate on paper. A competitive ratio does not buy a runnable cryptographic algorithm; Q9 claims a qubit-overhead composition, not algorithm runnability, and the band is memory-to-memory only. The separately reported decoder-channel advantage ≈8.4× (95% CI 4.2–16.9×) is **not** a memory-overhead ratio and is never conflated with the band.
>
> Q9 is claimed as a **system composition** — the specific four-layer / five-lever architecture on the (12,3) coset geometry — novel regardless of where within the band the overhead is realized; even at the ≈28× ceiling the layered architecture is the claimed invention.

The two-stage construction, decomposed (connective; the structure behind the floor framing). The two-stage construction the banner invokes resolves into three ordered operations (DERIVED methodology): **Stage-1-part-1 — the admissibility gate** (the curvature-positivity / triangle-inequality test of §H.2.2, i.e. gate G0 of the §H.9.1a compiler) restricts the four moduli to the admissible region $\mathcal{A}$; **Stage-1-part-2 — the scale-invariance**

restriction further admits, out of $\mathcal{A}$, only those candidate configurations whose projected physical-per-logical overhead is substantially independent of the logical-qubit count, the count-independence being recorded as a *constructed property of the restricted sub-space* (with a sub-space identifier) before any code is chosen; **Stage-2 — the topological search within** then selects the geometry-derived code (its logical-sector structure, check structure, connectivity, and asymmetric (d_x, d_z)) from inside that restricted sub-space. **Why the order is load-bearing:** because Stage 1 isolates the count-independent regime *before* Stage 2 runs, every configuration the search can reach already carries the floor-at-scale property, so the search selects only *where within* the count-independent band the design lands, not *whether* it is count-independent — which is precisely why the reported overhead is defensible as a floor that holds at scale (evaluated at 10,000 logical qubits, nominally unchanged at larger counts) rather than as a small-scale figure that would climb with the logical-qubit count under conventional constructions.

H.9.1 The locked geometry and the mission framing

Element	Value	Status
Outer code	asymmetric rotated XZZX, **(d_x, d_z) = (12, 3)** (binary-lever-locked)	DERIVED
Substrate	K_6 = SU(3)/U(1)² flag manifold	DERIVED
Frozen moduli (R, x_1, x_2, x_3)	(0.880, 0.750, 1.087, 1.087)	DERIVED (pinned; the iter-4e squashed basin — see the §H.0 moduli-point disclosure)
Holonomy phase / gap	$\Phi = -0.12827$ rad / $\Delta = 0.2589200$ (prose "≈0.259")	DERIVED (exact, landed run S1; pinned to the SM flavor calibration)
Topological offset	g(Δ) = 18.305 at (12,3)	DERIVED (exact, landed run S1)
Fabric	**150 modules × 67 logical = 10,050 logical**	DERIVED (the 10,000+-logical mission shape)

The (12,3) geometry is the promoted lock (minimum defensible d_z = 3; the relaxed Z-distance demand enabled jointly by the GKP inner Z-suppression of §H.1 and the Q8 orthogonal routing of §H.8); the (8,3) X-verified configuration is the lower endpoint. The four layers (FIG03): (1) GKP inner — exponential suppression in $1/\Delta\_GKP^2$, relaxing the Z-distance demand; (2) K_6 chamber — runtime admissibility projector + null-space firewall + geometric T-state factory replacing distillation (§H.2); (3) asymmetric XZZX outer at the (12,3) lock (removable per Q6 at the G5-bounded squeezing point; thinnable to d_outer = 2); (4) photonic modular fabric (Q5), with the ≈4.3–5× link rate deficit of §H.5.4 disclosed wherever its multiplier is used.

H.9.1a The staged geometry-to-code compiler (G0→G6) — the named derivation pipeline

The locked geometry of §H.9.1 is not asserted by fiat: it is the output of an **ordered seven-stage compiler** that maps the four scalar moduli (R, x_1, x_2, x_3) — and *no additional free parameters* — to the error-correction figures of merit and then to the (d_x, d_z) code specification. The pipeline is the single named enablement object behind the geometry-derived code; it is **DERIVED as an algorithm**, with any numeric output it produces carrying the SIMULATED label of the run that produced it. The stage list (DERIVED, ordered — the construction is invalid if reversed; a backend-adapter stage G7 is the only hardware-facing step):

- **G0 — Admissibility gate.** $\mathbb{R}^4\_\{>0\}$ → the admissible region $\mathcal{A}$, i.e. the positive-Ricci subset defined by the §H.2.2 triangle inequality x_k < x_i + x_j (equivalently max_i(x_i)/min_i(x_i) < 2.0), plus a non-zero Kaluza–Klein mass-gap requirement, plus a re-parameterizable search window (R ∈ [0.80, 1.25], x_i ∈ [0.75, 1.35] — explicitly a sampling window, NOT formalism) (DERIVED). A modulus point failing the triangle test is REJECTED as a class.
- **G1 — Tensor-stack construction ("parameter collapse").** The four scalars are expanded into the full internal tensor stack with **no new free parameters introduced** — the load-bearing fact: the metric g_{MN} and inverse, the distance/overlap kernels, the 6×6 Hermitian internal Hamiltonian H_int, the logical projector

P_log, the Berry connection, the rank-4 curvature tensor, the geometric and logical Choi noise superoperators, the leakage operator, plus the canonical invariant tuple (Δ, $\bar{\kappa}$, S_sec/$\hbar$, σ_eff) and the metric-sensitivity tensors $\partial g/\partial \theta_a$, a ∈ {R, x_1, x_2, x_3} (DERIVED). The entire internal Hamiltonian is therefore a function of four numbers.

- **G2 — Holonomy / logical-sector extraction.** The spin-c (1,0) bundle connection on K_6 is parallel-transported around the generation 1-cycle; its holonomy operator U_γ = diag(1, $e^{i\Phi}$, $e^{-i\Phi}$) on the fundamental-3 is diagonalized. Its eigenstructure fixes which modes become **logical qubits** (P_log, the orthogonal projector onto a chosen eigenspace, the eigenpair with maximal separation) versus syndrome ancillas, and the spectral gap $\Delta = min_{j\neq k}|\lambda_j - \lambda_k| = 2|\sin(\Phi/2)| + O(\delta x_i)$ is read off (DERIVED). At the frozen point Δ = 0.2589200 and Φ = −0.12827 rad (the exact landed values of record, run S1, §H.9.8). With the geometry removed (no moduli seed), G2 is **mathematically undefined** — the bundle holonomy is non-existent as an object, so G3–G6 cascade-fail; this is a definitional property of the compiler, not a claim about nature.
- **G3 — Figures of merit (the geometry-derived control quality).** From P_log and the Berry-connection/holonomy curvature, the **noise bias η and the common-mode rejection r_rej are computed as DERIVED outputs, not free knobs** (alongside the control correction γ_ctrl). Conformed to the Tahoe calibration: **η ≈ 3.1 central** (BOUNDED — the package's bias-creation central estimate, distinct from any higher full-bias bound), and r_rej ≈ 18.3–19.3 (SIMULATED), the latter held **strictly distinct from η** — r_rej is a correlated-error suppression ratio, η a Pauli-channel asymmetry, and the two are never conflated.
- **G4 — Noise-channel compilation.** The Choi noise channels are derived from the metric sensitivities $\partial g/\partial \theta_a$ — how a modulus perturbation propagates through H_int into logical error rates — yielding the effective logical-noise scale σ_eff (DERIVED).
- **G5 — Asymmetric code distances.** The gap-to-noise protection figure Δ/σ_eff is formed; it fixes the **asymmetric code distances (d_x, d_z)** and assembles the canonical invariant tuple (DERIVED). At the frozen point this selects the (d_x, d_z) = (12, 3) lock (the binary-lever-locked anchor of §H.9.1), the relaxed d_z = 3 being the minimum defensible value enabled jointly by §H.1 and §H.8.
- **G6 — Hardware projection.** Only at this final stage do hardware-facing parameters (transmon frequencies, couplings, pair pitch, routing, readout taper) appear; the projection is late and fully determined by the upstream geometric outputs (DERIVED).

Why this is the load-bearing enablement object. The compiler is the explicit, replayable bridge from "four numbers" to "a code spec and an overhead," ordered and parameter-conserving end-to-end (G1 introduces no parameter not already implied by the moduli; G3's η and r_rej are read off, not fit). The §H.3 additive offset g(Δ) = 18.305 is the G5-stage consumption of the G2 gap; the §H.9.2 composition algebra consumes the G5 (d_x, d_z) = (12, 3) output. The geometric necessity is internal to this compiler family — the flag-manifold structure supplies the *language* (Hamiltonian, holonomy, curvature) within which the moduli are optimized; physical correctness of the underlying geometry is neither claimed nor required for the engineering result, which is independently replayable from the four moduli (DERIVED / SIMULATED). The pipeline is drafted as a computer-implemented design-method dependent under the CLM-Q9 family (§J).

H.9.2 The composition algebra (the step table, with the clamp audit)

The canonical chain (`compute_composition_algebra()`, the closed-form engine of record — SIMULATED) at the (12,3) geometry. At d_qudit = 4, K_{12} chamber multiplier 0.30 lo-bound (**Gap-C-contingent**, §H.2.5), Path-4 delta −0.08 (the only validated cross-layer figure, §H.4.6), modular overhead 0.04, exact offset g(Δ) = 18.305:

Step	Formula	Value	Clamp status
gkp_base_raw	$2 \cdot d\_x \cdot d\_z - g(\Delta) = 72 - 18.305$	53.695	—
qudit_factor	$1/\log_2(4)$	0.5	—
gkp_base	max(1.0, 26.848)	26.848	clamp present, **NOT binding**
factory_reduction	0.30 · 0.20 / 0.60	0.10	Gap-C-contingent input
non_factory / factory_new	chamber-credit-applied	—	—
chamber_out → topo_out	with the chamber and topo credits firing	≈6.7	—
single_module	with Path-4 −0.08	**≈6.6** (= the ≈14.5× anchor before the M18-Majorana topo credit)	clamp present, **NOT binding**
10K fabric	≈6.6 × 1.04 × 1.005	**≈6.9**	—

Clamp audit (verified). Both max() gates are non-binding on this chain and the Path-4 −0.08 is fully expressed, not absorbed (verified step-by-step; DERIVED). The ≈6.6/6.9 floor is the (12,3) position with every accounting-gated credit firing; the ≈28× ceiling is the position with the chamber 0.60 credit failing its lo-bound. **The ≈6.6 floor is not a floor-clamp artifact** (it exceeds 1.0 by a wide margin before any gate) — the retired square-geometry sub-unity number, whose pre-clamp value was physically meaningless, stays retired.

H.9.3 The levers, with Lever-0 (exactification) added

Lever effect arithmetic is SIMULATED (composition campaign, real engine runs); the 24-month probabilities are unchanged:

Lever	Mechanism	Effect	24-mo P	Status
0 — exactification sweep (LANDED; runs S1/S2)	Kostant-exact Δ = 0.2589200, exact N_sites,gap and g(Δ) = 18.305 (settled the §H.3.5 g(Δ) question), exact K_6/K_{12} multiplicity counts (settles Gap-C), Gap-A holonomy re-derivation	anchored the ≈9.6–14.5× floor band at exact status (the posture is now defending an exact floor, not reaching a target)	n/a — a verification, not a success lever; **LANDED**	**DONE (machine-lane); the highest information-per-cost action in the program, now executed**
1 — qudit GKP (d = 4)	2 logical/mode on the coset-resonant cavity	reduces the GKP inner raw cost by $1/\log_2(4)$ within the floor band	~15%	SIMULATED effect; MEASURED at break-even (Brock 2025); FT demo OPEN
2 — modular 1.05	5-round BBPSSW + post-select (§H.5)	fabric ×1.075 → ×1.05	~70%	SIMULATED effect; **the ≈4.3–5× rate deficit of §H.5.4 is disclosed and travels with this row**
3 — cross-layer cancellation	3-layer joint decoder (§H.4)	Path-4 −0.08 only validated; Paths 1+2 CONDITIONAL on the g4 disjoint-factor chamber (G3); −0.46/−0.59 NOT ACHIEVABLE today	~25%	SIMULATED (Path 4); CONDITIONAL (Paths 1+2)
4 — outer-code removal @ 20 dB	d_outer 3 → 1 (§H.6)	toward ~1.0 single-module	~10%	BOUNDED by G5 ($Q \gtrsim 3\times10^8$ or ≳5.4 kHz stabilizer; fallback d_outer = 2)
5 — cavity Q	≈2×10^8 compound path (§H.7)	enables Lever 4	~30%	OPEN at 7–8 GHz; filter window characterized (G5)

Compound outcomes (SIMULATED waterfall; FIG20/FIG21/FIG22): the (8,3) X-verified frontier lands at the ≈9.6× lower end; the (12,3) lock with full accounting-gated credits lands at the ≈14.5× anchor (single-module ≈6.6 / fabric ≈6.9); the conservative ≈28× ceiling (chamber 0.60 credit fails its lo-bound) bounds the band from above and is carried at all times.

H.9.3a Sharp-optimum search and ablation (the mechanism evidence for the geometry-derived basin)

The Stage-2 topological search is not a smooth optimization that drifts onto a plateau — it lands on an **isolated optimum** over the admissible squashing space, and a leave-one-out ablation localizes *which* geometry-derived levers carry the saving. Both are reported here as **qualitative mechanism evidence, status SIMULATED**, and both are evaluated over a **bare-surface-code candidate space — i.e. before the GKP-inner + K_6-chamber + modular-outer levers of §H.1/§H.2/§H.5 are applied.** The bare-surface-code axis values produced by that search are accordingly the **bare-surface-code geometry-derived baseline (before the GKP/chamber/modular levers)** and are NOT a competing overhead figure; the levers then improve that bare basin into the two-stage scale-invariant floor band (the ≈9.6×–14.5× band, with the ≈28× ceiling, of §H.9.0). They must never be read as a headline ratio.

The isolated-optimum result (SIMULATED). A search over **360 candidate moduli** spanning squashing ratios ≈1.33 → 2.0 (sampling both the original territory and the territory newly opened by the §H.9.4a bound correction) yields a single narrow winning basin requiring the specific combination (R = 0.88, $x_1 = 0.75$, $x_2 = x_3 = 1.087$) — the frozen moduli of §H.9.1 — that **does not lie on the smooth grid-search curve**: every other promotable candidate sits at the bare-surface-code geometry-derived baseline level, and the next-best candidate is

well-separated from the winner by a wide margin within that bare-surface-code search (SIMULATED). Pushing to higher anisotropy raises Δ by up to ≈43% at ratio 2.0 yet *worsens* the bare-baseline overhead, because η and r_rej degrade faster than Δ improves — the implicit objective the geometry makes expressible is "maximize the coupled product of η, r_rej, and Δ subject to their geometric coupling," which has no representation without the prior. This wide, well-isolated separation is the load-bearing mechanism fact: the optimum is sharp and geometry-selected, not a broad accident of the sampling grid.

The ablation (leave-one-out; SIMULATED, qualitative). Removing each of the three geometry-derived levers in turn, holding the other two, degrades the bare-surface-code geometry-derived baseline by an isolable factor (SIMULATED):

Lever removed (other two kept)	Effect on the bare-surface-code baseline	Interpretation
Noise bias η removed	overhead worsens by ≈4×	bias alone is worth ≈4×
Common-mode rejection r_rej removed	overhead worsens by ≈2×	CMR alone is worth ≈2×
Spectral gap Δ removed	overhead worsens by ≈1.8×	gap alone is worth ≈1.8×

Each lever is individually worth ≈1.8×–4×; their **combined** effect is *less than the product* of the individual factors because the three share a single physical error budget (SIMULATED). The flag-manifold prior with its moduli seed removed produces no design at all (the §H.9.1a G2 stage is undefined), and a random positive-definite metric seed lands far from the basin — confirming the saving is the specific geometry-derived coupling, not a generic anisotropy. This ablation is the mechanism evidence for the §H.9.1a compiler's η/r_rej/Δ outputs; it is **not** a restatement of the floor band, which is the levered §H.9.2 composition result.

H.9.4 The floor-band honesty and the Shor orthogonality (both load-bearing, both reported)

The floor is reported as the value of a scale-invariant floor (not a probabilistic single point), with both endpoints feasibility-verified at SIMULATED design-stage grade. The reality-check budget run (SIMULATED, unfavorable, reported): with total_error = n_T·p_L·n_L and Shor-2048 requiring $p_L < 10^{-11}$, the design FAILS at every scenario (best case $p_L \approx 10^{-10} \rightarrow$ total $\approx 2.2\times10^{5}$); Grover-10^{6} also fails the total budget though the per-gate threshold is met. Closing this gap needs ~10 dB more squeezing or deeper chamber selectivity — a separate axis from the qubit-count ratio Q9 claims. The logical clock rate (~5,000 Hz at d = 4) clears the throughput thresholds on paper; the binding constraint is the logical error rate.

H.9.4a Robustness to a solution outside the identified search space (claim-scope disclosure)

The (12,3) lock and its frozen moduli are the *found* optimum, not a ceiling: the ideal or good-enough code may lie **outside the admissible search space identified so far**, and the architecture is disclosed to support that case without losing its held invariants. Three scope-broadening facts are recited (DERIVED methodology; the held invariants — Atiyah–Singer index = 3, Δ = 0.2589200, the Standard-Model calibration — re-checked against any candidate reached this way):

- **(a) The admissible region is a correctable, expandable object.** The §H.2.2 curvature-positivity bound is a *computed* set, not a fiat boundary: a prior over-conservative triangle bound (ratio < 1.45) was corrected to the true bound max_i(x_i)/min_i(x_i) < 2.0 (ratio < 2.0, equivalently |log(x_i/x_j)| < ln 2), and that single curvature-bound relaxation **opened additional design space** that the §H.9.3a search then sampled (the "newly opened territory" at squashing ratios above the prior bound). The admissibility gate G0 of the §H.9.1a compiler can therefore be re-derived, relaxed, expanded, or re-parameterized, and Stages 1 and 2 re-run over the expanded region (DERIVED).
- **(b) A good or ideal code may sit OUTSIDE the initially-identified admissible region.** Because the found optimum is an isolated narrow basin off the smooth grid curve (§H.9.3a), an admissible-region grid search need not *contain* the best solution: the disclosed method identifies a candidate found *anywhere* — reached,

for example, by a Stage-1 widening or an alternative prior — with the selected structure recorded under a sub-space identifier and re-checked against the held invariants (index = 3, $\Delta \approx 0.2589200$, the SM calibration). The (12,3) winner is recited only as a preferred embodiment; a candidate identified outside the initially-identified admissible region is within disclosed scope (DERIVED).

- **(c) Alternative priors, parameterizations, encodings, and modalities.** For QC purposes the substrate-selection constraints may be relaxed to admit alternative homogeneous-space priors beyond K_6 (e.g. other rank-controlled sub-family members) and random positive-definite metric seeds; alternative overhead cost models (a closed-form expression with a regime guard, and an empirical fit) selected by the regime guard; alternative code menus (the (12,2)-plus-post-select branch alongside (12,3)); representation-independent specifications of the same prior (root-and-weight, Weyl-chamber, or Borel–Weil–Bott bundle, the construction recovering the same mode-to-stabilizer mapping from any one); and alternative qubit modalities (the geometry-derived advantages being properties of the *code*, not the substrate, so they transfer in principle to trapped-ion, photonic, and neutral-atom realizations) (DERIVED / CONDITIONAL). These establish that the broad claim recites the *general* two-stage construction and the count-independence property, the specific found winner and its search space being quarantined to dependent embodiments.

H.9.5 Sensitivity honesty — the soft spots inside the un-clamped chain (disclosed)

1. **The g(Δ) calibration (§H.3.5)** — formerly the highest verification demand on the floor, now settled by run S1 (exact $g(\Delta) = 18.305$ at (12,3); $\Delta = 0.2589200$); the chamber-site-count reading is confirmed and the band anchored at exact status.
2. **The factory step rests on the Gap-C credit (§H.2.5).** The 0.30 lo-bound sits below even the 0.375 bare geometry cannot reproduce; if the chamber 0.60 credit fails its lo-bound, the position moves to the ≈28× ceiling. Gap-C is therefore load-bearing for where within the band the floor lands; resolver S2 (the Kostant exact counts Gap-C lacked) has landed and re-anchors the chamber-multiplicity dependent.

H.9.6 The system claim anchor (CLM-Q9)

The independent system claim (**CLM-Q9**; §J) recites: (a) the d = 4 GKP bosonic inner code providing two logical encodings per mode and Z-suppression; (b) the asymmetric outer code at the (12,3) geometry, the relaxed Z-distance enabled by (a); (c) the coset chamber projector applied to the magic-state-factory portion to admit chamber-class syndromes and reject the null-space complement, replacing a distillation cascade; (d) the photonic modular fabric; (e) arrangement on the K_6 substrate with the holonomy phase pinned to a Standard-Model flavor-calibration value; wherein the projected overhead is computed by a composition algebra subtracting the topological additive offset $g(\Delta) = 18.305$, the result being a **two-stage scale-invariant floor (≈9.6–14.5×, conservative ≈28× ceiling) substantially independent of the logical-qubit count**. Dependents recite: the offset coupling at the gkp_base_raw step with the exact $g(\Delta)$; the five levers; the (12,3) anchor and (8,3) frontier as band endpoints; evaluation at 10,000 logical with nominal invariance at larger counts; and chamber-coset selection with selectivity recited as **exact Kostant weight-multiplicity counts** at the recited rank, the two senses disambiguated per §H.2.5 (the re-anchored claim formerly carrying verbatim percentage bands — absent from all claim language).

H.9.7 Substrate enablement and the Gap-04 closure (the hook; verification-gated)

The K_6 substrate underlying claims CLM-Q9 and CLM-Q1Q2 is enabled not by calibration fiat but by the theory-side modulus-stability discharge: **the S-7 modulus-stability premise is discharged by the native scaffold (H04_DISCHARGED_AT_DECISION_GRADE, countersigned 2026-06-12).** The native scaffold $V(\sigma)$ of §H.2.3 cleared its falsifier-attached bound by 35× in the weakest accepted variant (central clearance 2.4×10^3). The historical QC-side blocker **B-UQFC-14-NP-1 is DISSOLVED, not satisfied**: no non-perturbative ΔV_{NP} wall was ever supplied; the apparent runaway was a truncation artifact of a superseded expansion layer, and the native three-layer scaffold stabilizes without any such wall — the QC harness's fail-closed refusal was the correct

behavior, resolved by upstream re-diagnosis. **QC-side consumption is verification-gated (V0–V4: hash-pin → adoption ruling → re-freeze → fail-closed re-run with old runs preserved → countersign; zero auto-promotions), and the full narrative with its residuals (loop-coefficient value, magnitude refinement, σ-kinetic normalization, the NT-1 transverse residue) is Appendix §K.App-B.** The §H.0 moduli-point disclosure applies in full. The closure also carries method evidence (a fail-closed blocker named honestly by the harness → native re-diagnosis → exact-coefficient scaffold → machine clearance → countersigned discharge, with four engine generations halted by their own gates and five self-caught corrections); the companion credibility narrative is carried in Appendix §K.App-C.

H.9.8 The exactification posture, the corrected-invariant discipline, and the named runs

Exactification-landed statement. The sweep has landed: every occurrence of Δ, Φ, g(Δ), N_sites,gap, and the chamber multiplicity counts carries its exact value of record. The two executed machine-lane runs: **S1** — Kostant-exact Δ = 0.2589200 plus exact N_sites,gap, yielding g(Δ) = 18.305 at (12,3) (settled the §H.3.5 chamber-site-count question); **S2** — exact K_6/K_{12} chamber multiplicity counts (re-anchored the Gap-C factory step and the CLM-Q9 selectivity dependent). The result re-anchors the **two-stage scale-invariant ≈9.6–14.5× floor band at exact status** — the posture is now **defending an exact floor**, with the ≈28× ceiling carried if the chamber 0.60 credit fails its lo-bound. NF-4 is the dependency chart.

The corrected-invariant discipline (credibility support, not a claim). The substrate's curvature-invariant registry was upgraded during this program by a **4.6× correction of the dominant quadratic invariant to the banked $|\mathbf{Riem}|^2 = 276$** (with $|\mathrm{Weyl}|^2 = 216$ — the substrate is not conformally flat; the superseded prior value is preserved only in the Appendix §K.App-A supersession record, with the UV-evidence computations that consumed it quarantined there). Three facts make this evidence *for* the filing: (i) the suspicion was **flagged by the corpus's own QC-side harness** before the theory side banked the correction; (ii) the corrected registry does **not** move Δ, Φ, g(Δ), or the overhead chain (none consume $|\mathrm{Riem}|^2$); (iii) the same self-correction discipline produced the decoded-versus-undecoded XZZX re-measurement and the sub-Kitaev retirement (§H.9.0). Appendix §K.App-C carries the three arcs together.

H.10 INVENTION Q10 — THE PRIOR 28-FAMILY PORTFOLIO, INLINED AT CLAIM-CHART LEVEL

H.10.1 Role and counting note

Q10 is the prior QC IP portfolio of the same sole inventor, carried both **by incorporation by reference** (the underlying specifications and the 2026-05-04 handoff remain the controlling detail documents) and **inlined at claim-chart level** so the filing is self-contained on the portfolio's strongest assets. Two bodies of work: (1) the **AEGIS spectral architecture** (the 28th family — correlated-error spectral-certification with four hardware mechanisms); (2) the **27 lettered claim families** on the K_6 substrate (Stage-B converged, Stage-C backed). Counting note: the handoff labels the lettered corpus "27 families"; this package counts 28 by including AEGIS as the architecture the iterations layer onto. Status: **9/9 Stage-C regressions PASS** (SIMULATED); of the 27 families, 20 carry direct Monte-Carlo verification, 5 are STRUCTURAL, 2 METHODOLOGICAL; AEGIS is independently SIMULATED (9/9 adversarial tests PASS). The new Q1–Q9 claims are drafted as improvements over, and compositions onto, this portfolio.

H.10.2 AEGIS — the spectral-certification architecture (inlined)

Thesis. Treat all pairwise correlated multi-qubit error as a single symmetric kernel and certify the spectrum, rather than attacking channels in isolation:

> **$K_{total}(i,j) = K_{EM}(i,j) + K_{ctrl}(i,j) + K_{QP}(i,j)$**, additive because the three channels live on distinct timescales (GHz always-on cavity EM; gate-time control crosstalk; sub-MHz burst-like quasiparticle diffusion);

> cross-channel terms second-order.

Certification criterion: $\lambda_1(K_{total}) \leq 0.01$ — the dominant-eigenvalue threshold above which a surface-code decoder is overwhelmed along the dominant correlated eigenmode, a *certification* criterion (worst-case over error patterns), not merely a design metric (DERIVED). Baseline: **$\lambda_1 = 0.005105$ (50% margin)**; IPR 0.018; $\xi \approx 35$ qubits (< 50 required); scaling exponent $d(\log \lambda_1)/d(\log N) = \mathbf{-0.0014}$ (negative — certification *improves* with scale); fabrication-disorder failure 0.0% over 200 MC samples at ±3–20% variation; **9/9 adversarial tests PASS** (all SIMULATED).

The four hardware mechanisms (the claim-bearing apparatus):

1. **M1 — Modular EM firewalls.** Superconducting trenches partition the chip into 25-qubit modules; a 5 mm firewall penalty suppresses inter-module EM coupling → $\exp(-12.5) \approx 3.7\times10^{-6}$; alone moves λ_1 0.035 → 0.0095 (3.7×). Novelty: confining cavity modes to modules so the kernel's long-range block is bounded by construction.
2. **M2 — Orthogonal control routing.** Control wiring perpendicular to the EM mode pattern so K_EM and K_ctrl have orthogonal dominant **eigenvectors** — λ_1 of the sum behaves as max() not sum (~30% additional λ_1 reduction; eigenvector overlap 0.8 → <0.2). Novelty: eigenvector orthogonality as an explicit layout target — matrix structure, not scalar crosstalk.
3. **M3 — Ultra-low coupling + hard 500 µm cutoff.** Control coupling 36× lower (0.0018 → 0.00005) AND exactly zero beyond 500 µm. The cutoff is load-bearing: without it the residual $1/r^2$ accumulation drives λ_1 to 4522 (450,000× over threshold) — the catastrophic mode exponential-decay models miss. Novelty: the hard cutoff as a certification-enabling element.
4. **M4 — Quasiparticle trap network.** A dense normal-metal trap grid (5 µm spacing, ~40,000 traps/mm², 92% efficiency) collapses K_QP by $(1-0.92)^2 = 0.0064$ (156×; diffusion length 80 µm → 2.5 µm). **Honesty note on its face:** unmitigated bursts fail certification above **0.25 MeV**; full-range resilience is **mitigation-operational** ("all combined" passes through 10 MeV at >800× reduction). All SIMULATED.

AEGIS claim families (chart): A — core spectral architecture (one broad independent claim); B — M1–M4 embodiments; C — the design workflow (build kernel → decompose → spectral criterion → select parameters → adversarially validate); D — certification endpoint (deterministic acceptance matrix: fixed seeds, source hashes, pass/fail gates); E — burst-resilience orchestration. The Kostant generator claim (**CLM-KOST-1**, §H.2.6) is a dependent under this lineage — the $\lambda_1 \leq 0.01$ kernel certification and the integer-gated substrate-spectrum certification are the same certify-the-spectrum move applied to noise and geometry respectively. AEGIS is the 2D-lattice spectral-hardware line, independent of K_6; the two meet only at the Q9 stack as different layers.

H.10.3 The 27 lettered families — claim cores by group (chart level)

Group A — iter-4 mechanisms (#1–#16, all SIMULATED). Claim cores: **#1** Z_2 chirality syndrome pre-classifier reading the bundle's forced Dirac grading. **#2** chirality parity-protection on a DFS transmon pair. **#3/#6/#9/#12** composability matrices (subsumed into #16). **#4** Parisi-RSB ultrametric cascade decoder (per-subclass error spans 3.8× at d = 11). **#5** iter-2-as-Level-0 verbatim predicate preservation under a build-failing HARD GATE. **#7** cross-layer PI controller modulating DD rate, decoder/cascade/re-read thresholds (Bode-stable at 20% noise). **#8** runtime admissibility-range enforcement on controller setpoints. **#10** cascade-of-ensembles decoder with MIMO-tunable vote weights. **#11** L0 single-vote degeneracy preserving the verbatim gate through the ensemble layer. **#13** cycle-frequency adaptive scheduling (QEC period as fifth controlled variable). **#15** gain-matrix-extension discipline (zero-initialized row+column).

The two strongest, expanded:

- **#14 — FOUR-LAYERED HARD-COMPOSITION-GATE preservation discipline (STRONGEST META).** Four preservation tiers — code-time (#5), runtime (#8), ensemble (#11), cycle-period (#13) — each with its own build-failing hard gate, composing multiplicatively so flat-predicate semantics, the SM calibration, index = 3, and Δ are preserved across all four regimes simultaneously (no prior art uses a four-layered QEC-evolution preservation discipline; verified end-to-end across Stage-C, SIMULATED). This governs how the new Q1–Q9 layers were added without corrupting the substrate guarantees.
- **#16 — the cumulative composability matrix (STRONGEST SINGLE CLAIM).** The K_6 substrate + a nine-slot composability matrix (five base mechanisms + cascade/PI/ensemble/scheduling) + the #14 gate + the #15 gain-matrix discipline, all preserving SM-calibration, index = 3, and $\Delta \approx 0.259$, with cumulative ≥60% EM-external overhead reduction vs iter-0 (SIMULATED, all five Stage-C campaigns PASS). The natural carrier into which the new four layers slot.

Group B — prior-iter mechanisms (#17–#20). #17 substrate-anchored dynamical-decoupling (CPMG + XY-4, locks anchored to substrate properties). **#18** EM-gap-closure admission-predicate API (three calibrated thresholds, ~70% EM-gap closure in its original regime; the predicate *logic* survives, the original d_z = 2 floor is superseded below). **#19** XY-surface-code drop-in at bias $\eta = 4$ (threshold gain library-derived — disclosed; superseded below). **#20** multi-stage suspect-gated re-read (suspect-only second reads + hybrid disagreement handling; f^n scaling verified across four decades). All SIMULATED.

Group C — geometry-side foundations (#21–#25; STRUCTURAL). #21 — the Article-A six-constraint substrate admissibility filter (HIGHEST geometry-side): admit only compact 6-manifolds with transitive Lie action satisfying all six constraints (dimension 6; compactness; SU(3) isometry; chiral Dirac index = 3 via the spin-c (1,0) bundle; tripartite tangent split; consistent spin-c structure) — a substrate-*selection* discipline whose unique solution among 62+ candidates is K_6 (Stage-A search). The entire Rung-3 stack is built on this output. **#22 — the Φ↔Δ bridge calibration anchor:** $\Delta = 2|\sin(\Phi/2)|$ ties QEC spectral-gap optimization to Standard-Model observables (CKM/Jarlskog), so any Δ change forces SM-fit re-validation ($R^2 \geq 0.9957$); the Q3 additive offset (§H.3) is its direct descendant. **#23** index = 3 preservation as a code-time hard gate (falsified by a fourth generation). **#24** spin-c chirality-grading READ discipline (no competing Z_2). **#25** the find-N-instances-then-extract-invariants substrate-search methodology (METHODOLOGICAL; process-claim eligibility risk flagged).

Group D — methodology (#26–#27; METHODOLOGICAL, process-claim risk flagged). #26 iteration governance (per-iteration cognitive-stack caps + composite-degradation HALT + zero-new-blocker checks). **#27** the record-decision override pattern for value-driven cap extensions.

H.10.4 The superseded / still-stands / composes map (binding on claim construction)

Relation	Items	Statement
Superseded	#19 (XY overlay); #18's prior-iter floor; AEGIS-as-headline; "27-as-the-portfolio"	The locked asymmetric XZZX (12,3) geometry replaces the XY overlay on the critical path; GKP inner suppression + Q8 orthogonal routing relax the Z-distance demand below the iter-2 floor (the predicate API survives as fallback logic); AEGIS remains the correlated-error layer but the Q9 composition is top-line; this package's Rung-3 specs supersede/extend the 28-family spec.
Still stands	#21–#25; all AEGIS families; #1, #2, #17, #20, #26, #27	Substrate-selection, calibration-bridge, and invariant-preservation disciplines are the foundations the new stack is built on; AEGIS is an orthogonal standalone defensive asset; the mechanism/process families compose onto any K_6 decoder/scheduler regardless of inner code.
Composes	#16 + #14; #4 + #10 (→ Q4 decoder); #7 + #13 + #15 (→ Q5 control); #18's API (→ Q5 link post-selection)	The new layers slot into #16's matrix under #14's four-layer gate; the Q4 joint decoder builds on the cascade/ensemble substrate; the PI/scheduling/gain-matrix spine drives the new levers' runtime setpoints; the admission-predicate pattern reappears at the modular link.

H.11 SECTION-H CROSS-REFERENCE SPINE (for A-PASS)

Claims: CLM-Q1Q2 (§H.1.7, §H.2.7), CLM-Q9 + dependents (§H.3.6, §H.5.3, §H.8.4, §H.9.6), CLM-Q4-N1…N9 (§H.4.3–H.4.5), CLM-KOST-1 (§H.2.6, §H.10.2). Figures: FIG03, FIG20, FIG21, FIG22, FIG23 (G4 selectivity correction), FIG33, FIG34, NF-4, NF-6, NF-7, NF-8. Appendix hooks: §K.App-A (invariant registry + supersession + quarantine), §K.App-B (Gap-04 V0–V4 trail), §K.App-C (status-label legend + the three self-correction arcs), §K.App-D (banned-numbers ledger, exactification dependency, pre-filing gates: the landed S1/S2, G5-EM, g4 Gaps A/B/C, the Path-4 scheduler re-confirmation, the (12,3) Stim check). Simulation evidence: §SE (SE.2 GKP/bosonic campaign, SE.3 composition + the two-stage floor-band derivation, SE.5 gate board g1–g5).

H.12 WA-3 SELF-CERTIFICATION (blueprint §2 rules → compliance)

This Section was drafted against `C7_PROVISIONAL_100P_BLUEPRINT.md` §2 (RULE-1…RULE-11, plus §3.3 / §1-row-H / label discipline) and is compliant with every rule. The load-bearing ones, with their controlling sites: **RULE-1/6/7** — the UV-evidence layer, the disputed CERT, and the stale curvature value are referenced only as quarantined/descriptive (§H.0, §H.9.8), never relied upon; the no-consumption blast-radius statement (Δ/Φ/g(Δ)/the ≈9.6–14.5× band do not consume $|\mathrm{Riem}|^2$) appears in §H.0 and §H.9.8. **RULE-2/8/10/11** — modulus stability is grounded only in the native scaffold V(σ), B-UQFC-14-NP-1 is framed "DISSOLVED, not satisfied," and every theory-side consumption states its moduli point and propagates Gap-04 only via the V0–V4 trail (§H.2.3, §H.9.7). **RULE-3** — claim language recites exact Kostant weight-multiplicity counts (run S2); the admitted-fraction vs firewall-rejection senses are disambiguated at every use; 0.375/0.30 always carry the Gap-C flag; historical bands appear once as ESTIMATED background with the G4 correction adjacent (§H.2.5). **RULE-4** — $\Delta = 0.2589200$, $g(\Delta) = 18.305$ at (12,3), $\Phi = -0.12827$ rad carried as exact at every load-bearing occurrence; the floor is the two-stage scale-invariant ≈9.6–14.5× band with the conservative ≈28× ceiling. **RULE-5/9(c)/9(h)** — the aggressive Stage-B figure, the retired sub-unity framing, and Q8's 11.4-without-SPECULATIVE are each banned/fenced and appear nowhere as live values. **RULE-9(a/b/d/i/j)** — the floor-band/count-independence framing governs via the §H.9.0 banner; −0.46/−0.59 are NOT-ACHIEVABLE-fenced everywhere; the ≈4.3–5× Q5 rate deficit travels with every 1.05 appearance; the Shor-2048-fails-at-every-scenario orthogonality caveat travels with every ratio. **§3.3** — the Kostant generator is described as claimable tooling (§H.2.6). Every quantitative statement carries exactly one master-vocabulary status label.

Approximate section word counts (post-trim): the Section runs ≈12,800 words ≈ 23 pages at the blueprint's 1 page ≈ 500–550 words with tables counting ~double their line count.

STATUS-PRESERVATION FOOTER. This file is WRITE-NEW-ONLY: it edits no stub, spec, gate result, narrowed-claims file, simulator, runtime tensor, FINAL_WINNER_RECORD, or frozen artifact. Status invariants preserved by reference: SoT hash `88e5903fe1c3d17655acef7e4fb45078e3168add25362040e21598b3b161ff17`; hard-gate vector `D, D-local, D, D, D`; BG-10 Open/Diagnostic; Stage-B iter-4e FINAL-WINNER-SCIENCE-PASS untouched (the (12,3) geometry is an Option-C patent-context designation only); pinned invariants $\Phi = -0.12827$ rad, $\Delta = 0.2589200$ (exact, run S1; prose "≈0.259"), $g(\Delta) = 18.305$ at (12,3) per RULE-4. Sole inventor: Chris Bergstrom. Assignee: Tahoe Labs (entity TBD). AI used as a tool, not a co-inventor (*Thaler v. Vidal*; 2024 USPTO guidance). Drafted 2026-06-12 by WA-3; scrubbed and rebuilt to the two-stage scale-invariant frame 2026-06-17.

End of §H (PART_C_RUNG_3_QC.md).

PART D — SIMULATION EVIDENCE AND APPENDICES

Provisional Patent Application — Tahoe Labs (Sections §I and §K of the assembled 100-page draft). Sole inventor: Chris Bergstrom. **Assignee:** Tahoe Labs (entity TBD). **Status-label convention:** every quantitative statement carries exactly one of MEASURED / SIMULATED / DERIVED / PROJECTED-TARGET / BOUNDED / CONDITIONAL / OPEN / SPECULATIVE (legend in Appendix K3). No number is presented as achieved unless so labeled. The exactification sweep (runs S1/S2) has landed: $\Delta = 0.2589200$ and $g(\Delta) = 18.305$ at the (12,3) geometry (Atiyah–Singer index = 3) are exact values of record; σ_eff remains the one ESTIMATED input. **SoT hash `88e5903fe1c3d17655acef7e4fb45078e3168add25362040e21598b3b161ff17` preserved by reference.** AI is a tool, not a co-inventor (*Thaler v. Vidal*; 2024 USPTO inventorship guidance).

> **Attorney work-product pointer.** Detailed §101 eligibility and §102/§103 prior-art analysis is maintained as a separate attorney work-product companion and is available on request.

---

§I — SIMULATION EVIDENCE

Results are from REAL runs only, on commodity hardware with free open-source tooling (versions recorded in the reproducer of §I.6). Unfavorable results are reported alongside the claim they qualify. Analytic-only quantities are labeled as such. The single-command reproducer of §I.6 re-derives every load-bearing number from a clean checkout.

§I.1 — The decoder-channel advantage (load-bearing for E3, Q9)

High-statistics Monte Carlo on the locked asymmetric rotated XZZX surface code ($d_x = 12$, $d_z \in \{1,2,3\}$) under a device-created biased-noise channel (local-Clifford Hadamard deformation; Bonilla Ataides et al., *Nat. Commun.* 12, 2172, 2021), decoded by MWPM (Stim 1.15.0 + PyMatching 2.3.1), 4×10^6 shots/point at the headline:

Bias point	Metric	Unbiased LER	Biased LER	Reduction	95% CI
usable $\eta = 3.10$	**MWPM-decoded**	1.90×10^{-5}	2.25×10^{-6}	**8.44×**	**[4.2, 16.9]**

The single filable decoder statement carried through §F (E3) and §H (Q9) is: **the MWPM-decoded same-distance logical-Z reduction at the device-central usable bias is 8.4× (95% CI 4.2–16.9)** [SIMULATED]. **This is a decoder-channel figure — NOT a memory-overhead ratio** — and is never conflated with the composition floor of §I.2. The earlier undecoded raw-syndrome-flip figure is retired (named only in Appendix K4). Below-threshold operability holds at the design point: at $p_{phys} = 3.9\times10^{-4}$, $d_x = 12$, the dominant X-axis returns $0 / 3\times10^6 \rightarrow$ **$p_{L,X} < 1.0\times10^{-6}$** (rule-of-three 95% upper bound) [SIMULATED], deeply below threshold. The minimum defensible Z-axis distance is 3 (the degenerate unencoded small-axis line is retired as a code; Appendix K4 R6).

§I.2 — The two-stage scale-invariant composition floor (load-bearing for Q9)

Overhead is computed as a **two-stage, scale-invariant floor** — Stage 1 restricts the design search to the count-independent regime; Stage 2 performs a topological search within it — evaluated at **10,000 logical qubits** and reported as nominally unchanged at larger counts (e.g. 20,000 logical). The reported floor is the **SIMULATED ≈9.6×–14.5× band** of physical(bosonic-mode) per logical, carried with a **conservative ≈28× ceiling** (DERIVED) if the chamber 0.60 credit fails its lo-bound. Each endpoint preserves all three held invariants (Atiyah–Singer index = 3, spectral gap $\Delta = 0.2589200$, Standard-Model calibration):

Configuration	fabric @ 10K	role in the band
(12,3) binary-lever-locked (canonical lock / headline anchor)	**≈6.9**	the **≈14.5×** anchor (all credits firing)
(12,3) credit-stripped	**≈28**	the conservative **≈28×** ceiling
(8,3) X-verified frontier (dependent geometry)	**≈9.6**	the **≈9.6×** lower frontier

The floor is **substantially independent of the logical-qubit count** (a floor at scale, not a quantity that climbs with the count): the band is nominally unchanged from 10K to 20K logical (≤1% drift). The (8,3) row is DERIVED + EXTRAPOLATED-X with dependent-geometry caveats (binary lever forfeited, X-axis extrapolated). The only validated cross-layer reduction applied is **−0.08 (Path 4; Litinski ancilla-sharing, SIMULATED)**; deeper Paths-1+2 figures are quarantined (the SU(3) commutativity check returned BOTH_FAIL; Appendix K4). The decoder-channel advantage of §I.1 (8.4×) is a different quantity from this floor and is never conflated with it.

The stack separately FAILS the Shor-2048 logical-error budget at every scenario including the R = 1.0 best case [SIMULATED]: the physical-to-logical floor ratio and the logical-error budget are **orthogonal** — a competitive

overhead ratio is a qubit-count win, not a runnable cryptographic algorithm. This caveat travels wherever the ratio appears (Appendix K4, RULE-9(j)).

§I.6 — Reproducibility: the one-command reproducer (enablement centerpiece)

A skeptic can re-derive the load-bearing simulation evidence from a clean checkout with free, open-source tools in well under a minute, via a single command:

```
python SIMULATION_CAMPAIGN/verify_campaign.py
```

Exit code 0 if and only if every core numeric reproduction passed. The verifier runs from any working directory, needs no internet and no GPU, prints the tool versions found (Python 3.12, NumPy 2.4.2, SciPy 1.16.3, Stim 1.15.0, PyMatching 2.3.1, QuTiP 5.2.3, Matplotlib 3.10.8), imports the corpus engines unmodified (falling back to committed CSVs if absent and saying so), and exits 0 only if all core checks reproduced. Core checks include the (12,3) lock floor @ 10K (≈6.9 fabric, the ≈14.5× anchor), the (8,3) frontier (≈9.6×), the scale-invariance check (floor @ 10K vs 20K, ≤1% drift), and the decoder biased-advantage @ η≈3.1 (8.44×, a decoder-channel figure NOT a memory-overhead ratio). A healthy run ends with `RESULT: CORE REPRODUCED` and exit 0. Items no open-source laptop stack can compute are **named as gates, not failed**: the heavy electromagnetic binaries (openEMS / FastHenry2 / FasterCap), the thermal conjugate-FEM (named gate ANSYS Icepak / Cadence Celsius), and the deep ideal GKP tails (analytic, not finite-Fock-simulable).

---

§J — CLAIMS (PART E OF THE 100-PAGE PROVISIONAL)

Writing agent: WA-5 (claims assignment per `C7_PROVISIONAL_100P_BLUEPRINT.md` §1 row J, §4.5) **Date:** 2026-06-12 **Assignee:** Tahoe Labs (entity TBD) · **Sole inventor:** Chris Bergstrom (AI used as a drafting tool only — *Thaler v. Vidal*, Fed. Cir. 2022; 2024 USPTO AI-inventorship guidance) **Controlling language:** `GATE_CLOSURE/narrowed_claims/NARROWED_CLAIMS_NVIDIA.md`, `NARROWED_CLAIMS_E3.md`, `NARROWED_CLAIMS_Q4.md` (deviations only to tighten); base stub §9; A7 §5.2 (6-layer claims); A6 §4.1.4 + R5 (Kostant tooling; Q9-5 re-anchor). **Numbering:** all claim identifiers are SYMBOLIC (`CL-…`); A-PASS assigns final numbers and resolves antecedent basis. Registry cross-walk: `CL-` IDs map 1:1 onto the blueprint §4.0 registry (`CL-F1/F2` = `CLM-FOUND-1/2`; `CL-Q1-*/Q2-*` = `CLM-Q1Q2`; `CL-N3/N4/N5` = `CLM-N3N4N5`). Bracketed *[italic]* passages are attorney drafting notes, not claim text.

> **Re-anchoring self-certification (RULE compliance).** No claim below recites: the historical selectivity percentage bands (RULE-3, replaced by exact-Kostant-pending language); any quarantined UV-evidence quantity (RULE-1); any superseded-coefficient value of the retired flow lane (RULE-2); the banned aggressive compiler figure (RULE-5); a stale curvature invariant (RULE-6); or a retired cross-layer aggregate, retired floor/reduction artifact, local thermal figure as a system number, or linear-code-kernel non-membership assertion for a classical amplitude channel (RULE-9). The QC value-driver claims recite the vetted **two-stage, scale-invariant physical-to-logical overhead floor — a count-independent band of approximately 9.6× to 14.5×, anchored at (12,3) (≈14.5×) and reaching the (8,3) X-verified frontier (≈9.6×)** — with the held-fixed invariants (Atiyah–Singer index = 3, $\Delta \approx 0.260$, Standard-Model calibration) carried as constraints; every recited overhead value is an EXTERNAL / SIMULATED projection, not a measured device property. The retired single-module headline ratios and provisional swap-slot markers are removed; no claim recites a degenerate d_z = 1 line as a code. Conditional claims recite their contingencies in the claim text itself.

J.9A — FOUNDATIONAL INDEPENDENT CLAIMS (the coset-admittance principle, classical system + quantum chamber)

> *[Scope discipline (F-05): the classical instance of the coset thesis is carried by the SYSTEM trade (bandwidth densification conditioning thermal-resource reallocation), never by the signaling alphabet; the "not a linear-code kernel" clause is asserted ONLY for the coherent quantum embodiment.]*

CL-F1 (independent — system, spanning classical and quantum embodiments). An information-processing system comprising a physical channel and an admittance subsystem, wherein: (a) the admittance subsystem restricts an admissible set of channel states to the simultaneous +1 eigenspace of a maximally-abelian stabilizer subgroup S of a Heisenberg–Weyl or holonomy algebra of a compact Lie group, said eigenspace being a coset of said algebra; (b) the admittance subsystem rejects, at a channel input and/or output, every state lying in the complementary null space of a projector P_S onto said eigenspace; (c) **in a coherent quantum embodiment**, the channel carries a momentum-conjugate (X-basis) degree of freedom in addition to an amplitude (Z-basis) degree of freedom, and the admissible set is not the kernel of any linear code over a finite field or real vector space of equivalent rate; and (d) **in a classical embodiment**, the multi-level admitted alphabet carries a specified aggregate bandwidth on a reduced physical-channel count relative to a binary-signaling baseline of equal aggregate bandwidth, and the physical-channel resources thereby freed are reallocated to a thermal-extraction function in number determined by the reduced channel count, the system-level reallocation trade — and not the admitted alphabet itself — constituting the classical subject matter.

CL-F2 (independent — method). A method comprising: defining a coset of a compact-Lie-group Heisenberg–Weyl or holonomy algebra as the simultaneous +1 eigenspace of a maximally-abelian stabilizer subgroup;

admitting only states, signal levels, or syndromes whose representation lies in said coset; and firewalling the complementary null space by suppression at a transmitter or rejection at a receiver or decoder prior to downstream processing; wherein the same coset-admittance structure is instantiated (a) classically as a d-level signaling alphabet on a through-silicon via with $d \geq 4$, the classically-claimed benefit being the system-level thermal-via reallocation conditioned on the resulting bandwidth densification, and (b) quantum-mechanically as a self-adjoint idempotent projector onto an admissible-state subspace of a flag manifold SU(N)/$U(1)^{N-1}$, the coherent instance alone carrying signaling-alphabet novelty.

J.9B-I — RUNG 2 (NVIDIA / 3D-CHIP)

J.9B-I.1 The strongest standalone system claim and dependents (controlling language: NARROWED_CLAIMS_NVIDIA §1)

> **Claim-hierarchy note.** The NVIDIA-rung headline is the 6-layer integrated-system claim **CL-6L-1** (§J.9B-I.2). **CL-N2-1** is the strongest *standalone, prior-art-defensible* claim (it also reads on the 2-die HBM-PHY and the wakeup-chip RTP) and is incorporated by reference into the 6-layer claims (CL-6L-1(d), CL-6L-2(e)).

CL-N2-1 (system, independent — the strongest STANDALONE, prior-art-defensible claim of the classical rungs; incorporated by reference into the 6-layer headline claims). A three-dimensional stacked integrated circuit comprising: (a) a plurality of inter-layer signaling through-silicon vias driven with a multi-level admitted signal alphabet, the admitted alphabet being the set of signal levels admitted by membership in a coset of a Heisenberg–Weyl group HW(d) of dimension $d \geq 4$, such that the plurality of inter-layer signaling through-silicon vias carries a specified aggregate inter-layer bandwidth using a **signal-through-silicon-via count that is reduced relative to a binary-signaling baseline carrying the same specified aggregate inter-layer bandwidth**; (b) a plurality of dedicated thermal vias occupying through-silicon-via slots that are **freed by, and in number determined by, the reduced signal-through-silicon-via count of (a)**; wherein (c) a density of the dedicated thermal vias of (b) is at least 30% greater than a thermal-via density of said binary-signaling baseline carrying the same specified aggregate inter-layer bandwidth; and (d) the dedicated thermal vias of (b), by reason of their occupying the slots freed in (a), provide a reduction in vertical thermal resistance relative to said binary-signaling baseline; wherein the **freed-slot causality of (b)** — that the dedicated thermal vias occupy slots freed by, and in number determined by, the reduced signal-through-silicon-via count of (a) — the **equal-aggregate-bandwidth baseline** against which the signal-via count and thermal-via density of (a) and (c) are defined, and the **quantified thermal-via-density coupling of (c)** are each **ESSENTIAL AND NON-SEVERABLE limitations of this claim**, none of which may be read out without removing the claim from its scope, the combination of these three limitations being what distinguishes the circuit from the mere addition of a thermal via to a stacked integrated circuit. *[The d = 4 / PAM-8-plus-parity case is intentionally within scope, consumed by the system framing. The numeric R_th range is held severable in CL-N2-2 until the conjugate-thermal FEA gate closes. The freed-slot causality, equal-aggregate-bandwidth baseline, and quantified coupling are recited as the essential, non-severable anti-KSR triple in the claim language itself; making them express limitations tightens (does not broaden) the claim, per spec §D.]*

CL-N2-2 (dependent — quantified R_th, FEA-gated; severable numeric anchor). The circuit of CL-N2-1, wherein the reduction in vertical thermal resistance of (d) is in the range of 15% to 25% as determined by a conjugate-thermal finite-element analysis of the stacked integrated circuit at a compute-die heat flux of approximately 1 W/mm². *[Status: DERIVED + FEA-ESTIMATED, density-dependent. The G2 conduction FEM measured ~6.5% (SIMULATED) at 4,800 vias / 50 mm²; the 15–25% range requires ~2–3.7× densification and is licensed only by the ANSYS Icepak / Cadence Celsius conjugate-thermal FEA (filing-grade gate). Keep the range OUT of CL-N2-1 until that gate closes.]*

CL-N2-3 (dependent — freed-slot causality made explicit / anti-obviousness). The circuit of CL-N2-1, wherein the number of dedicated thermal vias equals approximately the number of signal through-silicon vias eliminated by the multi-level admitted signal alphabet relative to the binary-signaling baseline of equal aggregate bandwidth, such that substantially every through-silicon-via slot freed by the reduced signal count is occupied by a dedicated thermal via.

CL-N2-4 (dependent — bond-interface bypass mechanism). The circuit of CL-N2-1, wherein the dedicated thermal vias provide a direct copper conduction path that bypasses a silicon-dioxide hybrid-bond interface, said hybrid-bond interface otherwise constituting a majority of the vertical thermal-resistance path between a compute die and an adjacent memory die. *[Mechanism-tied; no local conductance figure recited as a system number (RULE-9(f)).]*

CL-N2-5 (dependent — counts / geometry, single-process). The circuit of CL-N2-1, wherein a total through-silicon-via count is approximately 16,000, of which approximately 8,000 are coset-signaling vias and approximately 4,800 are dedicated thermal vias, and wherein each dedicated thermal via is approximately 3.2 μm in diameter and 40 μm deep and is filled with electroplated copper formed by the same via-last process as the signaling through-silicon vias, adding no additional process step.

CL-N2-6 (dependent — application boundary). The circuit of CL-N2-1, configured as a memory-to-logic physical-layer interface between a compute die and a high-bandwidth-memory base die joined by a copper-to-copper hybrid bond.

CL-N2-7 (dependent — crosstalk subordinated to the thermal architecture). The circuit of CL-N2-1, wherein the multi-level admitted signal alphabet of (a) suppresses intermediate-level signal transitions at the transmitter so as to reduce capacitive coupling to adjacent through-silicon vias, and wherein the dedicated thermal vias of (b) are interleaved among the signaling through-silicon vias so as to maximize vertical heat extraction from a buried logic die.

J.9B-I.2 The 6-layer hierarchical-stack claims — THE NVIDIA HEADLINE (A7 §5.2; honest-magnitude bindings of blueprint §3.4)

CL-6L-1 (independent — heterogeneous-role multi-layer system; THE LEAD CLAIM OF THE NVIDIA 6-LAYER PRODUCT). A three-dimensional stacked integrated circuit comprising at least six vertically-bonded active device layers, wherein: (a) each of the at least six device layers carries a **distinct heterogeneous functional role**, the roles comprising a signal-distribution layer, an inter-column routing layer, an input layer, an output layer, a compute layer, and a feedback-cache layer; (b) the device layers are **thermally ordered** such that the layer having the highest power density is positioned adjacent to a primary heat-extraction path of the stack; (c) the layers are organized into vertical functional columns spanning all of the at least six device layers, inter-layer wiring within and among the columns being partitioned into a majority of intra-column local vias, a minority of lateral vias, and a smaller minority of global vias through a hierarchical return network; and (d) inter-layer signaling among the device layers employs a multi-level coset-admitted signal alphabet according to CL-N2-1(a), the through-silicon-via slots freed thereby being reallocated to dedicated thermal vias according to CL-N2-1(b). *[Novel surface unreachable by the 2-die rung: the heterogeneous-role partition (two roles is prior practice; six thermally-ordered roles is not). The thermal-via reallocation of (d) is a supporting knock-on (~6.5% SIMULATED), never the stack's thermal solution. The thermally-ordered assignment (compute@inlet → distribution → I/O → output → feedback-cache → NoC@top) is SIMULATED — coupled OpenFOAM CFD CONFIRMED, buried-die ΔT ≤11.1 K (`topo_search/integrated/INTEGRATED_CONFIRM_CFD.md`); ANSYS Icepak gate OPEN.]*

CL-6L-1.1 (dependent — NEW; guard-TSV signal-integrity allocation, the marginal→pass BER fix; full-wave-upgraded). The circuit of CL-6L-1 or CL-N2-1, wherein a **guard-through-silicon-via allocation of at least 16% of the inter-layer PHY via array**, each guard through-silicon via being a grounded copper via

interposed between an aggressor and a victim signal via and tied to a power-ground plane, reduces effective nearest-aggressor crosstalk coupling such that the multi-level coset-admitted inter-die channel achieves **a bit-error rate of at most 1×10^{-12} at 10 Gbaud**, the channel being marginal at a bit-error rate of approximately 1×10^{-11} at a 12.5% baseline guard allocation, the at-least-16% allocation being the minimal allocation that clears the target; in a further embodiment a forward-error-correction code is provided as an equivalent crosstalk-margin alternative. *[Status: SIMULATED — FULL-WAVE EM. A full-wave openEMS (FDTD) extraction (`NVIDIA_SIM_CAMPAIGN_2026-06-13/openems_crosstalk/OPENEMS_CROSSTALK_RESULTS.md`) pins per-guarded-pair shielding at ~96% (≈30 dB, 5–20 GHz); the 16%-guard case gives BER_SU8 = 1.4×10^{-19}, clearing 1×10^{-12} with margin; a FastHenry/FasterCap cross-check PASSES (~1%). The shielding ratio is robust; absolute mV stays foundry-PDK-gated (named gate: PDK-calibrated extraction + BERT). F-05: a guard-allocation signal-integrity claim, NOT signaling-alphabet novelty; the SU(8) alphabet remains PAM-8+parity on a classical wire.]*

CL-6L-2 (independent — integrated-stack composition; the Rung-2 analog of CL-Q9-1). A three-dimensional stacked integrated circuit comprising, in combination on a columnar modular fabric in which vertical functional columns constitute repeating modules (MFT-13 fabric): (a) a hierarchical return-network inter-layer wiring topology with a dedicated per-cluster return network bounding total inter-via coupling noise to scale linearly, O(N), in the number of device layers N, rather than quadratically; (b) a spare-bank (k-of-n) redundancy of N functional and $k \geq 1$ spare device layers bounded by four nested fault-isolation domains (transistor, block, die, package), each domain having a hardware isolation latch operable to disconnect power and signal across the domain boundary on fault detection; (c) for at least one critical inter-die channel, $k \geq 4$ vertex-disjoint via routes with per-route reinforcement-weighted selection and quorum arbitration; (d) an inter-die liquid-cooling subsystem comprising a manifold branching according to a radius-cubed-conserving, cost-minimizing branching law ($r_{parent}^3 = \Sigma\, r_{child}^3$), counterflow-interdigitated inter-die channel layers, and a fractal heat-spreader layer; and (e) coset-signaling-enabled thermal-via reallocation according to CL-N2-1; wherein the elements (a)–(e) are composed on the same stack such that the composition — independent of any single realized performance number — provides yield recovery, fault tolerance, bounded coupling noise, and parallelized heat extraction for a stack of at least six active device layers. *[The composition is the claim, exactly as CL-Q9-1 on the quantum rung. No single realized number is load-bearing.]*

CL-6L-2.1 (dependent — spare-array yield, iid-gated, with the heterogeneous-role and bond-sparing fixes). The circuit of CL-6L-2, wherein N = 6 and k = 2 or k = 3, the spare-bank (k-of-n) redundancy raising a six-layer stack yield from approximately 73.5% (no spares) to approximately 99.4% (k = 2) or approximately 99.9% (k = 3) under an independent-and-identically-distributed per-die yield model with per-die yield approximately 0.95; and wherein, accounting for the heterogeneous die roles, the spares are provisioned **per-role or as universal known-good-die spares**, and each inter-die interface connection is provided by **$r \geq 2$ redundant copper-to-copper bond pillars with k-of-n acceptance** (CL-N4-8/CL-N4-9), the resulting post-fix system yield being approximately **99.37% for an 8-die universal-KGD (6+2, r=2) configuration**, approximately **99.73%** for a per-role (kA=2, kB=2, r=3) 10-die configuration, and approximately **99.00%** for a per-role (kA=1, kB=2, r=2) 9-die configuration. *[Status: bare single-class figures DERIVED (binomial); the post-fix heterogeneous-role + bond-sparing figures (99.37% / 99.73% / 99.00%) are SIMULATED by a 10^7-trial Monte-Carlo (`topo_search/INTEGRATED_CHIP_TOPO_SEARCH.md`) — recite these. The bond iid caveat is addressed by r≥2 bond-sparing; correlated die-failure modes remain; RTL fault-injection is the named gate.]*

CL-6L-3 (independent — inter-die microfluidic parallelization; force at ≥3 active layers). A cooling apparatus for a three-dimensional stacked integrated circuit having at least three active device layers, comprising: (a) an inter-die microfluidic coolant layer disposed between every adjacent pair of active device layers in the stack; and (b) a coolant distribution manifold supplying said inter-die layers in parallel; wherein each active device layer rejects heat through at most one die thickness and one bond interface into an adjacent coolant layer, such that the vertical heat-extraction topology of the stack is converted from a **series** conduction chain traversing

every die and bond interface of the stack to a plurality of **parallel** single-die extraction paths, whereby the number of active logic layers operable within a fixed per-interface bond-conductance ceiling is increased relative to a lidded series-conduction stack of equal layer count. *[Honest-magnitude binding (A7): the Cu–Cu bond ceiling (~10–50 W/cm² per interface) is parallelized, not removed; thermal is MITIGATED, not structurally solved. Force only at ≥3 layers. Numeric targets severable into CL-6L-3.1 pending the conjugate-CFD gate.]*

CL-6L-3.1 (dependent — quantified parallelization, CFD-gated; severable numeric anchor). The apparatus of CL-6L-3, wherein the series conduction chain of a six-layer lidded stack presents a vertical thermal resistance of approximately 11 K·mm²/W and each parallel single-die path presents approximately 1.5 K·mm²/W, supporting a sustained per-die power density in the range of approximately 1 to 3 W/mm² across four to six active logic layers. *[Status: DERIVED → SIMULATED by a coupled OpenFOAM 11 CFD: dense-array per-die R_th 1.45 K·mm²/W (= design central; ≈22 conservative floor carried), buried-die ΔT ≤11.1 K vs a series control's 144.6 K (×13),* `topo_search/integrated/INTEGRATED_CONFIRM_CFD.md`*. The 1–3 W/mm²/die target stays PROJECTED-TARGET, licensed for independent-claim assertion only by the ANSYS Icepak / Cadence Celsius filing-grade conjugate-CFD (OPEN); both R_th regimes travel; the bond ceiling is parallelized, not beaten.]*

CL-6L-3.2 (dependent — NEW; the dedicated coolant layer the TSVs traverse, MORE defensible than co-location). The apparatus of CL-6L-3, wherein each inter-die microfluidic coolant layer is a **distinct dedicated coolant layer interposed between an adjacent pair of device dies and separate from the inter-die signal-routing plane**, the inter-die through-silicon vias **traversing the channel walls of said dedicated coolant layer** to connect the adjacent dies; in a further embodiment the channels of the dedicated coolant layer are radius-cubed-graded ($r_{parent}^3 = \Sigma\, r_{child}^3$) and counterflow-interdigitated in a countercurrent heat-exchange arrangement at approximately 25 µm width × 100 µm height and approximately 32.5 µm pitch. *[Status: DERIVED (shared-space collapse: ~25 µm channels at ~32.5 µm pitch collide with ~35 µm TSV columns, so the dedicated layer is mandatory) → SIMULATED (bounds buried-die ΔT to ≤11.1 K vs series 144.6 K). A separate claimed coolant layer is less obvious than a co-located overlay; the true cost is Z-height (~+721 µm for an 8-die stack), not in-plane area or pump power.* `topo_search/INTEGRATED_CHIP_TOPO_SEARCH.md`*.]*

CL-6L-3.3 (dependent — NEW; TSV-through-channel-wall thermal assist). The apparatus of CL-6L-3.2, wherein a copper through-silicon via traversing a channel wall of the dedicated coolant layer acts as a **high-thermal-conductivity path conducting heat from a die to the coolant through a thin dielectric liner**, the via thereby reducing rather than increasing the per-die thermal resistance. *[Status: SIMULATED (per-die R_th 1.45 → 0.73 with the through-wall Cu via, ~301 K near the coolant; both ≪ the 2.5 falsify threshold). Thermal-only; the Cu/Si thermomechanical (CTE) stress is a separate named gate.]*

CL-6L-4 (dependent — column-as-module mapping). The circuit of CL-6L-1 or CL-6L-2, wherein each vertical functional column spanning the at least six device layers is mapped one-to-one onto a module of a modular fault-tolerant fabric (MFT-13), the column constituting the unit of yield sparing, of routing quorum, and of thermal budgeting for the stack.

J.9B-I.3 N1 signaling — coherent-only fallback + defensive posture (controlling language: NARROWED_CLAIMS_NVIDIA §2)

CL-N1-A (apparatus, independent — coherent embodiment ONLY; cross-referenced to the quantum rung). An inter-layer quantum-information signaling apparatus comprising: (a) a coherent qudit channel of dimension $d \geq 8$ that carries, in addition to an amplitude (Z-basis) degree of freedom, a **momentum-conjugate (X-basis) degree of freedom**; (b) an encoder constraining transmitted states to the simultaneous +1 eigenspace of a maximally abelian stabilizer subgroup $S = \langle X^2, Z^4 \rangle$ of the Heisenberg–Weyl group HW(d) defined by $ZX = \omega XZ$, $\omega = e^{2\pi i/d}$, wherein both the amplitude generator Z^4 and the momentum generator X^2 impose an observable constraint on the channel; (c) a projector P_S onto said eigenspace whose null space is rejected at the

transmitter; and (d) a receiver that measures **both** the amplitude and the momentum-conjugate observable to verify membership in said eigenspace and to detect a shift (X) or phase (Z) error that maps an admitted state into the null space of P_S. *[The coherent limitation of (a)/(d) is ESSENTIAL AND NON-SEVERABLE — the sole feature taking this claim out of PAM-8+parity. Enablement boundary (§2a.4): NOT enabled on a classical TSV; enabled only on the quantum rung's Hilbert-space substrate (K_6 chamber, §J.9B-II). Prosecute under the quantum rung; never broaden (a) to a single-real-amplitude channel.]*

CL-N1-A.1 (dependent — d = 8 instance). The apparatus of CL-N1-A, wherein d = 8, S = $\langle X^2, Z^4 \rangle$, the admitted eigenspace is spanned by the even-index states $\{|0\rangle, |2\rangle, |4\rangle, |6\rangle\}$, and a shift or phase error maps an admitted state into the odd-index null space $\{|1\rangle, |3\rangle, |5\rangle, |7\rangle\}$ detected at the receiver.

CL-N1-A.2 (dependent — quantum-rung firewall embodiment). The apparatus of CL-N1-A, wherein the coherent qudit channel is a cavity or bosonic mode of a quantum-error-correction substrate, and the projector P_S of (c) acts as a syndrome-admissibility firewall rejecting null-space states as a class.

J.9B-I.4 Standalone Rung-2 apparatus / architecture claims (scoped per A7 §5.3)

CL-N6 (apparatus, independent — branching-manifold counterflow liquid cooling; strongest standalone Rung-2 apparatus). A 3D-stacked-IC cooling apparatus comprising: a plurality of inter-die microfluidic channel layers each comprising interdigitated inlet and outlet channels of mutually opposite flow direction such that an outlet stream transfers heat to an adjacent inlet stream through an intervening die wall in a counterflow (countercurrent) heat-exchange arrangement; and a coolant distribution manifold supplying said inter-die channel layers; wherein the counterflow-interdigitated inter-die channel arrangement is the inter-die heat-extraction element of the apparatus. *[The counterflow-interdigitated arrangement is the fresh lead element; the radius-cubed branching law and fractal spreader are demoted to CL-N6-1 / CL-N6-2. Prosecute first among the N-series apparatus claims; conjugate-CFD + reliability gates named in App D. SIMULATED: a 7,200-topology calibrated search returns the winner geometry (25 µm × 100 µm, pitch 32.5 µm; manufacturing-limited), confirmed by a coupled OpenFOAM CFD at per-die R_{th} 1.45 K·mm²/W (≈22 conservative floor carried); ANSYS Icepak gate OPEN; the bond ceiling is parallelized, not beaten.]*

CL-N6-1 (dependent — the radius-cubed-conserving branching manifold; the demoted optimum, equation retained). The apparatus of CL-N6, wherein the coolant distribution manifold branches in a tree such that at each bifurcation the cube of the parent-channel radius equals the sum of the cubes of the daughter-channel radii ($r_{parent}^3 = \Sigma\, r_{child}^3$), said radius-cubed-conserving relation being a cost-minimizing optimum of the branching network for the stated flow-resistance cost functional. *[Recited as a provable optimum for the stated cost functional, demoted from the independent claim per the §D re-lead; equation retained verbatim, attribution dropped.]*

CL-N6-2 (dependent — fractal heat-spreader layer). The apparatus of CL-N6, further comprising a fractal (high-surface-area) heat-spreader layer of effective fractal dimension $D_f \approx 2.8–2.9$ between a topmost die and a lid.

CL-N3 (architecture, independent — return-network exponent bound). A 3D-stacked-die integrated circuit comprising $N \geq 4$ device layers organized into vertical functional columns, the through-silicon vias partitioned into a majority of intra-column local vias, a minority of lateral vias, and a smaller minority of global vias, with inter-layer returns routed through a hierarchical return network having dedicated per-cluster trunks, **such that total inter-via coupling noise scales linearly, O(N), rather than quadratically, O(N²), in N**. *[The exponent bound is the independent claim; specific wiring ratios and layer-role assignments are dependent optimizations of known hierarchical-return / NoC practice. No sub-linear scaling claimed.]*

CL-N4 (architecture, independent — four-scale fault-isolation hierarchy). A 3D-stacked-die integrated circuit comprising N functional and $k \geq 1$ spare die positions, the dies bounded by **four nested fault-isolation**

domains — transistor, block, die, and package — each domain having a hardware isolation latch operable to disconnect power and signal across the domain boundary on fault detection, the four-scale nested-isolation hierarchy with a hard isolation boundary at each scale confining a fault at the smallest enclosing scale and activating a spare at the die scale only when lower-scale confinement fails. *[k-of-n sparing alone is prior art (DRAM/FPGA); the four-scale nested hierarchy is the independent kernel.]*

CL-N4-8 (dependent — NEW; bond-sparing, the demonstrated fix for the bond-yield-dominates falsifier; believed novel — counsel [DB-CHECK]). The circuit of CL-N4, wherein each logical inter-die interface connection is provided by **r ≥ 2 redundant copper-to-copper bond pillars with k-of-n acceptance at the bond-pillar scale**, the spare-bank k-of-n redundancy thereby being applied at the bond interface as well as at the die, the interface remaining functional while at least the acceptance count of the r pillars bonds successfully. *[Status: DERIVED → SIMULATED (carried at r=2 in the 99.37% post-fix 10^7-trial MC, `topo_search/INTEGRATED_CHIP_TOPO_SEARCH.md`). Closes the §G.N4 falsifier that Cu–Cu bond yield, not die yield, dominates the stack. Believed novel over hybrid-bonding redundancy art — [DB-CHECK].]*

CL-N4-9 (dependent — NEW; design-class spare fix, making the dependent 99% yield real). The circuit of CL-N4, wherein, the N functional die positions carrying heterogeneous functional roles realized by at least two distinct die designs, the k spare positions are provisioned either **per-role** (a count of each die design, each spare able to substitute only its own design class) or as **universal known-good-die spares** each physically able to substitute either die design, such that the spare-array yield is achieved against the heterogeneous-role stack rather than a single-class binomial idealization. *[Status: SIMULATED post-fix (99.37% universal-KGD 8-die / 99.73% per-role 10-die). The bare single-class binomial over-states yield by assuming any spare backfills any role; this dependent is the THREAT-1 fix.]*

CL-N5 (architecture, independent — disjoint routing + reinforcement + quorum combination). A 3D-stacked-die integrated circuit comprising, for at least one critical inter-die channel, k ≥ 4 vertex-disjoint through-silicon-via routes; a per-route flux-reinforcement weighting that strengthens routes in proportion to carried traffic and decays unused routes; and a quorum arbitration that requires agreement of a majority of surviving routes before committing a critical transaction. *[No single primitive survives alone; the three-primitive combination is the claim.]*

CL-N5-8 (dependent — NEW; thermal-disjoint routing, demonstrated; the most defensible N5 narrowing). The circuit of CL-N5, wherein the k ≥ 4 vertex-disjoint routes are placed through **thermally-separated regions of the stack such that no single thermal-hotspot event disables more than ⌊k/2⌋ of the routes**, the channel-survival independence 1−(1−p)^k thereby being preserved as a demonstrated structural property rather than an assumption; in a further embodiment the thermally-separated regions are the separately-cooled channel lanes of an inter-die coolant manifold (per CL-N6 / CL-6L-3), which physically realizes the thermal disjointness. *[Status: SIMULATED — routing survival 0.99990 (thermal-disjoint) vs 0.9974 (clustered) by a 5×10^7-trial fault-injection + MC campaign (`topo_search/ROUTING_TOPO_SEARCH.md`); the N6 manifold's 435 cooled lanes per face (≫ k=4) supply the regions. Converts an asserted independence assumption into a demonstrated structural property, non-obvious over graph-disjoint multipath (Pavlidis–Friedman). Named gate: gate-level RTL fault-injection on a real floorplan.]*

J.9B-II — RUNG 3 (QUANTUM COMPUTING)

J.9B-II.1 The integrated-stack system claim (Q9) and its dependents

CL-Q9-1 (system, independent — the integrated stack). A fault-tolerant quantum-computing system comprising: (a) a bosonic inner code encoding each physical cavity mode in a Gottesman–Kitaev–Preskill (GKP) code of dimension d = 4, providing two binary logical encodings per mode and suppressing Z-type noise; (b) an outer asymmetric rotated surface code with X-distance $d_x = 12$ and Z-distance $d_z = 3$, the asymmetric

(d_x, d_z) = (12, 3) "binary-lever-locked" code distance being enabled by the inner-code Z-suppression of (a), the Z-distance being held at a minimum defensible value of d_z = 3; (c) a chamber projector defined on a coset $SU(N)/U(1)^{N-1}$, applied to a magic-state-factory portion of the outer code so as to admit syndromes whose pre-image lies in the chamber and reject the null-space complement, thereby replacing a magic-state distillation cascade; (d) a photonic modular fabric coupling a plurality of modules by inter-module Bell pairs; and (e) the codes of (a)–(d) arranged on a flag-manifold $K_6 = SU(3)/U(1)^2$ substrate having a holonomy phase pinned to a Standard-Model flavor-calibration value; wherein a physical-to-logical qubit ratio is computed by a composition algebra that subtracts a topological additive offset g(Δ) derived from the substrate holonomy gap. *[The composition is patentable independently of the realized overhead number. The physical-to-logical overhead is recited as a two-stage, scale-invariant floor — an EXTERNAL / SIMULATED band of ~9.6× to 14.5× (anchor ≈14.5× at (12,3); frontier ≈9.6× at (8,3)), carried alongside a conservative ceiling of ~28× should the chamber credit fail its lower bound — detailed in CL-Q9-2 and the two-stage family CL-Q9-SI-1/2/3. The floor is a memory-overhead qubit-count metric, not a runnable algorithm (the Shor-2048 budget fails at every scenario). Substrate enablement: K_6 moduli stabilized at decision grade theory-side (2026-06-12); QC-side consumption verification-gated V0–V4 (App B).]*

CL-Q9-2 (dependent — the offset coupling and the scale-invariant overhead band; severable numeric anchor). The system of CL-Q9-1, wherein the topological additive offset g(Δ) is subtracted at a gkp_base_raw step of the composition algebra, the substrate holonomy phase Φ and the spectral gap Δ being held fixed at Δ ≈ 0.260 with g(Δ) ≈ 18.305 at the (12,3) configuration; and wherein the physical-to-logical overhead computed by the composition algebra is a **scale-invariant floor in a band of approximately 9.6× to 14.5×** physical (bosonic-mode) qubits per logical qubit, the **upper anchor of approximately 14.5× corresponding to the (12,3) binary-lever-locked configuration (d_x = 12, d_z = 3)** and the **lower frontier of approximately 9.6× corresponding to the (8,3) X-verified configuration (d_x = 8, d_z = 3)**, the band being carried together with a **conservative ceiling of approximately 28×** as the design-stage position should the chamber credit fail its lower bound, and the held-fixed invariants — Atiyah–Singer index = 3, spectral gap Δ ≈ 0.260, and Standard-Model calibration — being preserved across the configuration selection. *[The recited band, anchor/frontier values, and ≈28× ceiling are EXTERNAL / SIMULATED design-stage projections of a memory-overhead floor, not measured device properties. Minimum defensible Z-distance is d_z = 3; a degenerate unencoded d_z = 1 line is not a code and is not recited. The two-stage construction is the subject of CL-Q9-SI-1. The named Kostant-exact computation (S1/S2) refines g(Δ) within the band without changing the configuration lock.]*

CL-Q9-SI-1 (method, independent — the two-stage scale-invariant overhead-floor construction; THE LEAD QC VALUE-DRIVER CLAIM, mirroring the standalone two-stage family). A computer-implemented method of provisioning physical-qubit resources for a fault-tolerant quantum-error-correction architecture derived from a fixed internal geometry, the method comprising: (a) receiving the fixed internal geometry as configuration data and holding fixed a plurality of geometry-derived structural invariants comprising at least an Atiyah–Singer index equal to 3, a spectral gap Δ ≈ 0.260, and a Standard-Model calibration; (b) computing a physical-to-logical overhead ratio of the architecture as a **scale-invariant floor that is substantially invariant with respect to the number of protected logical qubits over a range of logical-qubit counts**, by a **two-stage construction** comprising (i) a **first stage** that restricts a design search to a sub-space of candidate constructions in which the computed overhead ratio is substantially independent of the protected-logical-qubit count, and (ii) a **second stage** that performs a topological search within the restricted sub-space to select a geometry-derived code structure, the first stage being performed before the second stage such that every configuration reachable by the second-stage topological search already carries the scale-invariant-floor property; and (c) provisioning a number of physical qubits according to the said scale-invariant overhead floor and recording the floor as a quantity that does not increase with the protected-logical-qubit count across the range, wherein the recited scale-invariance is an EXTERNAL / SIMULATED design-stage property of the architecture

and not a measured property of any device, and wherein the held-fixed invariants of (a) are preserved by the two-stage construction. *[The single most valuable QC claim: the count-independent two-stage construction. The kernel is the change in functional dependence — the construction removes count-dependence and produces a floor at scale, the inverse of conventional overhead that grows with distance/count. The invariance of (b) is a constructed property of the restricted sub-space, established before any code structure is selected. Mirrors the standalone two-stage family (standalone Claims 35–37 + [0026a]); pull no scope beyond it. Eligibility: rebut "optimizing a number" with the functional-dependence change (Berkheimer non-conventionality).]*

CL-Q9-SI-2 (dependent — the evaluation point and the floor-at-scale property). The method of CL-Q9-SI-1, wherein the scale-invariant overhead floor is recited as a floor at scale and the range of logical-qubit counts over which the ratio is recorded as substantially invariant includes, in a representative embodiment, a count on the order of **ten thousand logical qubits**, the architecture recording that the ratio does not increase as the number of logical qubits grows beyond said representative count, for example to twenty thousand logical qubits, the recited invariance being an EXTERNAL / SIMULATED design-stage property of the architecture and not a measured property of any device.

CL-Q9-SI-3 (dependent — the two endpoints of the second-stage topological search). The method of CL-Q9-SI-1, wherein the second-stage topological search selects between at least a first configuration recorded as a **(12,3) binary-lever-locked configuration** corresponding to a physical-to-logical overhead anchor of approximately **14.5×** and a second configuration recorded as an **(8,3) X-verified configuration** corresponding to a physical-to-logical overhead frontier of approximately **9.6×**, each said configuration preserving the held Atiyah–Singer index equal to 3, the held spectral gap $\Delta \approx 0.260$, and the held Standard-Model calibration, the two endpoints defining a scale-invariant overhead band of approximately 9.6× to 14.5× carried together with a conservative ceiling of approximately 28×, each said overhead value being an EXTERNAL / SIMULATED design-stage projection recorded in the architecture and not a measured device property. *[The decoder-channel advantage of ~8.4× (95% CI 4.2×–16.9×), where recited elsewhere, is a decoder-channel figure, NOT a memory-overhead ratio, and is not conflated with this claim's band.]*

> *[Search-space-scope additions (ported from the standalone two-stage family, standalone §XVIII + Claims 4/5/6/3, and conformed to the LIVE Tahoe scheme: the held invariants are Atiyah–Singer index = 3, spectral gap $\Delta \approx 0.260$ (recited consistently with the existing CL-Q9-SI-1/2/3 sibling claims; the exact compiler value is 0.2589200, of which 0.260 is the documented rounding), holonomy phase $\Phi = -0.12827$ rad with $g(\Delta) = 18.305$ at the locked (12,3) geometry, and the Standard-Model calibration; every overhead recitation is the two-stage scale-invariant floor band ≈9.6×–14.5× (anchor ≈14.5× at (12,3); frontier ≈9.6× at (8,3); conservative ceiling ≈28×), evaluated SIMULATED at 10,000 logical qubits. No retired single-module point ratio and no bare overhead number is recited. F-05/RULE discipline preserved.]*

CL-Q9-SI-4 (dependent — non-convergence: admissible-region expansion and two-stage re-run). The method of CL-Q9-SI-1, further comprising, upon a non-convergence condition or a separation criterion being met in the second-stage topological search, **expanding, iterating, or re-parameterizing** the admissible region defining the said restricted sub-space — including by relaxing a curvature-positivity or triangle-inequality bound so as to open additional design space — and re-running both the first stage and the second stage over the expanded admissible region, the held-fixed invariants of CL-Q9-SI-1(a) (Atiyah–Singer index = 3, spectral gap $\Delta \approx 0.260$, and Standard-Model calibration) being re-checked and preserved across the expansion such that no overhead reduction is obtained by relaxing any held invariant; the scale-invariant overhead floor recorded after re-running remaining the SIMULATED design-stage band of approximately 9.6× to 14.5× and not a measured device property. *[Status: the admissible region is a DERIVED, re-derivable set; the standalone records that the curvature bound was itself corrected by a factor of two, opening additional design space [standalone §V, §XVIII(a)]. No retired figure recited; the expansion is recited functionally, not by a superseded percentage.]*

CL-Q9-SI-5 (dependent — selected structure outside the initially-identified sub-space). The method of CL-Q9-SI-1 or CL-Q9-SI-4, wherein the selected geometry-derived code structure **lies outside the initially-identified admissible search sub-space** and is reached by the said expansion or re-parameterization or by an alternative prior, the selected structure being **recorded with a sub-space identifier** and **re-checked against the held invariants** of CL-Q9-SI-1(a) (Atiyah–Singer index = 3, spectral gap $\Delta \approx 0.260$, and Standard-Model calibration) before acceptance, the overhead of the selected structure being recorded as the SIMULATED scale-invariant floor band of approximately 9.6× to 14.5× (conservative ceiling approximately 28×) and not a measured device property. *[Status: SIMULATED / DERIVED — supports a claim to a candidate identified anywhere, with the (12,3) lock recited only as a preferred embodiment; mirrors standalone Claim 5 / §XVIII(b). The sub-space identifier makes count-independence a constructed, recorded property re-validated after relocation.]*

CL-Q9-SI-6 (dependent — isolated optimum, not a grid-search point). The method of CL-Q9-SI-1, wherein the configuration selected by the second-stage topological search is an **isolated optimum** that does not lie on a smooth grid-search curve over the admissible region, the second stage identifying the said configuration by an overhead figure of merit notwithstanding that a uniform grid search over the admissible region does not contain it, each said configuration preserving the held Atiyah–Singer index = 3, the held spectral gap $\Delta \approx 0.260$, and the held Standard-Model calibration, and the recorded overhead being the SIMULATED scale-invariant floor band of approximately 9.6× to 14.5× and not a measured device property. *[Status: SIMULATED — the standalone records the winner as an isolated narrow basin separated from the next-best grid candidate and not on the smooth search curve [standalone §X, §XVIII(b)]; the separation supports the isolated-optimum recitation. The standalone's bare-surface-code separation figures are the bare-surface-code geometry-derived baseline (BEFORE the GKP-inner + K_6-chamber + modular-outer levers) and are SUPERSEDED — not recited here as a live ratio.]*

CL-Q9-SI-7 (dependent — the staged compiler that maps moduli to the code without additional free parameters). The method of CL-Q9-SI-1 or CL-Q3-1, wherein the geometry-derived code structure is produced by a **staged compiler** that maps the geometry moduli, **without introducing any additional free parameter**, through, in order: (i) an **admissibility gate** restricting the moduli to a curvature-positive admissible region; (ii) a **tensor-stack construction** assembling, from four scalar moduli, the metric, inverse-metric, distance, overlap, internal-Hamiltonian, and logical-projector tensors of the substrate; (iii) a **holonomy / logical-sector extraction** that diagonalizes a spin-c bundle holonomy operator to yield a logical projector P_log and a spectral gap $\Delta \approx 0.260$ ($\Phi = -0.12827$ rad at the locked geometry); (iv) **figures of merit** comprising a noise bias η and a common-mode-rejection figure r_rej, the two being **distinct recited quantities** in which $\eta \approx 3.1$ is the usable design-central noise bias and r_rej (≈18.3 to 19.3) is the rejection ratio, η and r_rej never being conflated; and (v) **asymmetric code distances** (d_x, d_z) = (12, 3) at the locked geometry; wherein every quantity of (iii)–(v) is a compiler output rather than a free parameter, and the resulting overhead is the SIMULATED scale-invariant floor band of approximately 9.6× to 14.5× and not a measured device property. *[Status: SIMULATED / DERIVED. $\eta \approx 3.1$ central conforms to Tahoe (the standalone's $\eta = 4.0$ frozen-anchor is NOT recited as the live bias); r_rej kept strictly distinct from η per RULE-9(g). Mirrors standalone Claim 3(b); the "no additional free parameters" recitation is the parameter-collapse property of the G1 stage. The standalone's bare-surface-code overhead point values are the bare-surface-code geometry-derived baseline BEFORE the GKP-inner + K_6-chamber + modular-outer levers and are SUPERSEDED — not recited.]*

CL-Q9-SI-1-BROAD (broadening recitation appended to the CL-Q9-SI-1 independent body; tethered to the held invariants). The method of CL-Q9-SI-1, the second-stage topological search of CL-Q9-SI-1(b)(ii) being further recited to cover **a configuration identified anywhere within or outside the initially-identified restricted sub-space** — including a configuration reached by the admissible-region expansion of CL-Q9-SI-4, by an alternative prior, or by a re-parameterization — provided that the configuration so identified **preserves the held-fixed invariants of CL-Q9-SI-1(a)**, namely the Atiyah–Singer index = 3, the spectral gap $\Delta \approx 0.260$, and

the Standard-Model calibration, and is recorded with a sub-space identifier; the recited scale-invariance and the recorded overhead floor of approximately 9.6× to 14.5× remaining EXTERNAL / SIMULATED design-stage properties of the architecture and not measured properties of any device. *[Drafting note: this broadening recitation closes the CL-Q9-SI-1 independent with the standalone's "anywhere within or outside the initially-identified restricted sub-space" reach (standalone Claims 1–3 closing clause, Claim 5, §XVIII), but is TETHERED to the three held invariants so it does not over-reach past the locked geometry. Added as a pure-insertion broadening clause referencing CL-Q9-SI-1(b)(ii); A-PASS may fold it into the CL-Q9-SI-1 body verbatim at numbering. No bare overhead number recited; band only, SIMULATED.]*

CL-Q9-3 (dependent — the lever set). The system of CL-Q9-1, further configured to reduce overhead by one or more of: a qudit ququart encoding per CL-Q1-1; a five-round entanglement-distillation recurrence with real-time fidelity post-selection per CL-Q5-1; a cross-layer joint decoder over the inner, chamber, and outer layers per CL-Q4-N1; outer-code removal at a squeezing level satisfying the crossover condition of CL-Q6-1; and a cavity quality factor of at least 10^8 per CL-Q7-1.

CL-Q9-4 (dependent — cross-layer element with the F1/F2 condition embedded). The system of CL-Q9-1, further comprising the joint decoder of CL-Q4-N1, wherein at least one physical ancilla is time-shared between a chamber magic-state preparation and an outer-code stabilizer cycle by a scheduling amortization that does not require any operator-commutativity condition. *[Path-4 scheduling is the presently-operative embodiment (SIMULATED); commuting-projector embodiments are expressly conditional per CL-Q4-N6/N7 — F1/F2 returned BOTH_FAIL, basis-invariant, under SU(3)/K_6.]*

CL-Q9-5 (dependent — chamber selectivity, RE-ANCHORED per RULE-3). The system of CL-Q9-1, wherein the chamber coset is selected from SU(3)/U(1)², SU(4)/U(1)³, and SU(5)/U(1)⁴, and wherein a chamber selectivity at the selected coset rank is **determined by exact Kostant weight-multiplicity counts (computation pending) at the recited coset rank**, the selectivity being characterized in two distinct senses each derived from said counts: (i) an **admitted-fraction** sense, the fraction of the syndrome space admitted by the chamber projector, equal to 1/(N+1) and **decreasing** with coset rank (DERIVED at K_6; G4); and (ii) a **firewall-rejection** sense, the fraction of the syndrome space rejected as a class by the null-space firewall, equal to N/(N+1) and **increasing** with coset rank; the decoder-workload benefit of increasing coset rank being claimed under the firewall-rejection sense only. *[RULE-3: the historical selectivity percentage bands are removed from claim language, surviving only as ESTIMATED §H background. Any chamber floor multiplier below the bare-geometry value is Gap-C-contingent and is NOT recited as a claim value. The named derivation apparatus is the Kostant generator CL-KOST-1; pending computations are the S1/S2 exactification runs.]*

J.9B-II.2 Q1 / Q2 method claims + the Kostant-generator tooling claim

CL-Q1-1 (method, independent — ququart GKP on the coset substrate). A method of encoding quantum information, comprising: providing a bosonic cavity mode resonant with a squashing-asymmetric tangent-generator root pair of a coset SU(N)/$U(1)^{N-1}$ carrying a non-zero signed-area topological invariant; encoding two logical qubits in said single mode as a d = 4 ququart GKP code whose codespace is the simultaneous eigenspace of commuting displacement stabilizers forming a Z_4 Heisenberg–Weyl group ($X_4Z_4 = iZ_4X_4$) on a lattice of primitive symplectic-cell area 8π; and performing error correction by modular-quadrature measurement, rounding each Z_4 syndrome to the nearest lattice coset and applying an inverse displacement, such that a per-cycle logical error rate is suppressed approximately as $\exp(-\pi/4\Delta\_GKP^2)$, wherein Δ_GKP^2 denotes the quadrature variance 10^(−s/10)/2 at squeezing s expressed in decibels. *[Z_4 algebra MACHINE-PRECISION verified; variance convention recited so the suppression figure is auditable.]*

CL-Q1-2 (dependent — operating point and mitigation). The method of CL-Q1-1, wherein the bosonic mode is a transmon-coupled superconducting cavity stabilized to a squeezing of at least 15 dB and targeted to 20 dB,

and wherein cross-Kerr dephasing of the ququart codewords is mitigated by Carr–Purcell–Meiboom–Gill dynamical decoupling applied to the dispersively-coupled ancilla.

CL-Q2-1 (method, independent — chamber projector on the flag manifold). A method of operating a geometric magic-state resource for a fault-tolerant quantum processor, comprising: defining, on a flag manifold $K_6 = SU(3)/U(1)^2$ at an Einstein point of the substrate metric, a Haar-average chamber projector P_chamber^(d) = (1/d) $\Sigma_{k=0}^{d-1}$ C_d^k over a Z_d center subgroup of the U(1)² toric directions, said projector being idempotent, self-adjoint, and commuting with the coset holonomy operator U_Φ; applying P_chamber to a syndrome stream so as to admit the trivial-holonomy sector and reject the null-space complement as a class; and preparing a magic state approximating a T-state by accumulating a coset-holonomy phase over d_chamber windings of a magic angle Φ that is independent of the code dimension d, the accumulated phase d_chamber × |Φ| approximating π/4 corrected by a single local Clifford.

CL-Q2-2 (dependent — higher-SU(N) extension, selectivity per RULE-3). The method of CL-Q2-1, wherein the coset is extended to $SU(4)/U(1)^3$ or $SU(5)/U(1)^4$, the null-space firewall-rejection fraction N/(N+1) thereby increasing with coset rank, and wherein the chamber admitted-fraction and chamber-multiplicity counts at the extended rank are **determined by exact Kostant weight-multiplicity counts (computation pending)**. *[CONDITIONAL: the $SU(4)/U(1)^3$ Einstein-metric derivation is the named gate (g4 family; the program's longest pole). No extended-rank selectivity percentage recited. Four SU(4) Einstein metrics PASS at G4; chamber realization remains gated.]*

CL-KOST-1 (dependent method — the Kostant weight-multiplicity generator as design tooling; the named derivation engine behind CL-Q2-1/CL-Q2-2/CL-Q9-5). The method of CL-Q2-1, further comprising computing exact spectral data of the flag-manifold substrate by an integer-gated Kostant weight-multiplicity generator, the computing comprising: (a) enumerating weight multiplicities of irreducible representations of the compact Lie group by Kostant's multiplicity formula evaluated in exact integer arithmetic, each multiplicity gated on integrality; (b) assembling therefrom exact Laplace–Beltrami spectral towers on the coset in scalar, vector, Dirac, and tensor channels to eigenvalue λ ≈ 4052, including an explicit 63-degree-of-freedom graviton-internal mode table; (c) deriving from the exact towers (i) an exact chamber-multiplicity count and admitted-fraction at a recited coset rank, (ii) an exact protected-site count N_sites,gap, and (iii) an exact spectral gap Δ; and (d) supplying the exact quantities of (c) to the selection, calibration, or verification of a chamber projector, a topological additive offset, or a physical-to-logical provisioning ratio of a quantum-error-correction stack. *[Tooling lineage: extends the prior portfolio's spectral-certification family (AEGIS $\lambda_1 \leq 0.01$ Rayleigh-certificate discipline, §J.9B-II.5) from bound-certification to exact-spectrum generation. Status: banked theory-side tooling; QC-side consumptions are the named S1/S2 exactification runs (pending — RULE-4 labels travel until landed).]*

J.9B-II.3 Q3 additive-offset method (with the built-in Δ→0 falsifier recited)

CL-Q3-1 (method, independent — topological additive offset in the qubit budget). A method of estimating and provisioning physical-qubit resources for a fault-tolerant quantum memory or processor encoded on a compact-Lie-group coset substrate, the method comprising: (a) determining a holonomy phase Φ of a spin-c connection around a generation 1-cycle of a toric base of a flag manifold $K_6 = SU(3)/U(1)^2$; (b) computing a spectral gap $\Delta = 2|\sin(\Phi/2)|$ of a holonomy operator acting on a fundamental representation of SU(3) at an Einstein point of the substrate metric; (c) computing a topological additive offset $g(\Delta) = (N_{sites,gap} / N_{log}) \cdot 2 \cdot d_x \cdot d_z \cdot (1 - e^{-\Delta/\sigma_{eff}})$, where N_sites,gap is a count of logical sites within a $\Phi \leftrightarrow \Delta$-protected chamber, N_log is a logical-qubit count, d_x and d_z are X- and Z-code distances, and σ_eff is an effective logical-noise scale; (d) computing a physical-to-logical qubit ratio as $2 \cdot d_x \cdot d_z \cdot \gamma_{ro} \cdot \gamma_{ctrl} - g(\Delta)$; and (e) provisioning a number of physical qubits according to said ratio. *[No prior-art surface-code resource formula contains an additive spectral-gap offset; the additive entry is forced by, and inseparable from, the calibration lock of CL-Q3-*

2. Operating-point numerics carried at CL-Q9-2 (SIMULATED 9.6×–14.5× floor; Δ ≈ 0.260, index = 3, SM held fixed); no swap-slot marker recited.]

CL-Q3-2 (dependent — the Standard-Model calibration lock). The method of CL-Q3-1, wherein Φ is constrained to equal a Standard-Model CP-violating holonomy phase such that increasing Δ beyond the value fixed by said phase would violate a flavor-sector calibration, whereby g(Δ) is treated as a fixed floor offset rather than a tunable parameter.

CL-Q3-3 (dependent — the Δ→0 falsifier recited as a verification step). The method of CL-Q3-1, further comprising verifying the offset by the limit checks that (i) setting Δ = 0 yields g(0) = 0 identically, the provisioning ratio of step (d) reducing exactly to the prior-art formula 2·d_x·d_z·γ_ro·γ_ctrl, and (ii) for Δ → ∞ the offset saturates at (N_sites,gap/N_log) · 2·d_x·d_z, and rejecting the offset computation as erroneous if a residual offset persists at Δ = 0. *[The built-in falsifier converts auditability into a claim limitation: if the term were a disguised constant, the recited verification fails. Spec 05 Q3.10 F1.]*

J.9B-II.4 Q4 three-layer joint decoder (controlling language: NARROWED_CLAIMS_Q4) + Q5/Q6/Q7

CL-Q4-N1 (system, independent — narrowed three-layer joint decoder). A quantum-error-correction system comprising: (a) a **bosonic inner code** (GKP) defined on a cavity mode and producing displacement-parity syndromes; (b) a **geometric admissibility chamber**, interposed between the inner code and the outer code and constituting a **distinct third code layer**, defined by a projector P_chamber onto an admissible subspace of a compact-Lie-group coset and having a non-trivial null-space N_chamber of dimension ≥ 1; (c) a **topological outer code** (surface or XZZX) producing stabilizer syndromes decoded by a minimum-weight-perfect-matching (MWPM) or union-find decoder; and (d) a **joint decoder** configured to construct a joint syndrome hypergraph whose vertices comprise syndrome bits drawn from all three of the inner code, the chamber layer, and the outer code — including at least one vertex that is a chamber-admission bit derived from P_chamber — and whose hyperedges represent error mechanisms spanning two or more of said three layers; wherein the joint decoder is configured to identify at least one hyperedge whose vertex set is explained by a layer-spanning correlated event comprising a GKP displacement error whose syndrome contribution lies in the chamber null-space N_chamber, and to cancel the contribution of that event from the outer-code syndrome — by computing an induced-syndrome estimate $\hat{\sigma}$_induced from the hyperedge structure and supplying the outer decoder a residual syndrome σ_outer ⊖ $\hat{\sigma}$_induced — **prior to** invoking said MWPM or union-find decoder; and wherein the chamber-admission information represented by said null-space vertex is co-measured as an object of the joint hypergraph, such that the cancellation distinguishes the system from a two-layer concatenation of a bosonic inner code and a topological outer code in which only classical soft information of the inner code is passed to the outer decoder and no geometric-coset chamber null-space is co-measured. **Express disclaimer (recited):** the system of CL-Q4-N1 is not a two-layer GKP–surface concatenated decoder; the third, geometric-coset chamber layer and the co-measurement and cancellation of its null-space N_chamber are essential limitations, and a decoder lacking a co-measured coset-chamber null-space is outside the scope of this claim. *[No cross-layer reduction number is asserted. Required citations (NARROWED_CLAIMS_Q4 §3): Vuillot 2019 PRA 99, 032344; Noh–Chamberland 2020 PRA 101, 012316; Poulin 2006 PRL 97, 230501 (the three closest); plus Litinski 2019; Dennis 2002; Haah 2017; GKP 2001 / Campagne-Ibarcq 2020 / Brock 2025; optionally Higgott belief-matching.]*

CL-Q4-N2 (dependent — physical ancilla sharing). The system of CL-Q4-N1, wherein at least one physical ancilla qubit is time-shared between a syndrome-extraction operation of a first one of said three code layers and a syndrome-extraction operation of a second one of said three code layers in adjacent timeslots, such that the total number of physical ancilla preparation/measurement cycles per error-correction round is reduced relative to independent per-layer extraction — said sharing being of physical measurement hardware and thereby distinct from the sharing of classical soft information alone.

CL-Q4-N3 (dependent — Path 4 / Litinski scheduling; THE CURRENTLY-ENFORCEABLE EMBODIMENT). The system of CL-Q4-N2, wherein a final-round magic-state distillation operation of a chamber-based magic-state factory and a stabilizer-extraction cycle of the topological outer code are scheduled to coincide such that a shared physical ancilla simultaneously completes the final distillation round and measures an outer-code stabilizer, thereby reducing physical-to-logical overhead by a scheduling amortization that does **not** require any operator-commutativity condition between P_chamber and the inner or outer syndrome operators. *[Status: SIMULATED; Δ = −0.08 single-module; the only currently-validated reduction-bearing embodiment. Pre-filing gate: re-confirm −0.08 net-of-correlation against a runnable scheduler artifact (RED_TEAM F-14).]*

CL-Q4-N4 (dependent — cancellation pre-pass / effective distance). The system of CL-Q4-N1, wherein, by removing the induced component $\hat{\sigma}$_induced upstream, the outer decoder solves a residual matching problem at an effective distance d_eff ≤ d_outer for the residual outer problem. *[No numeric d_outer − d_eff benefit is recited; magnitude is CONDITIONAL on the CL-Q4-N6/N7 algebra.]*

CL-Q4-N5 (dependent — correlation-cost guard / accounting-artifact falsifier). The system of CL-Q4-N4, further comprising a calibration in which a deliberately non-correlated or non-commuting configuration is required to yield zero overhead reduction, thereby certifying that any reported reduction is attributable to a genuine layer-spanning mechanism and not to an accounting artifact.

CL-Q4-N6 (dependent — EXPRESSLY CONDITIONAL; Path 1, co-measured GKP+chamber). The system of CL-Q4-N1, wherein the admissibility chamber is constructed from a coset of SU(N) with N ≥ 4 **such that the chamber projector P_chamber commutes with the GKP displacement-parity operators ([S_q, P_chamber] = [S_p, P_chamber] = 0)**, and wherein a single joint ancilla measurement reports both a GKP displacement-parity syndrome and a chamber-admission result; **the operability of this claim being contingent on a chamber construction satisfying the recited commutator condition**, which condition is not satisfied by the SU(3)/K_6 construction (measured residuals 1.18/1.32, basis-invariant). *[CONDITIONAL — gate: the SU(4) Einstein-metric chamber derivation. The projected reduction is NOT ACHIEVABLE today and not recited. Hold for continuation/CIP once an SU(4) chamber passes F1.]*

CL-Q4-N7 (dependent — EXPRESSLY CONDITIONAL; Path 2, chamber test as outer post-processing). The system of CL-Q4-N1, wherein the chamber is constructed **such that P_chamber commutes with every outer-code stabilizer generator S_k ([P_chamber, S_k] = 0 for all k)**, and wherein the chamber-admission test is recovered as a linear post-processing function of the outer stabilizer syndrome stream without a dedicated chamber-syndrome measurement round; **the operability of this claim being contingent on the recited commutator condition**, which is not satisfied under SU(3)/K_6 (measured residuals 1.00–2.06, basis-invariant). *[CONDITIONAL — same gate and posture as CL-Q4-N6.]*

CL-Q5-1 (method, independent — modular distillation with real-time post-selection). A method of operating a modular quantum computer comprising a plurality of modules connected by heralded photonic links, the method comprising: heralding raw entangled pairs across a link; performing a fixed five-round 2→1 entanglement-distillation recurrence on the heralded pairs to produce distilled pairs; estimating, in real time from the per-pair heralding record, a fidelity of each distilled pair; admitting for inter-module operations only those distilled pairs whose estimated fidelity exceeds an acceptance threshold and discarding the remainder; and consuming an admitted distilled pair in an inter-module logical operation whose boundary error-detection depth is reduced relative to the bulk error-detection depth in dependence on the reduced inter-module error rate so produced.

CL-Q5-2 (dependent — thresholds and the modular multiplier; rate-deficit disclosed). The method of CL-Q5-1, wherein the acceptance threshold is F > 0.9995, the recurrence is seeded by a heralded fidelity F_heralded = $1 − (1 − \eta)^2$ with link efficiency $\eta \geq 0.95$, and admitting only pairs above the threshold reduces a modular overhead multiplier applied to the physical-to-logical ratio to approximately 1.05; and further comprising re-

enabling bulk-depth boundary correction whenever the link efficiency falls below 0.95. *[Honesty binding (RULE-9(i)): wherever the 1.05 multiplier appears, the disclosed rate deficit travels — at 10K-scale the distillation+post-selection chain is ~4.3–5× short of the required Bell-pair throughput absent a photonic raw-rate uplift; the overhead reduction is real but rate-infeasible today.]*

CL-Q6-1 (method, independent — outer-code removal at the crossover). A method of reducing physical-qubit overhead in a fault-tolerant quantum computer having logical qubits each encoded by an inner GKP bosonic code wrapped by an outer topological code, the method comprising: stabilizing the GKP code at a squeezing level at which the inner-code logical error rate per cycle, computed as $\frac{1}{2}\cdot\mathrm{erfc}(\sqrt{\pi}/2\sqrt{(\Delta\_GKP^2)})$ in the variance convention $\Delta\_GKP^2 = 10^{\wedge}(-s/10)/2$, is below a gate-injection error rate of the outer code's syndrome extraction; and, upon satisfaction of said crossover condition, reducing the outer code distance such that each logical qubit is carried by the GKP mode with a reduced or removed outer topological code, residual logical-sector tracking being performed by a coset-admittance chamber projector in place of the removed outer code.

CL-Q6-2 (dependent — the BOUNDED operating window and the d_outer ≥ 2 fallback). The method of CL-Q6-1, wherein the squeezing level is approximately 20 dB sustained by a cavity having internal quality factor $Q \gtrsim 3\times10^8$, or $Q = 2\times10^8$ with a thermal-reservoir engineering bandwidth of approximately 5.4 kHz; and wherein, at a sustained squeezing of approximately 17.6 dB corresponding to $Q = 2\times10^8$ at a 5 kHz reservoir, the outer code distance is reduced to $d\_outer \geq 2$ rather than removed. *[Status: BOUNDED (g5). Full outer-code removal is licensed only inside the recited Q/reservoir window; the d_outer ≥ 2 fallback is the recited honest floor.]*

CL-Q7-1 (apparatus, independent — high-Q cavity with multi-pole Purcell co-design). A microwave cavity apparatus for hosting a bosonic error-correction code, comprising: a niobium cavity body defining a re-entrant mode geometry that concentrates the cavity electric field in a sub-millimeter gap region spaced from the cavity walls; an interior surface prepared by buffered-chemical-polish or electropolish followed by a nitrogen-infusion bake; and a transmon ancilla dispersively coupled to the re-entrant mode through a **multi-pole Purcell filter of at least two poles providing at least 45 dB rejection at the storage-cavity frequency**, at a filter bandwidth of approximately 50–100 MHz and a cavity–filter detuning of approximately 0.3–0.8 GHz; the apparatus exhibiting an internal quality factor $Q \geq 10^8$ in the 7–8 GHz band while maintaining a dispersive coupling $\chi/2\pi \geq 2$ MHz between the ancilla and the mode.

CL-Q7-2 (dependent — the co-design and the squeezing service). The apparatus of CL-Q7-1, wherein the at-least-two-pole filter is integrated within the re-entrant cavity geometry as a co-designed structure, the apparatus sustaining approximately 20 dB squeezing of a GKP code state by maintaining cavity internal loss $\kappa\_int/2\pi < 1$ kHz, thereby supplying the crossover condition of CL-Q6-1. *[The Purcell EM solve is the named filing-grade gate; unfiltered collapse is to $Q \sim 10^4$ — approximately four orders below the recited Q.]*

J.9B-II.5 Prior-portfolio claims incorporated at claim level (compact recitation of the four strongest families)

> *[The 28-family portfolio is inlined at claim-chart level in §H.Q10; the four strongest are recited here compactly as claims of THIS filing. All four are SIMULATED (Stage-C 9/9 regressions PASS) or STRUCTURAL as labeled.]*

CL-P21 (method, independent — Article-A six-constraint substrate-admissibility filter; family #21, highest geometry-side). A method of selecting a geometric substrate for a quantum-error-correction architecture, comprising admitting a candidate manifold only if it simultaneously satisfies all six of: (i) real dimension six; (ii) compactness; (iii) a transitive SU(3) isometry action; (iv) a chiral Dirac index of 3 realized via a spin-c (1,0) bundle; (v) a tripartite tangent-space splitting; and (vi) a consistent spin-c structure; and instantiating the quantum-error-correction substrate on the admitted manifold, whereby the substrate inherits a Standard-Model flavor calibration, a three-family index count, and a natural Z_2 grading. *[STRUCTURAL — Stage-A enumeration*

of 62+ candidates admits uniquely $K_6 = SU(3)/U(1)^2$. The kernel is the substrate-selection discipline (physics-grounded filter vs. hardware-convenience choice).]

CL-P22 (method, independent — Φ↔Δ bridge calibration anchor; family #22). A method of calibrating a quantum-error-correction design parameter against a particle-physics observable, comprising: relating a substrate holonomy phase Φ to a code spectral gap Δ by the bridge identity $\Delta = 2|\sin(\Phi/2)|$; pinning Φ to a Standard-Model CP-violation phase; optimizing a quantum-error-correction resource quantity as a function of Δ; and, upon any proposed change to Δ, re-validating the substrate against the Standard-Model flavor-sector observables before accepting the change. *[STRUCTURAL + indirect-MC. Ancestor of CL-Q3-1/Q3-2; recited independently because the bridge discipline is broader than the offset formula.]*

CL-P14 (method, independent — FOUR-LAYERED hard-composition-gate preservation discipline; family #14, strongest META). A method of composing error-correction mechanisms on a calibrated geometric substrate, comprising enforcing four composition-preservation tiers — a code-time gate, a runtime gate, an ensemble gate forcing a bottom ensemble to collapse to its original flat predicate, and a cycle-period gate — each tier having its own hard gate that must pass before a mechanism is admitted, the four gates composing multiplicatively such that a flat-predicate semantics, a Standard-Model calibration, a chirality index of 3, and a substrate spectral gap Δ are preserved verbatim across all four regimes. *[SIMULATED — end-to-end PASS across the composing regression suites. The preservation guarantee the CL-Q9-1 multi-layer stack consumes.]*

CL-P16 (system, independent — cumulative composability matrix; family #16, strongest single prior claim). A quantum-error-correction system comprising a topologically-grounded $K_6 = SU(3)/U(1)^2$ substrate; a composability matrix of at least nine mechanism slots comprising at least five base mechanisms and cascade, proportional-integral, ensemble, and scheduling extensions, the mechanisms composing multiplicatively as $\Pi(1-\delta_i)$; the four-tier preservation gate of CL-P14; and a gain-matrix extension discipline whereby a new controlled variable is added to a multi-input controller by appending a zero-initialized row and column and tuning without redesign; the system preserving a Standard-Model calibration, an index of 3, and a substrate gap $\Delta \approx 0.260$ across all slots. *[SIMULATED — all five composing regression suites PASS. The carrier into which the GKP/chamber/outer/modular layers of CL-Q9-1 slot as additional orthogonal mechanisms.]*

J.9B-III — RUNG 1 (EASY / EVIDENTIARY)

J.9B-III.1 E2 apparatus

CL-E2-1 (apparatus, independent — paired-transmon differential qubit site). A superconducting quantum-processor qubit site comprising a first and a second transmon coupled by an exchange interaction, wherein quantum information is encoded in the antisymmetric (differential) superposition of the two transmons and a readout/control subsystem addresses the differential mode while perturbations common to both transmons are rejected into the symmetric (common) mode at a common-mode rejection ratio $r_{rej} > 5$, such that correlated dephasing is suppressed relative to uncorrelated relaxation and the site exhibits a structural noise bias $\eta = p_X/p_Z > 2$ arising from the pair coupling and the matching of the two transmons. *[r_{rej} (rejection ratio, ~19; dual provenance 19.25 geometry-path / 18.27 adapter-path) and η (usable bias, ~3.1) are distinct recited quantities, never conflated (RULE-9(g)). All values SIMULATED/DERIVED, not MEASURED; the on-device measurement is the named post-filing validation.]*

CL-E2-2 (dependent — rejection mechanism detail). The site of CL-E2-1, wherein $r_{rej} > 18$ and is given by $r_{rej} = 1/\Sigma_k c_k|\delta_k|$ in which δ_k are normalized mismatch parameters between the first and second transmons (including at least a frequency mismatch δω and a coupling mismatch δg) and c_k are sensitivity coefficients of the pair Hamiltonian, the bias η being computed from the pair-Hamiltonian parameters with the effective dephasing rate reduced by r_{rej} on its correlated fraction only, and the two transmons sharing a common readout resonator and local flux environment within a site extent ≤ 40 μm.

J.9B-III.2 E3 device-coupled method (controlling language: NARROWED_CLAIMS_E3 Claims 1–6, near-verbatim)

CL-E3-1 (method, independent — configuring and operating a biased-noise code from a device-derived bias). A method of configuring and operating quantum error correction for a superconducting quantum processor, the method comprising: **(a)** providing the processor with a plurality of qubit sites, **each qubit site comprising a first transmon and a second transmon coupled by an exchange interaction and configured to encode quantum information of the site in an antisymmetric differential mode of the two transmons**, such that perturbations common to both transmons are rejected into a symmetric common mode at a common-mode rejection ratio r_rej; *[ESSENTIAL — names the hardware bias source]* **(b) measuring, on the processor, the common-mode rejection ratio r_rej** of the qubit sites by injecting a correlated perturbation on the first and second transmons of at least one site and quantifying the suppressed response in the differential mode, and characterizing a correlated-error fraction of the site dephasing; *[ESSENTIAL — the bias is measured, not assumed]* **(c) deriving a noise-bias parameter η = p_X_eff / p_Z_eff from the measured r_rej**, wherein an effective dephasing rate p_Z_eff is obtained by reducing a correlated fraction of a physical dephasing rate by the measured r_rej while an uncorrelated fraction passes through, and an effective bit-flip rate p_X_eff is obtained from a physical relaxation rate, such that η is computed from the said measured hardware common-mode rejection rather than supplied as an external assumption; *[ESSENTIAL]* **(d)** selecting, for an XZZX surface code, a first code distance d_x along a first axis and a second code distance along a second axis as **d_z = max(d_z_floor, ⌈d_x/η⌉)**, wherein d_z_floor ≥ 3 is a safety floor and η is the device-derived parameter of step (c); *[ESSENTIAL — distances selected from the device-derived η, under a floor]* **(e)** laying out and operating the XZZX surface code on the processor with d_x data qubits along the first axis and d_z data qubits along the second axis, performing stabilizer measurement, and decoding the syndromes along each axis under a biased-noise model in which the Z-axis error rate is the said reduced p_Z_eff; and **(f)** thereby achieving a reduced physical-to-logical qubit overhead at a fixed target logical error rate relative to a symmetric XZZX code operated on a processor without the said common-mode rejection; **wherein the noise-bias parameter η used in step (d) is the parameter derived in step (c) from the measured common-mode rejection ratio r_rej of the differential transmon-pair qubit sites, and is not a value assumed independently of the said qubit sites.** *[ESSENTIALITY ANCHOR: the closing wherein clause is the firewall. Bonilla Ataides 2021 / Tuckett 2020-2021 teach (d) for an ASSUMED η, none of (a)–(c). If (a) is broadened to "a hardware mechanism," or (b)–(c) deleted, the claim reads on that art and MUST NOT be filed in that form. §101: physical apparatus + syndrome operations recited. §112: the on-device r_rej measurement of (b) is the named OPEN enablement gate.]*

CL-E3-2 (dependent — the r_rej mechanism pinned). The method of CL-E3-1, wherein the common-mode rejection ratio is given by $r_{rej} = 1/\Sigma_k c_k|\delta_k|$, in which δ_k are normalized mismatch parameters between the first and second transmons of a site (including at least a frequency mismatch $\delta\omega$ and a coupling mismatch δg) and c_k are sensitivity coefficients of the pair Hamiltonian, and wherein the measured r_rej is at least 10 (in a further embodiment at least 18), such that the derived noise bias η is at least 3 at a correlated-dephasing fraction of at least 0.8.

CL-E3-3 (dependent — safety floor + bias-headroom margin). The method of CL-E3-1, wherein d_z_floor = 3, and wherein the selecting step further applies a bias-headroom margin such that the selected d_z is held above the minimum value permitted by η by a headroom fraction (approximately 0.30, corresponding to a d_z-aggression factor of approximately 1.5), whereby the method refuses to select an under-protected distance $d_z \in \{1, 2\}$ when the realistic device-derived η is small and retains a safety margin against bias-estimation error.

CL-E3-4 (dependent — re-measurement / closed calibration loop). The method of CL-E3-1, further comprising re-measuring the common-mode rejection ratio r_rej of the qubit sites during operation of the processor on a recurring calibration schedule, re-deriving the noise-bias parameter η from each re-measurement, and re-selecting $d_z = \max(d_{z\_floor}, \lceil d_x/\eta \rceil)$ when the re-derived η crosses a threshold, whereby the code

geometry tracks drift in the device-measured bias. *[Closed-loop hardware-software co-design with no antecedent in the bias-code literature; continuation seed for a "biased-code controller" apparatus claim.]*

CL-E3-5 (dependent — essentiality validation; the decoded anchor). The method of CL-E3-1, further comprising validating that the said device-derived bias is essential to the distance reduction by simulating, at equal code distance d_z, the minimum-weight-perfect-matching-decoded logical-Z error rate with and without the said common-mode rejection, and confirming the with-bias logical-Z error rate is lower by at least a factor of four at a device-central bias $\eta \approx 3.1$, **the said factor being a same-distance logical-Z error reduction of approximately 8.4× (95% confidence interval 4.2× to 16.9×) at d_z = 3 and $\eta \approx 3.1$**, as established by a Monte-Carlo decoding campaign of approximately 1×10^8 decoded shots, rising toward and past two orders of magnitude as the bias approaches the full common-mode rejection ratio. *[The decoded figure; the prior undecoded raw-syndrome statistic is retired and never recited.]*

CL-E3-6 (dependent — the overhead-reduction band). The method of CL-E3-1, wherein d_x = 13 (or, in a further embodiment, d_x = 12) and the selected $d_z \in \{3, 4, 5\}$ corresponds to a device-derived η of approximately $\{\geq 4.3, 4, 3.1\}$ respectively, yielding a physical-to-logical overhead reduction of approximately **37% to 62%** (overhead approximately 269 reduced to approximately 168, 134, or 101 physical qubits per logical qubit) relative to a symmetric d_x = d_z = 8 XZZX baseline at equal logical error rate. *[Keep the BAND, never the single best case — written-description support spans the η band.]*

J.9B-III.3 E1 wakeup-chip embodiment (reduction-to-practice dependent)

CL-E1-RTP (dependent — production-silicon embodiment of the lead system claim). The circuit of CL-N2-1, embodied as a two-die sensor-wake integrated circuit comprising a first die carrying wake-detection logic and a second die carrying an analog sensor frontend, the dies joined by a copper-to-copper hybrid bond at approximately 8 µm pitch and a 256-via through-silicon-via array in a 16×16 arrangement allocated as approximately 128 signal vias (50%) carrying the d = 8 $\langle X^2, Z^4\rangle$ coset-admitted alphabet on 64 differential pairs, 32 dedicated thermal vias (12.5%), 32 power vias, 32 ground-return vias, and 32 guard vias; wherein the 32 dedicated thermal vias occupy slots freed by the coset alphabet's reduced signal-via count relative to a binary-signaling interface of equal aggregate bandwidth — a thermal allocation of 12.5% versus 0% for the equal-bandwidth binary baseline — thereby practicing, in production silicon, the conditioned thermal-via reallocation of CL-N2-1. *[IP-role lock (A8): purely EVIDENTIARY — reduction-to-practice of the CL-N2-1 system trade plus commercial proof — NEVER a signaling-novelty claim. The prior d = 8 "minimum modulus for linear-code-kernel non-membership" recital is WITHDRAWN per F-05 (STMT-N1-B). Measurement package (silicon-pending): KPI-1 BER ≤ 10^{-9} at 1 Gbps/pair over 10^{12} bits; KPI-2 ≥ 40% coset-detector false-codeword rejection vs an 8-PAM baseline; KPI-3 IR-thermography within ±30%.]*

J.9D — DEFENSIVE STATEMENTS

STMT-N1-B (defensive acknowledgement — classical amplitude embodiment; NO novelty asserted; controlling language verbatim). Applicant acknowledges that, in the **classical, amplitude-only embodiment** — wherein the through-silicon via transmits a single real voltage level per symbol (the Z/diagonal basis only) — the admitted level set {0, 2, 4, 6} of the d = 8 stabilizer $\langle X^2, Z^4\rangle$ is **identical to the set of eight-level pulse-amplitude (PAM-8) symbols having a least-significant bit of zero**, i.e. PAM-8 gated by a single parity bit; that this set **is** a binary linear code, being exactly the kernel of the single GF(2) parity check H = [0 0 1]; and that, in this amplitude-only embodiment, the second stabilizer generator X^2 imposes no observable constraint and the construction is therefore **anticipated by, and is not asserted to be novel over, PAM-8 + parity** and the multi-level signaling references of record. Applicant claims no patentable distinction for the classical amplitude-only embodiment; the coset framing in that embodiment is a derivation-only relabeling of PAM-8 + parity. *[Purpose:*

forecloses inequitable-conduct/§112 exposure from the executed F-05 retraction; bounds CL-N1-A by explicit contrast (amplitude-only excluded); executes the spec-03 line-465/504 reconciliation. A statement, not a claim.]

J.9E — DEFENSIBILITY MAP (for counsel)

A per-claim defensibility map (exposure rating, load-bearing distinguishing limitation, and FILE-NOW/CONDITIONAL/DEFENSIVE/EVIDENTIARY status for each claim) is maintained as a separate attorney work-product companion.

J.9E.1 Claim-count summary

Category	Independent	Dependent	Defensive statements	Total
Foundational (CL-F)	2	0	0	2
Rung 2 — N2 lead chain	1	6	0	7
Rung 2 — 6-layer	3 (6L-1, 6L-2, 6L-3)	6 (6L-1.1, 6L-2.1, 6L-3.1, 6L-3.2, 6L-3.3, 6L-4)	0	9
Rung 2 — N3/N4/N5/N6 + dependents	4	5 (N4-8, N4-9, N5-8, N6-1, N6-2)	0	9
Rung 2 — N1 coherent fallback	1	2	1 (STMT-N1-B)	3 + 1 stmt
Rung 3 — Q9 system chain (+ two-stage scale-invariant family)	2 (Q9-1, Q9-SI-1)	11 (Q9-2…Q9-5, Q9-SI-2, Q9-SI-3, Q9-SI-4, Q9-SI-5, Q9-SI-6, Q9-SI-7, Q9-SI-1-BROAD)	0	13
Rung 3 — Q1/Q2 + Kostant	2	3 (Q1-2, Q2-2, KOST-1)	0	5
Rung 3 — Q3 offset chain	1	2	0	3
Rung 3 — Q4 narrowed chain	1	6 (incl. 2 expressly conditional)	0	7
Rung 3 — Q5/Q6/Q7	3	3	0	6
Rung 3 — prior portfolio (compact)	4	0	0	4
Rung 1 — E2/E3/E1	2	7 (E2-2, E3-2…6, E1-RTP)	0	9
TOTAL	**26**	**51**	**1**	**77 claims + 1 statement**

> *[A-PASS note: the WA-5 self-certification table and status-preservation footer of this Part are retained in the source file `PART_E_CLAIMS.md` and in `A_PASS_COMPLIANCE_REPORT.md`.]*

End of §J.

---

§K — APPENDICES

Appendix K1 — Load-bearing curvature invariant of record

The K_6 = SU(3)/U(1)² Riemann-squared invariant of record is **$|Riem|^2 = 276$** (DERIVED, banked first-principles), with $|Weyl|^2 = 216$ (non-zero — K_6 is **not** conformally flat) and scalar curvature R = 30. The value supersedes a prior **$|Riem|^2(K_6) = 60$** (conformally-flat / maximally-symmetric, implying $|Weyl|^2 = 0$ — qualitatively wrong for SU(3)/T²); the superseded value is retained only as a retirement record (Appendix K4 R1). The corrected invariant does **not** mechanically move the overhead chain: $\Delta = 0.2589200$, $\Phi \approx -0.12827$ rad, and $g(\Delta) = 18.305$ are spectral/holonomy/counting objects that consume no $|Riem|^2$; the blast radius of any change to $|Riem|^2$ is confined to the UV/heat-kernel evidence layer.

Appendix K2 (App B) — Substrate-enablement V0–V4 verification trail

This appendix records the **Gap-04 substrate-enablement gate**: the named **V0→V4 verification trail** by which the theory-side substrate result (the K_6 = SU(3)/U(1)² moduli, stabilized at decision grade theory-side, 2026-06-12) is propagated into the QC-side overhead and selectivity claims of §H. It is recorded here as a **named open gate, not a closed result** [OPEN].

- **Status [CONDITIONAL — verification-gated].** QC-side consumption of the substrate result is gated through V0→V4; no §H quantitative statement is promoted above its label by this trail until the corresponding stage lands. The exactification dependency is tracked in Appendix K4.3: **S1/S2 are LANDED** (exact $\Delta = 0.2589200$ and the exact protected-sector mode count → $g(\Delta) = 18.305$ at (12,3); exact chamber multiplicity counts at K_6/K_{12}), and **S3 is pending** (the heat-kernel re-assembly with $|Riem|^2 = 276$, which re-opens R1/R2/R3 in the UV-evidence layer only and does **not** touch the overhead floor).
- **Named residuals to close [OPEN].** The trail's remaining residuals are the **loop-coefficient value**, the **magnitude refinement**, the **σ-kinetic normalization**, and the **NT-1 transverse residue**. Each is a named gap (Appendix K3 legend, label OPEN) requiring a specific derivation or simulation; none is asserted as DERIVED here.
- **Re-pointing discipline.** The native-scaffold repoint of the superseded modulus-potential inputs (Appendix K4 R4: $c_{KK} = -8.892\times10^{1}/R_Y^4$ negative, $c_{bdry} = -1.08\times10^{-2}$, Wilson sums $S(0) = +4$ / $S(\pi) = +92$) is performed **per V0→V4**; the superseded inputs survive only as retirement records in Appendix K4.
- **Blast-radius bound.** Closing V0→V4 does **not** change the operative overhead posture: the two-stage scale-invariant **≈9.6×–14.5×** floor band (anchor (12,3) ≈ 14.5×, frontier (8,3) ≈ 9.6×, conservative ≈ 28× ceiling) is computed from spectral/holonomy/counting objects (Δ, Φ, g(Δ)) that consume no UV-evidence-layer quantity (Appendix K1 blast-radius note).

Appendix K3 (App C) — Status-label legend and the three self-correction arcs

Every quantitative statement carries exactly one label:

Label	Meaning
MEASURED	Demonstrated in hardware (ours or cited literature).
SIMULATED	Produced by a runnable simulation in this corpus (open-source or first-principles).
DERIVED	Follows analytically from stated assumptions / geometry / Hamiltonian.
PROJECTED-TARGET	A design goal or extrapolation, not yet simulated or measured at the stated scale.
BOUNDED	A constraint characterized within stated limits (the window exists at a named price; not yet delivered).
CONDITIONAL	Unblockable only under a named, not-yet-realized condition.
OPEN	A named gap requiring a specific external demonstration or derivation.
SPECULATIVE	A hypothesized refinement not yet supported by derivation or simulation.

K3.2 — The three self-correction arcs (the credibility narrative). This application records, rather than hides, three places where an earlier internal value was corrected. Each arc is carried at its primary site of record and is consolidated here:

1. **The curvature-invariant correction arc.** $|\mathrm{Riem}|^2(K_6)$ was corrected from the stale **60** (conformally-flat, implying $|\mathrm{Weyl}|^2 = 0$ — qualitatively wrong for $SU(3)/T^2$) to the banked **276** ($|\mathrm{Weyl}|^2 = 216 \neq 0$; DERIVED). The corrected value is the invariant of record (Appendix K1); the stale value survives only as a retirement record (Appendix K4 R1), and the change is confined to the UV/heat-kernel evidence layer — it does **not** move Δ, Φ, or $g(\Delta)$.
2. **The decoded-versus-undecoded re-measurement arc.** The earlier **undecoded raw-syndrome-flip ratio is retired** as an undecoded artifact; the load-bearing figure is the **MWPM-decoded same-distance logical-Z reduction of 8.4× (95% CI 4.2–16.9)** [SIMULATED] (§I.1), held strictly distinct from the memory-overhead floor (Appendix K4 RULE-9(d)).
3. **The sub-Kitaev retirement arc.** The earlier **"sub-Kitaev" (single-module < 1.0) framing is retired** as a floor-clamp/algebra-clamp artifact, and its pre-clamp value is not reproduced (Appendix K4 RULE-9(c)). The operative posture is the two-stage scale-invariant **≈9.6×–14.5×** floor band with the conservative **≈28×** ceiling (Appendix K4 RULE-9(a)).

Appendix K4 — The quarantine / banned-number retirement ledger

The **single place** in this application where quarantined items may be named, and only as supersession/quarantine records. For each: what is quarantined, why, and what re-opens it.

K4.1 — The quarantine ledger (R1–R7).

Row	Quarantined item(s)	Why quarantined	What re-opens / resolves it
R1	The stale curvature-contraction artifacts: **$\|Riem\|^2(K_6) = 60$**, block-additive **64**, cubic **2048**, and the implied **$\|Weyl\|^2 = 0$**	Stale invariant; the conformal-flatness implication is qualitatively wrong for $SU(3)/T^2$	Recompute with the banked **276** and the banked curvature code (run S3); keep the old file as frozen history with a supersession note (K1)
R2	The UV-evidence obstruction numbers: **$\|B1/B2\| = 14.2489$** (and "14.25"), "obstruction present," "n_positive_g_NGFP = 0," "NGFP not found," and the "13.0 → 14.25 away-from-saturation" language	These consumed the stale $\|Riem\|^2 = 60$; the dominant quadratic invariant changed by ≈4.6×	**BANNED from all filed text** until the S3 re-run lands; re-opened pending re-run (RULE-1)
R3	The single **disputed certificate** whose truncation basis includes the superseded invariant class	Same invariant class in its operator flow	Carry as **exactly one disputed certificate** — neither dropped nor multiplied — pending its invariant-value audit; **not cited as support anywhere** (RULE-7)
R4	The superseded modulus-potential inputs: **$c_{KK} = +0.0095$** (wrong sign vs the native $-8.892\times10^{1}/R_{Y}^{4}$), $c_{bdry} = -1.267\times10^{-2}$, $c_{Wilson} = -4.88\times10^{-5}$, and any sentence grounding modulus stability in "FRG-4 V_eff"	Superseded by the native scaffold (the positive c_KK sign is the load-bearing error)	Rewrite to the **native scaffold** with **$c_{KK} = -8.892\times10^{1}/R_{Y}^{4}$ (negative)**, $c_{bdry} = -1.08\times10^{-2}$, Wilson sums $S(0) = +4$ / $S(\pi) = +92$; repoint per V0–V4; preserve old certificates (RULE-2)
R5	The verbatim selectivity percentages **"approximately 8–15%, 25–50%, and 50–75%"** in claim language, and the **0.375 / 0.30** chamber multipliers asserted as derived	Admitted-fraction = 1/(N+1) **decreases** with rank (20% at K_{12}); the sub-0.60 multiplier needs the un-derived Gap-C credit (bare geometry computes ≥1.1–1.5)	**Removed from claim language**, replaced by *"chamber selectivity determined by exact Kostant weight-multiplicity counts at the recited coset rank"* (counts from landed run S2); 0.375/0.30 always flagged **Gap-C-contingent** (RULE-3)
R6 (RESOLVED)	The retired estimate-fed forks: the constancy-reading **$g(\Delta) = 12.45$ → fabric 1.25–1.43**, the formula-scaled-reading **$g(\Delta) \approx 6.23$ → ≈1.88**, the **F1 fallback 2.175 fabric**, and the **$d_z = 1$** degenerate line	Both readings were estimate-fed pending the exact protected-sector mode count; now superseded by the landed exactification	**RESOLVED by S1/S2 (landed).** Values of record: **$\Delta = 0.2589200$, $g(\Delta) = 18.305$ at (12,3)** (constancy/chamber-site reading confirmed); the floor is the two-stage ≈9.6×–14.5× band. The left-column numbers appear **only here, as a retirement record** (RULE-4). σ_eff remains the one ESTIMATED input. d_z = 1 is retired as a code (min defensible d_z = 3).
R7	The cascade-gate prose ("gated on S-7" / "blocked awaiting ΔV_NP")	The gate premise changed from "blocked awaiting a non-perturbative wall" to "discharged by the native scaffold"	Update to: *"the S-7 modulus-stability premise is discharged by the native scaffold (countersigned 2026-06-12); per-gap statuses change only by their own re-runs and countersigns"*; blocker **B-UQFC-14-NP-1 is DISSOLVED, not satisfied** (RULE-8)

K4.2 — Carried bans and caveats (RULE-5 and RULE-9). Banned from the operative 100-page text except where this appendix names them as quarantine records; the assembly pass greps for each:

- **The aggressive compiler headline.** The conservative Stage-B figure is **≈22–30 qubits/logical**; the aggressive figure is **BANNED from the entire text**, including the Q8 baseline comparison (restated against ≈22–30) and any provenance line (RULE-5). The residual +26.6% curvature-richness factor is disclosed with its ruling pending.
- **The cross-layer reductions. −0.46 and −0.59 are never quoted as achievable** — CONDITIONAL on a physically-realized disjoint-factor SU(4) chamber; **Path-4 −0.08 is the only validated figure** (RULE-9(b)).
- **"Sub-Kitaev" stays retired** (and the 0.614 pre-clamp figure never appears) — a floor-clamp artifact (RULE-9(c)).
- **The undecoded raw-flip ratio stays retired** (an undecoded artifact); the decoded figure is **8.4× (CI 4.2–16.9)** (RULE-9(d); §I.1).
- **The 15–25% R_th never appears without its correction** (≈6.5% SIMULATED at 4,800 vias / 50 mm²; density-dependent; ANSYS Icepak is the filing-grade gate); claim language keeps the numeric range severable (RULE-9(e)).
- **The local thermal figures (213×, 0.372 W/K, 275×) are never system numbers**, and the τ_RC = 1 ns error stays corrected (1 ps; RC not binding) (RULE-9(f)).
- **The bias-rejection ratio (≈19) and the usable bias (≈3.1) are never conflated** (RULE-9(g)).
- **The retired composition-floor framing stays retired (RULE-9(a)).** The lever-success bin-artifact probability (formerly read against a 1.73 "TARGET"), the figures **1.7286 / 1.73-as-target / 1.58 nominal**, the probability **12.6%** (and the earlier 91%), and the **3.24:1 "Robust fallback"** are all superseded and appear nowhere as live values; named only here as retirement records. The operative posture is the **two-stage scale-invariant ≈9.6×–14.5× floor band** (anchor (12,3)≈14.5×, frontier (8,3)≈9.6×) with the **conservative ≈28× ceiling**.
- **The orthogonality caveat travels** wherever the ratio appears — the floor ratio is a qubit-count metric and the Shor-2048 error budget FAILS at every scenario including R = 1.0 (RULE-9(j)).
- **No standalone classical d ≥ 8 signaling-novelty claim** — it is OPEN/weak; the coset does genuine work only in the coherent quantum embodiment (RULE-9(k)).

K4.3 — Exactification dependency (S1/S2 LANDED; S3 pending). The formerly estimate-fed floor (R6) and selectivity claims (R5) became exact via the banked Kostant generator: **S1** (exact $\Delta = 0.2589200$ plus the exact protected-sector mode count — settled $g(\Delta) = 18.305$ at (12,3)) and **S2** (exact chamber multiplicity counts at K_6 and K_{12}). The landed posture is "defending an exact floor" (the two-stage ≈9.6×–14.5× band), not "reaching a target." One item remains pending: **S3** (the heat-kernel re-assembly with $|\mathrm{Riem}|^2 = 276$ — re-opens R1/R2/R3 in the UV-evidence layer only), which does not touch the overhead floor.

Status-preservation footer. This section (§I + §K) edits no pre-existing artifact and downgrades no frozen status. Preserved by reference: the F1/F2 BOTH_FAIL basis-invariant negative, the Path-4-only −0.08 figure, the **two-stage scale-invariant ≈9.6×–14.5× overhead floor band** (anchor (12,3)≈14.5×, frontier (8,3)≈9.6×, conservative ≈28× ceiling), the exact landed invariants ($\Delta = 0.2589200$, $g(\Delta) = 18.305$ at (12,3), Atiyah–Singer index = 3, $|\mathrm{Riem}|^2 = 276$), and SoT hash `88e5903fe1c3d17655acef7e4fb45078e3168add25362040e21598b3b161ff17`. All "no prior art teaches X" assertions and patent numbers remain pending a live database search. Sole inventor: Chris Bergstrom. Assignee: Tahoe Labs (entity TBD). AI is a tool, not a co-inventor.

End of Part D (§I Simulation Evidence; §K Appendices).

CONSOLIDATED STATUS-PRESERVATION FOOTER (assembled document)

This assembled document integrates the five corrected source part files. Status invariants preserved by reference: SoT hash `88e5903fe1c3d17655acef7e4fb45078e3168add25362040e21598b3b161ff17`; hard-gate vector `D, D-local, D, D, D`; BG-10 Open/Diagnostic. The QC rung is recited on the current two-stage scale-invariant (12,3)/(8,3) floor frame (anchor ≈14.5×, frontier ≈9.6×, conservative ceiling ≈28×); held-fixed invariants index = 3, Δ = 0.2589200, and Standard-Model calibration are unchanged. Superseded QC figures (1.197 / 1.25–1.43 / d_z=1-as-a-code / 1.73 target / 12.6% / the 0.30 multiplier / 3.24×) appear in this document only inside retirement/supersession records, never as a live recited value. Every quantitative figure is an EXTERNAL / SIMULATED design-stage projection, not an independently benchmarked hardware result; σ_eff is an ESTIMATED input; the inventor countersign is the named gating step for external consumption. Sole inventor: Chris Bergstrom. Assignee: Tahoe Labs (entity TBD). AI used as a tool, not a co-inventor (*Thaler v. Vidal*; 2024 USPTO guidance). Assembled 2026-06-17 (rev 2) by A-PASS v2.

End of assembled provisional draft.