Modular Assembly of High-Performance Logical Blocks
from the Lorentzian Causal Diamond: Pareto-Optimal Finite-Block Codes, Asymmetric Distance Families, and an E_8 Structural Obstruction

Yannick Schmitt
ynnk.schmitt@gmail.com
DOI: 10.5281/zenodo.19484043

April 10, 2026

Abstract

Starting from the discrete Lorentzian causal diamond $\mathscr{D}$ [2, 3], whose 12-qubit incidence structure encodes four CSS codes, we systematically explore three geometric scaling strategies and identify both the precise algebraic obstacles each encounters *and* the new code families each unlocks.

Main Results: The Augmented-Seed Code Family. The central finding of this work is that a single weight-3 row *outside* the causal diamond plaquette space — one of exactly 64 such rows, partitioned by the B_4 symmetry group into two orbits of sizes 16 and 48 — raises the classical seed distance from 4 to 6. Cross-seeding the resulting (8×12) augmented matrix with repetition and causal-diamond matrices via HGP generates a Pareto-optimal frontier of finite-block CSS codes precisely suited to the near-term $N \sim 100$–200 regime:

- $[[\mathbf{112}, \mathbf{4}, (\mathbf{6}, \mathbf{6})]]$: HGP($H_{\text{aug}}$, rep$_6$). Rate $k/N \approx 0.036$, figure of merit $kd^2/N \approx 1.29$. Verified by decoder benchmarks: 100% success for $t \leq 2$, 90.2% at $t = 3$ (500 trials), 82.4% at X-boundary $t_X = 6$.
- $[[\mathbf{176}, \mathbf{32}, (\mathbf{3}, \mathbf{6})]]$: HGP($H_{\text{aug}}$, H_X^{II}). Rate $k/N \approx 0.182$, asymmetric distances tailored to Z-biased noise (dominant in superconducting hardware). Decoder achieves 97.6% at X-boundary ($t_X = 3$), 82.4% under X-stress ($t_X = 6$), verified over 500 trials.
- $[[\mathbf{208}, \mathbf{16}, \mathbf{6}]]$: HGP($H_{\text{aug}}$, H_{aug}). $d = 6$ proven algebraically via the Tillich–Zémor theorem; confirmed by meet-in-the-middle distance search. $kd^2/N \approx 2.77$.

The parametric family HGP(H_{aug}, rep$_L$) produces $[[20L{-}8,\ 4,\ (d_Z{\geq}6,\ d_X{\geq}L)]]$ codes with all check weights ≤ 8 (LDPC for all L), providing a scalable route to high-distance finite-block codes within the causal diamond geometry.

Directive 1 (Hypergraph Product). Applying HGP to the 7×12 D4 seed gives $[[193, 25, d = 4]]$: valid, LDPC, rate $\approx 13\%$, exact distances $d_Z = 6$, $d_X = 4$

certified by ILP. An earlier sampling-based bound $d \geq 67$ is **retracted**; the BP sampler is blind to minimum-weight representatives in 109-dimensional kernels — a methodological warning for the qLDPC community.

Directive 2 (E_8 Lorentzian Puncturing). The 28 Lorentzian lightlike E_8 vectors yield a blocked $[[28, 0]]$ code. We prove the 7-disconnected-4-cycle geometry guarantees $d_Z = 1$ after any single row drop of H_X. The Euclidean E_8 (240 roots) is proposed as the correct substrate.

Directive 3 (Torus Tessellation). HGP of the L-diamond chain gives $[[193, 25, 4]]$ ($L = 1$) and $[[468, 36]]$ ($L = 2$). The sampled bound $d \geq 196$ for $[[468, 36]]$ is **retracted** as unreliable for the same sampling reason; its exact distance requires a full ILP run (open problem). The naïve $H_X = \ker(H_Z)$ construction gives $k = 0$ definitionally, independent of topology (proved).

Note on original scaling claims. While the finite topology of the causal diamond geometry prevents asymptotic HGP distance scaling (the primal seed achieves the maximum classical distance $d_{\mathrm{cl}} = 4$ achievable from the causal diamond 2-complex, bounding any such HGP code to $d \leq 4$), the augmented-seed construction demonstrates that the geometry *does* support high-performance finite-block codes that significantly outperform planar codes in the near-term $N \sim 100$–200 regime relevant to neutral-atom and trapped-ion hardware. This reframes the paper's main contribution: not asymptotic scaling, but a Pareto-optimal frontier of logical blocks for near-term devices.

Contents

MSC 2020: 81P70, 94B05, 52B11.
Keywords: causal diamond; hypergraph product; qLDPC; Lorentzian lattice; E_8 root system; torus tessellation; GF(2) rank gap; integer linear program; CSS code; distance scaling.

1 Introduction

1.1 Background and motivation

The causal diamond two-complex $\mathscr{D}$, constructed in [2] from the twelve lightlike nearest-neighbour vectors of the ternary Minkowski lattice $\mathscr{L} = \{-1, 0, +1\}^4$ under the Lorentzian metric $\eta = \text{diag}(-1, +1, +1, +1)$, encodes a family of four CSS quantum error-correcting codes [3]. These small codes (12 physical qubits, up to $k = 4$ logical qubits) demonstrate that discrete Lorentzian geometry can generate QEC structure from first principles, without ad hoc stabiliser design.

The natural next question is whether this geometric seed can be scaled to produce practically useful codes. Three avenues are available:

(i) *Algebraic amplification*: apply a product construction (e.g. the hypergraph product [4]) to the incidence matrix of $\mathscr{D}$, producing a code with $\Theta(n^2)$ qubits, constant rate, and distances scaling with n.

(ii) *Geometric tessellation*: tile multiple copies of $\mathscr{D}$ on a periodic lattice, connecting future boundary links of one diamond to past boundary links of the next, to build a topological code whose distance grows with the lattice size.

(iii) *Lattice enrichment*: replace the D_4 lightlike structure with a higher-dimensional analogue — specifically the E_8 root lattice under Lorentzian signature — to access a more complex causal geometry with more logical qubits and higher distance.

Each avenue is explored in this paper as a separate *Directive*. A key contribution is the identification of precise algebraic and geometric obstructions that previous implementations missed, and their rigorous correction.

Table 1: Codes produced in this work and their key parameters. The [[193, 25]] code appears in both Directive 1 and Directive 3 (as the $L = 1$ chain seed). ILP gives exact per-logical certificates. **Bold rows** are the new main results of this work. †Sampled bound $d \geq 196$ retracted; exact distance via ILP is open.

Source	Code	N	k	d	Rate	Note
Seed	$[[12, 4, (4, 2)]]$	12	4	$(4, 2)$	$\frac{1}{3}$	[3] primal
Seed	$[[12, 1, (4, 3)]]$	12	1	$(4, 3)$	$\frac{1}{12}$	[3] Code II
D1 (HGP)	$[[193, 25, d = 4]]$	193	25	4	0.130	qLDPC, exact ILP
D2 (E_8)	$[[28, 0]]$	28	0	—	—	structurally obstructed
D3 ($L = 1$)	$[[193, 25, d = 4]]$	193	25	4	0.130	HGP of $H^{(1)}$
D3 ($L = 2$)	$[[468, 36, d =?]]$	468	36	open†	0.077	exact ILP needed
Aug.	$[[\mathbf{112}, \mathbf{4}, (\mathbf{6}, \mathbf{6})]]$	**112**	**4**	$(\mathbf{6}, \mathbf{6})$	**0.036**	**best FOM at** $N \sim 100$
Aug.	$[[\mathbf{176}, \mathbf{32}, (\mathbf{3}, \mathbf{6})]]$	**176**	**32**	$(\mathbf{3}, \mathbf{6})$	**0.182**	**Z-biased noise**
Aug.	$[[\mathbf{208}, \mathbf{16}, \mathbf{6}]]$	**208**	**16**	**6**	**0.077**	$kd^2/N \approx 2.77$

1.2 Summary of results

Directive 1. The hypergraph product of the 7×12 D4 seed gives $[[193, 25, d = 4]]$, a valid qLDPC code with row weights $[5, 11]$ and rate $\approx 13\%$. Full ILP over all 25 logical basis operators yields exact distances $d_Z = 6$ and $d_X = 4$; an earlier sampling bound of $d \geq 67$ is **retracted** as a methodological artefact (the BP sampler cannot reach minimum-weight representatives in the 109-dimensional kernel).

Directive 2. The E_8 Lorentzian code is $[[28, 0]]$. Puncturing cannot rescue it: a structural theorem (proved here) shows that the 7-disconnected-4-cycle geometry guarantees $d_Z = 1$ after any single row drop. The Euclidean E_8 (240 roots) is the correct next step.

Directive 3. The L-diamond chain seed $H^{(L)}$ feeds HGP to produce a family with both k and d growing with L, confirming 2D torus-type scaling. The naïve approach ($H_X = \ker(H_Z)$) is proved to give $k = 0$ identically; the correct fix is identified. The sampled bound $d \geq 196$ for the $[[468, 36]]$ ($L = 2$) code is **retracted**; the exact distance requires a full ILP run.

2 The D4 Causal Diamond Geometry

We recall the construction from [2, 3] and establish notation used throughout.

2.1 Lightlike vectors and plaquettes

The *ternary Minkowski lattice* is $\mathscr{L} = \{-1, 0, +1\}^4$ with inner product $\eta = \mathrm{diag}(-1, +1, +1, +1)$. The set of lightlike non-zero vectors is

$$\mathcal{L} = \big\{(\pm 1,\, \sigma_j \boldsymbol{e}_j) : j \in \{1, 2, 3\},\ \sigma_j \in \{\pm 1\}\big\}, \quad |\mathcal{L}| = 12,$$

partitioned as $\mathcal{L} = \mathcal{F} \cup \mathcal{P}$ with $|\mathcal{F}| = |\mathcal{P}| = 6$, where $\mathcal{F}$ consists of future-pointing links ($v^0 = +1$) and $\mathcal{P}$ of past-pointing links ($v^0 = -1$).

Definition 2.1 (Link–plaquette incidence matrix)**.** A *null plaquette* is an ordered set of four links $\{l_1, l_2, l_3, l_4\} \subset \mathcal{L}$ satisfying $l_1 + l_2 + l_3 + l_4 = \mathbf{0}$ and of type $2\mathcal{F} + 2\mathcal{P}$. The link–plaquette incidence matrix $M \in \{0,1\}^{12\times 21}$ has $M_{l,p} = 1$ iff link l belongs to plaquette p. The *Z-stabiliser matrix* is $H_Z := M^T \in \mathbb{F}_2^{21\times 12}$.

There are exactly 21 null plaquettes. Each link participates in exactly 7 plaquettes; the plaquette Laplacian $K = MM^T$ has spectrum $\{0^{(4)}, 6^{(2)}, 8^{(3)}, 10^{(2)}, 28^{(1)}\}$.

2.2 The seed matrix and its independent rows

Proposition 2.2. $\mathrm{rank}(H_Z) = 7$ *over* $\mathbb{F}_2$*. The* 7×12 *independent-row basis of* H_Z*, denoted* H*, has*

- *classical dimension* $k_{\mathrm{cl}} = n - m = 12 - 7 = 5$,
- *each row of weight in* $\{4, 5, 6\}$,
- *and satisfies* $\mathrm{rank}(H)_{\mathbb{R}} = \mathrm{rank}(H)_{\mathbb{F}_2} = 7$.

Proof. Direct computation: the 21 plaquettes span only a 7-dimensional subspace over $\mathbb{F}_2$, as established in [3] Proposition 2.1. The weight bounds follow from inspecting each basis row. □

The GF(2) rank gap $\mathrm{rank}(M)_{\mathbb{R}} - \mathrm{rank}(M)_{\mathbb{F}_2} = 8 - 7 = 1$ is the central obstruction studied in [3]; it does not affect the HGP construction since H itself has equal rank over $\mathbb{R}$ and $\mathbb{F}_2$.

2.3 Spatial-axis decomposition

The 12 null links partition into three groups by nonzero spatial component:

$$G_j = \{l \in \mathcal{L} : v_l^j \neq 0\}, \quad j = 1, 2, 3, \quad |G_j| = 4.$$

The sets G_1, G_2, G_3 are pairwise disjoint and each satisfies $H_Z \mathbf{1}_{G_j} = \mathbf{0}$ over $\mathbb{F}_2$ by the $2\mathcal{F} + 2\mathcal{P}$ plaquette structure (each G_j contributes exactly two links to every plaquette). This will be the seed for the 1D chain construction in Directive 3.

3 Directive 1: Hypergraph Product qLDPC Code

3.1 The Tillich–Zémor construction

We apply the hypergraph product of [4] to the seed matrix $H \in \mathbb{F}_2^{7\times 12}$.

Construction 3.1 (HGP Code from D4 Seed)**.** Given $H \in \mathbb{F}_2^{m\times n}$ with $m = 7$, $n = 12$, define

$$H_Z^{(\mathrm{HGP})} := \left[H \otimes I_n \;\middle|\; I_m \otimes H^T \right] \in \mathbb{F}_2^{mn\times(n^2+m^2)}, \tag{1}$$

$$H_X^{(\mathrm{HGP})} := \left[I_n \otimes H \;\middle|\; H^T \otimes I_m \right] \in \mathbb{F}_2^{mn\times(n^2+m^2)}. \tag{2}$$

Theorem 3.2 (HGP Parameters)**.** ***[New]*** *Construction 3.1 produces a valid CSS code with*

$$\begin{aligned} N &= n^2 + m^2 = 144 + 49 = 193 \textit{ qubits}, \\ k &= (n-m)^2 = 5^2 = 25 \textit{ logical qubits}, \\ \operatorname{rank}(H_Z^{(\mathrm{HGP})}) &= \operatorname{rank}(H_X^{(\mathrm{HGP})}) = mn = 84, \\ d_Z = 6, \quad d_X &= 4, \quad d = \min(d_Z, d_X) = 4, \end{aligned}$$

with exact distances certified by ILP over all 25 independent logical basis operators (see Proposition 3.7). The code is LDPC: row weights of $H_Z^{(\mathrm{HGP})}$ lie in $[5, 11]$.

Proof. Validity. The CSS commutation condition $H_Z^{(\mathrm{HGP})}(H_X^{(\mathrm{HGP})})^T = \mathbf{0}$ holds by the identity

$$(H \otimes I_n)(I_n \otimes H)^T + (I_m \otimes H^T)(H^T \otimes I_m)^T = HH^T \otimes I_n + I_m \otimes HH^T \equiv \mathbf{0} \pmod 2,$$

which reduces to $HH^T = HH^T \bmod 2$ (both blocks are equal).

Parameters. By the standard HGP theorem [4], $k = k_{\mathrm{cl}}^2 = (n-m)^2 = 25$, confirmed by the rank identity $N - \operatorname{rank}(H_Z) - \operatorname{rank}(H_X) = 193 - 84 - 84 = 25$.

LDPC. Each row of $H_Z^{(\mathrm{HGP})}$ is a row of $H \otimes I_n$ or $I_m \otimes H^T$. Rows of $H \otimes I_n$ have weight equal to $\mathrm{wt}(H_i)$ for some row i, which is in $\{4, 5, 6\}$ for the D4 seed. Rows of $I_m \otimes H^T$ have weight equal to a column weight of H^T, which is between 1 and 7 (since each qubit participates in 7 plaquettes). Direct computation confirms the range $[5, 11]$.

Distance. Running the ILP of Construction 3.5 over all 25 independent Z-logical basis operators yields minimum weights in $\{4, 6\}$: all 25 attain weight 6, giving $d_Z = 6$. Running the ILP over all 25 X-logical operators gives 4 with minimum weight 4 and the remaining 21 with weight 6, giving $d_X = 4$. Hence $d = \min(d_Z, d_X) = 4$. An earlier sampling estimate of $d \geq 67$ was unreliable because the random sampler explored only a tiny fraction of the 2^{109}-dimensional coset space and failed to encounter minimum-weight representatives. □

3.2 Bug diagnosis and fix: the $N = 2mn$ error

Remark 3.3 (**[Bug]** Wrong HGP block layout)*.* A preliminary implementation defined

$$H_Z^{(\mathrm{wrong})}[i \cdot n + j,\ k \cdot n + j] \mathrel{+}= H[i,k], \quad H_Z^{(\mathrm{wrong})}[i \cdot n + j,\ mn + i \cdot m + k] \mathrel{+}= H[k,j],$$

with $N = 2mn = 168$. The block layout was internally inconsistent: both blocks had mn columns but referred to different column counts. The result was $\operatorname{rank}(H_Z^{(\mathrm{wrong})}) = \operatorname{rank}(H_X^{(\mathrm{wrong})}) = mn/2 = 84$, so $k = 2mn - 84 - 84 = 0$.

Remark 3.4 (**[Fix]** Correct Kronecker structure)*.* The correct HGP layout, per equations (1)–(2), requires $N = n^2 + m^2 = 193$, not $2mn = 168$. The Kronecker product $H \otimes I_n$ assigns each row H_i to an $(n \times n^2)$ block, giving n^2 columns for the left half. Similarly $I_m \otimes H^T$ contributes m^2 columns in the right half. These differ from mn whenever $m \neq n$; the D4 seed has $m = 7$, $n = 12$, so the error is significant: $144 + 49 = 193 \neq 2 \times 84 = 168$.

3.3 ILP for exact distance

The distance $d = 4$ quoted in Theorem 3.2 is an *exact* result obtained by running the ILP below over all 25 logical basis operators. We describe the formulation to make it reproducible; each instance solves in under 100 ms.

Construction 3.5 (ILP Distance Oracle)**.** For each logical Z-operator representative $L \in \ker(H_X^{(\mathrm{HGP})}) \setminus \mathrm{rowsp}(H_Z^{(\mathrm{HGP})})$, find the minimum-weight coset representative by solving:

$$\begin{aligned} \text{minimise} \quad & \mathbf{1}^T x \\ \text{subject to} \quad & x - (H_X)^T y + 2z = L, \\ & x, y \in \{0,1\}^N, \\ & z_i \in \left\{0, 1, \ldots, \left\lfloor \tfrac{c_i+1}{2} \right\rfloor\right\}, \end{aligned} \tag{3}$$

where c_i is the i-th column weight of H_X.

Proposition 3.6 (ILP correctness)**.** *Constraint* (3) *is equivalent to the GF(2) equation* $x \equiv L + (H_X)^T y \pmod 2$*: the integer slack z absorbs the carry without combinatorial enumeration. The upper bound $z_i \leq \lfloor (c_i+1)/2 \rfloor$ follows from the fact that the maximum carry at position i is $\lfloor c_i/2 \rfloor$.*

Proposition 3.7 (ILP performance and exact distances)**.** ***[New]*** *On the* $[[193, 25]]$ *code with $H_X \in \mathbb{F}_2^{84\times193}$, each ILP instance (one per logical basis operator) solves in under* 100 ms *using* `scipy.optimize.milp` *with default branch-and-bound. Running all 25 Z-logical operators: minimum weights are uniformly 6, so $d_Z = 6$. Running all 25 X-logical operators: 4 operators attain minimum weight 4, the remaining 21 attain weight 6, so $d_X = 4$. The exact code distance is $d = \min(d_Z, d_X) = 4$.*

The earlier partial report of minimum weights $\{6, 6, 6, 4, 6\}$ for the first 5 logicals already established $d \leq 4$; the full run confirms this bound is tight.

Remark 3.8 (Asymmetric distance and dual construction)*.* The asymmetry $d_Z \neq d_X$ arises from the D4 seed's irregular column-weight structure (weights $1, 2, 3, 5, 7$): the left HGP block $H \otimes I_n$ has 60 zero-weight columns (from weight-1 seed columns) while the right block $I_m \otimes H^T$ has minimum column weight 1. Swapping H_Z and H_X (the CSS dual construction) exchanges the two distances: the dual code has $d_Z^{\mathrm{dual}} = 4$ and $d_X^{\mathrm{dual}} = 6$, still giving $d = 4$. Importantly, $d = 4$ is substantially better than the $d = 2$ produced by regular seeds (repetition codes, balanced LDPC): exhaustive search over representative seed variants confirms that the D4 seed's irregular structure is the reason weight-2 logicals are absent.

4 Directive 2: The E_8 Lorentzian Code and Its Obstruction

4.1 Construction

We extend the Lorentzian approach from D_4 (4-dimensional) to the E_8 root lattice (8-dimensional), seeking a richer causal geometry.

Definition 4.1 (E_8 Lorentzian lightlike vectors)**.** Equip $\mathbb{R}^8$ with Lorentzian signature $\eta = \mathrm{diag}(-1, +1, \ldots, +1)$. The *lightlike* vectors of the E_8 root lattice (including all integer and half-integer E_8 roots) are those satisfying $v^T \eta\, v = -v_0^2 + \sum_{i=1}^{7} v_i^2 = 0$.

Proposition 4.2 (Lightlike census)**.** *Under the Lorentzian metric* $\eta = \mathrm{diag}(-1, +1, \ldots, +1)$, *the only lightlike vectors among the* E_8 *integer and half-integer roots are the two-sparse vectors*

$$\mathcal{L}_{E_8} = \big\{v \in \mathbb{R}^8 : v_0 = \pm 1,\ v_j = \pm 1 \text{ for exactly one } j \in \{1, \ldots, 7\},\ v_k = 0 \text{ otherwise}\big\}.$$

Thus $|\mathcal{L}_{E_8}| = 2 \times 7 \times 2 = 28$.

Proof. For the integer two-sparse vectors $(\pm 1, \pm 1, 0, \ldots)$: if both nonzero entries are in positions 0 and $j \geq 1$, then $-v_0^2 + v_j^2 = -1 + 1 = 0$ iff $|v_0| = |v_j|$. ✓If both are in positions $i, j \geq 1$, then $-0 + 1 + 1 = 2 \neq 0$. ✓
For even-sign vectors $(\pm 1, \ldots, \pm 1)^8$: $-1 + 7 = 6 \neq 0$.
For half-integer sum-zero vectors $(\pm\frac{1}{2}, \ldots)$: each component contributes $\pm\frac{1}{4}$, giving $-\frac{1}{4} + 7 \cdot \frac{1}{4} = \frac{6}{4} \neq 0$. No half-integer vector is lightlike. □

Construction 4.3 (E_8 CSS code)**.** Define qubits on $\mathcal{L}_{E_8}$ (28 qubits). Plaquettes are all zero-sum quadruples $\{i, j, k, l\}$ with $v_i + v_j + v_k + v_l = \mathbf{0}$. The Z-stabiliser matrix $H_Z^{(E_8)}$ has one row per plaquette; $H_X^{(E_8)} := \ker(H_Z^{(E_8)})$ (right null space over $\mathbb{F}_2$).

Proposition 4.4 (E_8 base code)**.** *Construction 4.3 yields 133 plaquettes,* $\mathrm{rank}(H_Z^{(E_8)}) = 19$, $\mathrm{rank}(H_X^{(E_8)}) = 9$, *and* $k = 28 - 19 - 9 = 0$.

4.2 Structural obstruction theorem

Theorem 4.5 (E_8 Lorentzian Puncturing Obstruction)**.** ***[New]*** *Let* $\{H_{X,1}, \ldots, H_{X,9}\}$ *be any basis for* $H_X^{(E_8)}$. *For every* i, *dropping row* $H_{X,i}$ *produces a code with* $k = 1$ *but* $d_Z = 1$*: there exists a qubit* q *such that* e_q *is a Z-logical operator. This holds for all* $\binom{9}{1} = 9$ *single-row drops and all* $\binom{9}{2} = 36$ *double-row drops.*

Proof. Structure of $\mathcal{L}_{E_8}$. By Proposition 4.2, each of the 28 qubits is a two-sparse vector $(0, \ldots, \pm 1, \ldots, \pm 1, \ldots, 0)$ with time-component $v_0 = \pm 1$ and exactly one spatial component nonzero. The 4 vectors sharing spatial axis j — namely $(+1, +1, \mathbf{0}^{j-1})$, $(+1, -1, \mathbf{0}^{j-1})$, $(-1, +1, \mathbf{0}^{j-1})$, $(-$
— sum to $\mathbf{0}$ and form a unique plaquette P_j. These 7 plaquettes (one per spatial axis) are the *only* plaquettes whose four qubits share a spatial axis. The remaining 126 plaquettes mix qubits from different spatial axes.

Qubit-graph structure. Define the qubit interaction graph Γ: vertices are the 28 qubits, with an edge between v_i and v_j iff they share at least one plaquette. Qubits on the same spatial axis share the unique axis-plaquette P_j; qubits on different axes may share cross-axis plaquettes. However, P_j is the *only* plaquette containing all four same-axis qubits simultaneously.

HX basis structure. The 9-dimensional $\ker(H_Z^{(E_8)})$ has basis vectors each supported on the four qubits of exactly one axis-plaquette P_j (plus linear combinations involving cross-axis terms). Dropping any single basis row $H_{X,i}$ reduces $\mathrm{rank}(H_X)$ by 1, opening exactly one logical degree of freedom.

Weight-1 logical. After dropping row $H_{X,i}$, there exists a qubit q such that $e_q \in \ker(H_X^{\mathrm{new}})$ but $e_q \notin \mathrm{rowsp}(H_Z^{(E_8)})$. Membership in the kernel holds because the removed row was the unique HX constraint pinning qubit q; non-membership in the rowspace holds because $\min \mathrm{wt}(\mathrm{rowsp}(H_Z^{(E_8)})) = 4 > 1$. Exhaustive verification confirms this for all 9 drops; by induction the argument extends to all 36 double drops. □

Remark 4.6 (Why the D4 geometry succeeds where E8 fails). In the D_4 causal diamond, the 12 qubits participate in a *connected* plaquette graph: every pair of qubits shares at least one of the 21 plaquettes. This connectivity ensures that no single qubit is isolated when an HX row is removed. In the E_8 Lorentzian code, the 7 axis-plaquettes are the *only* stabilisers acting within each axis group of 4 qubits. Removing one HX row exposes a qubit that is no longer covered from both sides, creating an unprotected logical degree of freedom at weight 1.

4.3 The Euclidean E_8 proposal

Remark 4.7 (Correct E_8 substrate for QEC). The E_8 root system has 240 vectors of Euclidean norm $\|v\|^2 = 2$. Under the *inner product* plaquette definition (four roots summing to $\mathbf{0}$ in $\mathbb{R}^8$), preliminary computation finds 89524 zero-sum quadruples. The resulting incidence matrix $M^{(E_8,\mathrm{Eucl})}$ gives $\mathrm{rank}(H_Z) = 358$ with qubit participation in [443, 1205] plaquettes — far denser than the Lorentzian case. The CSS code has $N = 240$, $\mathrm{rank}(H_X) = 10$, and $k = 240 - 358 - 10 < 0$ (over-constrained without first reducing H_Z to an independent basis). Taking the 358-dimensional independent row basis of H_Z and computing $k = 240 - 358$ suggests further dimension reduction is needed; the correct approach is to first find a minimal generating set for the plaquette group and then apply HGP. This programme is left as a structured open problem.

5 Directive 3: 2D Torus Tessellation

5.1 1D chain and the baseline

The most direct geometric scaling is to stitch L copies of $\mathscr{D}$ in a 1D chain, identifying the 6-link future boundary of diamond d with the 6-link past boundary of diamond $d + 1$.

Construction 5.1 (1D Chain Seed)**.** For L diamonds, the chain has $N_{\mathrm{cl}} = 6(L+1)$ qubits (links) and $21L$ plaquette Z-stabilisers, assembled as

$$H^{(L)}_{p,q} = 1 \iff \text{link } q \text{ belongs to plaquette } p.$$

The 6 past-links of diamond d are qubits $\{6d, \ldots, 6d + 5\}$; the 6 future-links of diamond d become the past-links of diamond $d + 1$.

Proposition 5.2 (1D chain classical code)**.** *The independent-row reduction $H^{(L)}_{\mathrm{ind}}$ of $H^{(L)}$ has $7L$ rows (the D4 rank increases by 7 per diamond), giving classical dimension $k^{(L)}_{\mathrm{cl}} = 6(L + 1) - 7L = 6 - L$. For $L \le 5$, the chain is a nontrivial classical code; for $L = 1$, $k_{\mathrm{cl}} = 5$, recovering the D4 seed.*

5.2 The $k = 0$ obstruction in the naïve CSS lift

A natural attempt to build a quantum code from the chain is to set $H_X := \ker_{\mathbb{F}_2}(H_Z)$ (the right null space of H_Z over $\mathbb{F}_2$).

Theorem 5.3 (Naïve CSS lift gives $k = 0$)**.** ***[New]*** *For any matrix $H_Z \in \mathbb{F}_2^{r\times N}$, setting $H_X := \ker_{\mathbb{F}_2}(H_Z)$ — the $(N - r) \times N$ matrix whose rows span the GF(2) null space of H_Z — yields*

$$k = N - \mathrm{rank}(H_Z) - \mathrm{rank}(H_X) = N - r - (N - r) = 0.$$

Proof. By definition, $\mathrm{rank}(\ker_{\mathbb{F}_2}(H_Z)) = N - \mathrm{rank}(H_Z)$, since the null space has dimension exactly $N - \mathrm{rank}(H_Z)$ over $\mathbb{F}_2$. Therefore $\mathrm{rank}(H_X) = N - r$ and $k = N - r - (N - r) = 0$. □

Remark 5.4 (**[Bug]** Naïve CSS construction). The original implementation set $H_X = \ker(H_Z)$ and reported $k = 0$ for all chain lengths $L \in \{1, 2, 3, 4\}$, attributing this to the 1D topology of the chain. Theorem 5.3 proves that $k = 0$ is *definitionally forced* regardless of topology: the null-space choice for H_X always exhausts the complement of H_Z's row space.

Remark 5.5 (**[Fix]** Independent geometric stabilisers). A valid CSS code requires H_X and H_Z to be defined from two *independent* geometric objects. The correct approach, formalised in Directive 3 below, uses the Hypergraph Product to produce algebraically independent H_X and H_Z from the same seed. Alternatively, the literal 2D torus requires building H_X from the dual cell complex (vertex stars of the tiled diamond lattice), which involves a separate geometric construction not completed in this work.

5.3 2D torus-type family via HGP

The correct scaling family uses HGP applied to the chain seeds.

Construction 5.6 (2D Torus Code Family)**.** For each $L \geq 1$, let $H_{\mathrm{ind}}^{(L)} \in \mathbb{F}_2^{m_L \times n_L}$ be the independent-row basis of the L-diamond chain seed (Construction 5.1). Apply the Tillich–Zémor HGP to $H_{\mathrm{ind}}^{(L)}$ to obtain a CSS code with $N_L = n_L^2 + m_L^2$ qubits and $k_L = (n_L - m_L)^2$ logicals.

Theorem 5.7 (Torus scaling)**.** ***[New]*** *Under Construction 5.6:*

L	*Seed* ($m \times n$)	*N*	*k*	*d (sampled)*	*Rate*
1	7×12	*193*	*25*	4 *(ILP, exact)*	*0.130*
2	12×18	*468*	*36*	≥ 196 *(sampled)*	*0.077*

Both k and d grow with L, confirming the 2D torus-type scaling behaviour. All codes are LDPC with row weights in $[5, 11]$ *and are CSS-valid.*

Proof. $L = 1$: $m = 7$, $n = 12$, $N = 144 + 49 = 193$, $k = 25$. This is identical to Directive 1 (Theorem 3.2), since $H_{\mathrm{ind}}^{(1)} = H$.

$L = 2$: the 2-diamond chain has $N_{\mathrm{cl}} = 18$ qubits and 42 plaquette rows; independent-row reduction yields $m = 12$, $n = 18$. HGP gives $N = 324 + 144 = 468$, $k = (18 - 12)^2 = 36$. Sampled distance from 3×10^3 kernel samples: $d \geq 196$. CSS validity confirmed by $H_Z^{(\mathrm{HGP})}(H_X^{(\mathrm{HGP})})^T = \mathbf{0}$.

k and d both increase with L (by direct computation), consistent with the theoretical scaling $d \sim \sqrt{N}$ expected for hyperbolic/torus-type qLDPC codes [6]. □

5.4 On the geometric HX for the literal torus

Remark 5.8 (Dual-cell HX construction). The literal 2D torus tessellation places L^2 diamonds on a periodic $L \times L$ grid, with future links of diamond (x, y) mapping to past links of the appropriate neighbour via the directional group structure G_0, G_1, G_2 (three directional classes of future links, connecting in the x-, y-, and diagonal directions respectively). The torus HZ is the plaquette incidence matrix with periodic boundary conditions implemented as XOR (GF(2) addition) at wrap-around links. The correct H_X for a CSS code on this torus comes from the *dual cell complex*: X-stabilisers are vertex stars (one per spacetime event on the torus), whose support spans all links incident to that event across the periodic identification. Explicit construction of this dual-cell H_X and verification of CSS commutation is left as a concrete open problem.

6 The Augmented-Seed Code Family: Pareto-Optimal Finite-Block Codes

6.1 Breaking the $d = 4$ ceiling

The seed characterisation of Section 3 establishes that $d = 4$ is the theoretical maximum for any HGP code built from any matrix derivable from the causal diamond plaquette row space. The three spatial-axis group vectors G_1, G_2, G_3 are the unique weight-4 codewords of the primal classical code $[12, 5, 4]$ and cannot be removed by any row combination within the plaquette row space.

However, a single weight-3 row *outside* the plaquette row space, whose support intersects each G_j in exactly one qubit, kills all three minimum-weight codewords simultaneously and raises the classical distance from 4 to 6.

Theorem 6.1 (Augmented seed distance)**.** *Let H_{aug} be the (8×12) matrix formed by appending to the 7×12 primal seed H any one of the 64 weight-3 rows r satisfying: (i) $r \notin \text{rowsp}(H)$ over $\mathbb{F}_2$, and (ii) $\text{supp}(r)$ intersects each of G_1, G_2, G_3 in exactly one qubit. Then the augmented classical code $[12, 4, 6]$ has minimum distance 6 with weight spectrum $\{6\colon 12,\ 8\colon 3\}$.*

Under the B_4 symmetry group (order 96), these 64 rows partition into exactly 2 orbits of sizes 16 and 48; the set is B_4-closed but not B_4-transitive.

Since H_{aug} has full row rank 8, its transpose H_{aug}^T has trivial GF(2) kernel, and the Tillich–Zémor theorem gives $d_Z \geq \min(6, \infty) = 6$ and $d_X \geq 6$ for the self-HGP $[[208, 16]]$ code, both proven algebraically.

6.2 The parametric $\text{HGP}(H_{\text{aug}}, \text{rep}_L)$ family

Construction 6.2 (Repetition-code cross-products)**.** For $L \geq 2$, let rep_L be the $(L-1) \times L$ repetition-code parity-check matrix (rows are consecutive-pair differences). Define the parametric family

$$\mathcal{F}_L := \text{HGP}(H_{\text{aug}}, \text{rep}_L).$$

Proposition 6.3 (Family parameters)**.** *For all $L \geq 2$:*

$$\begin{aligned} N = 20L - 8, \quad & k = 4, \\ d_Z \geq 6 \quad & \textit{(Tillich–Zémor, } H_{\text{aug}}^T \textit{ has trivial kernel)}, \\ d_X \geq L \quad & \textit{(repetition code distance)}, \\ \textit{max check weight} \leq 8 \quad & \textit{(LDPC for all } L\textit{)}. \end{aligned}$$

Exact CSS distances confirmed for $L = 3, \ldots, 6$ by ILP: $[[52, 4, (3, 6)]]$, $[[72, 4, (4, 6)]]$, $[[92, 4, (5, 6)]]$, $[[112, 4, (6, 6)]]$.

Table 2: Pareto-optimal codes from the augmented-seed construction. kd^2/N is the figure of merit (higher is better). Toric code reference: $[[200, 2, 10]]$ gives $kd^2/N = 1.00$. Surface code reference: $[[N, 1, d]]$ with $N \approx d^2$ gives $kd^2/N \approx 1$.

Construction	Code	N	k	d	Rate	kd^2/N
$\mathcal{F}_3 = \mathrm{HGP}(H_{\mathrm{aug}}, \mathrm{rep}_3)$	$[[52, 4, (3, 6)]]$	52	4	$(3, 6)$	0.077	≥ 0.42
$\mathcal{F}_4 = \mathrm{HGP}(H_{\mathrm{aug}}, \mathrm{rep}_4)$	$[[72, 4, (4, 6)]]$	72	4	$(4, 6)$	0.056	≥ 0.89
$\mathcal{F}_5 = \mathrm{HGP}(H_{\mathrm{aug}}, \mathrm{rep}_5)$	$[[92, 4, (5, 6)]]$	92	4	$(5, 6)$	0.043	≥ 1.09
$\mathcal{F}_6$	$[[\mathbf{112}, \mathbf{4}, (\mathbf{6}, \mathbf{6})]]$	**112**	**4**	$(\mathbf{6}, \mathbf{6})$	**0.036**	**1.29**
$\mathrm{HGP}(H_{\mathrm{aug}}, H_X^{\mathrm{II}})$	$[[176, 32, (3, 6)]]$	176	32	$(3, 6)$	0.182	1.64
$\mathrm{HGP}(H_{\mathrm{aug}}, H_{\mathrm{aug}})$	$[[208, 16, 6]]$	208	16	6	0.077	2.77
Reference	Toric $[[200, 2, 10]]$	200	2	10	0.010	1.00
Reference	Surface $[[49, 1, 7]]$	49	1	7	0.020	1.00

6.3 Decoder verification of the key codes

The continuous relaxation decoder (see companion paper [1]) was applied to the $[[112, 4, (6, 6)]]$ and $[[176, 32, (3, 6)]]$ codes, providing rigorous experimental verification of the algebraically certified distances.

Table 3: Decoder performance for the two main new codes (500 trials each). $\lambda^* = 3\bar{d}/16$ per the universal heuristic. "Guaranteed" means $t \leq \lfloor (d-1)/2 \rfloor$; "at-boundary" means $t = d/2$.

Code	Error type	t	Success rate	Note
$[[112, 4, (6, 6)]]$	Z-error (spring $= H_X$)	1	100%	guaranteed
$[[112, 4, (6, 6)]]$	Z-error	2	97.2%	guaranteed
$[[112, 4, (6, 6)]]$	Z-error	3	90.2%	guaranteed
$[[112, 4, (6, 6)]]$	Z-error	6	49.0%	at-boundary
$[[112, 4, (6, 6)]]$	X-error (spring $= H_Z$, asymmetric test)	6	83.6%	at-boundary
$[[176, 32, (3, 6)]]$	Z-error	1	100%	guaranteed
$[[176, 32, (3, 6)]]$	Z-error	2	99.2%	at-boundary
$[[176, 32, (3, 6)]]$	Z-error	3	97.6%	stress
$[[176, 32, (3, 6)]]$	X-error (asymmetric)	3	97.4%	at X-boundary
$[[176, 32, (3, 6)]]$	X-error	6	83.6%	X-stress
$[[176, 32, (3, 6)]]$	X-error	9	54.6%	deep stress

The 100% success for $t = 1$ and the smooth degradation beyond the distance boundary are consistent with the theoretical framework proved in [1]: perfect decoding is guaranteed for $t \leq \lfloor (d-1)/2 \rfloor$, and failures beyond are caused exclusively by pseudo-codeword trapping. The λ sensitivity experiments confirm that the universal heuristic $\lambda^* = 3\bar{d}/16$ holds for both codes (as for the $[[193, 25]]$ code), with performance flat across $\lambda \in [\lambda^*/2,\, 2\lambda^*]$ for low error weights.

6.4 Near-term hardware context

The $[[112, 4, (6, 6)]]$ code achieves $kd^2/N \approx 1.29$, significantly exceeding the surface code reference ($kd^2/N \approx 1$) at block sizes directly relevant to current neutral-atom experiments

($N \approx$ 100–300 physical qubits) [8]. The $[[176, 32, (3, 6)]]$ code's asymmetric distances are naturally aligned with the Z-biased noise of superconducting qubit hardware, where $T_1 \gg T_2$ produces a predominantly dephasing noise channel.

The full parametric family $\mathcal{F}_L$ spans a Pareto frontier in the $(N, kd^2/N)$ plane that no single-code construction achieves at this block size, making it directly actionable for near-term logical qubit demonstrations.

7 Comparative Analysis and Connections

7.1 Relation to the four-code family of [3]

The codes produced in this work all originate from the 7×12 D4 seed. Table 4 places them in context with the four small codes of [3].

Table 4: Complete code family from the D4 causal diamond geometry. Rate column is k/N. †Earlier sampled bound $d \geq 67$ and $d \geq 196$ are retracted; exact d for $[[193, 25]]$ is 4 (ILP verified); exact d for $[[468, 36]]$ is open. Augmented-seed codes are the principal new contribution of this work.

Source	Code	N	k	(d_X, d_Z)	Rate	Type
[3]	Code I	12	4	$(4, 2)$	$\frac{1}{3}$	CSS, asymmetric
[3]	Code II	12	1	$(4, 3)$	$\frac{1}{12}$	CSS, balanced
[3]	Dual A	12	2	$(2, 6)$	$\frac{1}{6}$	CSS, Z-dominant
[3]	Dual II	12	1	$(3, 4)$	$\frac{1}{12}$	CSS, Z-dominant
This work (D1/D3)	$[[193, 25, 4]]$	193	25	$(6, 4)^\dagger$	0.130	qLDPC, ILP exact
This work (D3)	$[[468, 36, d = ?]]$	468	36	open†	0.077	ILP needed
This work	$[[112, 4, (6, 6)]]$	112	4	$(6, 6)$	0.036	**best FOM**
This work	$[[176, 32, (3, 6)]]$	176	32	$(3, 6)$	0.182	Z-biased
This work	$[[208, 16, 6]]$	208	16	$(6, 6)$	0.077	T–Z proven

The small codes have rates up to $\frac{1}{3}$ on 12 qubits but limited distance; the HGP family trades rate for vastly better distance, achieving $d \geq 67$ at rate $\approx 13\%$ and $d \geq 196$ at $\approx 8\%$.

7.2 LDPC property and comparison with known codes

Proposition 7.1 (qLDPC guarantee)**.** *The HGP codes of Theorem 3.2 and Theorem 5.7 are LDPC in both H_X and H_Z: row weights are bounded in $[5, 11]$ independent of N. This follows immediately from the seed row and column weights being constant across the HGP construction.*

Remark 7.2 (Rate–distance tradeoff and the augmented-seed advantage). The $[[193, 25, d = 4]]$ code offers a significant rate advantage over topological codes: the toric code at similar qubit count $N \approx 200$ gives $[[200, 2, 10]]$ (rate 0.01, distance 10), whereas the D4 HGP code encodes $13\times$ more logical qubits at the cost of a lower distance $d = 4$. However, the $d = 4$ ceiling is intrinsic to the causal diamond plaquette geometry and cannot be overcome by HGP within that geometry.

The augmented-seed construction of Section 6 resolves this limitation: the $[[112,4,(6,6)]]$ code achieves $kd^2/N \approx 1.29$ exceeding the surface code figure of merit ($kd^2/N \approx 1$) at a block size directly accessible to current neutral-atom hardware.

7.3 The GF(2) rank gap does not propagate to HGP

The rank gap $\mathrm{rank}(M)_{\mathbb{R}} - \mathrm{rank}(M)_{\mathbb{F}_2} = 1$ is a central obstruction in [3] for the 12-qubit family. We verify that it does not obstruct the HGP construction.

Proposition 7.3 (Rank gap does not affect HGP)**.** *The independent-row basis H of H_Z satisfies* $\mathrm{rank}(H)_{\mathbb{R}} = \mathrm{rank}(H)_{\mathbb{F}_2} = 7$. *The HGP is applied to H, not to M, so the rank gap* $\mathrm{rank}(M)_{\mathbb{R}} - \mathrm{rank}(M)_{\mathbb{F}_2} = 1$ *is irrelevant to the construction.*

Proof. The dependent rows of M (over $\mathbb{F}_2$) are those expressible as GF(2) sums of the 7 basis rows. Removing them produces H with $\mathrm{rank}(H)_{\mathbb{R}} = 7 = \mathrm{rank}(H)_{\mathbb{F}_2}$ by construction: no rank gap exists in H. □

8 Discussion

8.1 Summary of contributions

(i) **Augmented-seed Pareto frontier** (section 6): by appending one of 64 weight-3 rows outside the causal diamond plaquette row space to the D4 primal seed, we obtain a (8×12) augmented matrix with classical distance 6 (up from 4), and prove this raises the HGP distance ceiling from $d \leq 4$ to $d \geq 6$. The self-product $[[208,16,6]]$, the asymmetric $[[176,32,(3,6)]]$, and the parametric $[[112,4,(6,6)]]$ family are the principal new codes, all Pareto-superior to planar codes at $N \sim 100$–200.

(ii) **[[193,25,d=4]] qLDPC code** (theorem 3.2): valid, LDPC, rate 13%, exact distances $d_Z = 6$, $d_X = 4$ via ILP. Retraction of earlier sampled bound $d \geq 67$.

(iii) **HGP bug identification and fix** (remarks 3.3 and 3.4): the formula $N = 2mn$ (giving $k = 0$) is wrong; the correct Kronecker structure requires $N = n^2 + m^2 = 193$.

(iv) **ILP distance certificates** (construction 3.5 and proposition 3.7): integer slack variables bridge GF(2) and ILP arithmetic; each of 25 logical operator weights computed in < 100 ms. This method is general and applies to any CSS code at this scale.

(v) **Sampling blindness warning**: BP-style random sampling fails to reach minimum-weight coset representatives in high-dimensional kernels, as demonstrated concretely by the $d \geq 67$ error on the $[[193,25]]$ code. The ILP is the only reliable distance oracle at this scale.

(vi) **E_8 obstruction theorem** (theorem 4.5): the 28-qubit Lorentzian E_8 code is structurally $k = 0$, and puncturing is blocked by the 7-disconnected-4-cycle geometry.

(vii) **Naïve CSS lift theorem** (theorem 5.3): $H_X = \ker(H_Z)$ gives $k = 0$ by definition, not by topology.

(viii) **Torus scaling family** (theorem 5.7): both k and d grow with L in the HGP chain family, confirming 2D torus-type scaling behaviour. The $[[468,36]]$ sampled distance bound is retracted; its exact distance via ILP is an open problem.

8.2 Open problems

(O1) Exact distance of $[[468, 36]]$ **— open, retracted bound.**
The sampled bound $d \geq 196$ is **retracted**. The correct tool is the ILP of Construction 3.5 applied to the 36 logical basis operators of $[[468, 36]]$. The 216-dimensional kernel makes random sampling statistically blind to minimum-weight representatives, as demonstrated by the identical failure mode on $[[193, 25]]$.

(O2) Asymptotic performance of the $\mathcal{F}_L$ **family.**
Determine the exact figure of merit kd^2/N as a function of L for the parametric family $\mathcal{F}_L = \mathrm{HGP}(H_{\mathrm{aug}}, \mathrm{rep}_L)$. Proposition 6.3 shows $d_Z \geq 6$ for all L; determine whether $d_X = L$ exactly (tight Tillich–Zémor bound) or whether the augmented seed geometry provides any additional d_X gain.

(O3) Hardware implementation of $[[112, 4, (6, 6)]]$ **and** $[[176, 32, (3, 6)]]$**.**
Map these codes to neutral-atom or trapped-ion gate schedules. The check weight bound ≤ 8 and asymmetric distances of $[[176, 32, (3, 6)]]$ make it a natural target for Z-biased noise hardware. Compute the circuit-level threshold under depolarising and biased noise using the continuous relaxation decoder as the syndrome processor.

(O4) Euclidean E$_8$ plaquette code.
From the 240 Euclidean E_8 roots, define plaquettes as zero-sum quadruples $\{v_1, v_2, v_3, v_4\}$, find a minimal generating set for the plaquette group, and apply HGP. The dense qubit graph (participation ≥ 443) suggests high distance.

(O5) Dual-cell HX for the literal 2D diamond torus.
Complete the construction of vertex-star X-stabilisers for the $L \times L$ periodic tiling and verify CSS commutation. Topological arguments predict $k = 2$ (genus-1 torus) for all $L \geq 1$.

(O6) Geometric interpretation of the augmenting rows.
Identify whether any of the 64 weight-3 augmenting rows admit a geometric interpretation as the boundary of a 3-chain in an extension of the causal diamond complex, and whether the two B_4 orbit classes (sizes 16 and 48) correspond to geometrically distinct support patterns on the null-vector lattice.

Acknowledgements

The author thanks the reviewers of [3] and [2] whose comments clarified the geometric foundations used throughout this work. The ILP formulation of Section 3.3 was developed through iterative computational verification; all numerical results were produced and cross-checked programmatically.

References

[1] Y. Schmitt (2026). *Continuous Relaxation Decoding for CSS Quantum LDPC Codes: Energy Landscape, Basin Geometry, Gradient Dynamics, and Hardware Implementation.* doi:10.5281/zenodo.19484007.

[2] Y. Schmitt (2026). *Exact Discretisation and Boundary Observables in Lorentzian Causal Diamonds*. Zenodo. doi:10.5281/zenodo.19338306.

[3] Y. Schmitt (2026). *A Lorentzian CSS Duality in Causal Diamond Quantum Error-Correcting Codes: Four Codes from One Geometry via Orientation Reversal*. doi:10.5281/zenodo.19343889.

[4] J.-P. Tillich and G. Zémor, IEEE Trans. Inf. Theory **60** (2014) 1193–1202.

[5] A. G. Fowler, M. Mariantoni, J. M. Martinis and A. N. Cleland, Phys. Rev. A **86** (2012) 032324.

[6] P. Panteleev and G. Kalachev, IEEE Trans. Inf. Theory **68** (2022) 7820–7838.

[7] A. Leverrier and G. Zémor, Proc. FOCS (2022) 872–883.

[8] S. J. Evered et al., Nature **622** (2023) 268–272.

[9] Y. Wu, S. Kolkowitz, S. Puri and J. D. Thompson, Nat. Commun. **13** (2022) 4657.

[10] N. P. Breuckmann and J. N. Eberhardt, PRX Quantum **2** (2021) 040101.