

Relations Among Listed Pauli Terms Bound Distance in Qubit Stabilizer qLTCs

Kun Zhu^{1*}, Zihao Chen¹ and Langtao Hu¹

¹School of Electronic Information and Integrated Circuits, Anqing
Normal University, Anqing, Anhui, China.

*Corresponding author(s). E-mail(s): 082404@aqnu.edu.cn;

Abstract

How many linear relations must a quantum locally testable stabilizer Hamiltonian contain? We prove a finite relation–distance law for qubit stabilizer presentations that bounds code distance in terms of local-term weight, code rate, normalized soundness, and the dimension of the relation space among listed terms. Thus, under bounded term weight, a family cannot simultaneously have constant rate, linear distance, and constant normalized soundness without linearly many independent relations. With bounded qubit incidence, this remains true after duplicate terms are removed. For Calderbank–Shor–Steane presentations with constant normalized soundness in both check sectors, the requirement separates by sector, so redundancy cannot be concentrated in one sector. The result provides a rank-computable necessary test for candidate locally testable stabilizer Hamiltonians, a soundness ceiling for natural rank-minimal hypergraph-product presentations, and sectorwise lower bounds for term-appending amplifiers.

Keywords: quantum locally testable codes, quantum LDPC codes, stabilizer codes, quantum error correction, local testability, stabilizer Hamiltonians

1 Introduction

Recent quantum low-density parity-check code constructions achieve constant rate and linear distance through stabilizer presentations with bounded generator weight and bounded qubit incidence [1–3]. A quantum locally testable code (qLTC) further requires the fraction of violated local terms to be at least a constant times an error’s relative Pauli distance from the stabilizer centralizer [4]. This requirement is closely connected to the quantum PCP conjecture [5]. Earlier high-dimensional and hemicubic

constructions established nontrivial qLTC soundness [6, 7], whereas recent cubical, check-product, and transformation-based constructions have advanced different parts of the rate, distance, soundness, and locality tradeoff [8–11]. They do not, however, determine how much redundancy a local term presentation must contain.

The code space alone does not determine local testability. Adding a Pauli term that is dependent on the listed generators imposes no new noiseless constraint, because its outcome is fixed by the independent generators, yet it can change the fraction of local terms violated by an error. This distinction is illustrated by soundness amplification, which can preserve the number of qubits, distance, and dimension while changing locality [11]. Soundness and locality are therefore properties of an explicitly listed local term presentation rather than of the code space alone. For such a list, the vector space of linear dependencies among its Pauli terms is the relation space, and its dimension is the relation dimension. The question is whether a bounded-weight, constant-rate stabilizer qLTC with linear distance and constant normalized soundness can be tested by a local term list with sublinearly many independent relations.

The BSGKSV redundant-tester theorem states that a positive-rate, constant-query locally testable code whose rejection probability is at least a constant times relative distance requires linearly many redundant low-weight dual words [12]; good classical locally testable codes are now known [1, 13]. The quantum setting uses a different metric. Stabilizer distance is measured in the centralizer modulo the stabilizer group, whereas the classical theorem controls ordinary Hamming distance to a kernel. Low-weight stabilizers can keep that ordinary distance bounded even when the logical quotient distance grows linearly. For Calderbank–Shor–Steane (CSS) presentations, each check kernel contains the low-weight stabilizers generated by the opposite sector. A direct sectorwise use of the classical theorem therefore cannot yield the linear relation lower bound required under the quantum quotient metric.

To retain the quotient metric, we encode an arbitrary qubit stabilizer presentation in a three-term binary complex and we construct a sparse retraction onto the undetectable subspace. Its induced quotient map is onto the logical Pauli space. The resulting finite law shows that local-term weight, rate, normalized soundness, and relation dimension bound the distance. We find that a bounded-weight family with positive rate, linear distance, and constant normalized soundness must therefore have a relation space of linear dimension. When qubit incidence is bounded, we show that the same requirement applies to the deduplicated nonidentity term list. For a finite presentation, its term count and matrix rank fix the relation dimension, so we can evaluate the resulting distance ceiling directly from known parameter bounds. For a CSS presentation, the two check counts and two ranks likewise determine the sectorwise relation dimensions, and we derive a separate bound for each sector. A bounded-weight, positive-rate CSS family with linear distance and constant normalized soundness in both check sectors must therefore meet linearly many independent relations in each sector, so redundancy in one sector cannot compensate for a deficit in the other. We apply these constraints to exclude positive-rate, bounded-degree rank-minimal hypergraph-product families from the constant-normalized-soundness, growing-distance regime and to bound the sectorwise terms required by row-space-preserving, bounded-locality soundness amplifiers.

2 Results

2.1 A relation–distance law for stabilizer Hamiltonians

For a specified stabilizer Hamiltonian, the relation dimension counts the independent linear dependencies among its listed local terms. It is a resource of the presentation, rather than an invariant of the code. The following finite law connects that resource to logical distance.

Theorem 1 (Relation bound for qubit stabilizer qLTCs). *Consider $m > 0$ commuting nonidentity Hermitian Pauli observables whose common $+1$ eigenspace defines an $[[n, K, d]]$ stabilizer code. Suppose every listed term has weight at most w , and let $q = m - (n - K)$ be the dimension of the relations among their binary symplectic vectors. Let $G \in \mathbb{F}_2^{m \times 2n}$ contain those vectors as rows and set $J = \begin{pmatrix} 0 & I_n \\ I_n & 0 \end{pmatrix}$. Suppose the presentation has normalized linear qLTC soundness $\rho > 0$: for every Pauli error represented by $x \in \mathbb{F}_2^{2n}$,*

$$\frac{|GJx|}{m} \geq \rho \frac{\text{dist}_P(x, \ker(GJ))}{n}, \quad (1)$$

where dist_P is Pauli weight modulo the centralizer. If $K \geq Rn > 0$, then, with $r = \text{rank } G = n - K$,

$$d \leq 1 + \left\lfloor \frac{wrn(q+1)}{\rho m K} \right\rfloor \leq 1 + \left\lfloor \frac{w(q+1)}{\rho R} \right\rfloor. \quad (2)$$

The proof lifts systematic syndrome columns through the qLTC inequality and constructs a sparse retraction onto the undetectable subspace. Its induced quotient map is onto the logical Pauli space, so averaging identifies a short nontrivial logical Pauli.

In particular, an independent term list has $q = 0$ and therefore obeys

$$d \leq 1 + \left\lfloor \frac{w}{\rho R} \right\rfloor. \quad (3)$$

Thus a positive-rate family presented by independent bounded-weight terms cannot have both growing distance and constant normalized soundness.

2.2 General and sector-resolved consequences

The finite law first constrains every specified stabilizer term list, including one that repeats existing terms. Under bounded incidence, the same conclusion applies after deduplication. The CSS specialization is stronger in a different way: relations among X -checks control the Z -logical distance, and relations among Z -checks control the X -logical distance. A relation surplus in one sector therefore cannot compensate for a deficit in the other.

Corollary 2 (Good qubit stabilizer qLTCs require extensive term relations). *Consider a family of qubit stabilizer codes with rate at least $R > 0$, relative distance at least*

$\delta > 0$, bounded term weight, and normalized linear qLTC soundness at least $\rho > 0$. The relation dimension q of the specified term list satisfies

$$q \geq \frac{\rho R}{w}(d-1) - 1 = \Omega(n), \quad (4)$$

where w is a common upper bound on term weight.

Rearranging (2) gives the finite lower bound, which is linear for families with $d \geq \delta n$.

Corollary 3 (Good qubit stabilizer qLTCs require extensively many distinct terms). *Under the hypotheses of Corollary 2, assume also that a family-wide constant g bounds the number of listed terms incident on each qubit. Delete repeated copies of the same nonidentity Hermitian Pauli observable, and let $\tilde{m}$ be the remaining term count. Then*

$$\tilde{m} - (n - K) \geq \frac{\rho R}{wg}(d-1) - 1 = \Omega(n). \quad (5)$$

Deduplication preserves the code and distance, while bounded term incidence reduces normalized soundness by at most a constant factor.

Corollary 4 (Finite CSS relation–distance tradeoff). *Consider an $[[n, k, d]]$ CSS code specified by $\mathbf{H}_X \in \mathbb{F}_2^{m_X \times n}$ and $\mathbf{H}_Z \in \mathbb{F}_2^{m_Z \times n}$, with $\mathbf{H}_X \mathbf{H}_Z^T = 0$ and $k \geq Rn > 0$. Let d_Z be the minimum weight in $\ker \mathbf{H}_X / \text{im } \mathbf{H}_Z^T$, and define d_X by exchanging the two sectors. If every check has weight at most w , and the normalized soundness constants of $\mathbf{H}_X$ and $\mathbf{H}_Z$ are both at least $\rho > 0$, then, with $q_X = m_X - \text{rank } \mathbf{H}_X$, $q_Z = m_Z - \text{rank } \mathbf{H}_Z$, and $d = \min\{d_X, d_Z\}$,*

$$d_Z \leq 1 + \left\lfloor \frac{w(q_X + 1)}{\rho R} \right\rfloor, \quad d_X \leq 1 + \left\lfloor \frac{w(q_Z + 1)}{\rho R} \right\rfloor. \quad (6)$$

In particular,

$$d \leq 1 + \left\lfloor \frac{w(\min\{q_X, q_Z\} + 1)}{\rho R} \right\rfloor. \quad (7)$$

Consequently,

$$q_X, q_Z \geq \frac{\rho R}{w}(d-1) - 1, \quad Q \geq \frac{2\rho R}{w}(d-1) - 2, \quad (8)$$

where $Q = q_X + q_Z$.

The two inequalities apply the same sparse retraction to the primal and reversed CSS complexes; the normalized conversions appear in (34) and (35).

For later presentation-level calculations, the two sector counts obey the exact identity

$$m_X + m_Z = n - k + q_X + q_Z. \quad (9)$$

For a CSS presentation, $m = m_X + m_Z$, $K = k$, and $q = q_X + q_Z = Q$. The general relation dimension is therefore the sum of the two sector relation dimensions. Likewise, (5) measures the excess of distinct X - and Z -checks over $n - k$.

Remark 5 (Uniform repetition). For a binary $m \times n$ matrix M , let $q(M) := m - \text{rank } M$. If $M^{[c]}$ is obtained by stacking c copies of every row, its kernel and normalized soundness are unchanged, whereas $q(M^{[c]}) = q(M) + (c - 1)m$. Repeating both CSS sectors therefore leaves the code and its normalized tester unchanged but gives

$$Q^{[c]} = Q + (c - 1)(m_X + m_Z) = cQ + (c - 1)(n - k). \quad (10)$$

For fixed c , locality remains bounded while the incident check degree grows by c . Thus Q depends on multiplicity, whereas Corollary 3 identifies what repetition cannot manufacture.

2.3 A normalized CSS relation budget

The sector-resolved bound turns a target rate, distance, and soundness into a minimum relation budget. To display that dependence, consider a balanced asymptotic slice. Write $\chi = Q/n$ for the endpoint-relation density and $\eta = \rho\delta/\Delta$ for the normalized demand. Set

$$k = Rn, \quad q_X = q_Z = \frac{Q}{2}, \quad m_X = m_Z = \frac{(1 - R)n + Q}{2},$$

both normalized sector soundness constants equal ρ , both sector distances are at least δn , and all row and column weights are at most Δ . Here $\text{rank } \mathbf{H}_X = \text{rank } \mathbf{H}_Z = (1 - R)n/2$. Substitution into the sector bound gives

$$\eta \leq \frac{(1 - R)\chi}{2R(1 - R + \chi)} + o(1). \quad (11)$$

For $0 \leq \eta < (1 - R)/(2R)$, the balanced specialization requires, in the asymptotic limit,

$$\chi \geq \chi_{\min}(R, \eta) := \frac{2\eta R(1 - R)}{1 - R - 2\eta R}. \quad (12)$$

Fig. 1 displays the independent sector budgets and their balanced specialization. Its first panel uses $\chi_X = q_X/n$ and $\chi_Z = q_Z/n$, while the second uses the balanced total density χ . In the balanced slice, any presentation below (12) is ruled out. A one-coordinate error violates at most Δ checks, and the asymptotic quantum Singleton bound gives $\delta \leq (1 - R)/2 + o(1)$ [14]. Together they imply

$$\eta \leq \frac{1 - R}{1 - R + Q/n} + o(1). \quad (13)$$




Fig. 1 Independent relation budgets for CSS qLTC presentations. (a) At rate $R = 1/2$ and equal normalized demands of 0.10, each shaded strip marks a sector deficit excluded by the corresponding sector inequality, up to the vanishing finite-length correction. The upper-right region is not excluded, rather than an achievability claim. (b) The balanced specialization gives the required total relation density. Colour shows $\chi_{\min}$, and hatching marks demands ruled out for every relation density by the combined necessary bounds. The marked target $R = 1/2$, $\eta = 0.10$ requires $Q/n \geq 1/8$. Panel (a) uses the presentation-uniform rate-only bound; panel (b) retains the exact rank-to-row factor fixed by balanced accounting.

Combining this with (12) gives the relation-independent ceiling

$$\eta \leq \eta_{\max}(R) := \frac{1 - R}{1 + R}. \quad (14)$$

Thus every balanced good CSS qLTC obeys $\eta \leq \eta_{\max}(R)$, regardless of its relation density; at rate $R = 1/2$, the ceiling is $1/3$. As a direct budget lookup, a target with $R = 1/2$ and $\eta = 0.1$ requires $Q/n \geq 1/8$ in this slice.

2.4 Consequences for code presentations

The following calculations use the general relation dimension q , denoted $Q = q_X + q_Z$ in the CSS cases. They show how a rank calculation screens a natural presentation, quantifies the cost of repairing one, and distinguishes several ways that constructions supply relation space. Extensive relation dimension is a necessary resource, not a certificate of soundness. The deduplicated form in Corollary 3 removes bounded row multiplicity from this test, and Table 1 collects the resulting classifications.

The standard hypergraph-product (HGP) construction gives an immediate exclusion. Let $\mathbf{H}_i \in \mathbb{F}_2^{\mu_i \times \nu_i}$ have rank r_i , and set $k_i = \nu_i - r_i$ and $s_i = \mu_i - r_i$. Its check matrices are

$$\mathbf{H}_X^{\text{HGP}} = [I_{\nu_1} \otimes \mathbf{H}_2 \mid \mathbf{H}_1^T \otimes I_{\mu_2}], \quad (15)$$

$$\mathbf{H}_Z^{\text{HGP}} = [\mathbf{H}_1 \otimes I_{\nu_2} \mid I_{\mu_1} \otimes \mathbf{H}_2^T] \quad (16)$$

and encode $K = k_1 k_2 + s_1 s_2$ qubits [15]. Direct rank calculation gives

$$q_X = k_1 s_2, \quad q_Z = s_1 k_2. \quad (17)$$

In particular, full-row-rank seeds give $q_X = q_Z = 0$. Their rows are already distinct and nonzero, so deduplication leaves this relation dimension unchanged. Thus the natural HGP term list cannot support a positive-rate, growing-distance family with constant normalized soundness.

Corollary 6 (Soundness ceiling for rank-minimal HGP presentations). *For full-row-rank seeds of constant aspect ratio and bounded degree, assume the standard HGP satisfies $K \geq Rn > 0$ and $d_X, d_Z > 1$. Write $d_{\text{HGP}} := \min\{d_X, d_Z\}$, and let w_X and w_Z bound the corresponding check weights. Let ρ_X^{HGP} and ρ_Z^{HGP} denote the normalized soundness constants of its natural CSS presentation. Then*

$$\rho_X^{\text{HGP}} \leq \frac{w_X}{R(d_Z - 1)}, \quad \rho_Z^{\text{HGP}} \leq \frac{w_Z}{R(d_X - 1)}. \quad (18)$$

In particular, both normalized sector soundness constants are $O(d_{\text{HGP}}^{-1})$.

We obtain the result by substituting $q_X = q_Z = 0$ into Corollary 4 and using $d_X, d_Z \geq d_{\text{HGP}}$.

For good classical seeds, $d_{\text{HGP}} = \Theta(\sqrt{n})$, so both normalized sector soundness constants are at most $O(n^{-1/2})$. This is the inverse-square-root necessary scale for these natural presentations [7, 16].

The same calculation gives a repair cost for a term-appending soundness amplifier. Suppose it appends t_X and t_Z nonzero rows in the existing stabilizer row spaces. The code space, dimension, distance, and rank remain unchanged, and

$$q'_X = q_X + t_X, \quad q'_Z = q_Z + t_Z. \quad (19)$$

If the bounded-locality output has positive rate, growing sector distances, and the displayed normalized sector soundness constants, then

$$t_X \geq \frac{\rho_X R}{w_X}(d_Z - 1) - 1 - q_X, \quad t_Z \geq \frac{\rho_Z R}{w_Z}(d_X - 1) - 1 - q_Z. \quad (20)$$

Here the soundness, weights, and distances are output parameters. Under common bounds ρ, w and quantum distance d , summing gives

$$t_X + t_Z \geq \frac{2\rho R}{w}(d - 1) - 2 - Q_{\text{in}}. \quad (21)$$

Here $Q_{\text{in}} = q_X + q_Z$ is the input relation dimension. Thus this form of amplifier cannot repair a sector with sublinear relation supply using only a sublinear number of added local terms when the output has constant rate, soundness, and linear distance [11].

The asymmetric CSS family of Cross et al. calibrates the sector separation. Starting from a good classical locally testable code $C = \ker H$, its presentation

$$H_X = [H \mid H], \quad H_Z = [I \mid I]$$

has constant rate, locality, and soundness, with $d_Z = d(C) = \Theta(n)$, $d_X = 2$, $q_X = \Theta(n)$, and $q_Z = 0$ [9]. The linear-distance sector carries linear relation dimension, whereas the independent opposite sector has constant distance. Exchanging the sectors gives the dual calibration. Because the overall quantum distance remains two, this does not settle simultaneous sharpness for good qLTCs.

Bivariate-bicycle (BB) presentations supply relation space internally. For cyclic dimensions L and M , set $s = LM$ and consider the standard BB presentation

$$\mathbf{H}_X = [\mathbf{A} \mid \mathbf{B}], \quad \mathbf{H}_Z = [\mathbf{B}^T \mid \mathbf{A}^T], \quad (22)$$

where both matrices have s rows and $2s$ columns [17]. If $h = \dim(\ker \mathbf{A} \cap \ker \mathbf{B})$, the BB dimension formula gives $k = 2h$, while the exact identity between the left kernel of $\mathbf{H}_Z$ and $\ker \mathbf{A} \cap \ker \mathbf{B}$ gives $\text{rank } \mathbf{H}_Z = s - h$. Since the CSS dimension identity gives the same rank for $\mathbf{H}_X$, we obtain

$$q_X = q_Z = h = \frac{k}{2}, \quad q_X + q_Z = k. \quad (23)$$

The common kernel supplies both endpoint relation spaces and the logical subspace, explaining $Q = k$. Thus a positive-rate BB family clears this necessary relation test without importing relations from a separate tester. A fixed- k sequence instead has vanishing rate and relation density.

A check product instead transports a classical relation space across a CSS base. Let the full-row-rank base have n_q qubits and r_X and r_Z check rows. Let $\mathbf{H} \in \mathbb{F}_2^{m \times n_c}$ be a classical tester of rank r , kernel distance d_c , and relation dimension $q_c = m - r$. The product uses $\mathbf{H}_X \otimes \mathbf{H}$ and $\mathbf{H}_Z \otimes \mathbf{H}$ [9], so multiplicativity of rank gives

$$q_X = r_X q_c, \quad q_Z = r_Z q_c. \quad (24)$$

When the tester has positive rate and relative distance, constant query size, and constant soundness, the classical redundant-tester bound gives $q_c = \Theta(n_c)$ [12]. If the base also has a nonvanishing stabilizer-rank density, then $r_X + r_Z = n_q - k = \Theta(n_q)$. The product consequently supplies $\Theta(n_c n_q)$ relations. In the parameterization of Cross et al., the distance is $\Theta(\min\{d_c, n_q\})$ while locality is $O(n_q)$. Fixed locality therefore keeps distance bounded, so this example explains a consistency mechanism rather than the target parameter regime.

The native middle-level presentation of the Dinur–Lin–Vidick (DLV) cubical qLTC provides a complementary benchmark [8]. Writing N_{phys} for its physical block length, the chain-dimension count gives $Q = \Theta(N_{\text{phys}})$. The published family has constant rate and check weight, distance $\Theta(N_{\text{phys}} / \log^3 N_{\text{phys}})$, and soundness at least a constant multiple of $\log^{-3} N_{\text{phys}}$. At that guaranteed soundness scale, Theorem 1 certifies

$Q = \Omega(N_{\text{phys}}/\log^6 N_{\text{phys}})$. The certified lower scale is within an $O(\log^6 N_{\text{phys}})$ factor of the native count. This comparison does not determine whether the bound or the native presentation is optimal.

Table 1 Screening outcomes from the relation-dimension test. Extensive relation dimension leaves a presentation eligible for further qLTC analysis, whereas a sublinear value excludes the target parameter set.

Presentation	Relation dimension or scaling	Screening outcome
Arbitrary stabilizer presentation	$q = m - (n - K)$, or $\tilde{m} - (n - K)$ after deduplication	A sublinear value excludes the simultaneous good-qLTC parameter set; bounded term incidence leaves the asymptotic decision unchanged.
Natural full-row-rank HGP	$Q = 0$	Excludes the natural term list from constant normalized soundness with growing distance; gives an $O(1/d_{\text{HGP}})$ ceiling.
Standard bivariate bicycle	$Q = k$, hence $Q/n = k/n$	Positive-rate families pass the relation test; fixed- k sequences lie outside the positive-rate target regime.
DLV cubical qLTC	$Q = \Theta(N_{\text{phys}})$ in the native middle-level presentation	Realizes the extensive relation scale; its proved relative distance and soundness are inverse-polylogarithmic.
Cross asymmetric qLTC	$q_X = \Theta(n)$, $q_Z = 0$	Calibrates the sector laws separately, with $d_Z = \Theta(n)$ and $d_X = 2$.
Cross check product	$Q = (r_X + r_Z)q_c = \Theta(n_c n_q)$	Transports structured relations from the classical tester; fixed locality gives bounded distance in the proved parameter regime.
Bounded-locality amplifier	$q'_X = q_X + t_X$, $q'_Z = q_Z + t_Z$	Each sector has its own linear deficit; under common bounds, $t_X + t_Z \geq 2\rho R(d-1)/w - 2 - Q_{\text{in}}$.

3 Discussion

The stabilizer theorem is a construction filter for local commuting-Pauli Hamiltonians, since a bounded-weight term list with sublinear relation dimension cannot simultaneously realize constant rate, linear distance, and constant soundness. Given the term list, we answer this first design question with one rank computation; for CSS presentations, we use two sector ranks to retain the finer endpoint accounting. The answer concerns the chosen local terms rather than the underlying code, and Corollary 3 prevents bounded term multiplicity from supplying the required distinct terms.

For a CSS list, $q_X + q_Z = m_X + m_Z - (n - k)$. This identity sharpens a check-count criterion, because linear check counts in both sectors can still leave zero relation density when their leading density equals the stabilizer-rank density $1 - R$, as in a full-row-rank HGP presentation; the corollary instead demands a positive linear excess over $n - k$.

We interpret this resource as a property of the Hamiltonian term list, rather than of noiseless syndrome information. Although an independent generating set fixes every

dependent-check outcome, each dependent projector still contributes to the violated-term count and hence to tester energy; measuring or inferring those terms under noisy readout is a separate operational question.

The closest classical antecedent is the theorem of Ben-Sasson, Guruswami, Kaufman, Sudan, and Viderman, which requires redundant low-weight tests for constant-query testers with constant rejection at the relative-distance threshold [12]. Direct application to either CSS sector does not give our bound. The ordinary distance of $\ker \mathbf{H}_X$ is bounded by any nonzero row of $\mathbf{H}_Z$, and conversely, so it may stay $O(1)$ while the logical quotient distances grow linearly. Constant qLTC soundness then gives only inverse-linear rejection at the classical theorem’s threshold and hence only a constant-scale redundancy guarantee. Pooling the sectors has the same defect because the ordinary centralizer code contains bounded-weight stabilizers. The sparse retraction avoids this loss by acting directly on Pauli quotient distance; Methods gives the exact parameter conversion.

Campbell’s Theorem 3 and Lemma 4 construct a special diagonalized, redundancy-free generating family with $(\infty, f(x) = x)$ -soundness, and Claim 1 bounds its distance when that particular family is LDPC [18]. Our result instead starts from an arbitrary listed qubit stabilizer presentation, needs only constant normalized Pauli soundness, and interpolates quantitatively through a nonzero relation space. Aharonov and Eldar’s Definition 15 and Claim 3 identify this Pauli syndrome inequality with stabilizer quantum local testability [4]. The qualitative observation that soundness entails redundant checks also appears in the single-shot setting [19]; here it becomes a finite, presentation-uniform law in the quotient metric and a separate requirement on each CSS sector.

Hong, Guo, and Lucas derive an extensive relation requirement from the different premise of a finite-temperature phase transition in a classical linear-code Hamiltonian [20]. The two results constrain the same presentation resource under different physical or coding assumptions. Gopalan, Vadhan, and Zhou’s coset-leader interpretation of soundness provides another classical precursor to the quotient language used here [21].

Our symplectic embedding covers arbitrary qubit stabilizer presentations, while the CSS decomposition adds sector resolution and shows that relations cannot be concentrated on one side. We do not treat subsystem codes, for which an extension would require a quotient distinguishing measured gauge operators from the stabilizer Hamiltonian.

The asymmetric Cross family proves that each sectorwise dependence is order-sharp when considered separately, but its overall quantum distance is two. Distance balancing can raise the weak sector distance with a controlled soundness cost, yet the published parameter statements do not determine the relation dimensions required for a simultaneous saturation claim [10]. The DLV comparison likewise benchmarks a native presentation rather than proving either the theorem or that presentation optimal. Sharpness with growing overall quantum distance therefore remains open.

We therefore arrive at the construction problem of building bounded-weight commuting Pauli term lists whose distinct terms have a linear-dimensional relation space while preserving constant rate, linear distance, and constant soundness. For CSS codes, we obtain the more explicit task of finding bounded-degree orthogonal checks with

extensive relations in the two sectors and with linear primal and dual distance. DLV and check products exhibit two mechanisms that create the required relation space, whereas Corollary 6 rules out the natural rank-minimal HGP presentation. We are left with the question of which relation geometries can also support sound testing.

4 Methods

4.1 Quotient homology and the one-endpoint retraction

We work over $\mathbb{F}_2$ and consider a three-term complex

$$C_2 \xrightarrow{B} C_1 \xrightarrow{A} C_0, \quad AB = 0. \quad (25)$$

Write $n_i = \dim C_i$, $H_1 = \ker A / \operatorname{im} B$, $k = \dim H_1$, and

$$d_1 = \min\{|x| : x \in \ker A \setminus \operatorname{im} B\}. \quad (26)$$

The lower endpoint relation dimension is $q_0 = \dim \operatorname{coker} A = n_0 - \operatorname{rank} A$, while for a binary matrix M we define its unnormalized Hamming soundness by

$$\alpha(M) = \min_{x \notin \ker M} \frac{|Mx|}{\operatorname{dist}(x, \ker M)}. \quad (27)$$

Lemma 7 (One-endpoint retraction). *Let (25) have $k > 0$, let $r_0 = \operatorname{rank} A$, and suppose every row of A has weight at most Γ_A . If $\alpha(A) \geq a > 0$, then*

$$d_1 \leq 1 + \left\lceil \frac{\Gamma_A r_0 (q_0 + 1)}{ak} \right\rceil. \quad (28)$$

Proof Let $S = \operatorname{im} A$, choose an information set I of size r_0 such that coordinate projection $Q : S \rightarrow \mathbb{F}_2^{r_0}$ is an isomorphism, and write $T = Q^{-1}$. Every systematic column t_i of T has weight at most $q_0 + 1$. Soundness supplies $u_i \in C_1$ with $Au_i = t_i$ and $|u_i| \leq (q_0 + 1)/a$. Let $U : \mathbb{F}_2^{r_0} \rightarrow C_1$ have columns u_i , extend Q by reading the coordinates in I , and set $h = UQ$ and $P = I - hA$. Then $AhA = A$, so $AP = 0$, and P restricts to the identity on $\ker A$. The quotient map $\pi : \ker A \rightarrow H_1$ therefore makes πP onto, so at least k of its coordinate columns are nonzero. Moreover,

$$\sum_{j=1}^{n_1} |hAe_j| \leq \sum_{i=1}^{r_0} |u_i| |\operatorname{row}_i(QA)| \leq \frac{\Gamma_A r_0 (q_0 + 1)}{a}. \quad (29)$$

Averaging over the nonzero quotient columns gives one whose correction weight is at most the floor in (28). Adding the original coordinate proves the claim. $\square$

4.2 Symplectic specialization

Let $J = \begin{pmatrix} 0 & I_n \\ I_n & 0 \end{pmatrix}$. Represent the listed commuting Pauli observables by $G \in \mathbb{F}_2^{m \times 2n}$, with $GJG^T = 0$, and write $r = \operatorname{rank} G = n - K$ and $q = m - r$. The associated

$m_X - \text{rank } \mathbf{H}_X$. The reversed complex exchanges the sectors. For an $m \times n$ matrix, define

$$\rho(M) = \frac{n}{m} \alpha(M). \quad (33)$$

The stabilizer qLTC inequality restricted to pure Pauli errors gives these two sector inequalities directly [4]. If $r_X = \text{rank } \mathbf{H}_X$, the primal substitution in Lemma 7 gives

$$d_Z \leq 1 + \left\lfloor \frac{w_X r_X n (q_X + 1)}{\rho_X m_X k} \right\rfloor \leq 1 + \left\lfloor \frac{w_X (q_X + 1)}{\rho_X R} \right\rfloor. \quad (34)$$

The reversed complex similarly gives

$$d_X \leq 1 + \left\lfloor \frac{w_Z (q_Z + 1)}{\rho_Z R} \right\rfloor. \quad (35)$$

For a pure Z -type error e , minimization over centralizer Paulis is exactly $\text{dist}(e, \ker \mathbf{H}_X)$, while its violated-term count is $|\mathbf{H}_X e|$. This proves the X -check sector inequality without a converse reduction; exchanging X and Z gives the other sector. Common bounds on the weights and soundness constants now prove Corollary 4.

For Corollary 3, let $c_{\max}$ be the largest multiplicity of a listed observable. A nonempty common $+1$ eigenspace excludes a pair that differs only by sign. Deduplication preserves the code and distance, while the number of violated distinct terms is at least $1/c_{\max}$ times the original count. Bounded term incidence gives $c_{\max} \leq g$, so the deduplicated normalized soundness is at least ρ/g . Applying (4) proves (5).

4.4 Why direct classical reductions lose the quotient distance

Let $C_X = \ker \mathbf{H}_X$, and let d_X^{cl} be its ordinary classical distance. Let w bound the query weight of the classical tester. A fixed normalized sector coefficient $\bar{\rho}_X > 0$ gives rejection probability only

$$\rho_B = \bar{\rho}_X \frac{d_X^{\text{cl}}}{3n} \quad (36)$$

at the relative-distance threshold used in the classical redundant-tester corollary [12]. Because C_X contains every row of $\mathbf{H}_Z$, bounded-weight opposite-sector stabilizers can make $d_X^{\text{cl}} = O(1)$ even when the quotient distance is linear. The resulting classical guarantee is therefore only

$$\frac{\rho_B \dim C_X}{w} - 1 = O(1). \quad (37)$$

The Z sector is identical. Pooling the two sectors into the ordinary centralizer code does not remove the loss: its dimension is $n+k$, but it still contains bounded-weight stabilizers, so the rejection probability at its minimum-distance threshold is $O(1/n)$. The sparse retraction avoids this threshold by acting directly on the quotient $\ker \mathbf{H}_X / \text{im } \mathbf{H}_Z^T$ and its dual.

4.5 Cubical-family relation accounting

For completeness, we record the dimension calculation used in the DLV benchmark. Let $\mathcal{G}$ denote the finite group underlying their cubical complex and write $M = |\mathcal{G}|$. Let ℓ be the fixed local degree, let q_{fld} denote the fixed field size, and let D_j be the dimension of the level- j chain space. Lemma 3.4 of arXiv:2402.07476v3 gives

$$D_j = M\ell^j 2^{4-j} \sum_{\substack{T \subseteq [4] \\ |T|=4-j}} \prod_{r \in T} m_r. \quad (38)$$

For the two-low, two-high split $m_1 = m_2 = \vartheta\ell$ and $m_3 = m_4 = (1 - \vartheta)\ell$, substitution yields

$$D_1 + D_3 - D_2 = 8M\ell^4\vartheta(1 - \vartheta) = \Theta(D_2) \quad (39)$$

for every fixed $\vartheta \in (0, 1)$. Since the binary physical block length is $N_{\text{phys}} = (\log_2 q_{\text{fld}})D_2$, the native presentation obeys

$$Q = (\log_2 q_{\text{fld}})(D_1 + D_3 - D_2) + K = \Theta(N_{\text{phys}}). \quad (40)$$

Statements and Declarations

Acknowledgements

This work was supported by the Youth Project of the Anhui Provincial Department of Education, “Research on Semantic-Aware Quantum LDPC Code Construction and Decoding Mechanisms for Fault-Tolerant Quantum Computing”, Grant No. 2025AHGXZK40236.

Author contributions

K.Z. conceived the project and developed the main argument. K.Z., Z.C., and L.H. analysed the consequences, checked the derivations, and revised the manuscript. K.Z. prepared the initial draft and verification package. All authors approved the submitted version.

Competing interests

The authors declare no competing interests.

Data availability

All finite-diagnostic data generated in this study are included in the Supplementary Information source archive.

Code availability

The exact finite-field verifiers and plotting scripts used for the reported checks are included in the Supplementary Information source archive.

References

- [1] Panteleev, P., Kalachev, G.: Asymptotically good quantum and locally testable classical LDPC codes. In: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pp. 375–388 (2022). <https://doi.org/10.1145/3519935.3520017>
- [2] Leverrier, A., Zémor, G.: Quantum tanner codes. In: 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science, pp. 872–883 (2022). <https://doi.org/10.1109/FOCS54457.2022.00117>
- [3] Dinur, I., Hsieh, M.-H., Lin, T.-C., Vidick, T.: Good quantum LDPC codes with linear time decoders. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing, pp. 905–918 (2023). <https://doi.org/10.1145/3564246.3585101>
- [4] Aharonov, D., Eldar, L.: Quantum locally testable codes. *SIAM Journal on Computing* **44**(5), 1230–1262 (2015) <https://doi.org/10.1137/140975498>
- [5] Aharonov, D., Arad, I., Vidick, T.: Guest column: The quantum PCP conjecture. *SIGACT News* **44**(2), 47–79 (2013) <https://doi.org/10.1145/2491533.2491549>
- [6] Hastings, M.B.: Quantum codes from high-dimensional manifolds. In: 8th Innovations in Theoretical Computer Science Conference (ITCS 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 67, pp. 25–12526 (2017). <https://doi.org/10.4230/LIPIcs.ITCS.2017.25>
- [7] Leverrier, A., Londe, V., Zémor, G.: Towards local testability for quantum coding. *Quantum* **6**, 661 (2022) <https://doi.org/10.22331/q-2022-02-24-661>
- [8] Dinur, I., Lin, T.-C., Vidick, T.: Expansion of high-dimensional cubical complexes: With application to quantum locally testable codes. In: 2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS), pp. 379–385 (2024). <https://doi.org/10.1109/FOCS61266.2024.00031>
- [9] Cross, A., He, Z., Natarajan, A., Szegedy, M., Zhu, G.: Quantum locally testable code with constant soundness. *Quantum* **8**, 1501 (2024) <https://doi.org/10.22331/q-2024-10-18-1501>
- [10] Wills, A., Lin, T.-C., Hsieh, M.-H.: Local testability of distance-balanced quantum codes. *npj Quantum Information* **10**, 120 (2024) <https://doi.org/10.1038/s41534-024-00908-8>
- [11] Wills, A., Lin, T.-C., Hsieh, M.-H.: Tradeoff constructions for quantum locally testable codes. *IEEE Transactions on Information Theory* **71**(1), 426–458 (2025) <https://doi.org/10.1109/TIT.2024.3503500>
- [12] Ben-Sasson, E., Guruswami, V., Kaufman, T., Sudan, M., Viderman, M.: Locally testable codes require redundant testers. *SIAM Journal on Computing* **39**(7), 3230–3247 (2010) <https://doi.org/10.1137/090779875>
- [13] Dinur, I., Evra, S., Livne, R., Lubotzky, A., Mozes, S.: Good locally testable codes.

Annals of Mathematics **203**(2), 511–553 (2026) <https://doi.org/10.4007/annals.2026.203.2.3>

- [14] Knill, E., Laflamme, R.: Theory of quantum error-correcting codes. Physical Review A **55**(2), 900–911 (1997) <https://doi.org/10.1103/PhysRevA.55.900>
- [15] Tillich, J.-P., Zémor, G.: Quantum LDPC codes with positive rate and minimum distance proportional to the square root of the blocklength. IEEE Transactions on Information Theory **60**(2), 1193–1202 (2014) <https://doi.org/10.1109/TIT.2013.2292061>
- [16] Leverrier, A., Tillich, J.-P., Zémor, G.: Quantum expander codes. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, pp. 810–824 (2015). <https://doi.org/10.1109/FOCS.2015.55>
- [17] Bravyi, S., Cross, A.W., Gambetta, J.M., Maslov, D., Rall, P., Yoder, T.J.: High-threshold and low-overhead fault-tolerant quantum memory. Nature **627**, 778–782 (2024) <https://doi.org/10.1038/s41586-024-07107-7>
- [18] Campbell, E.T.: A theory of single-shot error correction for adversarial noise. Quantum Science and Technology **4**(2), 025006 (2019) <https://doi.org/10.1088/2058-9565/aafc8f> . Result numbering follows arXiv:1805.09271v6
- [19] Hong, Y.: Single-shot preparation of hypergraph product codes via dimension jump. Quantum **9**, 1879 (2025) <https://doi.org/10.22331/q-2025-10-07-1879>
- [20] Hong, Y., Guo, J., Lucas, A.: Quantum memory at nonzero temperature in a thermodynamically trivial system. Nature Communications **16**, 316 (2025) <https://doi.org/10.1038/s41467-024-55570-7>
- [21] Gopalan, P., Vadhan, S., Zhou, Y.: Locally testable codes and Cayley graphs. In: Proceedings of the 5th Conference on Innovations in Theoretical Computer Science, pp. 325–336 (2014). <https://doi.org/10.1145/2554797.2554807>