

Algorithmic Graph-Search Design of Heralded Linear Optical Circuits for Multipartite Entanglement: A Reconciled Quantitative Assessment

Algorithmic Graph-Search Design of Heralded Linear Optical Circuits for Multipartite Entanglement: A Reconciled Quantitative Assessment

Abstract

Heralded multipartite photonic entanglement underpins quantum communication, distributed sensing, and fault-tolerant computation, yet the manual synthesis of linear-optical circuits that generate a prescribed target state remains a combinatorial bottleneck. A recent preprint [1,2] recasts this design task as an algorithmic graph-search problem in the linear quantum graph (LQG) picture, reporting automated discovery of heralded circuits for hypergraph magic states, quantum error-correcting code states, general three-qubit states, and length-1 caterpillar graph states. This paper provides a single reconciled synthesis and quantitative appraisal. We formalize the search-space reduction claimed by the LQG representation with explicit combinatorial counts: for a representative six-mode, three-beam-splitter budget, raw circuit enumeration (3,375 ordered-slot instances) reduces to 455 graph instances, a factor of ≈ 7.4 (up to ≈ 44.5 under a reordering-inclusive convention), growing to ≈ 151 at ten modes and five beam splitters. We derive heralding probabilities for canonical fusion-based constructions: a three-qubit GHZ state via type-II fusion succeeds with probability $p^2/2$ per trial (5×10^{-3} at source probability $p = 0.1$), and a four-qubit caterpillar projection spans $[3.2 \times 10^{-7}, 7.5 \times 10^{-3}]$ across plausible hardware ranges — showing that hardware parameters, not circuit topology, dominate success probability. We further derive multiplexing arithmetic (a $p = 0.01$ circuit needs ≈ 458 attempts for 99% cumulative success; 64-fold multiplexing converts this to 8 rounds of 64 executions (512 total, $\approx 12\%$ more than the unmultiplexed 458) and loss-limited fidelity bounds $F \approx \eta^{2q}$). We identify falsification criteria and the decisive open problems: graph-isomorphism overhead, ancilla-mode blowup, loss-aware search objectives, and scalability beyond the demonstrated $q \leq 4$ target classes.

1. Introduction

Photons are attractive carriers of quantum information: they decohere negligibly at room temperature and are naturally networked, but they do not interact strongly. Linear optics sidesteps this via measurement-induced nonlinearity — ancillary photons are injected, interfered with the signal, and detected, and a particular detection pattern *heralds* that the desired transformation succeeded. The price is probabilistic operation and a proliferation of ancillary modes, which makes manual design of heralded circuits for nontrivial multipartite states laborious and error-prone.

The preprint under analysis [1,2] proposes to replace expert intuition with search. Its central move is representational: instead of searching over circuits (arrays of beam splitters and phase shifters), one searches over graphs in the *linear quantum graph* (LQG) picture, in which passive linear optical networks acting on Fock-state inputs are encoded as graphs whose vertices are modes and whose edges carry interaction data. Because many distinct circuits correspond to the same graph, and because graph structure constrains what is physically reachable, the effective search space shrinks. The authors report automated discovery of heralded circuits for a broad menu of multipartite resource states.

This paper has three goals. First, to make the framework accessible to an adjacent-field expert (Sections 2–3). Second, to quantify the claims: what "substantially reducing the search space" means in numbers, and what the circuit structures imply experimentally (Section 4). Third, to assess honestly what would falsify or limit the approach (Section 6). Where the underlying source drafts disagreed on modeling conventions, we adopt one convention in the main text and document every conflict explicitly in Appendix A; no divergence is silently resolved.

2. Background and Related Work

The paper under review [1,2]. References [1] and [2] are the query record and paper record for *Algorithmic Design of Heralded Linear Optical Circuits for Multipartite Entanglement* (arXiv:2609.18002v1); we treat them as one work. Its core contribution is to reconstruct circuit structures as graphs in the LQG picture and to search over graphs rather than circuits, reporting efficient schemes for hypergraph magic states, QEC code states, general three-qubit states, and length-1 caterpillar graph states. The novel claim relative to prior art is the *search procedure* and its pruning power.

The direct predecessor [9]. Reference [9] (arXiv:2310.10291v3) introduced a graph approach for creating heralded multipartite entanglement via the graph picture of linear quantum networks, building on npj Quantum Information 10, 67 (2024), explicitly to manage the ancillary particles and modes that "amplify the circuit intricacy." The lineage matters for evaluation: [9] used graphs as a *description* language for hand-designed heralded protocols; [1,2] uses them as a *search space* for automated discovery. The representational machinery is inherited; the search is the new element.

Entanglement taxonomy [3]. Reference [3] (arXiv:2409.04566v1) reviews rigorous definitions of separability, entanglement classification, transformations, and measures across several subsystems. This matters because "design a circuit for state X " is well-posed only once X is specified up to well-defined equivalence classes — e.g., GHZ class versus W class under SLOCC — and any automated search must encode such classes as its acceptance predicate. The multipartite classification problem is far subtler than the bipartite one, and [3] makes that precise.

Entanglement quantification [4]. Reference [4] (arXiv:quant-ph/0603281v2) characterizes pure-state multipartite entanglement via the probability density function of bipartite entanglement across partitions. This is relevant twice over: as a *verification* tool for heralded outputs (sampling the entanglement PDF across cuts certifies genuine multipartite entanglement rather than a mixture of lower-order correlations), and as a candidate *scoring function* for the graph search itself, since an algorithmic designer needs a computable objective to rank candidate LQGs.

Measurement-conditioned entanglement dynamics [5]. Reference [5] (arXiv:2407.03206v4) studies how n -partite GHZ entanglement builds up in monitored random Clifford circuits, exhibiting a measurement-induced transition between volume-law and area-law phases. Although the physical setting differs from passive optics, the conceptual parallel is strong: heralded linear optics is also measurement-conditioned dynamics, and the phase-transition phenomenology suggests that the space of LQGs may have structured regions where GHZ-type outputs are dense — potentially exploitable by heuristic search.

Network distribution [6]. Reference [6] (arXiv:2103.14759v3) presents an algorithm for generating multipartite entanglement between distant nodes of a noisy quantum network. Heralded photonic resource states are exactly the objects such network algorithms wish to distribute, and [6]'s noise model supplies the loss channels relevant to fidelity analysis; the composition of network-distribution algorithms with algorithmically designed local heralded sources is an open systems-level question.

Hardware frontier [7,11]. Reference [7] (arXiv:2203.06515v1) discusses photovoltaic-ferroelectric materials for all-optical devices in which light controls light — the competing deterministic-interaction paradigm. If such devices mature, the heralded-postselection architecture becomes partially obsolete, which is precisely why quantitative cost analysis matters: the graph-search framework is valuable in proportion to how long the "photons don't interact" regime persists. Reference [11] (DOI 10.5281/zenodo.21515894), a due-diligence assessment of QuiX Quantum, grounds the hardware-side plausibility: large-mode-count integrated circuits of the kind LQG search would output are exactly the product class integrated-photonics vendors are engineering.

Non-qubit encodings [8]. Reference [8] (arXiv:quant-ph/0104011v2) generates multipartite entangled coherent states via entanglement swapping with an ion-trap realization, quantifying entanglement by concurrence and the N-tangle. It is an early example of scheme-level design in a non-qubit encoding, and the N-tangle is one measure applicable to the caterpillar graph states targeted by [1,2]; coherent-state encodings could in principle be treated as alternative LQG decorations.

Diagrammatic-methods caution [10]. Reference [10] (DOI 10.5281/zenodo.22018102) examines the "cafeteria problem" of cross-disciplinary imports of diagrammatic languages (ZX spiders, Pauli webs, gadgets), cautioning that importing a formalism into a new domain imports hidden assumptions along with its elegance. The LQG picture is exactly such an import: a graph formalism that prunes the search space does so by encoding a theory of *which circuits matter*, and if that theory is incomplete (blind to loss, mode mismatch, or adaptive strategies), the pruned space may exclude the best circuits.

Corpus context [12,13]. Reference [12] (DOI 10.5281/zenodo.18327721) supplies benchmarking-methodology context, and [13] (DOI 10.5281/zenodo.17955898) analyzes thermodynamic and informational bottlenecks of scalable fault-tolerant computation — pertinent because heralded resource-state generation sits inside the fault-tolerance overhead budget. We note the corpus contains no dedicated work on the computational complexity of circuit synthesis; our complexity remarks are our own analysis, not literature-grounded claims.

3. Methods

Our method is analytical reconstruction and closed-form derivation; no simulations were performed, and every number in Sections 4–5 is either derived with stated inputs or explicitly labeled a projection with stated assumptions.

3.1 Circuit-versus-graph enumeration (adopted convention). We model a raw circuit search as follows: a linear-optical circuit on m modes built from k two-mode beam splitters has, per ordered slot, $C(m,2) = m(m-1)/2$ choices of unordered mode pair, giving

$$N_{\text{circ}}(m,k) = C(m,2)^k.$$

In the LQG picture, the same family is a graph on m vertices with k edges, where edge ordering is absorbed into the graph structure:

$$N_{\text{graph}}(m,k) = C(C(m,2), k), R(m,k) = N_{\text{circ}}/N_{\text{graph}}.$$

This is our model of the pruning step, consistent with [1,2]'s claim that the LQG picture "substantially reduc[es] the search space"; the actual framework may prune further. (Alternative enumeration conventions used by the source drafts are documented in Appendix A.)

3.2 Heralding-probability model for fusion-based GHZ construction. As a topology-independent baseline, we compute the success probability of the standard type-II fusion construction of a three-qubit GHZ state from two heralded Bell pairs. Inputs: per-Bell-pair heralded probability p (source parameter), and per-photon output

transmission η . Fusion of one photon from each pair at a 50:50 beam splitter followed by photon-number-resolving detection succeeds with conditional probability $qf = 1/2$ (*the standard type-II fusion figure of merit: Hong–Ou–Mandel interference gives bunched outcomes $|2,0\rangle$ and $|0,2\rangle$ in 2 of 4 cases, and exactly one of the two bunched detector patterns heralds success, giving $qf = 1/2$ conditional on bunching, i.e. $(2/4) \times (1/2) = 1/4$ unconditionally but $1/2$ as the conditional figure of merit used here*). Total per-trial success: $P_{\text{3GHZ}} = p^2 \cdot (1/2) \cdot \eta^2$.

3.3 Chain-fusion projection for caterpillar states. A length-1 caterpillar graph state (four qubits) requires three successful fusion events from four Bell pairs under a linear schedule, with $qf = 1/2$ and independence assumed: $P_{\text{4cat}} = p^4 \cdot (1/8) \cdot \eta^4$. We propagate parameter ranges $p \in [0.05, 0.5]$, $\eta \in [0.8, 0.99]$.

3.4 Multiplexing and attempt-count model. For per-attempt success p , cumulative success after N attempts is $1 - (1-p)^N$; with s independent multiplexed copies, per-round success is $1 - (1-p)^s$. All inputs are stated assumptions, not measurements of $[1,2]$'s circuits.

3.5 Loss-fidelity model. With uniform per-mode transmissivity η , a q -qubit dual-rail output traversing $2q$ modes is loss-free with probability η^{2q} , bounding the heralded-state fidelity.

4. Analysis

4.1 Search-space reduction (computed).

For $m = 6$ modes, $k = 3$ beam splitters:

- $C(6,2) = 15$.
- $N_{\text{circ}} = 15^3 = 3,375$ (ordered-slot model). If one additionally counts each three-beam-splitter sequence together with its $3! = 6$ reorderings as distinct synthesis paths, the reordering-inclusive count is $3,375 \times 6 = 20,250$.
- $N_{\text{graph}} = C(15,3) = (15 \cdot 14 \cdot 13)/6 = 2730/6 = 455$.
- Reduction: $R = 3375/455 = 7.4176$; $R' = 20250/455 = 44.5055$.

Scaling check, $m = 10$, $k = 5$: $C(10,2) = 45$; $N_{\text{circ}} = 45^5 = 184,528,125$ ($45^2 = 2025$; $45^3 = 91,125$; $45^4 = 4,100,625$; $45^5 = 184,528,125$). $N_{\text{graph}} = C(45,5) = (45 \cdot 44 \cdot 43 \cdot 42 \cdot 41)/120 = 146,611,080/120 = 1,221,759$. $R = 184,528,125/1,221,759 \approx 151.03$. The reduction factor grows with budget size — the quantitative content of the claim that LQG pruning is the enabling step for automated design.

4.2 Heralding probability, three-qubit GHZ via fusion (computed).

$P_{\text{3GHZ}} = p^2/2$ (lossless outputs):

- $p = 0.1$: $P = 0.01/2 = 5 \times 10^{-3}$.
- $p = 0.25$: $P = 0.0625/2 = 3.125 \times 10^{-2}$.
- $p = 0.5$: $P = 0.25/2 = 0.125$.

With per-photon output loss η applied to the two surviving photons: $P_{\text{3GHZ}}(\eta) = p^2 \cdot (1/2) \cdot \eta^2$. At $p = 0.1$, $\eta = 0.9$: $0.01 \times 0.5 \times 0.81 = 4.05 \times 10^{-3}$.

4.3 Four-qubit caterpillar projection (labeled projection).

$P_{\text{4cat}} = p^4 \cdot (1/8) \cdot \eta^4$:

- $p = 0.1$, $\eta = 1$: $0.0001/8 = 1.25 \times 10^{-5}$.
- $p = 0.1$, $\eta = 0.9$: $0.0001 \times 0.125 \times 0.6561 = 8.20 \times 10^{-6}$ ($0.9^4 = 0.6561$).
- $p = 0.5$, $\eta = 0.95$: $0.0625 \times 0.125 \times 0.81450625 = 6.36 \times 10^{-3}$ ($0.95^4 = 0.81450625$).

Bounds over $p \in [0.05, 0.5]$, $\eta \in [0.8, 0.99]$: $P_{\min} = 0.05^4 \times 0.125 \times 0.8^4 = 6.25 \times 10^{-6} \times 0.125 \times 0.4096 = 3.2 \times 10^{-7}$. $P_{\max} = 0.5^4 \times 0.125 \times 0.99^4 = 0.0625 \times 0.125 \times 0.96059601 = 7.5 \times 10^{-3}$. The projected per-trial success spans over four orders of magnitude across plausible hardware — a key finding: the framework's value is gated almost entirely by source and loss parameters, not circuit topology.

4.4 Attempt-count and multiplexing arithmetic (labeled projection).

Assume $p = 0.01$ per attempt (representative of postselected linear-optical entangling operations, not a measurement of $[1,2]$). Expected attempts to first success: $1/p = 100$. For 99% cumulative success: $1 - (0.99)^N = 0.99 \rightarrow N = \ln(0.01)/\ln(0.99) = (-4.6052)/(-0.010050) \approx 458$ attempts. With $s = 64$ multiplexed copies: per-round success $= 1 - (0.99)^{64}$; $\ln(0.99) \times 64 = -0.64322$; $\exp(-0.64322) = 0.5256$; so 0.4744 per round. Rounds for 99%: per-round success is $1 - 0.5256 = 0.4744$, so $N = \ln(0.01)/\ln(0.5256) = (-4.6052)/(-0.6432) \approx 7.16 \rightarrow 8$ rounds; total executions $8 \times 64 = 512$ versus 458 unmultiplexed. Multiplexing converts latency (458 rounds $\rightarrow$ 8 rounds) at a total-execution overhead of $512/458 \approx 1.12$, i.e. $\approx 12\%$ more executions, so under these assumptions it reduces latency but not total work.

4.5 Loss-limited fidelity (labeled projection).

With $\eta = 0.99$ per mode: $q = 4$ output qubits (8 modes): $F_{\text{bound}} = 0.99^8 = \exp(-0.0804027) = 0.923$. $q = 8$: $0.99^{16} = \exp(-0.160805) = 0.852$. Each additional qubit costs a factor $\eta^2 \approx 0.98$ — the exponential in q is the fundamental enemy.

4.6 Hardware threshold (labeled projection).

For a fault-tolerant application demanding $R_{\text{target}} = 1$ successful four-qubit caterpillar state per second at trial rate $f = 10^6$ /s: $P_{\text{req}} = 10^{-6}$. Requiring $p^4 \eta^4/8 \geq 10^{-6}$ gives $(p\eta)^4 \geq 8 \times 10^{-6}$, so $p\eta \geq (8 \times 10^{-6})^{1/4} = 1.6818 \times 10^{-1.5} = 0.0532$. At $\eta = 0.9$: $p \geq 0.0591$. Any source-and-loss combination with $p\eta > 0.053$ meets a 1 Hz useful rate under these assumptions.

4.7 Illustrative concrete circuit (single-draft contribution, retained as an example). One source draft instantiated the framework as a four-mode network with two single-photon sources ($p_s = 0.80$ assumed), two 50:50 beam splitters, and two photon-number-resolving detectors ($\eta = 0.90$ assumed), computing $P_{\text{herald}} = p_s^2 \times \eta^2 \times (2RT)^2 = 0.64 \times 0.81 \times 0.25 = 0.1296$, with claimed unit fidelity. We retain this as an illustrative parameterized example, not as a verified property of $[1,2]$'s circuits (see Appendix A).

5. Results

All numbers are computed in Section 4; none are measured or simulated; projections are labeled.

- 1. Search-space reduction (computed).** $m = 6$, $k = 3$: 3,375 $\rightarrow$ 455 graphs, factor 7.4176 (44.5055 reordering-inclusive). $m = 10$, $k = 5$: 184,528,125 $\rightarrow$ 1,221,759, factor ≈ 151.03 . Reduction grows with system size, supporting the qualitative claim in $[1,2]$.
- 2. Three-qubit GHZ heralding (computed, model-based).** $P = p^2/2$: 5×10^{-3} ($p = 0.1$), 3.125×10^{-2} ($p = 0.25$), 0.125 ($p = 0.5$); 4.05×10^{-3} at $p = 0.1$, $\eta = 0.9$.
- 3. Four-qubit caterpillar (projection).** $P = p^4 \eta^4/8 \in [3.2 \times 10^{-7}, 7.5 \times 10^{-3}]$ over the stated hardware ranges.
- 4. Attempt counts (projection, $p = 0.01$).** 100 expected attempts; 458 for 99%; 8 rounds $\times$ 64 = 512 executions under 64-fold multiplexing, $\approx 12\%$ more than the 458 unmultiplexed attempts.
- 5. Loss-limited fidelity (projection, $\eta = 0.99$).** ≈ 0.923 ($q = 4$), ≈ 0.852 ($q = 8$).
- 6. Hardware threshold (projection).** $p\eta \geq 0.0532$ for a 1 Hz four-qubit caterpillar rate at 10^6 trials/s.

7. **Scale of demonstrated targets (from the abstract of [1,2]).** All four named target classes are consistent with $q \leq 4$; no scaling result to large mode count is claimed in the available material.
8. **Benchmarking implication.** Because topology-independent fusion baselines span four orders of magnitude across plausible (p, η) , any LQG-discovered circuit must be benchmarked against the fusion baseline at fixed (p, η) , not against other circuits at unspecified parameters.

6. Discussion

Limitation 1: representational pruning may hide the best circuits. The search-space compression is a virtue only if the equivalence classes preserve optimality. If the LQG picture, as imported from [9], cannot express certain heralding strategies — adaptive (feed-forward) circuits, higher-rank ancilla detections, non-Gaussian inputs — the search is over a truncated space, and the "optimal" graph may be optimal only within the picture. This is the cafeteria problem of [10]: the imported formalism's elegance may smuggle in its blind spots.

Limitation 2: our numbers are projections. The p and η inputs are assumptions, not measurements of [1,2]'s circuits. If algorithmically designed circuits achieve $p = 0.1$, the 99%-success attempt count drops from 458 to $\ln(0.01)/\ln(0.9) = 44$; at $p = 10^{-4}$ it rises to $\approx 46,000$. Our cost conclusions span two orders of magnitude and should be read as a framework, not a verdict.

Limitation 3: small-scale demonstration. All four target classes sit at $q \leq 4$. Whether structural constraints (e.g., degree bounds implied by sparse code-state entanglement structure) become load-bearing at larger m — where the graph space grows as $2^{\{m(m-1)/2\}}$ ($\approx 10^{131}$ at $m = 30$) — is the decisive open question. Search-space reduction is not search-tractability: even after a factor-151 reduction, the ten-mode space contains over a million graphs, and acceptance predicates (e.g., "is a QEC code state") may be expensive per candidate; graph-isomorphism testing adds further overhead we have not quantified.

Limitation 4: hardware parameters dominate. As Section 4.3 shows, hardware parameters shift success probability by orders of magnitude — an algorithmically optimal circuit at bad (p, η) is useless. The framework optimizes structure, not physics. Conversely, near-term integrated photonics [11] makes the multiplexing arithmetic of Section 4.4 realistic, converting the heralded scheme's probabilistic weakness into a latency cost only — the strongest argument for the framework's relevance. If deterministic photon-photon mediators of the type explored in [7] mature, the design problem changes character entirely.

Limitation 5: verification. The PDF-of-entanglement method of [4] and the N-tangle of [8] provide the verification toolkit, but heralded states are produced conditionally at low rate; accumulating enough samples for a statistically rigorous multipartite entanglement witness is itself expensive. A circuit optimal in silico but unverifiable in the laboratory is not yet a resource. Idealized-component assumptions (perfect indistinguishability, lossless optics, exact 50:50 splitting) further inflate model probabilities relative to real devices; partial distinguishability most plausibly degrades the conditional $q_f = 1/2$.

Falsification criteria. The central claim of [1,2] would be falsified if: (i) a target class outside the LQG-expressible set exists for which the search returns no circuit while a hand-designed one does; (ii) discovered circuits' success probabilities, measured at fixed (p, η) with realistic distinguishability, are systematically no better than the fusion baselines computed here (e.g., $P = p^2/2$ for three-qubit GHZ); (iii) search cost grows super-exponentially in mode count so that no target beyond ~ 6 qubits is reachable; or (iv) experimental heralding rates fall substantially below predictions after calibrated losses, indicating the interference model is insufficient.

Open questions. Can entanglement-PDF scoring [4] serve as the search objective, at what per-candidate cost? Do the measurement-induced transitions of [5] imply dense regions of LQG space for heuristic search? How do LQG-designed sources compose with network distribution [6] under realistic noise? Can non-qubit encodings [8] or all-optical device concepts [7] be expressed as LQG decorations? Can the search objective incorporate loss directly rather than only ideal-state overlap?

7. Conclusion

The algorithmic graph-search formulation of heralded circuit design [1,2] is a genuine methodological advance: it converts an artisanal design problem into a computational one, building on the graph picture of [9] and consistent with the broader diagrammatic-methods perspective of [10]. Our reconciled analysis quantifies the pruning (factor ≈ 7.4 – 44.5 at a minimal six-mode budget, ≈ 151 at ten modes, growing with size) and the experimental arithmetic (fusion baselines $P = p^2/2$ for three-qubit GHZ; four-qubit caterpillar projections spanning $[3.2 \times 10^{-7}, 7.5 \times 10^{-3}]$; ≈ 458 attempts for 99% success at $p = 0.01$, reducible to 8 multiplexed rounds (512 total executions, $\approx 12\%$ more than unmultiplexed); loss-limited fidelity ≈ 0.92 at four output qubits). The framework's future value hinges on three testable questions: whether the searched graph class contains the optimal circuits, whether structural pruning activates at scale, and whether the search objective can be made loss-aware. If those resolve favorably, automated discovery of heralded resource states could do for photonic entanglement what logic synthesis did for digital circuits.

References

[1] TITLE: arXiv Query: `searchquery=&idlist=2609.18002&start=0&max_results=1` [2] Algorithmic Design of Heralded Linear Optical Circuits for Multipartite Entanglement. arXiv:2609.18002v1. <https://arxiv.org/abs/2609.18002v1> [3] Multipartite entanglement. arXiv:2409.04566v1. <https://arxiv.org/abs/2409.04566v1> [4] Probability density function characterization of multipartite entanglement. arXiv:quant-ph/0603281v2. <https://arxiv.org/abs/quant-ph/0603281v2> [5] Multipartite Greenberger-Horne-Zeilinger Entanglement in Monitored Random Clifford Circuits. arXiv:2407.03206v4. <https://arxiv.org/abs/2407.03206v4> [6] Distributing Multipartite Entanglement over Noisy Quantum Networks. arXiv:2103.14759v3. <https://arxiv.org/abs/2103.14759v3> [7] Photovoltaic-ferroelectric materials for the realization of all-optical devices. arXiv:2203.06515v1. <https://arxiv.org/abs/2203.06515v1> [8] Multipartite entangled coherent states. arXiv:quant-ph/0104011v2. <https://arxiv.org/abs/quant-ph/0104011v2> [9] Heralded Optical Entanglement Generation via the Graph Picture of Linear Quantum Networks. arXiv:2310.10291v3. <https://arxiv.org/abs/2310.10291v3> [10] DOI 10.5281/zenodo.22018102. QNFO: ZX Diagrams at the Seam: Spiders, Pauli Webs, Gadgets, and the Cafeteria Problem of Cross-Disciplinary Imports. [11] DOI 10.5281/zenodo.21515894. QNFO: Due Diligence Report: QuiX Quantum. [12] DOI 10.5281/zenodo.18327721. QNFO: Spectral Benchmarking of Holographic Quantum Simulations. [13] DOI 10.5281/zenodo.17955898. QNFO: Thermodynamic and Informational Bottlenecks of Scalable Fault-Tolerant Quantum Computation.

Appendix A. Divergence report

D1. Quantitative search-space reduction (C2).

- Draft A: five-mode network, naïve simple-graph count $2^{\wedge C(5,2)} = 2^{10} = 1024$, pruned to 128 (factor 8) via photon-number conservation and degree bound $\Delta_{\text{max}} = 3$.
- Draft B: nine-mode case; discretized circuit space 10^{36} ($B = 10$ settings per beam splitter, Reck decomposition) versus $2^{36} \approx 6.87 \times 10^{10}$ simple graphs, a $\approx 1.46 \times 10^{25}$ representational compression; degree

constraints prune negligibly at this scale.

- Draft C: six-mode/three-beam-splitter enumeration, 3,375 (or 20,250) $\rightarrow$ 455, factor 7.4176–44.5055.
- *Convention behind the disagreement*: the drafts count fundamentally different objects — A counts simple graphs on few vertices with hard constraints; B counts continuous circuit manifolds versus graphs (representational compression); C counts discrete beam-splitter topology placements versus edge sets (structural compression). These are not contradictory but incommensurable.
- *Resolution*: the main text adopts Draft C's enumeration convention (Sections 3.1, 4.1) because it is the most explicit and reproducible arithmetic model of the pruning step, and reports its scaling. Draft A's 1024 $\rightarrow$ 128 figure and Draft B's $\approx 10^{25}$ representational-compression figure are documented here as alternative conventions; no attempt is made to average or reconcile them into a single number.

D2. GHZ heralding probability (C3).

- Draft A: concrete four-mode circuit with $p_s = 0.80$, $\eta = 0.90$, two 50:50 beam splitters; Pherald = $0.64 \times 0.81 \times 0.25 = 0.1296$. (Draft A's abstract stated ≈ 0.2592 , inconsistent with its own body derivation of 0.1296; the body derivation is internally explicit and is taken as A's position.)
- Draft B: assumes $p = 0.01$ per attempt as a literature-representative figure and computes attempt counts (458 for 99%), explicitly not a measurement of [1,2].
- Draft C: fusion model $P = p^2/2$, giving 5×10^{-3} at $p = 0.1$.
- *Convention behind the disagreement*: A models a specific circuit with optimistic source/detector parameters and a product-form interference factor (2RT per beam splitter); C models the canonical type-II fusion baseline with a parameterized source; B declines to commit to a circuit and treats p as a free assumption.
- *Resolution*: the main text adopts Draft C's parameterized fusion baseline as the primary quantitative convention (Sections 3.2, 4.2), because it is topology-independent and benchmarkable at fixed (p , η). Draft A's circuit-level computation (0.1296 under its stated parameters) is retained in Section 4.7 as an illustrative instantiation, clearly labeled as resting on A's assumed $p_s = 0.80$, $\eta = 0.90$, and product-form interference model. Draft B's attempt-count framework is retained in Section 4.4 with its assumption labeled.

D3. Fidelity of the generated GHZ state (C14).

- Draft A: claims $F = 1.0$ exactly, by analytical construction.
- Drafts B and C: no exact-fidelity claim; B derives loss-limited bounds $F \approx \eta^{\{2q\}} < 1$.
- *Resolution*: the main text does not assert $F = 1.0$ as a general result; A's claim is confined to the idealized, lossless instantiation of Section 4.7, and the loss-limited bound of Section 4.5 governs realistic settings.

D4. Internal inconsistency in Draft A. Draft A's abstract (≈ 0.2592) contradicts its own derivation (0.1296). Resolved in favor of the explicit derivation; documented here rather than silently corrected in the abstract-level claim.

Appendix B. Claim attribution

Claim	Description	Drafts	Status
C1	LQG graph-search substantially reduces the circuit design search space	A, B, C	CONVERGENT
C2	Quantitative search-space reduction factor (specific counts)	A: 1024 $\rightarrow$ 128; B: $\approx 10^{25}$; C: 7.4–44.5	DIVERGENT (D1)
C3	GHZ heralding probability (specific value/model)	A: 0.1296; B: $p = 0.01$ projection; C: $p^2/2$	DIVERGENT (D2)

Claim	Description	Drafts	Status
C4	Four-qubit caterpillar heralding projection $P = p^4 \eta^4 / 8 \in [3.2 \times 10^{-7}, 7.5 \times 10^{-3}]$	C	SINGLE-DRAFT (projection)
C5	Attempt-count model: $1/p = 100$ expected; 458 attempts for 99% at $p = 0.01$	B	SINGLE-DRAFT (projection)
C6	Multiplexing arithmetic: 64-fold multiplexing $\rightarrow$ 8 rounds, 512 total executions	B	SINGLE-DRAFT (projection; corrected per Section 4.4)
C7	Loss-limited fidelity bound $F \approx \eta^{\{2q\}}$	B	CONVERGENT with C14 resolution
C8	Hardware threshold $p\eta \geq 0.0532$ for 1 Hz four-qubit caterpillar rate at 10^6 trials/s	B	SINGLE-DRAFT (projection)
C9	Search-space reduction grows with system size (≈ 151 at $m = 10$, $k = 5$)	C	SINGLE-DRAFT (computed)
C10	All demonstrated target classes satisfy $q \leq 4$	A, B, C	CONVERGENT (from abstract of [1,2])
C11	Benchmark against fusion baseline at fixed (p, η)	B, C	CONVERGENT
C12	Verification via entanglement-PDF [4] and N-tangle [8]	B	SINGLE-DRAFT (proposed toolkit)
C13	Illustrative four-mode circuit with $P_{\text{herald}} = 0.1296$ and unit fidelity	A	SINGLE-DRAFT (retained as example, Section 4.7)
C14	Fidelity of generated GHZ state: A claims $F = 1.0$; B, C give loss-limited bounds	A vs. B, C	DIVERGENT (D3)