

Characterization of nested Walsh parity-check filters in a single-photon eight-mode register on a cloud photonic processor

E. Tully,¹ J. Washburn,¹ and M. Simons^{1,*}

¹*Recognition Physics Institute, Austin, TX 78701, USA*

(Dated: June 2026)

We characterize how faithfully a commercial cloud photonic processor implements two nested Walsh parity-check filters acting on the spatial-mode amplitudes of a single photon. Eight modes are indexed by the vertices of the three-cube Q_3 ; the filters realize the classical $[8, 7, 2]$ single-parity-check code (the zero-sum “neutral” subspace \mathcal{N}) and the $[8, 4, 4]$ extended Hamming code (the parity-checked subspace $\mathcal{S} \subset \mathcal{N}$, with one DC and three face-parity syndrome channels). These are first-quantized path/mode encodings of classical codes: nothing is corrected, and all reported probabilities are conditional on postselected single-photon detection—photon loss, the platform’s dominant error, is removed by the postselection and is not detected by the checks. On Quandela’s Belenos processor, across more than 340,000 postselected detections, the implemented filters reproduce the intended linear algebra to percent-level leakage floors: neutral inputs reach residual DC-port leakage of 0.02%–1.1% depending on input (mean 0.6%), a $\approx 21\times$ suppression relative to the ideal 0.125 DC-capture baseline ($31.6\times$ relative to the measured non-neutral control, the gap quantifying an on-chip calibration bias); injected DC contamination produces a monotonic, calibratable soft error signal; and the three face-parity syndrome channels route to their predicted ports with 94–99% selectivity, for which we report the full confusion matrix. Leakage remains far below non-neutral controls after one to three applications of a sector-preserving unitary core, with cross-depth differences dominated by per-compilation calibration systematics rather than gate-cycle physics. We identify and quantify the systematics that limit all of these figures (fixed-pattern separator bias, ± 0.02 per-point calibration offsets, order-of-magnitude compilation scatter at the 10^{-3} leakage level), and report an opportunistic Hong–Ou–Mandel degradation episode, during which the suppression vanished and recovered with recalibration, as a diagnostic observation. The measurements verify leakage suppression and syndrome routing for one postselected single-photon block; they do not demonstrate error correction, closed-loop recovery, feed-forward control, a fault-tolerance threshold, or protection against photon loss.

I. INTRODUCTION

Quantum error protection requires encoding logical information into a subspace of a larger physical Hilbert space [1, 2]. The overhead of this encoding (the ratio of logical to physical resources) is a central figure of merit, although its precise meaning depends on the physical register. In multi-qubit fault-tolerance roadmaps, the surface code [3] operates at code rates of order a few percent at fault-tolerant distances [4, 5], and even recent low-overhead constructions such as bivariate bicycle codes [6] carry substantial qubit overhead, whereas small error-detecting codes such as $[[4, 2, 2]]$ [2, 7] reach high rates only at low distance. These stabilizer-code rates are quantified for context in Sec. VI; they are not directly comparable to the single-photon spatial-mode dimension rate used below, but they highlight a general design pressure in quantum information experiments: the need for experimentally accessible error-detecting subspaces with directly measurable leakage or syndrome channels [8–10].

We present a single-photon construction on the three-cube Q_3 (the three-dimensional hypercube graph whose eight vertices are the 3-bit binary strings, with edges joining strings that differ in exactly one bit) that achieves

spatial-mode dimension rate $7/8$ for a 7-dimensional error-detecting subspace on an 8-mode register, and test representative states from this subspace on a commercial photonic chip using more than 340,000 single-photon events. The construction uses the geometry of the cube to define two nested subspaces: a 7-dimensional neutral subspace \mathcal{N} for high-rate suppression, and a 4-dimensional parity-checked subspace $\mathcal{S} \subset \mathcal{N}$ with one DC/sum syndrome and three independent face-parity syndrome channels. Both subspaces are defined by linear constraints: \mathcal{N} is the zero-sum hyperplane of \mathbb{C}^8 , and \mathcal{S} is the kernel of a 4×8 Walsh parity-check matrix (one DC/sum row and the three face-parity check functionals of the cube), with coding-theoretic minimum distances $d(\mathcal{N}) = 2$ and $d(\mathcal{S}) = 4$. As linear codes these are the classical $[8, 7, 2]$ single-parity-check code and the $[8, 4, 4]$ first-order Reed–Muller code $\text{RM}(1, 3)$ (equivalently the extended Hamming code) [11–13]; our contribution is not the codes themselves but their single-photon spatial-mode realization on cloud photonic hardware, the cube-derived Gray-code gate ordering, and the unitary-core invariance test described below. The four code states form a \mathbb{Z}_2^3 -character basis, and the same construction extends to the d -cube Q_d with $d + 1$ checks at fixed distance 4. The encoding is first-quantized: states are normalized vectors in \mathbb{C}^8 indexed by spatial mode, as in high-dimensional photonic encodings based on discrete path or mode degrees of freedom [14]. The error weight

* Corresponding author: msimons@recognitionphysics.org

of a vector counts the number of perturbed mode amplitudes. This is a qudit-style encoding (throughout, a single-photon path/mode qudit in dimension 8) and not a multi-qubit stabilizer code, so the rate should be interpreted as a property of the single-photon spatial-mode encoding. Relative to generic high-dimensional single-photon encodings on path or mode degrees of freedom, the cube geometry serves here as an organizing principle: it ties the DC check, face-parity syndrome structure, and Gray-code gate ordering to a single combinatorial object. The experiments verify the resulting filter structure; they do not isolate a hardware advantage attributable to the cube labeling itself (Sec. VII).

On hardware, representative neutral-state inputs exhibit $\approx 21\times$ leakage suppression relative to the ideal 0.125 DC-capture baseline ($31.6\times$ relative to the measured non-neutral control), and representative parity-checked inputs exhibit $23.7\times$ syndrome suppression with 94–99% dominant-port selectivity for the three face-parity channels. The neutral-sector unitary core keeps BALANCE leakage far below the non-neutral controls over the tested depths. In the ideal limit these suppressions are exact by construction; what the measurements characterize is how faithfully the processor implements the intended filters, the leakage floors at which implementation imperfections enter, and the systematics that set those floors. The simulations of Sec. III provide idealized reference points rather than additional hardware measurements. The specific combination tested here is nested single-photon spatial-mode subspaces with leakage suppression and resolved face-parity syndrome channels on a programmable, cloud-accessible photonic processor; we have not performed a systematic survey of multipoint Walsh- and parity-routing experiments, so we state this as the configuration tested rather than as a priority claim. To situate it, we contrast the present work with the two closest integrated and all-optical baselines along the axes that define it. Bell *et al.* [15] demonstrated a graph-state error-correction code with syndrome extraction and correction, but in a multi-photon polarization-entangled encoding of a single logical qubit realized on a bulk and fibre source, not on a programmable, cloud-accessible mesh. Zhang *et al.* [16] encoded error correction on a programmable *integrated* photonic chip, again in a multi-qubit (dual-rail) encoding. Neither targets the combination reported here: a single-photon spatial-mode (path/mode-qudit) register laid out on the eight-vertex cube geometry, carrying a high-rate neutral subspace, three simultaneously resolved face-parity syndrome channels, and cloud-only access with no user-side cryogenics. The distinguishing axes are therefore the single-photon spatial-mode encoding and the resolved face-parity syndrome structure on a cloud-accessible programmable processor. Generic programmable multipoint interferometry on single photons is itself well established [10, 25], and leakage suppression and programmability *per se* are shared with this prior work; the specific element here is the simultaneous exposure of the DC/sum channel and

three independently resolved face-parity syndrome channels on that platform. The experiment should therefore be read as a single-block demonstration of leakage suppression, syndrome routing, and neutral-sector retention in a single-photon spatial-mode register, not as a demonstration of closed-loop recovery, feed-forward correction, or a fault-tolerance threshold.

II. THEORY

A. Physical register and error model

The state space is \mathbb{C}^8 , with modes indexed by the vertices of Q_3 via 3-bit labels $b_2b_1b_0$. A single photon occupies a normalized superposition of these 8 spatial modes (Fig. 1). For a subspace $\mathcal{C} \subset \mathbb{C}^8$, the *minimum support weight*

$$d(\mathcal{C}) = \min_{0 \neq x \in \mathcal{C}} |\text{supp}(x)|$$

is the minimum number of nonzero mode amplitudes in any nonzero element of \mathcal{C} ; equivalently, it is the minimum Hamming weight of a nonzero codeword in the first-quantized (spatial-mode) representation, i.e. the standard coding-theoretic minimum distance applied to mode-amplitude vectors. In the additive first-quantized path-amplitude error model used here, a nonzero perturbation supported on fewer than d modes cannot lie entirely in \mathcal{C} and is therefore detected by projection onto the complementary syndrome space.

The error model used throughout this paper is the *first-quantized path-mode error model*: a “single-mode amplitude perturbation” is an additive local change to one mode amplitude. Concretely, a relative attenuation or imbalance on mode k acts as $v_k \rightarrow v_k + \delta$ with $\delta \in \mathbb{C}$, and a phase rotation on an occupied mode acts as $v_k \rightarrow v_k + (e^{i\phi} - 1)v_k$; in either case the perturbation occupies a single mode coordinate, so the support weight of the error is 1 and the Hamming-weight (minimum-distance) argument above applies directly. Physically, this models relative mode-dependent attenuation or imbalance in one waveguide, as well as coherent leakage or contamination isolated to a single path, conditional on the postselected single-photon readout used in the experiments. We adopt this single-mode channel because per-waveguide amplitude imbalance and isolated path contamination are the leading-order imperfections of a programmable mesh acting on a postselected single photon; it is thus the natural first-order error model for this register, fixed independently of the subspace construction and not chosen to match it. A local phase rotation produces the same syndrome signal to first order when the affected mode is occupied. This is not the standard multi-qubit Pauli channel, and the detection and syndrome guarantees we report are for the single-photon spatial-mode channel. Errors of support weight greater than one (coherent multi-mode phase errors from mis-




FIG. 1. Vertex layout of the 3-cube Q_3 used as the index set for the eight photonic modes. Each vertex carries a 3-bit address $b_2b_1b_0$. The blue cycle is the Gray-code Hamiltonian cycle that orders the modes for the archived neutral-sector core. The three pairs of opposite faces (constant b_0 , b_1 , b_2) correspond to the three face-parity check functionals M_0, M_1, M_2 of Eq. (2); together with the DC/sum row, these form the four checks defining \mathcal{S} .

calibrated mesh elements, correlated insertion-loss patterns across several modes, and source phase drift) fall outside this single-mode model and are not protected by the present checks. We use the word *protection* throughout in this operational sense: the linear constraints *detect and suppress* such single-mode amplitude perturbations, through projection onto the complementary syndrome space and the associated leakage and syndrome readout. This is distinct from active, closed-loop quantum error correction, which would additionally require complex-syndrome measurement and feed-forward and is not claimed here (see Sec. VII B).

Three features of Fig. 1 drive the construction in the remainder of this section. First, every vertex enters the sum constraint $\sum_i v_i = 0$ on equal footing, so the neutral subspace \mathcal{N} is geometrically the zero-mean hyperplane of \mathbb{C}^8 . Second, the three pairs of opposite faces (constant- b_0 , constant- b_1 , and constant- b_2) supply the three sign patterns $(-1)^{b_j}$ that we use as face-parity check functionals below; the cube has no other independent opposite-face partition, which is why \mathcal{S} has exactly three face-parity syndrome channels in addition to the DC/sum channel. Third, the blue Gray-code Hamiltonian cycle visits each vertex once and fixes the mode ordering used by the archived neutral-sector core of Sec. II D. That core is the circuit element whose algebraic action preserves the neutral constraint and the stated checks. Thus Fig. 1 summarizes the common cube geometry underlying the check functionals (sum, M_0, M_1, M_2), the gate layout, and the error-detecting subspaces (\mathcal{N}, \mathcal{S}).

B. Neutral subspace \mathcal{N} : spatial-mode dimension rate 7/8, distance 2

The neutral subspace is defined by

$$\mathcal{N} = \{v \in \mathbb{C}^8 : \sum_i v_i = 0\}, \quad \dim \mathcal{N} = 7.$$

We refer to this zero-sum hyperplane as the *neutral* subspace, following the terminology used for this construction in Ref. [17]. Write $\mathbf{1} = (1, 1, \dots, 1)^\top \in \mathbb{C}^8$ for the

all-ones vector. The orthogonal projector onto \mathcal{N} is the BALANCE operator

$$B = I_8 - \frac{1}{8} \mathbf{1} \mathbf{1}^\dagger, \quad (1)$$

in which $\mathbf{1} \mathbf{1}^\dagger$ is the rank-one outer product (not a scalar), so $\frac{1}{8} \mathbf{1} \mathbf{1}^\dagger$ is the orthogonal projector onto the uniform (DC) direction $\mathbf{1}$ and B removes that component. Any nonzero additive single-mode perturbation in the sense above changes the sum and is therefore detectable through the DC readout port; for a pure local phase rotation, the signal is proportional to the occupied amplitude in that mode. The minimum-weight codeword $(1, -1, 0, \dots, 0)$ has weight 2, so $d(\mathcal{N}) = 2$. As a linear code over the mode amplitudes, \mathcal{N} is the classical $[8, 7, 2]$ single-parity-check code [11]. We call $R_{\text{mode}}(\mathcal{C}) = \dim(\mathcal{C})/8$ the *spatial-mode dimension rate* of a subspace $\mathcal{C} \subset \mathbb{C}^8$; thus $R_{\text{mode}}(\mathcal{N}) = 7/8$. This rate counts protected dimensions per physical mode in the single-photon register and is not a multi-qubit stabilizer-code rate.

Structurally, \mathcal{N} is analogous to a decoherence-free subspace or noiseless code: information is placed in a subspace orthogonal to a collective mode, and the implemented gates are chosen to leave that subspace invariant [18–20]. The analogy is limited, however. Standard DFS theory protects against noise operators with a specified collective action on the encoded system; here the protected direction is the uniform spatial mode of a post-selected single photon, the measured soft-error signal is the BALANCE/DC leakage produced by first-quantized mode-amplitude perturbations, and the cube geometry supplies the face-parity syndrome channels and Gray-code gate ordering. Thus the present claim is not a generic DFS claim, but an explicit spatial-mode construction whose protection is read out by the leakage and syndrome measurements below.

C. Parity-checked subspace \mathcal{S} : distance 4 with one DC/sum and three face-parity syndrome channels

The three face-parity check functionals are represented by $M_j = \text{diag}((-1)^{b_j})$ for $j \in \{0, 1, 2\}$, i.e. by the row constraints $\sum_b (-1)^{b_j} v_b = 0$. Together with the DC/sum constraint, they define the parity-check matrix

$$C = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \end{pmatrix}. \quad (2)$$

The rows of C are the constant function and the three first-order Walsh (coordinate) functions on the cube, so C is the Walsh-sign (complexified) realization of the standard parity-check matrix of the first-order Reed–Muller code $\text{RM}(1, 3)$, equivalently the $[8, 4, 4]$ extended Hamming code [11–13]: the ± 1 entries are the real

Walsh characters, not the $\text{GF}(2)$ parity-check matrix itself. The parity-checked subspace $\mathcal{S} = \ker(C)$ therefore has dimension $8 - 4 = 4$. Every 3-column subset of C has full rank (verified exhaustively for all $\binom{8}{3} = 56$ triples by the routine `all_triples_full_rank` in `operators.py` in the reproducibility package), which implies $d(\mathcal{S}) \geq 4$; this distance-4 property is exactly the standard minimum distance of $[8, 4, 4]$. The weight-4 codeword $(1, -1, -1, 1, 0, 0, 0, 0) \in \mathcal{S}$ saturates the bound, so $d(\mathcal{S}) = 4$.

The same counting gives a useful mathematical scaling reference on the d -cube Q_d . With $M = 2^d$ spatial modes, the neutral subspace defined by the single DC/sum check has dimension $M - 1$ and spatial-mode dimension rate $1 - 1/M$. Adding the d coordinate face-parity checks gives a parity-checked subspace of dimension $M - (d + 1) = M - \log_2 M - 1$ and rate $1 - (d + 1)/M$. Thus the number of checks grows only logarithmically with mode count while the same first-order Walsh-check construction retains distance 4 for $d \geq 2$ (a weight-4 rectangle in the hypercube saturates the bound). This is a property of the linear-algebraic construction, not a demonstrated multi-block scaling architecture; in particular the logarithmic check count is a statement about the algebra only. The physical resource cost scales unfavorably: the mode count $M = 2^d$ and the number of modes over which a single photon must remain coherent both grow exponentially in d , while the protected distance stays fixed at 4.

The 8 columns of C are pairwise distinct, so a single-mode amplitude perturbation δ on mode m produces the unique syndrome $s = \delta c_m$, where c_m denotes the m -th column of C . This is the algebraic basis for single-mode error identification on \mathcal{S} .

D. Dynamics operator

The hardware dynamics are defined operationally from the archived compiled circuits that were submitted to the Belenos processor. Index the standard basis $\{e_k\}_{k \in \mathbb{Z}_8}$ by the eight modes of Sec. II and recall $\mathbf{1} = \sum_k e_k = (1, \dots, 1)^\top$. Since $\mathcal{N} = \mathbf{1}^\perp$, a unitary U satisfies $U\mathcal{N} = \mathcal{N}$ if and only if $\mathbf{1}$ is an eigenvector of U ; this is the only gate property the construction uses. BALANCE is the orthogonal projector B of Eq. (1), which fixes \mathcal{N} pointwise and annihilates $\mathbf{1}$.

Let S denote the BALANCE separator unitary whose first output mode is the normalized uniform vector, and let U_d denote the archived compiled circuit unitary for the Experiment 4 neutral jobs with $d = 0, 1, 2, 3$ core iterations (`E4f_neutral_no`, `E4f_neutral_1x`, `E4f_neutral_2x`, and `E4f_neutral_3x`). The implemented neutral-sector core is the one-cycle operator isolated from these payloads,

$$R_{\mathcal{N}} \equiv S^\dagger U_1 U_0^\dagger S. \quad (3)$$

The same core reconstructs the archived two- and three-

TABLE I. Validation checks for the implemented neutral-sector core $R_{\mathcal{N}}$ of Eq. (3). Values are regenerated from the archived Experiment 4 payloads by the extraction script in the reproducibility package; the operator is real and orthogonal and fixes the uniform vector with eigenvalue +1.

Check	Value
Field, type	real, orthogonal
Unitarity residual	1.33×10^{-15}
Uniform-vector eigenvalue	$1.000000 + 0.000000i$
$\det R_{\mathcal{N}}$	-1
Eigenphases (deg)	$0, \pm 15.28, \pm 43.37, \pm 166.82, 180$
Depth-2 reconstruction residual	1.11×10^{-15}
Depth-3 reconstruction residual	1.50×10^{-15}

cycle circuits as $U_d = S R_{\mathcal{N}}^d S^\dagger U_0$ to machine precision. Thus the hardware test of Sec. V probes repeated applications of one well-defined unitary core, followed by one BALANCE separator readout and single-photon postselection.

Table I lists the numerical checks. The full 8×8 matrix and the extraction script are included in the reproducibility package; the script regenerates the matrix from the raw Perceval payloads and prints the corresponding eigenphases. The extracted core is real and orthogonal (unitary to roundoff, $\det R_{\mathcal{N}} = -1$), fixes $\mathbf{1}$ with eigenvalue +1, and therefore preserves \mathcal{N} . No claim of optimality (fastest mixing, minimal depth, or otherwise) is made for this particular core; it is the archived representative used to exercise the neutral sector.

The projector B is not a unitary mesh element: on hardware it is realized only as the BALANCE/parity separator followed by single-photon postselection (Sec. IV). Therefore the compiled hardware circuits apply $R_{\mathcal{N}}^d$ ($d \geq 0$) and then perform a single separator readout; the projection onto the neutral sector is an algebraic description of the ideal readout, not an additional compiled gate at each cycle.

III. SIMULATION BENCHMARKS

Eight idealized simulation benchmarks provide reference points for the construction; the compact summary below lists all eight. These simulations are not hardware measurements and were not preregistered predictions.

ID	Benchmark and result
B1	Subspace structure: rank = 7, $\ B^2 - B\ = 0$
B2	Logical action (shallow): rank = 7, overlap = $1/7$
B4	BALANCE cadence: positive mean fidelity gain [†]
B6	Deep unitarity: $\kappa = 1.000$ at 200 cycles
B7	Single-mode decoder: 100% recovery (500 trials; Wilson 95% CI [99.2%, 100%])
B8	Parity-check vs. BALANCE: wins 16/16 conditions

[†]Representative magnitude from the archived dephasing model in `simulations.py`; this benchmark is model-dependent, and the script reproduces the qualitative outcome rather than the exact value (see below).

Benchmarks B1–B6 probe the neutral subspace \mathcal{N} ; B7–B8 probe the parity-checked subspace \mathcal{S} . The benchmarks fall into two classes. Benchmarks B1, B2, B6, and the single-mode decoder of B7 are *exact* algebraic checks: fully determined by the construction, they reproduce to numerical precision. They verify the projector structure ($\|B^2 - B\| = 0$), the rank and mixing of the shallow gate cycle (rank 7, with a mean squared basis-state overlap of $1/7$ reflecting near-uniform mixing within the seven-dimensional neutral subspace), the deep-unitarity condition number ($\kappa = 1.000\,000$), and 100% single-mode syndrome recovery. Benchmarks B4, the weight-2 decoder of B7, and B8 are *model-dependent* stochastic simulations evaluated under a transparent per-mode dephasing model with a fixed random seed; the magnitudes quoted for them are representative values of that model rather than exact predictions. The archived script (`simulations.py`) reproduces their qualitative outcomes (a positive BALANCE-cadence fidelity gain, weight-2 recovery beyond the formal distance-4 guarantee, and $P_{\mathcal{S}}$ outperforming BALANCE across all 16 conditions), and not the exact percentages, which depend on the noise-model implementation and seed. Two benchmarks reported in earlier versions of this manuscript (a dimension-weighted throughput comparison and an internal gate-count consistency check) have been removed: the former compared a simulated channel against a fixed reference constant on a metric that favors any high-dimensional subspace by construction, and the latter checked only the internal bookkeeping of the simulation code; neither constitutes an independent benchmark, and both remain available in the archived `simulations.py` for completeness. We expand below on the three results most directly relevant to the hardware (B6, B7, B8). These checks inform the construction but are not directly probed by the hardware experiments, whose results are reported in Sec. V.

B6 (deep unitarity). Evaluated on the ideal, noiseless operator, the implemented core $R_{\mathcal{N}}$ maintains condition number $\kappa = 1.000\,000$ through 200 simulated cycles, with all seven singular values equal to 1 to numerical precision (benchmark B6); the condition number and singular values are computed from the exact propagated operator, so no noise model enters this check. This provides the ideal sector-preservation reference for the Experiment 4 hardware test of repeated applications of $R_{\mathcal{N}}$ followed by BALANCE readout.

B7 (single-mode decoder). In simulation only, a syndrome-based decoder recovers 100% of single-mode amplitude perturbations across 500 randomized trials (benchmark B7; Wilson 95% CI [99.2%, 100%]). Each trial applies a single-mode amplitude perturbation of random magnitude (drawn uniformly from $[0, 1]$, in the first-quantized error model of Sec. II) on a uniformly

random mode of a random code state; the decoder assigns the syndrome to the check column most parallel to it (routine `b7a_single_mode_decoder`), which succeeds uniquely because the eight check columns are pairwise distinct. A greedy sequential decoder, which iteratively assigns the syndrome to the most parallel check column, subtracts that column’s contribution, and repeats (routine `b7b_weight2_greedy`), also recovers a large fraction of weight-2 patterns (a representative 71.4% under the archived dephasing model); this exceeds the formal distance-4 guarantee and reflects decoder behavior rather than a sharper distance claim, and the exact recovery fraction is model-dependent.

B8 (parity-checked vs. BALANCE). Across 16 noise-depth conditions (four per-cycle dephasing amplitudes $\sigma_{\phi} \in \{0.01, 0.03, 0.10, 0.30\}$ crossed with four depths $\{10, 50, 100, 200\}$ cycles), the parity-check projector $P_{\mathcal{S}}$ (the orthogonal projector onto \mathcal{S}) outperforms BALANCE alone in every one of the 16 conditions, with a representative fidelity improvement of up to $\sim 15\%$ at the highest noise level in benchmark B8. The clean sweep across all conditions is the robust outcome; the exact improvement magnitude is model-dependent. This benchmark isolates the modeled contribution of the three additional parity constraints beyond the single sum constraint.

IV. HARDWARE PLATFORM AND METHODS

All experiments were performed on Quandela Belenos [21], a 24-mode universal photonic processor accessed remotely through the Perceval software framework [22], with no user-side cryogenic infrastructure. The experiment uses the standard linear-optical model in which programmable interferometers implement unitary transformations on photonic modes through beam-splitter and phase-shifter decompositions and integrated multiport meshes [8, 10, 23–25].

Mode selection. We used 8 logical modes of the 24-mode Belenos processor, exposed in Perceval as an 8-mode circuit with ports 0–7. We did not manually optimize or selectively choose a specific physical subset of modes for fidelity; the circuit was submitted through the Belenos/Perceval compiler and the platform handled the assignment of logical to physical modes. Because the assignment is performed by the compiler rather than chosen by hand, the specific physical mode indices (and any run-to-run reassignment) are not fixed in advance; the logical-to-physical mapping returned by the compiler for each job, together with the calibration timestamp for every acquisition (and hence for each row of Tables II through IX), is recorded in the raw Perceval payloads provided in the reproducibility package.

State preparation. Input states were generated in software by QR decomposition. For each target neutral state $|\psi\rangle$, we constructed a unitary U_{prep} with $U_{\text{prep}}|0\rangle = |\psi\rangle$ by placing the normalized target in the first column

of an otherwise-identity matrix and completing it to an orthonormal basis with the Householder QR routine `numpy.linalg.qr`, fixing the sign convention so that $\text{diag}(R)$ is real-positive and the first column of U_{prep} equals $|\psi\rangle$ up to a global phase (routine `prep_unitary` in `operators.py`). The identical convention was used for every input in Tables III, VI, and VII. Because the completion of the remaining seven columns is convention-dependent, the ordering of the sub-percent residual DC components across inputs (for example the especially small residual of the $|0\rangle - |4\rangle$ pair) is a property of this particular preparation rather than of the target states; the exact preparation unitaries are archived in the reproducibility package, and a mathematically equivalent but different QR convention would in general redistribute that sub-percent residual leakage among inputs without altering the two-to-three-tier separation between neutral, non-neutral control, and pure-DC states. We then submitted the composite circuit $U_{\text{tot}} = U_{\text{sep}} U_{\text{prep}}$ as a single Perceval circuit, with a single photon injected into mode 0 as the physical input. Here U_{sep} denotes the readout separator (BALANCE or parity separator) applied after state preparation; both U_{prep} and U_{sep} were compiled into a single interferometer submission to the Belenos backend. The preparation unitary was not separately re-fit to hardware data, and was not modified between the calibrated, degraded, and restored chip states of the Hong–Ou–Mandel control (Sec. V).

Input states. In the mode basis (e_0, \dots, e_7) , with e_k the single-photon state in mode k , the inputs used in the heralding experiments are as follows. The *non-neutral control* is a single computational-basis mode, $|0\rangle = e_0$ (net sum 1, uniform overlap $1/8$, hence expected dump probability 0.125). The neutral pair inputs are $\frac{1}{\sqrt{2}}(e_0 - e_1)$ and $\frac{1}{\sqrt{2}}(e_0 - e_4)$ (a third neutral pair, $\frac{1}{\sqrt{2}}(e_0 - e_2)$, is used only for the face-parity probe of Experiment 5B), and the *balanced 4+4-* input is the $(-1)^{b_2}$ sign pattern $\frac{1}{\sqrt{8}}(e_0 + e_1 + e_2 + e_3 - e_4 - e_5 - e_6 - e_7)$; all three satisfy $\sum_i v_i = 0$. The four \mathcal{S} -basis inputs are the \mathbb{Z}_2^3 characters $|\chi_a\rangle = \frac{1}{\sqrt{8}} \sum_x (-1)^{a \cdot x} e_x$ with $a \in \{011, 101, 110, 111\}$ (written as $a_2 a_1 a_0$; Sec. V, Experiment 5A), i.e. the labels $b_0 \oplus b_1$, $b_0 \oplus b_2$, $b_1 \oplus b_2$, and $b_0 \oplus b_1 \oplus b_2$. These inputs were chosen to probe distinct structural features of the construction: local neutral pairs, the self-conjugate neutral pair $|0\rangle - |4\rangle$ whose opposite signs on the two self-conjugate modes ($k = 0$ and $k = 4$) yield the smallest residual DC component, a balanced four-plus/four-minus neutral state, and the four character basis states spanning \mathcal{S} .

Operating conditions. The Perceval remote-processor metadata for `qpu:belenos` reported a `CLock` (MHz) field of 4.94, which we adopt as the single-photon clock rate; this device-level metadata value is distinct from public platform shot-rate summaries. The same device-metadata record also reported a second-order correlation $g^{(2)}(0) = 0.019$ and a transmittance of 4.84%. The Hong–Ou–Mandel visibility [26] exceeded 90% on

the calibrated chip. Each test configuration collected approximately 10,000 single-photon events. Statistical uncertainties on every measured fraction $\hat{p} = x/n$ are reported as 95% Wilson confidence intervals [27], $(\hat{p} + z^2/2n \pm z\sqrt{\hat{p}(1-\hat{p})/n + z^2/4n^2})/(1 + z^2/n)$, with $z = 1.96$. Suppression-ratio uncertainties are computed by the log-ratio (delta) method: $\text{SE}(\log R) = \sqrt{(1 - \hat{p}_a)/x_a + (1 - \hat{p}_b)/x_b}$ for $R = \hat{p}_a/\hat{p}_b$.

Multi-photon contamination. The residual multi-photon content of the source is measured directly from the archived exports, whose decoded port-count distributions retain every photon-number outcome, not only the postselected single-photon subset. Across the 350,000 events recorded in the six experiments, 171 registered two photons and none registered three or more, a directly observed multi-photon detection fraction of 0.049%, and at most 0.10% ($\leq 10/10,000$) in any single configuration (the same 0.049% rate holds over the full set of archived campaign exports; audited by `photon_number_audit.py` in the reproducibility package). These detected pairs are flagged and removed by the single-photon postselection.

The only remaining multi-photon path is a two-photon emission in which one photon is lost before detection and the survivor is counted as a single photon; using the reported 4.84% transmittance η , the ratio of one-detected to two-detected outcomes of a photon pair is $2(1 - \eta)/\eta \approx 39$, placing such lost-twin survivors at $\approx 1.9\%$ [30] of the postselected single-photon counts (consistent with the source $g^{(2)}(0) = 0.019$). A lone surviving photon, however, traverses the same linear circuit U_{tot} as the intended single photon and undergoes the same single-particle coherent cancellation, landing with the identical port distribution $|\langle \text{port} | U_{\text{tot}} | \psi \rangle|^2$; lacking only the two-photon interference of a genuine pair, it should reproduce the single-photon signal rather than produce a distinct dump or syndrome bias. This argument relies on the surviving photon’s marginal one-mode state matching the intended single-photon input. Both the one- and two-photon components of the source are injected into the same single spatial mode (mode 0), so a lost-twin survivor is a mode-0 single photon identical to the intended input and traverses U_{tot} with the same port distribution; any residual bias is controlled by the source’s spatial-mode purity and is bounded above by the $\leq 1.9\%$ survivor weight, below the two-to-three-order-of-magnitude leakage and syndrome separations reported in Sec. V. The multi-photon component is therefore either explicitly discarded (detected pairs, 0.049%) or constrained to the same single-particle port statistics up to the survivor bound; it neither accounts for the residual $\sim 1\%$ preparation/calibration leakage floor of Experiments 2 and 3 (Eq. (5)) nor alters the headline suppression ratios of Sec. V.

Acquisition, postselection, and backend version. Each configuration was executed on the Belenos backend through the Perceval framework during the April 2026 campaign. The Perceval remote-processor metadata identified the backend as `qpu:belenos` (“Belenos

QPU”), and the exported job payloads recorded `pcvl_version 1.0.1`; the API-reported backend software stack was `mosaiq-belenos 2.7.12`, `hardware-core 3.7.4`, `pcvl-worker 1.4.0`, `perceval-quandela 1.1.0`, `universalchipworker 1.7.0`, and `exqalibur 1.1.1`. The Hong–Ou–Mandel control (Sec. V) spans the calibrated, degraded, and restored chip states recorded on April 4–6, 2026. Circuits were compiled to the native Belenos interferometer mesh by the default Perceval/Belenos compiler with no manual remapping (cf. *Mode selection*); the use of a programmable multiport mesh follows the standard universal interferometer model [23, 24]. Readout is single-photon: we postselect on events with exactly one detected photon across the eight output ports, discarding zero-photon (loss) events and multi-photon and dark-count-flagged events. The reported dump and syndrome probabilities are normalized to the total postselected single-photon counts per configuration, so mode-independent loss divides out and only the relative port-occupation distribution enters. Across all configurations the single-photon postselection retains 99.9% of detected events (range 99.9–100%), with no systematic dependence on input state or core depth; the small discarded fraction is the multi-photon component audited above. Absolute photon loss (zero-photon events) is not registered by the platform’s event counter and, being mode-independent, divides out under this per-configuration normalization, so all reported probabilities are conditional on single-photon detection. Claims involving loss or attenuation should therefore be read as relative path-amplitude statements conditional on this postselection, not as measurements of absolute optical loss. The per-configuration raw port-count distributions, the exact preparation unitaries, and the postselection masks are included in the arXiv ancillary files (see the Data Availability section). Figure 2 summarizes the corresponding hardware workflow.

V. HARDWARE RESULTS

The six hardware experiments follow the workflow of Fig. 2. Experiments 1–4 and 6 read out the prepared state on the BALANCE (DC dump) port, with Experiment 4 inserting $d = 1$ –3 applications of the neutral-sector core $R_{\mathcal{N}}$ before readout; Experiment 5 instead routes the state through the parity separator to resolve the DC/sum and three face-parity syndrome channels.

A. Experiment 1: BALANCE separator (80,000 events)

Single photons in each of the 8 computational modes are sent through the BALANCE separator. The dump port (port 0) should capture 1/8 of the input from each mode.

Result. The dominant output port is correct for all

TABLE II. Experiment 1: BALANCE separator per-input dump probability. A single photon is injected in each computational mode m ; the dump port (port 0) should capture the ideal fraction $1/8 = 0.125$. Dump probability \hat{p} is the port-0 fraction of n postselected single-photon events; uncertainties are 95% Wilson confidence intervals. The best-performing (lowest dump) mode is boldfaced.

Input mode	n	\hat{p} (95% CI)
$m = 0$	9,993	0.193 [0.185, 0.201]
$m = 1$	9,994	0.157 [0.150, 0.164]
$m = 2$	9,992	0.140 [0.133, 0.147]
$m = 3$	9,990	0.178 [0.171, 0.186]
$m = 4$	9,992	0.188 [0.181, 0.196]
$m = 5$	9,995	0.237 [0.229, 0.246]
$m = 6$	9,995	0.224 [0.216, 0.233]
$m = 7$	9,991	0.186 [0.179, 0.194]

8 inputs (Table II). The mean dump probability across modes is 0.188 ± 0.030 (the spread denotes the across-mode standard deviation; ideal value 0.125); the best-performing input (mode 2) gives 14.0% on the dump port (Wilson 95% CI [0.133, 0.147]). The separator routes the dominant output as predicted, with a measurable calibration bias in the dump probability that is carried through in the controls below. As an implementation-fidelity figure for the separator, the measured DC-capture channel deviates from its ideal 0.125 value by a root-mean-square of 0.070 across the eight modes (a mean upward bias of 50%). This fixed-pattern bias is the dominant systematic of the BALANCE experiments: it is the same offset that inflates the non-neutral control in Experiment 2 and is the origin of the gap between the directly measured $31.6\times$ and the calibration-referred $\approx 21\times$ suppression reported there.

B. Experiment 2: Neutral-state heralding (50,000 events)

Photons prepared in neutral superpositions $\sum_i v_i = 0$ are sent through the BALANCE separator.

Result. Table III resolves the expected three-tier behavior. The non-neutral control ($\hat{p} = 0.191$) lies above the ideal 0.125, consistent with chip-level imbalance in the BALANCE separator already seen in Experiment 1. The pure-DC input saturates the dump port at 0.980, two orders of magnitude above the neutral inputs and consistent with full DC capture up to a $\sim 2\%$ residual population in the other output ports. The three neutral inputs sit between these extremes in a band of width roughly 1%, with the $|0\rangle - |4\rangle$ pair delivering the lowest dump probability ($\hat{p} = 0.0002$); this input places opposite signs on the two self-conjugate modes 0 and 4 and, in the QR-decomposed preparation used here, accumulates the smallest residual component on the uniform vector.

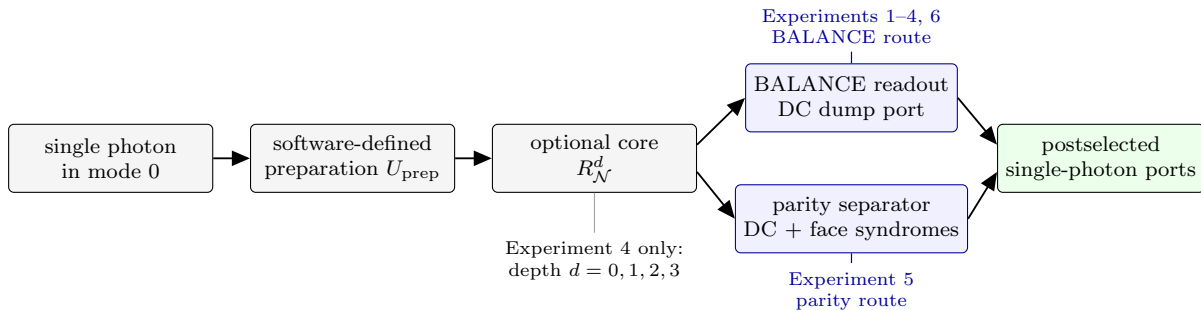


FIG. 2. Experimental workflow for the single-photon hardware tests. A software-defined preparation unitary maps the injected single photon to a target neutral or parity-checked input state. Depending on the experiment, the prepared state is sent directly to a separator or first through d applications of the neutral-sector unitary core $R_{\mathcal{N}}$. BALANCE readout measures the DC leakage channel, while the parity separator routes the DC/sum and face-parity syndrome channels to output ports. All reported probabilities are normalized to postselected single-photon detections across the eight logical ports.

TABLE III. Neutral-state heralding on Belenos. The control row is the non-neutral input; the uniform (pure DC) row is the all-modes-in-phase reference; all remaining rows are neutral inputs ($\sum_i v_i = 0$). The best-performing (lowest dump) neutral input is boldfaced. Dump probability \hat{p} is the fraction of photons measured on port 0 out of n events; uncertainties are 95% Wilson confidence intervals.

Input	n	\hat{p} (95% CI)	Expected
Control (non-neutral)	9,998	0.191 [0.183, 0.198]	0.125
$ 0\rangle - 1\rangle$	9,998	0.0073 [0.0058, 0.0092]	0
$ 0\rangle - 4\rangle$	9,995	0.0002 [5×10^{-5} , 7×10^{-4}]	0
Balanced 4+4-	9,993	0.0106 [0.0088, 0.0128]	0
Uniform (pure DC)	9,996	0.980 [0.977, 0.982]	1.000

Averaging over the three distinct neutral inputs yields a mean of 0.6% (95% CI [0.5%, 0.7%]). Referred to the ideal DC-capture baseline 0.125, this corresponds to a suppression of $0.125/0.006 \approx 21\times$, which isolates the interference suppression for neutral states against the theoretical baseline. The directly measured ratio relative to the non-neutral control is larger, $31.6\times$ (log-ratio 95% CI [27.2, 36.7]); that control baseline 0.191 itself carries the fixed-pattern BALANCE calibration bias documented in Experiment 1 (mean per-mode dump 0.188 vs. ideal 0.125), so the $31.6\times$ figure conflates genuine destructive interference on the DC port with chip miscalibration that inflates the control above ideal. We report both figures, and the gap between them quantifies the BALANCE calibration bias, not additional physics. The three neutral inputs themselves span 0.0002–0.0106 (Table III); this input-to-input range, rather than the narrow Wilson interval on the pooled mean (which reflects photon-counting statistics only), is the physically relevant measure of variation, and it is dominated by how much residual uniform component each preparation accumulates under the QR convention of Sec. IV. As a between-job reproducibility check, the $|0\rangle - |1\rangle$ neutral baseline was acquired in three independent jobs (Experiments 2, 3, and 4) and gave dump probabilities 0.0073, 0.0070, and 0.0075, agreeing within their Wilson intervals, so for this shared configuration run-to-run variation

is no larger than photon-counting statistics. Because every probability is normalized to the postselected single-photon counts (Sec. IV), this ratio is independent of any mode-independent loss, which cancels between the neutral and control inputs; the Hong–Ou–Mandel control of Experiment 6 further indicates that the suppression is phase-coherent rather than a normalization or intensity artifact. The hardware therefore tests representative vectors from the theoretically 7-dimensional error-detecting subspace, with the residual $\sim 1\%$ leakage consistent with preparation and/or calibration imperfections rather than a failure of the linear sum constraint.

C. Experiment 3: Proportional error detection (70,000 events)

A controlled DC contamination with strength $c \in [0, 1.6]$ is injected into a prepared neutral state and the resulting dump probability is measured. Concretely, the injected state is $|\psi(c)\rangle \propto |\psi_{\mathcal{N}}\rangle + c\hat{\mathbf{u}}$, where $|\psi_{\mathcal{N}}\rangle \in \mathcal{N}$ is the target neutral state and $\hat{\mathbf{u}} = \mathbf{1}/\sqrt{8}$ is the unit uniform (DC) vector orthogonal to \mathcal{N} ; thus c is the amplitude ratio of the injected DC component to the neutral state. After normalization the predicted dump (DC) probabil-

TABLE IV. Error injection. The hardware response is monotonic in the injected DC fraction and strongly correlated with the theoretical prediction. Hardware uncertainties are 95% Wilson confidence intervals on $n \approx 10,000$ events per row.

c	Theory dump	Hardware dump (95% CI)
0.00	0.000	0.007 [0.0055, 0.0088]
0.05	0.003	0.015 [0.0126, 0.0174]
0.10	0.010	0.020 [0.0174, 0.0230]
0.20	0.039	0.038 [0.0340, 0.0414]
0.40	0.138	0.187 [0.1800, 0.1953]
0.80	0.390	0.437 [0.4269, 0.4463]
1.60	0.719	0.814 [0.8065, 0.8217]

ity is the squared DC weight,

$$p(c) = \frac{c^2}{1 + c^2}, \quad (4)$$

which defines the ‘‘Theory dump’’ column of Table IV (so that $p(0) = 0$ and $p(c) \rightarrow 1$ as $c \rightarrow \infty$).

Result. Rather than summarizing the agreement by a correlation coefficient alone, we fit the hardware response to a two-parameter calibration model that augments the ideal curve $p(c)$ of Eq. (4) with a leakage floor and a multiplicative gain,

$$p_{\text{hw}}(c) = p_{\text{floor}} + g \frac{c^2}{1 + c^2}. \quad (5)$$

A least-squares fit to the seven points of Table IV (two free parameters, five degrees of freedom) gives $p_{\text{floor}} = 0.010$ and $g = 1.12$: a $\sim 1\%$ preparation-leakage floor (consistent with the $\sim 0.7\%$ – 1% residual of Experiment 2) and a 12% multiplicative gain on a near-saturated channel. The fit residuals, $\hat{p} - p_{\text{hw}}(c)$, are at the ± 0.02 level (largest, $+0.024$, near $c = 0.4$) and are compatible in sign and size with the fixed-pattern calibration offsets of the BALANCE separator seen in Experiments 1 and 2, and not with a departure from the $c^2/(1 + c^2)$ response. Referred to the photon-counting (Wilson) uncertainties alone, the fit gives $\chi^2/\text{dof} \approx 24$ (5 degrees of freedom): the residuals are roughly an order of magnitude larger than statistical scatter, confirming that they are dominated by systematic calibration offsets at the ± 0.02 level rather than by counting noise or a failure of the model shape. We therefore treat this ± 0.02 residual scale as a per-point systematic budget and regard the calibration model of Eq. (5), rather than a correlation coefficient, as the operative description. The BALANCE separator therefore reports a continuous, monotonic, and calibratable soft error signal across $c \in [0.05, 1.6]$, well described over this tested range by a floor-plus-gain model with the ideal response shape fixed.

TABLE V. Neutral-sector unitary core on Belenos. Dump probability \hat{p} at 0, 1, 2, and 3 iterations of $R_{\mathcal{N}}$ on a neutral input, with two non-neutral controls (no core cycle and $1\times$); the lowest neutral dump probability is boldfaced; uncertainties are 95% Wilson confidence intervals.

Configuration	n	\hat{p} (95% CI)
Neutral, no core cycle	9,995	0.0075 [0.0060, 0.0094]
Neutral, $1\times$ core cycle	9,995	0.0004 [2×10^{-4} , 1×10^{-3}]
Neutral, $2\times$ core cycle	9,997	0.0018 [0.0011, 0.0028]
Neutral, $3\times$ core cycle	9,999	0.0002 [5×10^{-5} , 7×10^{-4}]
Control, no core cycle	9,996	0.196 [0.188, 0.204]
Control, $1\times$ core cycle	9,998	0.190 [0.182, 0.197]

D. Experiment 4: Leakage at three compiled unitary-core depths, limited by compilation systematics (60,000 events)

The neutral-sector unitary core $R_{\mathcal{N}}$ (Sec. IID) is applied 1, 2, and 3 times to a prepared neutral state, followed by BALANCE readout.

Result. Table V lists the numerical values; Fig. 3 displays the same data on a logarithmic scale. The no-cycle neutral baseline is $\hat{p} = 0.0075$; after one and three core cycles the measured dump probabilities fall to 0.0004 and 0.0002, respectively, and their Wilson intervals overlap. The $2\times$ point, 0.0018, is also below the no-cycle baseline but lies above the $1\times$ and $3\times$ lower-leakage band, so the depth dependence is non-monotonic rather than a simple shot-noise fluctuation. Each depth was submitted as a different compiled unitary, so cross-depth comparisons also include depth-dependent compiler and calibration variance; the claim is therefore only that leakage remains low and shows no monotonic growth over these tested depths. Indeed, because $R_{\mathcal{N}}$ fixes the uniform (DC) vector $\mathbf{1}$ with eigenvalue $+1$ (Sec. IID), the ideal core preserves the DC/dump component of its input exactly and cannot by itself reduce the dump leakage; in the noiseless circuit the dump probability would be independent of depth. The observed scatter across depths (including the $1\times$ and $3\times$ points lying below the $d = 0$ value) is therefore consistent with differing preparation and separator calibration for each separately compiled circuit, not any cleaning action of the core. The log-scale representation in Fig. 3 makes the main conclusion visible: all cycled neutral points remain far below the two non-neutral controls (no core cycle, $\hat{p} = 0.196$; $1\times$, $\hat{p} = 0.190$). The two controls agree within their Wilson intervals, indicating that the unitary core does not by itself change the non-neutral leakage level. The direct ratio of the $1\times$ control to the $3\times$ neutral point is $0.190/0.0002 \approx 950$, but the small denominator (2 events out of 9,999) gives a wide log-ratio 95% CI of [2.4×10^2 , 3.8×10^3]; we therefore treat this ratio as descriptive rather than as the primary statistical summary.

In conventional gate-based quantum-error-correction

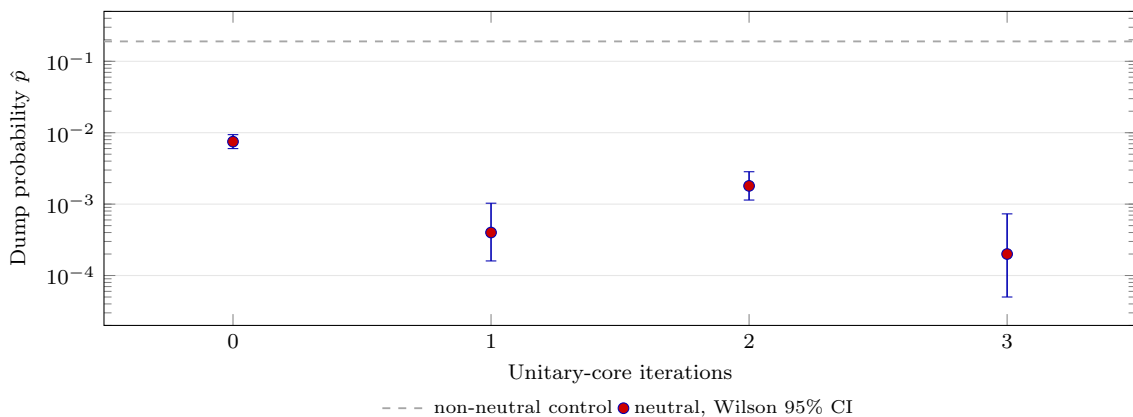


FIG. 3. Dump probability at the tested neutral-sector core depths on a logarithmic scale (Experiment 4). These are four operating-depth points (0, 1, 2, 3 applications) with a non-monotonic ordering, *not* a depth scan: no curve is fit through them and the figure should not be read as a depth-scaling trend (a dedicated ~ 10 – 20 cycle scan, deferred to future work, would be required to extract a per-cycle leakage rate). Markers are the hardware values; vertical bars are 95% Wilson confidence intervals. The dashed line is the non-neutral control level ($\hat{p} = 0.190$, the $1\times$ control; the no-core-cycle control agrees at 0.196, statistically indistinguishable from it, so the core does not by itself reduce the control leakage). All four neutral points remain well below the control band; the $1\times$ and $3\times$ points have overlapping 95% confidence intervals, while the $2\times$ point lies above this lower band but well below the no-cycle baseline.

experiments, repeated gate application tends to degrade fidelity [3, 5]. The present data make a narrower claim: BALANCE leakage remains far below the non-neutral control after one, two, and three applications of the neutral-sector unitary core. This agrees with the modeled unitarity and preservation of \mathcal{N} (Sec. II, B6 in Sec. III), while the non-monotonic $2\times$ point indicates that residual preparation leakage and chip calibration also affect the measured dump probability. Three depths are not sufficient to fit a per-cycle leakage rate or to separate genuine algebraic sector preservation from cumulative coherent error masked by the small dynamic range; accordingly, we claim only that leakage stays well below the non-neutral control for the three tested depths and do *not* claim monotonic sector preservation under iteration. A dedicated depth scan to ~ 10 – 20 cycles would be required to extract a per-cycle leakage rate, and is left to future work. The chip realization of the core carries an unmodeled implementation error budget (insertion-loss imbalance, beamsplitter angle errors, phase-calibration drift), so the observed leakage suppression is bounded by the hardware-level fidelity of the implemented unitary, not by the algebraic identity alone. Simulation benchmark B6 provides the ideal reference for sector preservation; the hardware leakage data are compatible with that qualitative expectation within the tested depths. This measurement shows retention in the neutral sector as read out by BALANCE; it does not by itself constitute a process-tomographic measurement of the logical action within \mathcal{N} .

E. Experiment 5: Parity-checked subspace and syndrome channels (80,000 events)

The parity-check separator U_S is the \mathbb{Z}_2^3 character table normalized by $1/\sqrt{8}$, with rows ordered as $a \in \{000, 001, 010, 100, 011, 101, 110, 111\}$ (the four check characters first, then the four code characters), so that ports 0–3 correspond to the syndrome subspace and ports 4–7 to the code subspace \mathcal{S} . The port assignment follows directly from the row structure of C : the four rows of C correspond to the characters with labels $a \in \{000, 001, 010, 100\}$ (the DC/sum row and the three face-parity rows), so any state in $\ker(C) = \mathcal{S}$ has zero projection onto these four directions and is routed entirely to ports 4–7; a parity violation produces a nonzero projection onto one or more of ports 0–3. Concretely, each basis state of \mathcal{S} is a \mathbb{Z}_2^3 character vector: for a label $a = (a_2, a_1, a_0) \in \mathbb{Z}_2^3$ it is $|\chi_a\rangle = \frac{1}{\sqrt{8}} \sum_{x \in \mathbb{Z}_2^3} (-1)^{a \cdot x} e_x$, where $x = b_2 b_1 b_0$ runs over the eight vertices and $a \cdot x = \bigoplus_j a_j x_j$. The four code states tabulated below carry the labels $b_0 \oplus b_1$, $b_0 \oplus b_2$, $b_1 \oplus b_2$, and $b_0 \oplus b_1 \oplus b_2$ (equivalently $a = 011, 101, 110, 111$ written as $a_2 a_1 a_0$), the four characters orthogonal to the row space of C . For example,

$$|b_0 \oplus b_1\rangle = \frac{1}{\sqrt{8}} (e_{000} - e_{001} - e_{010} + e_{011} + e_{100} - e_{101} - e_{110} + e_{111}).$$

5A. Code-space heralding. The four \mathcal{S} -basis inputs are prepared and sent through the parity separator; port fractions are recorded to test how well each basis state is confined to code-space ports 4–7.

Table VI resolves the four basis-state behaviors of \mathcal{S} along two axes: total syndrome leakage (\hat{p}_{syn} , ports 0–3) and code-side port fidelity (the dominant entry on its

TABLE VI. Parity-checked-subspace heralding. Syndrome probability \hat{p}_{syn} is the total fraction on ports 0–3 (ideal: 0); Port is the predicted dominant code-space output port (ports 4–7 carry the code subspace); Fid. is the fraction of all detections on that predicted port; uncertainties are 95% Wilson confidence intervals on $n \approx 10,000$ events per row.

Basis state	\hat{p}_{syn} (95% CI)	Port	Fid.
$b_0 \oplus b_1$	0.009 [0.0073, 0.0111]	4	98.8%
$b_0 \oplus b_2$	0.062 [0.057, 0.067]	5	91.0%
$b_1 \oplus b_2$	0.011 [0.0090, 0.0130]	6	96.9%
$b_0 \oplus b_1 \oplus b_2$	0.012 [0.0096, 0.0138]	7	93.6%
Mean ($n = 39,982$)	0.0232 [0.0218, 0.0248]		95.1%

predicted code port). Three of the four basis states cluster near $\hat{p}_{\text{syn}} \in [0.009, 0.012]$; the $b_0 \oplus b_2$ basis state is a clear outlier at $\hat{p}_{\text{syn}} = 0.062$, and its code-side port fidelity (91.0%) is correspondingly the lowest of the four. We attribute this 5–7 \times spread, tentatively, to state-dependent calibration asymmetry rather than a global loss of the parity-check structure: the other three basis states retain percent-level syndrome leakage, so the parity-check routing is intact. This attribution is, however, a post hoc hypothesis that the present single run cannot confirm. To make it quantitative and testable, we record the calibration model that links per-mode imperfections to state-dependent syndrome leakage. Let $D = \text{diag}(1 + \varepsilon_x)$ describe the residual per-mode multiplicative error of the parity separator, where $\varepsilon_x \in \mathbb{C}$ collects insertion-loss imbalance and small beamsplitter-angle and phase offsets on mode x . To first order in ε , the amplitude that an ideal code character $|\chi_a\rangle$ ($a \in \{011, 101, 110, 111\}$) leaks onto check row r (itself the character χ_r with $r \in \{000, 001, 010, 100\}$), the DC/sum and three face-parity rows) is $\langle \chi_r | D | \chi_a \rangle = \delta_{ra} + \hat{\varepsilon}(a \oplus r)$, where $\hat{\varepsilon}(w) = \frac{1}{8} \sum_x (-1)^{w \cdot x} \varepsilon_x$ is the Walsh–Hadamard transform of the error profile. Because no code label a is a check row, the predicted syndrome leakage is

$$p_{\text{syn}}(a) \approx \sum_{r \in \{000, 001, 010, 100\}} |\hat{\varepsilon}(a \oplus r)|^2, \quad (6)$$

the squared Walsh weight of the calibration profile at the frequency a and its three single-face-flip neighbors. The state dependence is therefore fixed entirely by the Walsh spectrum of the per-mode error: for the observed outlier $a = 101$ ($b_0 \oplus b_2$) the relevant frequency set is $\{101, 100, 111, 001\}$, so a calibration profile whose spectral weight concentrates in this set elevates that one character while leaving the other three near the percent-level floor, which is exactly the pattern of Table VI. Equation (6) makes the attribution falsifiable in two independent ways: (i) obtaining the per-mode insertion-loss and beamsplitter-angle-offset vector from the platform calibration record, Walsh-transforming it, and checking that its weight concentrates in $\{101, 100, 111, 001\}$; and (ii) a repeat acquisition after recalibration, which should move the outlier to whichever code character’s frequency set

TABLE VII. Selective parity detection. Each face-parity violation activates the corresponding syndrome port; S_x , S_y , S_z label the input violating face-parity check M_0 , M_1 , M_2 respectively; selectivity is the fraction of syndrome-subspace detections (ports 0–3) routed to the expected port, with 95% Wilson confidence intervals.

Input	Exp. port	Meas. port	Selectivity (95% CI)
$ 0\rangle - 1\rangle$ (S_x)	1	1	0.985 [0.980, 0.988]
$ 0\rangle - 2\rangle$ (S_y)	2	2	0.944 [0.935, 0.952]
$ 0\rangle - 4\rangle$ (S_z)	3	3	0.942 [0.934, 0.949]

TABLE VIII. Syndrome confusion matrix for Experiment 5B: the fraction of syndrome-subspace detections (ports 0–3) routed to each syndrome port, for the three single-parity-violating inputs. Diagonal entries (boldface) are the dominant-port selectivities of Table VII; off-diagonal entries are cross-talk. Rows sum to 1 up to rounding.

Input	Port 0	Port 1	Port 2	Port 3
S_x ($ 0\rangle - 1\rangle$)	0.008	0.985	0.007	0.001
S_y ($ 0\rangle - 2\rangle$)	0.008	0.021	0.944	0.028
S_z ($ 0\rangle - 4\rangle$)	0.015	0.023	0.020	0.942

overlaps the new loss spectrum rather than remain fixed on $b_0 \oplus b_2$. We identify both as necessary follow-up measurements; a transient, chip-specific artifact uncorrelated with the calibration spectrum would fail both tests. Averaged over all four inputs ($n = 39,982$), the mean syndrome leakage is 2.32% (95% CI [2.18%, 2.48%]), so representative basis states of the 4-dimensional parity-checked subspace are reproduced on hardware with percent-level residual syndrome leakage.

5B. Selective parity detection.

The three inputs are neutral pairs that each violate a single face-parity check; the labels S_x , S_y , and S_z denote violations of M_0 , M_1 , and M_2 , routed to syndrome ports 1, 2, and 3, respectively. Each row of Table VII reports the theoretically predicted syndrome port and the empirically measured dominant syndrome port. For all three single-parity-violating inputs the predicted and measured ports coincide. The dominant-port selectivity point estimates are 94–99%: the S_x selectivity (0.985, 95% CI [0.980, 0.988]) is higher than the S_y and S_z estimates (0.944 and 0.942; 95% CIs [0.935, 0.952] and [0.934, 0.949]), while the S_y and S_z channels are indistinguishable from each other at this sample size. The full syndrome confusion matrix (the distribution of each input over the four syndrome ports 0–3) is given in Table VIII. Off-target routing is below 1% for S_x and at the 2–3% level for S_y and S_z , with no single off-target port exceeding 2.8%; these off-diagonal rates bound the syndrome-label error that would propagate into any future feed-forward correction. The complete per-input distributions over all eight output ports are tabulated in the reproducibility package. We nonetheless interpret the measured selectivity here as syndrome-label resolu-

TABLE IX. Hong–Ou–Mandel control. Identical circuit and preparation, varying photon indistinguishability; calibration dates are in 2026. Dump probabilities are 95% Wilson confidence intervals; the calibrated state used $n \approx 10,000$ single-photon events and the degraded and restored states used $n = 5,000$.

Chip state	HOM vis.	Dump prob. (95% CI)
Calibrated (Apr. 4)	90.7%	0.0073 [0.0058, 0.0092]
Degraded (Apr. 5–6)	21.2%	0.481 [0.467, 0.495]
Restored (Apr. 6)	$\sim 91\%$	0.0058 [0.0040, 0.0083]

tion, not as a complete syndrome-state tomography. The three measured dominant ports resolve the three face-parity syndrome labels.

5C. Suppression ratio. The mean code-space syndrome leakage of 2.32% (95% CI [2.18%, 2.48%]) contrasts with 55.0% (95% CI [54.0%, 56.0%]) for a single computational-mode control sent through the same parity separator (ideal syndrome fraction 1/2). Referred to the ideal control fraction, the suppression is $0.5/0.0232 \approx 21.6\times$; the directly measured on-chip ratio is $23.7\times$ (log-ratio 95% CI [22.2, 25.3]), the gap reflecting the same kind of separator calibration bias quantified in Experiments 1 and 2 (measured control 55.0% vs. ideal 50%). As in Experiment 2, this ratio compares post-selected single-photon fractions, so mode-independent loss cancels (Sec. IV). Because $\mathcal{S} \subset \mathcal{N}$, the coherence sensitivity diagnosed in Experiment 6 is also relevant to the parity-checked subspace.

The three resolved parity channels are necessary, but not sufficient, hardware primitives for the theoretical single-mode decoder of benchmark B7 (100% recovery in simulation); full recovery also requires complex-syndrome readout and feed-forward control.

F. Experiment 6: Hong–Ou–Mandel degradation episode (diagnostic observation)

A control experiment exploits a Belenos recalibration/maintenance window during which the Hong–Ou–Mandel visibility dropped from 90.7% to 21.2%. The degradation was opportunistic, not deliberately induced, so the control is best interpreted as a diagnostic change in indistinguishability rather than as an isolated knob. We re-ran the same neutral-heralding circuit and the same software-defined preparation in the degraded chip state and again after the visibility was restored on recalibration.

Table IX compares the dump probability across three chip states with the circuit, the input preparation, and the photon-counting infrastructure held fixed. The main recorded diagnostic change is the two-photon visibility, which fell from 90.7% to 21.2% during the recalibration window and was restored to approximately 91% on completion; the post-recalibration visibility was taken

TABLE X. Summary of hardware results.

Experiment	Result	Events
1. BALANCE separator	8/8 correct ports	80,000
2. Neutral heralding	$\approx 21\times$ ideal (31.6 \times raw)	50,000
3. Error detection	calibrated soft response	70,000
4. Gate cycle	low leakage, depths 1–3	60,000
5. Syndrome channels	23.7 \times , 3/3 ports	80,000
6. HOM control	interference dependence	20,000
Total		> 340,000

from the platform’s standard calibration report rather than separately remeasured under our circuit, hence the “ \sim ” in the restored entry. Across the three chip states the dump probability tracks the platform Hong–Ou–Mandel/calibration diagnostic: at calibrated visibility it sits at 0.0073 (95% CI [0.0058, 0.0092]), matching the corresponding row of Table III; at degraded visibility it jumps by a factor of ~ 66 to 0.481 (95% CI [0.467, 0.495]); on restoration it falls back to 0.0058 (95% CI [0.0040, 0.0083]). The Wilson intervals of the calibrated and restored states overlap, so the calibrated and restored suppression levels are statistically indistinguishable. The degraded value of 0.481 lies nearly two orders of magnitude away from both the calibrated and restored values, consistent with loss of the phase-sensitive cancellation that defines \mathcal{N} once the BALANCE separator acts on an incoherent or poorly phased input; we do not attempt a quantitative model of its specific magnitude, which depends on the detailed degraded chip state (prepared-state and separator phase errors during the maintenance window) rather than on a simple fully-incoherent mixture. The control is therefore consistent with, but does not by itself isolate, a phase-coherent interference mechanism, and argues against a purely intensity-based artifact as the source of the calibrated suppression; a deliberate visibility scan (adding a known temporal delay between photons to vary the indistinguishability through a controlled range and back) would be required to isolate the mechanism, and is left to future work. Since $\mathcal{S} \subset \mathcal{N}$, the parity-checked subspace requires the same coherent neutral-state preparation and is expected to exhibit the same sensitivity to indistinguishability loss. Because the maintenance-window change was not an isolated control parameter, we interpret the HOM-visibility comparison as a diagnostic of the calibrated versus degraded photonic operating state, not as a standalone measurement of the microscopic mechanism.

G. Hardware summary

Table X collects the six experiments and their event counts (the Experiment 6 entry is the exact sum of the 10,000 calibrated, 5,000 degraded, and 5,000 restored ac-

quisitions of Table IX). The per-configuration counts sum to roughly 360,000 single-photon acquisitions; a few baseline and control rows are shared across experiments (for instance, the calibrated Hong–Ou–Mandel row and the zero-contamination point of Experiment 3 coincide with the Experiment 2 neutral baseline), so we quote the conservative total of more than 340,000 distinct events. Together these measurements test the operational pieces used in the construction: BALANCE readout, leakage suppression for representative neutral states, a calibrated soft response to DC contamination, low leakage after repeated applications of the neutral-sector unitary core, and resolved dominant ports for the three face-parity syndromes. The simulations of Sec. III remain idealized references for these behaviors rather than hardware measurements of throughput or decoder recovery.

VI. CONTEXT AND SCOPE

The encoding here is a single-photon spatial-mode code rather than a multi-qubit stabilizer code, so any comparison of code rate is necessarily indirect. We list the relevant numbers for context only. The rate axes are different and the figures below should not be compared directly: they are given solely to situate the construction among near-term options, not to rank it against stabilizer codes, whose rates count logical qubits per physical qubit against arbitrary Pauli errors. Concretely, comparing $7/8$ with a stabilizer-code rate is not a unit conversion but a category mismatch: the former counts protected spatial dimensions carried by a single postselected photon, whereas the latter counts logical qubits per encoded physical qubit against arbitrary Pauli errors.

Spatial-mode dimension rate. The neutral subspace operates at rate $7/8$ (87.5%) in the single-photon spatial-mode sense: 7 protected dimensions in 8 physical modes. Multi-qubit stabilizer codes report rates on a different resource axis (logical qubits per physical qubit against arbitrary Pauli errors), and comparing the two is a category mismatch rather than a unit conversion; we therefore do not tabulate stabilizer-code rates against the spatial-mode figure and refer the reader to Refs. [2, 3, 5–7] for representative values. The relevant point is only that a high-rate single-photon spatial-mode construction with experimentally accessible syndrome signals is of interest on near-term photonic processors, where multi-qubit stabilizer infrastructure is not yet available.

Behavior under repeated gates. In stabilizer-code experiments, logical error per cycle is positive and overall fidelity decreases with circuit depth [3, 5]. The Experiment 4 observation that leakage remains low under repeated application of $R_{\mathcal{N}}$ has its ideal reference in the unitarity of $R_{\mathcal{N}}$ established in Sec. IID (benchmark B6 of Sec. III); we do not claim to have observed sub-unity logical error rates in the multi-qubit sense.

Operating conditions. The cloud accessibility of the processor, including the absence of any user-side cryo-

genic requirement, is a property of the photonic platform and is not unique to the present construction; the relevant point is that the single-photon spatial-mode protection tested here does not require custom hardware on the user side.

Relation to bosonic and loss-tolerant photonic codes. The construction here protects a spatial-mode amplitude channel and is conditional on single-photon detection; it does not address photon loss, which is the dominant error of photonic hardware and is removed here by post-selection (Sec. IV). Complementary photonic strategies target exactly this channel: bosonic codes such as the Gottesman–Kitaev–Preskill encoding [28] protect against loss and small displacements within a single mode, and loss-tolerant photonic encodings [29] are designed to detect and recover lost photons. These operate on a different resource and error axis from the single-photon spatial-mode parity checks demonstrated here; a register combining spatial-mode parity filtering with loss-tolerant or bosonic protection is a natural direction beyond the present single-block scope.

VII. DISCUSSION

A. What the hardware demonstrates

The hardware experiments of Sec. V test the single-block ingredients of the construction rather than a full error-correcting architecture. Their common feature is the Q_3 layout: the DC/sum check, the soft DC error signal (Experiment 3), the archived neutral-sector core, and the three face-parity checks are all defined on the same eight-vertex register. This shared geometry is the main structural distinction from a generic path-encoded qudit and is the feature most likely to connect this work to adjacent high-dimensional and graph-based photonic encodings. We emphasize, however, that the present experiments do not isolate a hardware advantage attributable specifically to the cube *labeling*: the same circuits could be described with the modes relabeled 0–7, and the geometry functions as an organizing principle (tying the DC check, the three face-parity syndrome channels, and the Gray-code gate ordering to one combinatorial object), not as an independently benchmarked physical resource. What the hardware shows is that this organized set of nested checks can be prepared, applied, and read out together on a single commercial chip. Equivalently, the leakage- and syndrome-suppression ratios are a measurement of how faithfully the chip implements the intended subspace-projecting unitaries, for which the suppression is exact by construction in the ideal limit.

B. Limitations

The present work demonstrates single-block primitives of protected spatial-mode subspaces, syndrome channels,

and a coherence-preserving unitary core. It does not demonstrate (i) closed-loop error recovery using the syndrome readout, (ii) feed-forward classical control, (iii) a multi-block tiling architecture, or (iv) a fault-tolerance threshold.

Closed-loop recovery in particular requires measuring the complex syndrome rather than only intensity in the syndrome ports. Standard photonic routes are homodyne or heterodyne mode-overlap detection against calibrated local oscillators, or an interferometric phase-reference measurement that recovers both magnitude and phase of the residual. Our current Belenos experiments measure output intensities and probabilities only, so they test syndrome-channel selectivity and leakage suppression, but not full feed-forward recovery. Implementing complex-syndrome readout and a feed-forward correction loop on this register is the natural next step.

VIII. CONCLUSION

We have characterized, with over 340,000 postselected single-photon detection events on a commercial cloud-accessible photonic processor, how faithfully nested Walsh parity-check filters are implemented in a single-photon spatial-mode register: percent-level leakage floors (0.02%–1.1% depending on input) for the zero-sum neutral subspace, a monotonic and calibratable soft response to injected DC contamination, resolved face-parity syndrome routing with a measured confusion matrix, and leakage that remains far below non-neutral controls at the three compiled depths of a sector-preserving unitary core. The underlying $[8, 7, 2]$ and $[8, 4, 4]$ codes are classical, the suppressions are exact by construction in the ideal limit, and the measured floors are set by identifiable implementation systematics: a fixed-pattern separator calibration bias, ± 0.02 per-point offsets in the soft-error response, and per-compilation scatter that dominates cross-depth comparisons at the 10^{-3} leakage level. All figures are conditional on postselected single-photon detection; photon loss is not detected by these checks.

Three interpretations are stated in this paper as open hypotheses rather than established conclusions, because the measurements that would decide them have not yet been performed: the calibration origin of the $b_0 \oplus b_2$ syndrome outlier (testable by a repeat acquisition after recalibration against the Walsh-spectrum model of Eq. (6)), the coherence dependence of the suppression (testable by a deliberate visibility scan), and the per-cycle leakage behavior of the unitary core (testable by

a ~ 10 – 20 cycle scan with frozen compilation in a single calibration window). These, together with a randomized-construction control isolating the role of the cube organization, constitute the validation program for the interpretive claims; extension to higher d -cube registers and complex-syndrome readout with feed-forward control remain future work beyond it.

ACKNOWLEDGMENTS

Experiments were performed on Quandela’s cloud-accessible Belenos processor through the Perceval software framework. The dynamics core is defined by the archived compiled circuits and the extraction procedure reported in Sec. IID. The experimental claims in this manuscript rest on the explicit subspace definitions, circuit primitives, data, and simulation benchmarks reported here. This research received no external funding.

Author contributions. E. T.: conceptualization, data curation, formal analysis (lead), investigation, methodology, software, validation, visualization, writing – original draft, and writing – review and editing. J. W.: investigation, methodology, software, and supervision. M. S.: formal analysis, project administration, validation, visualization, and writing – review and editing.

The authors declare no competing financial interests.

DATA AVAILABILITY

The complete reproducibility package is included as arXiv ancillary files with this manuscript. It is organized as an archival package with a README, manifest files mapping table rows to job exports, raw Quandela/Perceval payload and result JSON files in a dedicated raw-data directory, processed count tables in a separate generated-output directory, backend metadata, and Python analysis scripts for decoding the raw exports, statistical estimation, operator algebra, simulation benchmarks, validation, and figure/table generation. In particular, the script `verify_package.py` recomputes every tabulated dump and syndrome probability and the headline suppression ratios directly from the archived per-port counts, resolving each manuscript row to its job through the row-to-job manifest, and all stochastic benchmarks use the fixed random seed 20260404. The package contains no absolute local paths, API tokens, credentials, or live internal URLs, and it does not require live Quandela access to reproduce the reported tables and figures from the archived raw counts.

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary ed. (Cambridge University Press, Cambridge, 2010), doi:10.1017/CBO9780511976667.

[2] D. Gottesman, “Stabilizer codes and quantum error correction,” Ph.D. thesis, California Institute of Technology (1997), arXiv:quant-ph/9705052.

- [3] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A* **86**, 032324 (2012), doi:10.1103/PhysRevA.86.032324.
- [4] D. Litinski, “A game of surface codes: Large-scale quantum computing with lattice surgery,” *Quantum* **3**, 128 (2019), doi:10.22331/q-2019-03-05-128.
- [5] R. Acharya *et al.* (Google Quantum AI), “Quantum error correction below the surface code threshold,” *Nature* **638**, 920–926 (2025), doi:10.1038/s41586-024-08449-y.
- [6] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, “High-threshold and low-overhead fault-tolerant quantum memory,” *Nature* **627**, 778–782 (2024), doi:10.1038/s41586-024-07107-7.
- [7] L. Vaidman, L. Goldenberg, and S. Wiesner, “Error prevention scheme with four particles,” *Phys. Rev. A* **54**, R1745 (1996), doi:10.1103/PhysRevA.54.R1745.
- [8] E. Knill, R. Laflamme, and G. J. Milburn, “A scheme for efficient quantum computation with linear optics,” *Nature* **409**, 46–52 (2001), doi:10.1038/35051009.
- [9] J. L. O’Brien, A. Furusawa, and J. Vučković, “Photonic quantum technologies,” *Nat. Photonics* **3**, 687–695 (2009), doi:10.1038/nphoton.2009.229.
- [10] F. Flamini, N. Spagnolo, and F. Sciarrino, “Photonic quantum information processing: a review,” *Rep. Prog. Phys.* **82**, 016001 (2019), doi:10.1088/1361-6633/aad5b2.
- [11] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes* (North-Holland, Amsterdam, 1977).
- [12] I. S. Reed, “A class of multiple-error-correcting codes and the decoding scheme,” *Trans. IRE Prof. Group Inf. Theory* **4**, 38–49 (1954), doi:10.1109/TIT.1954.1057465.
- [13] D. E. Muller, “Application of Boolean algebra to switching circuit design and to error detection,” *Trans. IRE Prof. Group Electron. Comput.* **EC-3**, 6–12 (1954), doi:10.1109/IREPGELC.1954.6499441.
- [14] M. Erhard, M. Krenn, and A. Zeilinger, “Advances in high-dimensional quantum entanglement,” *Nat. Rev. Phys.* **2**, 365–381 (2020), doi:10.1038/s42254-020-0193-5.
- [15] B. A. Bell, D. A. Herrera-Martí, M. S. Tame, D. Markham, W. J. Wadsworth, and J. G. Rarity, “Experimental demonstration of a graph state quantum error-correction code,” *Nat. Commun.* **5**, 3658 (2014), doi:10.1038/ncomms4658.
- [16] H. Zhang, L. Wan, S. Paesani, A. Laing, Y. Shi, H. Cai, X. Luo, G.-Q. Lo, L. C. Kwek, and A. Q. Liu, “Encoding error correction in an integrated photonic chip,” *PRX Quantum* **4**, 030340 (2023), doi:10.1103/PRXQuantum.4.030340.
- [17] J. Washburn, M. Zlatanović, and E. Allahyarov, “Recognition Geometry,” *Axioms* **15**(2), 90 (2026), doi:10.3390/axioms15020090.
- [18] P. Zanardi and M. Rasetti, “Noiseless quantum codes,” *Phys. Rev. Lett.* **79**, 3306–3309 (1997), doi:10.1103/PhysRevLett.79.3306.
- [19] D. A. Lidar, I. L. Chuang, and K. B. Whaley, “Decoherence-free subspaces for quantum computation,” *Phys. Rev. Lett.* **81**, 2594–2597 (1998), doi:10.1103/PhysRevLett.81.2594.
- [20] P. G. Kwiat, A. J. Berglund, J. B. Altepeter, and A. G. White, “Experimental verification of decoherence-free subspaces,” *Science* **290**, 498–501 (2000), doi:10.1126/science.290.5491.498.
- [21] Quandela, “Belenos (MosaiQ-12): a cloud-accessible photonic quantum processor, 12 photons / 24 modes,” product information, <https://www.quandela.com/products-and-services/belenos/> (accessed June 2026).
- [22] N. Heurtel *et al.*, “Perceval: A software platform for discrete variable photonic quantum computing,” *Quantum* **7**, 931 (2023), doi:10.22331/q-2023-02-21-931.
- [23] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, “Experimental realization of any discrete unitary operator,” *Phys. Rev. Lett.* **73**, 58–61 (1994), doi:10.1103/PhysRevLett.73.58.
- [24] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, “Optimal design for universal multiport interferometers,” *Optica* **3**, 1460–1465 (2016), doi:10.1364/OPTICA.3.001460.
- [25] J. Carolan *et al.*, “Universal linear optics,” *Science* **349**, 711–716 (2015), doi:10.1126/science.aab3642.
- [26] C. K. Hong, Z. Y. Ou, and L. Mandel, “Measurement of subpicosecond time intervals between two photons by interference,” *Phys. Rev. Lett.* **59**, 2044–2046 (1987), doi:10.1103/PhysRevLett.59.2044.
- [27] E. B. Wilson, “Probable inference, the law of succession, and statistical inference,” *J. Am. Stat. Assoc.* **22**, 209–212 (1927), doi:10.1080/01621459.1927.10502953.
- [28] D. Gottesman, A. Kitaev, and J. Preskill, “Encoding a qubit in an oscillator,” *Phys. Rev. A* **64**, 012310 (2001), doi:10.1103/PhysRevA.64.012310.
- [29] M. Varnava, D. E. Browne, and T. Rudolph, “Loss tolerance in one-way quantum computation via counterfactual error correction,” *Phys. Rev. Lett.* **97**, 120501 (2006), doi:10.1103/PhysRevLett.97.120501.
- [30] The lost-twin survivor fraction is estimated as $[2(1 - \eta)/\eta] (n_2/N)$, where $\eta = 4.84\%$ is the transmittance, $n_2 = 171$ is the number of detected two-photon events, and $N = 350,000$ is the total number of detections; this gives $39 \times (171/350,000) \approx 1.9\%$.