Near-Singleton Logical List Decoding of Quantum LDPC Codes from Faulty Syndromes

Kun Zhu[1]

[1] *School of Electronic Information and Integrated Circuits, Anqing Normal University, Anqing, Anhui, China*

(Dated: July 15, 2026)

Quantum low-density parity-check (QLDPC) codes promise scalable quantum memories, but their decoders receive only sparse check outcomes that may themselves be wrong. We prove that one such record supports logical list decoding near the Singleton radius in a fixed folded-block metric. For any fixed rate and positive gap below that radius, we construct bounded-locality Calderbank–Shor–Steane (CSS) QLDPC families with a randomized near-linear-time decoder. The decoder returns a constant number of Pauli candidates, independent of blocklength. An exact record gives zero-residual coverage of every compatible logical class within the target radius. If an adversarial fraction below a fixed threshold of outcomes is faulty, every compatible error within that radius has stabilizer-reduced relative distance at most a constant times that fraction from some returned candidate. The main advance is to reduce the high-weight logical ambiguity before residual correction begins. Each local syndrome outcome cuts the error space to an affine slice. We prove the quotient-compatibility conditions that let the published Alon–Edmonds–Luby (AEL) average-radius bound control distinct logical classes, leaving only a constant number of alternatives. We then turn the observed slices into the agreement graphs used by fast regularity enumeration. An existing single-shot quantum-Tanner decoder handles the small remainder. The resulting front end preserves the sparse measurement record, and its residual grows linearly with syndrome corruption below the stated threshold.

I. INTRODUCTION

Quantum low-density parity-check (QLDPC) codes now combine constant rate, growing distance, sparse checks, and efficient decoding in several asymptotic constructions [12–16]. These advances make large quantum memories plausible at the coding level. They also expose a harder decoding regime. Suppose many data blocks are wrong and some check outcomes are wrong as well. What can a decoder infer from a single sparse syndrome record?

At high error weight, one syndrome may remain compatible with several logical errors. Stabilizer-related Pauli representatives describe the same action on the encoded information, so physical explanations should be grouped by logical class. Near the Singleton radius, the natural decoder output is a short list of these classes. Quantum list coding formalizes this task and separates list generation from any later information used to select a candidate [1–3]. We ask whether a QLDPC device can produce that list from the syndrome record it actually measures.

Existing quantum results cover neighboring regimes. Quantum Alon–Edmonds–Luby (AEL) constructions give exact-syndrome logical-coset lists for QLDPC codes at a Johnson-type radius [5]. Folded-code and quantum-list-recovery constructions approach the Singleton radius modulo stabilizers [3, 4]. Taken together, these cited results stop short of combining QLDPC locality, a blocklength-independent list size, and stability from one faulty sparse record.

Classical coding and single-shot quantum decoding provide complementary tools. Erased-center AEL transfer and folded Reed–Solomon geometry control classical lists near capacity [6, 7]. Regularity methods can enumerate such lists in randomized near-linear time when given a received word or coordinate-wise candidates [8]. A raw syndrome has a different form. It specifies affine sets of compatible errors. Single-shot quantum-Tanner decoders enter at the other end of the problem, where a small residual is already available [9–11]. A high-radius QLDPC list decoder must turn the affine sets into a short logical list and then into a repairable residual.

The mismatch appears at the first decoding step. A syndrome reports local check constraints while the underlying Pauli error remains hidden. Each local outcome cuts the Pauli error space to an affine slice. Classical list algorithms compare candidates with an observed word, and single-shot repair expects a small residual. Replacing the fibers by one Pauli representative hides their local structure. Separately reducing the redundant record to an independent syndrome basis can remove check coordinates required by the final repair.

Our decoder keeps these affine slices intact and compresses their logical ambiguity before repair. We first prove a quotient-compatibility corollary that lets the published erased-center AEL bound separate distinct logical classes. Only a constant number of classes can then remain near the folded-block Singleton radius. Choosing an origin in each affine fiber makes zero coordinates coincide with the agreement pattern used by regularity enumeration. The resulting candidates differ from every compatible hidden error by a small outer residual, which an existing single-shot quantum-Tanner decoder repairs. This ordering keeps that decoder inside its promised small-residual regime. The raw record supplies the local coordinates used by enumeration and the unreduced outer coordinates used by quantum-Tanner repair.

For every fixed target rate and positive folded-block gap below the Singleton limit, the construction retains bounded check weight and bounded qudit degree. Its randomized near-linear-time decoder returns a constant

TABLE I. How published high-radius and single-shot decoding results relate to the interface proved here. The final row separates the new compatibility steps from the cited ingredients.

Prior result	Established capability and remaining interface
Quantum list coding and candidate selection [1, 2]	Logical list output and auxiliary information for choosing among candidates beyond unique recovery. The cited work stops short of a QLDPC decoder from one faulty sparse record.
Folded quantum list codes near the Singleton radius [3, 4]	Exact-syndrome stabilizer-distinct lists near the quantum Singleton radius, including quantum-AEL and quantum-list-recovery constructions. The direct high-radius construction has growing list size and lacks QLDPC locality; folded Hermitian codes leave the QLDPC and faulty-record interfaces unresolved.
Quantum Alon–Edmonds–Luby QLDPC list decoding [5]	Exact-syndrome logical-coset lists for QLDPC codes obtained with quantum-AEL amplification and a Sum-of-Squares decoder. The list radius is Johnson-type, and the theorem assumes an exact syndrome.
Classical erased-center average-radius transfer [6, 7]	Jeronimo et al. prove the erased-center local-to-global transfer, while Chen and Zhang provide the folded Reed–Solomon average-radius input used there. Applying this geometry to stabilizer-quotient logical classes requires an additional compatibility check.
Regularity-based expander decoding [8]	Randomized near-linear classical list decoding and list recovery from a received word or coordinate-wise candidate lists. A raw syndrome instead specifies affine fibers of possible errors.
Single-shot correction and quantum-Tanner repair [9–11]	Measurement-noise control after the data residual lies inside a fixed constant-relative-radius regime. Near-Singleton logical ambiguity remains upstream of this interface.
Present work	**A constant-size logical list from one faulty sparse QLDPC record. A quotient-compatibility corollary lets the published erased-center transfer bound stabilizer-quotient classes. Fixed affine origins turn local zero coordinates into the agreement relations used by fast enumeration. The unreduced outer coordinates then support exact correction and faulty-record quantum-Tanner repair.**

number of Pauli centers from one sparse record. An exact record gives zero-residual coverage throughout the target folded-block ball. With adversarial record corruption below a fixed threshold, the stabilizer-reduced relative distance from every compatible error to some returned center is at most a constant times the corrupted fraction. This front end extends the useful data-error range of small-residual repair; the repair decoder keeps its original unique-decoding radius. The theorem is stated in the fixed folded-block metric. Unfolded qubit-Hamming distance, unique candidate selection, and circuit-level fault tolerance lie outside its scope. Table I compares these capabilities with the adjacent literature.

Figure 1 maps the proof to the decoder. Section II fixes the observation model and the list-valued guarantee. Sections III–V bound the logical ambiguity, recover candidates from affine syndrome fibers, and control the residual created by faulty outcomes. Section VI gives the information-theoretic reasons for the folded alphabet and list output.

II. OBSERVATION MODEL AND MAIN RESULT

We now specify what the decoder observes and what it must return. Figure 1 separates the relevant objects. The red points are physical errors that satisfy the data and record budgets, and the gray points form one stabilizer orbit within that region. Passing to the quotient turns each orbit into one logical alternative. The orange block is the intermediate outer residual $\boldsymbol{v}$ passed to the single-shot interface. The green bound at the output is the final quotient distance from a hidden error to its witnessing center. Different hidden errors may have different witnesses.

Consider a vector-space Calderbank–Shor–Steane (CSS) stabilizer code [17, 18] whose coordinates are grouped into n folded blocks of fixed dimension. For a Pauli error $\boldsymbol{e} = (\boldsymbol{x}_{\mathrm{X}}, \boldsymbol{x}_{\mathrm{Z}})$, $\mathrm{wt}_{\mathrm{R}}(\boldsymbol{e})$ counts the right blocks touched by either CSS sector. The subscript R labels right blocks and is unrelated to the rate parameter used below. Let $\mathcal{S}$ denote the full Pauli stabilizer space. Representatives that differ by an element of $\mathcal{S}$ have the same syndrome and logical action, so their separation is measured in the quotient

$$d_{\mathcal{S}}(\boldsymbol{e}, \boldsymbol{e}') = \min_{\boldsymbol{s} \in \mathcal{S}} \mathrm{wt}_{\mathrm{R}}(\boldsymbol{e} - \boldsymbol{e}' - \boldsymbol{s}). \tag{1}$$

This distance removes physical degeneracy before any list is counted.

The ideal raw record $\Sigma_{raw}(\boldsymbol{e})$ contains the literal bounded-weight check outcomes, including the unreduced outer generator rows. The measured record $\widetilde{\boldsymbol{\sigma}}$ may differ from it. For m_n scalar outcomes, we write $\tau_{\widetilde{\boldsymbol{\sigma}}}(\boldsymbol{e}) = d_H(\widetilde{\boldsymbol{\sigma}}, \Sigma_{\mathrm{raw}}(\boldsymbol{e}))/m_n$ and call $\boldsymbol{e}$ τ-compatible when this fraction is at most τ. The data weight and record discrepancy are separate adversarial budgets. Their dependence is arbitrary, and the analysis retains the literal sparse coordinates throughout.

The decoder returns one small set of centers for the observed record. Every error meeting both budgets must lie near at least one center, but different hidden errors may be covered by different centers. We use the term *residual logical cover* for this list-valued output.

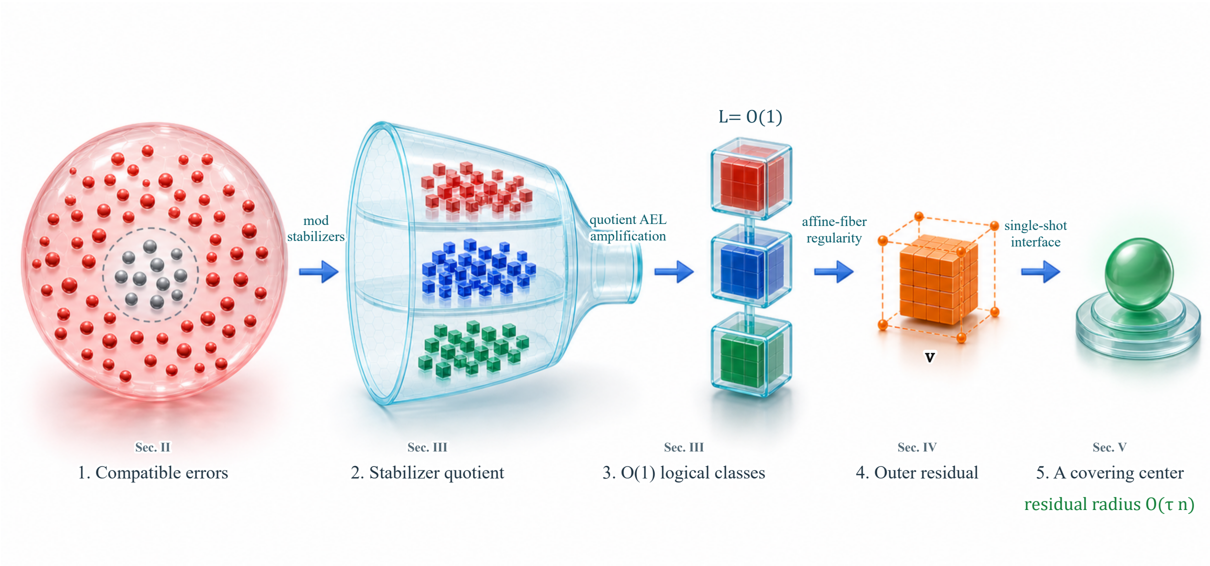

FIG. 1. **Logical ambiguity is compressed before residual repair.** The red region contains physical errors compatible with the data and record budgets, while the gray cluster marks one stabilizer orbit. Quotienting identifies stabilizer-equivalent representatives. The colored stacks are schematic logical alternatives, whose number is bounded by $L = O(1)$. For each compatible hidden error, syndrome-native enumeration supplies an assignment with small outer residual $\boldsymbol{v}$, which the single-shot interface turns into a covering center. The green $O(\tau n)$ label is the final quotient-distance bound, where τ is the raw-record discrepancy; the witnessing center may depend on the error.

Definition 1 (Residual logical cover)**.** For a record $\widetilde{\boldsymbol{\sigma}}$, a set $\widehat{\mathcal{E}} = \{\widehat{\boldsymbol{e}}_1, \dots, \widehat{\boldsymbol{e}}_L\}$ is a (t, τ, r) residual logical cover when

$$\begin{aligned} &\forall \boldsymbol{e} \quad \left[\mathrm{wt}_{\mathrm{R}}(\boldsymbol{e}) \le t,\ \tau_{\widetilde{\boldsymbol{\sigma}}}(\boldsymbol{e}) \le \tau\right] \\ &\implies \quad \exists j \le L \left[d_{\mathcal{S}}(\boldsymbol{e}, \widehat{\boldsymbol{e}}_j) \le r\right]. \end{aligned} \tag{2}$$

When $r = 0$, the cover is an exact list of logical classes. The definition also separates three objects that should not be conflated. A geometric list bound limits how many classes can exist, an algorithm may return a larger constant superset, and selecting one recovery requires information not present in the record.

The main theorem realizes this target at any fixed rate and fixed gap below the folded-block Singleton radius. The formal statement records the uniformity of the returned cover and the dependence on record corruption.

Theorem 1 (Residual-stable QLDPC list decoding)**.** *Fix $R \in (0, 1)$, $0 < \gamma < (1 - R)/4$, and $\eta_{\mathrm{alg}} \in (0, 1)$. There is an infinite family of constant-alphabet vector-space CSS QLDPC codes of asymptotic rate at least R, bounded check weight, and bounded qudit degree. For blocklength n, its raw record has $m_n = O(n)$ scalar coordinates. There are constants $\tau_0 > 0$, $\kappa < \infty$, and $L < \infty$, depending only on the fixed parameters. Given any raw record $\widetilde{\boldsymbol{\sigma}}$, a randomized $\widetilde{O}(n)$-time algorithm returns at most L Pauli centers. With probability at least $1 - \eta_{\mathrm{alg}}$, the same returned set simultaneously satisfies, for every compatible error $\boldsymbol{e}$,*

$$\begin{aligned} &\mathrm{wt}_{\mathrm{R}}(\boldsymbol{e}) \le \left(\frac{1-R}{2} - \gamma\right) n, \qquad \tau_{\widetilde{\boldsymbol{\sigma}}}(\boldsymbol{e}) \le \tau_0, \\ &\implies \qquad \min_{j \le L} d_{\mathcal{S}}(\boldsymbol{e}, \widehat{\boldsymbol{e}}_j) \le \kappa \tau_{\widetilde{\boldsymbol{\sigma}}}(\boldsymbol{e}) n. \end{aligned} \tag{3}$$

The probability is attached to one run on the observed record. On the success event, the returned centers are fixed and cover all errors that satisfy both budgets. An exact record gives zero residual and hence an exact logical list. A faulty record leaves a residual whose size vanishes with its discrepancy. Returned centers are judged by logical distance, so matching corrupted check outcomes is irrelevant. The theorem measures radius on fixed folded blocks and makes no corresponding near-Singleton claim in unfolded qubit-Hamming distance. Appendix D gives the complete order in which the constants are chosen.

The proof now separates into a geometric problem and an algorithmic one. We first bound how many logical classes can fit near one center. Candidate generation from the measured record comes afterward.

III. COMPRESSING LOGICAL AMBIGUITY NEAR THE SINGLETON RADIUS

The first task is geometric. A near-Singleton folded-block ball can contain many physical errors while representing only a constant number of logical classes. This bound is established before a measured record or a candidate-generation algorithm enters the argument.

Pairwise distance alone cannot control a list. Several logical classes might each be far from one another and still cluster around the same comparison center. Average radius rules out this collective clustering. The comparison must also allow some blocks to be omitted while local physical representatives are aligned across the expander. Here an erasure marks a block omitted during that alignment. The physical noise model and the syndrome record remain unchanged, and all distances are normalized by the full number of folded blocks.

For an erased center $\bar{\boldsymbol{g}}$, let s be the erased fraction and let $\Delta_P(\bar{\boldsymbol{g}}, \Gamma)$ denote its minimum relative distance to a logical class Γ in CSS sector P. The property needed below is the following.

Definition 2 (Quotient average radius with erasures)**.** A CSS code has $\mathrm{QAR}(\delta, k, \varepsilon)$ when, in either sector, every erased center $\bar{\boldsymbol{g}}$ with erased-block fraction s and every set $\mathcal{A}$ of at most k distinct logical classes obey

$$\sum_{\Gamma \in \mathcal{A}} \Delta_P(\bar{\boldsymbol{g}}, \Gamma) \geq (|\mathcal{A}| - 1)(\delta - s - \varepsilon). \tag{4}$$

The right side charges nearly one full distance for every additional class. Consequently, many distinct classes cannot all remain close to one center. With no erased blocks, the inequality bounds a quotient ball of radius $(k-1)(\delta - \varepsilon)/k$ by $k-1$ classes. The short sum-of-distances argument and the closed-ball margin appear in Appendix A.

The local and global parts of the construction have different jobs. A fixed folded Reed–Solomon CSS pair prevents many local candidates from clustering. A balanced bipartite expander spreads any persistent outer logical difference across many folded right blocks. The outer quantum-Tanner pair supplies the sparse QLDPC scaffold and positive quotient distance. Its published decoding interfaces will be used only after this geometric compression.

The d edges incident on each left vertex carry one fixed inner block, and the d edges incident on each right vertex form one folded output coordinate. Duality-preserving quotient sections identify the folded logical quotient with the outer quotient. The resulting rate is $r_{\mathrm{out}} r'_{\mathrm{in}}$, while fixed d preserves bounded check weight and bounded qudit degree. Figure 2 shows this propagation. Purple vertices are the left side $\mathcal{L}$ of the expander, blue cubes are its right side $\mathcal{R}$, and each right vertex defines one folded block. The red and blue inset assignments are local representatives of two logical classes; expansion repeats their disagreement across a linear number of right blocks. Appendix A gives the formal nested-pair construction.

The folded Reed–Solomon average-radius theorem is due to Chen and Zhang [7]. Its erased-center corollary is given by Jeronimo et al. [6]. We invoke both results as stated. The remaining check is that two such components can be chosen with explicit dual containment and compatible quotient maps. The field calculation and trace-dual presentation are given in Appendix B. We write $\mathrm{AR}(\delta, k, \xi)$ for this published local guarantee with erasures.

Lemma 3 in Appendix B constructs the required explicit dual-containing pair at any fixed inner rate and list slack. Its two components inherit the published folded Reed–Solomon average-radius guarantee. Admissible inner lengths can be arbitrarily large, so we choose one sufficiently large value of d and then hold it fixed. This meets the mixing requirement while preserving asymptotic locality.

The following corollary records the quotient compatibility needed to invoke the published erased-center transfer. It quantifies over every erased center and every collection of at most k logical classes.

Corollary 1 (Quotient compatibility for average-radius transfer)**.** *Fix $k \geq 2$ and $\zeta > 0$. Suppose both inner components have the average-radius guarantee with erasures $\mathrm{AR}(\delta_0, k, \zeta/2)$. Let δ_{out} be the smaller relative quotient distance of the two outer sectors, and let λ be the normalized second singular value of the expander. If $\lambda \leq \delta_{\mathrm{out}}\zeta/(6k^k)$, then the folded CSS code has $\mathrm{QAR}(\delta_0, k, \zeta)$ in both sectors.*

The corollary verifies the extra quotient compatibility required by the published transfer theorem. For each logical class, choose a representative nearest to the erased center. Distinct quotient classes map to outer words separated by the outer quotient distance, and distinct outer symbols give distinct local inner words. The published erased-center AEL theorem then applies to this finite set of representatives. Appendix C gives this reduction and checks that minimizing over representatives preserves the required distance sum.

One fixed quantum-Tanner outer family is used throughout. Published results supply its positive quotient distance and single-shot repair [10, 11]; setting the measurement-error input of that same decoder to zero gives exact small-residual correction. The paper-specific check is that one component pair and one unreduced generator matrix satisfy all three interfaces. Appendix D proves this common contract. The component pair is selected nonuniformly from a finite good set. The resulting family theorem is existential; a uniform construction algorithm remains open.

Combining the local geometry, quotient compatibility, and common outer family gives the geometric result used by the decoder.

Theorem 2 (Folded-block QLDPC logical lists near the Singleton radius)**.** *Fix $R \in (0,1)$ and $0 < \gamma < (1-R)/4$. There is an infinite family of constant-alphabet vector-space CSS QLDPC codes of asymptotic rate at least R, bounded check weight, and bounded qudit degree. For all sufficiently large blocklengths, every folded-block quotient ball in either CSS sector of radius $((1-R)/2 - \gamma)n$ contains at most $\lceil 4/\gamma \rceil - 1$ logical classes.*

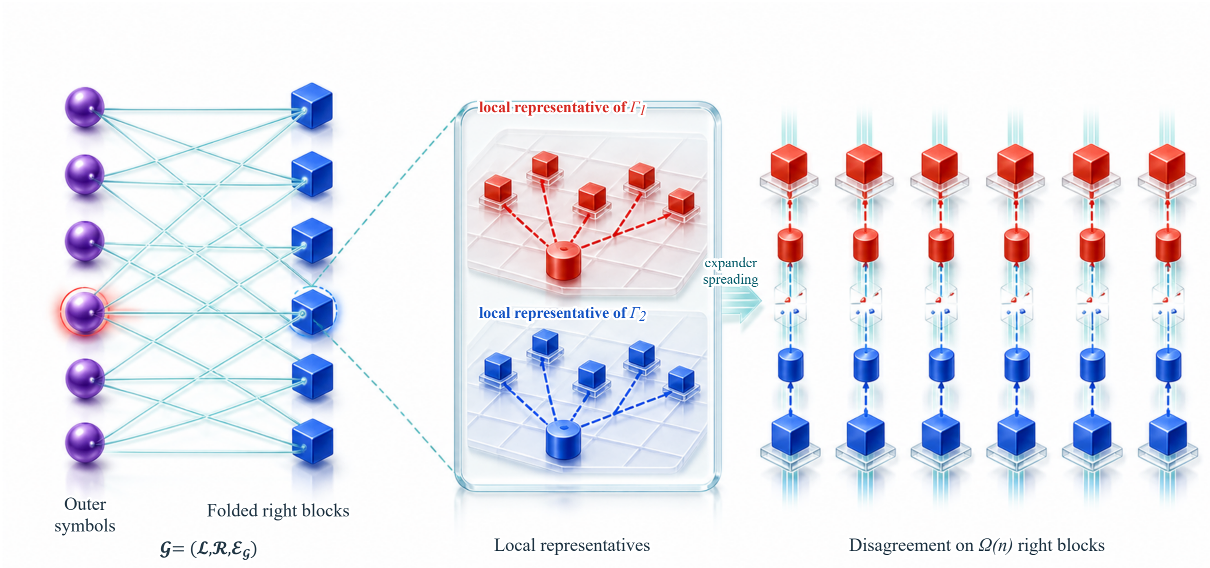

FIG. 2. **Expander amplification of logical separation.** Purple vertices and blue cubes are the two sides $\mathcal{L}$ and $\mathcal{R}$ of the balanced expander $\mathcal{G} = (\mathcal{L}, \mathcal{R}, \mathcal{E}_{\mathcal{G}})$; every $\mathcal{R}$-vertex defines one folded right block. Red and blue in the inset distinguish local representatives of Γ_1 and Γ_2. A persistent outer difference is exposed on a linear number of right blocks. This pairwise view shows the mechanism, while the average-radius transfer controls any fixed bounded collection of classes.

Taking the Cartesian product of the two sector lists gives at most $(\lceil 4/\gamma \rceil - 1)^2$ Pauli logical classes. The outer rate is chosen close enough to one to retain a fixed distance and the two decoding interfaces. The inner rate then supplies the target overall rate and the near-Singleton radius. Appendix D verifies the remaining margins.

Theorem 2 reduces the near-Singleton ball to a constant logical search space. Constructing that search space from a syndrome is a different problem because the record supplies affine constraints in place of a received word. Section IV develops the required conversion to agreement graphs.

IV. GENERATING CANDIDATES FROM A SPARSE SYNDROME

The geometric bound from Section III becomes useful only if the alternatives can be found from measured data. We begin with an exact raw record. A fixed linear-time normalization repackages its literal outcomes as $\boldsymbol{\sigma} = (\boldsymbol{y}, \boldsymbol{z})$ while retaining every local outcome and the redundant outer coordinates. This is a coordinate change on the existing sparse record.

Each local record $\boldsymbol{y}_\ell$ first selects an affine set of errors. Choosing a fixed origin turns that set into an ordinary inner-code list, and a zero coordinate becomes exactly the agreement event needed by the published regularity routine. The routine combines these local tests into a constant family of coherent assignments. Only then does the outer part $\boldsymbol{z}$ enter, supplying the syndrome used to correct the small difference between an assignment and the hidden error. Figure 3 contrasts this route with classical received-word decoding. The gray upper lane stops because the received word is unavailable. In the lower lane, green tests add edges to $\mathcal{H}_i$, red tests reject them, and $\boldsymbol{z}$ bypasses the local graph construction for the later outer-syndrome call.

Only four maps are needed in the main text. The local check matrix $\mathbf{J}$ defines each affine fiber, and its fixed right inverse ι chooses the origin used for translation. The map π reads the outer logical coordinate, while $\mathbf{H}$ is the sparse outer check matrix. Extending π to $\bar{\pi} = \pi(I - \iota\mathbf{J})$ keeps these two parts compatible. The canonical sparse syndrome is therefore

$$\Sigma(\boldsymbol{x}) = \left((\mathbf{J}\boldsymbol{x}_\ell)_{\ell\in\mathcal{L}}, \mathbf{H}\bar{\pi}^L\boldsymbol{x}\right). \tag{5}$$

Its kernel is the amplified sector code. Because the inner block and outer check degree are fixed, the map remains sparse and can be evaluated in linear time. The nested spaces, quotient sections, and kernel identity are verified in Appendix E. These coordinate checks establish the interface used here.

For each left vertex $\ell \in \mathcal{L}$ and observed local outcome $\boldsymbol{y}_\ell$, enumerate the affine list $\mathcal{E}(\boldsymbol{y}_\ell, \rho) = \{\boldsymbol{e} : \mathbf{J}\boldsymbol{e} = \boldsymbol{y}_\ell,\ \Delta(\boldsymbol{e}, 0) \le \rho\}$. Here Δ is local relative Hamming distance and ρ is the fixed local radius inherited from the target folded-block ball. If the hidden error vanishes on a right block b, agreement on the incident edge (ℓ, b) is equivalent to the zero-coordinate test $(\boldsymbol{e}_{\ell,i})_{(\ell,b)} = 0$. Candidate index i therefore defines the computable graph $\mathcal{H}_i$ by including precisely the edges that pass this test.

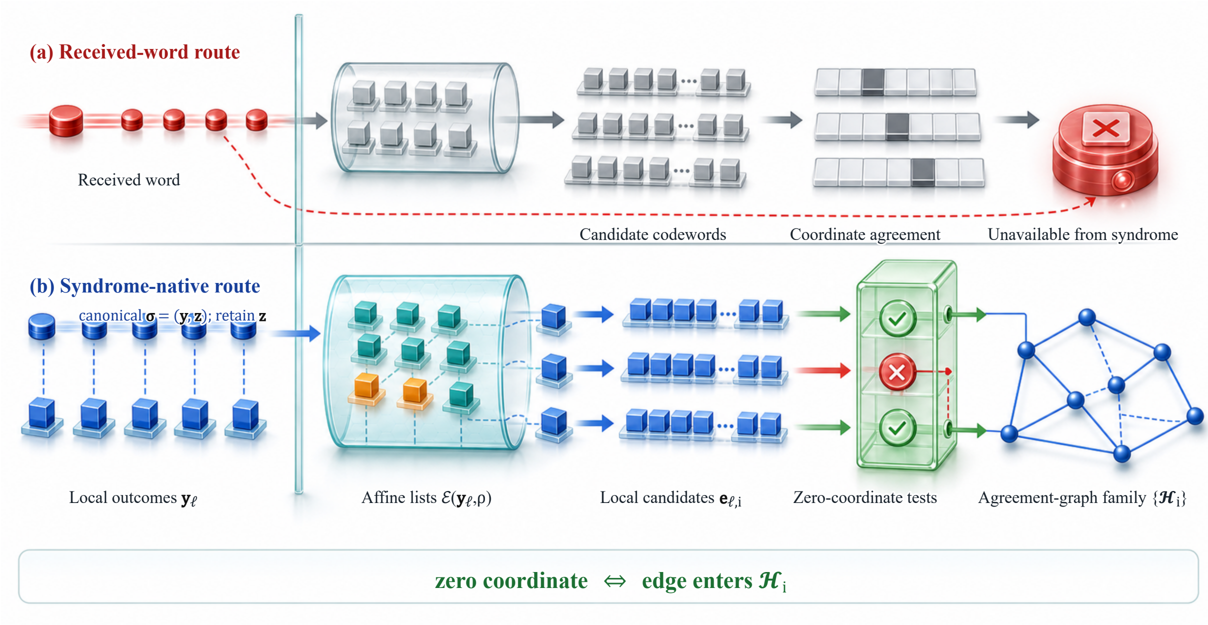

FIG. 3. **Candidate generation without a received word.** Classical list decoding compares codewords with an observed word. Here each local outcome $\boldsymbol{y}_\ell$ defines an affine error list, and a candidate contributes an edge to its agreement graph exactly when the corresponding coordinate is zero. Regularity combines these local graphs into coherent assignments. Green tests contribute edges to $\mathcal{H}_i$, whereas red tests do not. The outer record $\boldsymbol{z}$ bypasses the local construction and is retained for the subsequent quotient-syndrome call.

Because the inner blocklength is fixed, the entire graph family is built in constant work per left vertex.

The two agreement graphs coincide edge by edge. To see this, choose the fixed affine origin $\iota(\boldsymbol{y}_\ell)$. Every local candidate has a unique form $\boldsymbol{e}_{\ell,i} = \iota(\boldsymbol{y}_\ell) + \boldsymbol{c}_{\ell,i}$ with $\boldsymbol{c}_{\ell,i} \in \mathcal{C}_{\text{in}}$. Its local weight is the distance of $\boldsymbol{c}_{\ell,i}$ from the synthetic block $-\iota(\boldsymbol{y}_\ell)$, and the zero-coordinate test is ordinary agreement with that block. This synthetic block is computed from the observed affine origin and serves only as the reference used to build the graph.

We invoke the candidate-enumeration stage of the regularity method of Srivastava and Tulsiani [8]. The new reduction is the equality between its received-word agreement graph and the zero-coordinate graph obtained from each observed affine fiber. Appendix E proves this equality, checks the regularity hypotheses, and records the generic parameter form. After enumeration, the zero-measurement-noise specialization of the quantum-Tanner decoder [11] is used only within its promised small-residual radius.

The resulting decoder is concrete. It enumerates the low-weight affine list for every $\boldsymbol{y}_\ell$, builds the corresponding zero-agreement graphs, and applies regularity to obtain a constant family of assignments $\boldsymbol{e}'$. For each assignment it subtracts the outer syndrome contributed by $\boldsymbol{e}'$ from $\boldsymbol{z}$ and decodes that residual modulo the outer stabilizer. A final consistency test retains only representatives with canonical syndrome $\boldsymbol{\sigma}$. Every step uses the sparse record and the stated local and outer decoding interfaces.

Theorem 3 (Syndrome-native enumeration)**.** *Fix the amplified family and the parameter margins specified by Proposition 1 in Appendix E. For any exact canonical record $\boldsymbol{\sigma} \in \operatorname{im}\Sigma$, fixed target radius β, and fixed $\eta_{\text{reg}} \in (0,1)$, a randomized algorithm returns at most a constant number M of representatives. With probability at least $1-\eta_{\text{reg}}$, their quotient classes simultaneously contain every $\boldsymbol{x}$ satisfying $\Sigma(\boldsymbol{x}) = \boldsymbol{\sigma}$ and $\Delta_{\text{R}}(\boldsymbol{x}, 0) \le \beta$. The runtime is $\widetilde{O}(n) + MT_{syn}(n)$, where T_{syn} is the outer quotient-decoder runtime. A deterministic polynomial-time version gives the same coverage.*

To see why the enumeration is complete, fix any hidden error in the exact fiber and translate it by the same local origins. On an uncorrupted right block, a true local candidate produces the expected zero-agreement edges. Inner distance makes a false match lose many such edges, while expander mixing prevents these losses from being hidden in a small set of neighborhoods. The regularity output therefore contains an assignment $\boldsymbol{e}'$ whose difference from the hidden error lies inside the outer decoding radius.

The random event depends only on the observed graph family. Once it holds, the same returned assignments cover every hidden error in the exact fiber. For the assignment matched to a particular error, the remaining

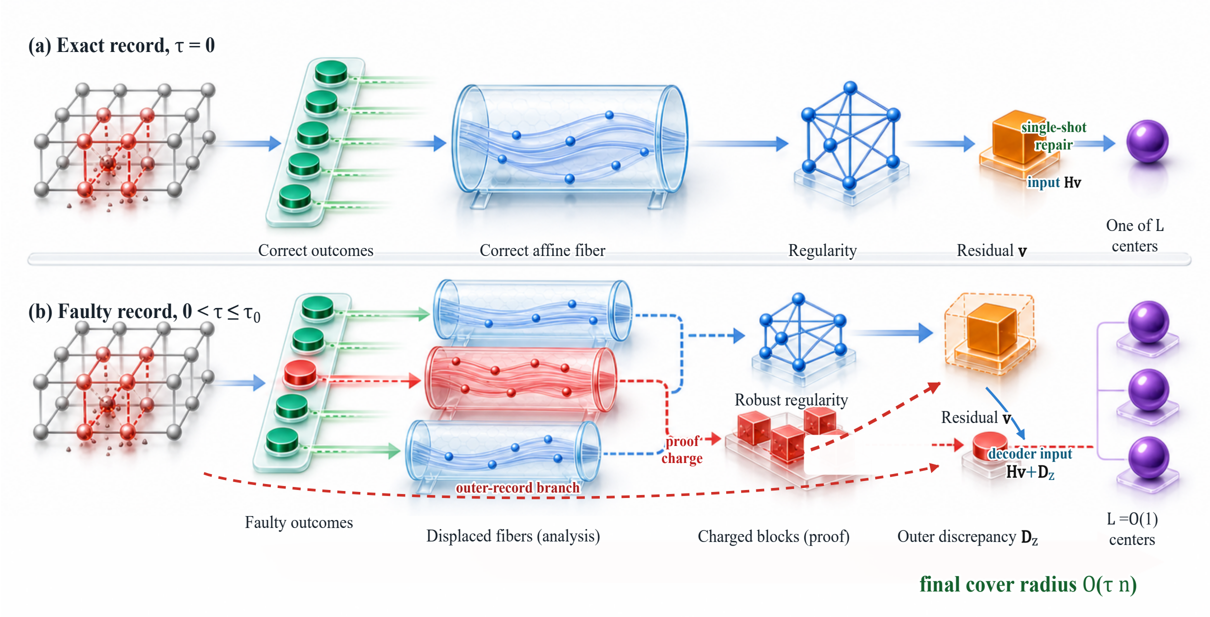

FIG. 4. **Faulty outcomes leave a bounded residual.** Panel (a) follows an exact record. Panel (b) separates the two effects of a faulty record. Red fibers and charged blocks are proof annotations defined relative to a fixed compatible hidden error; their incidence contributes to the support bound on $\boldsymbol{v}$. The term $\boldsymbol{D}_z$ comes separately from the normalized outer-record discrepancy. The single-shot decoder receives the combined input $\mathbf{H}\boldsymbol{v} + \boldsymbol{D}_z$, and the green $O(\tau n)$ label refers to the final quotient distance from the hidden error to a returned center.

relation is

$$\boldsymbol{v} = \pi^L(\boldsymbol{x}-\boldsymbol{e}'), \qquad \mathrm{wt}(\boldsymbol{v}) \le \delta_{\mathrm{dec}} n, \qquad \boldsymbol{z}-\mathbf{H}\bar{\pi}^L\boldsymbol{e}' = \mathbf{H}\boldsymbol{v}. \tag{6}$$

The quotient decoder repairs $\boldsymbol{v}$ modulo $\mathcal{T}_{\mathrm{out}}$. Its input is the small residual in Eq. (6); the original near-Singleton error has already been removed. This exact decoder is the zero-measurement-noise specialization of the single-shot map used below. Local enumeration costs linear total work, regularity is near-linear for fixed parameters, and a constant number of outer calls are made. The full coverage and runtime proof is in Appendix E.

Figure 4 places this exact branch beside the faulty-record branch analyzed next. Local-fiber faults enter through the support bound on the outer residual $\boldsymbol{v}$, while a separate outer-record discrepancy contributes $\boldsymbol{D}_z$. Their sum $\mathbf{H}\boldsymbol{v} + \boldsymbol{D}_z$ is the corrupted syndrome passed to the final repair stage.

Corollary 2 (Exact-record Pauli lists)**.** *For the family of Theorem 2, the two canonical sparse CSS records determine a constant number of Pauli representatives in randomized $\widetilde{O}(n)$ time. For every fixed $\eta_{\mathrm{reg}} \in (0,1)$, with probability at least $1-\eta_{\mathrm{reg}}$ their stabilizer classes simultaneously contain every Pauli error compatible with the exact two-sector canonical record and of folded-block weight at most $((1-R)/2-\gamma)n$.*

The algorithm may return a larger constant superset than the geometric list because it does not perform quotient-distance pruning. The linear-time sparse normalization from literal raw outcomes to the two canonical records places the corollary directly in the observation model of Section II.

An exact sparse record therefore yields a computable constant covering set. The next section quantifies the two fault contributions shown in Fig. 4.

V. STABILITY UNDER FAULTY SYNDROME MEASUREMENTS

We now make the two branches in Fig. 4 quantitative. For a fixed compatible hidden error, the analysis marks the local fibers whose observed outcomes displace that error and charges their incidence to the support of $\boldsymbol{v}$. The decoder itself continues to enumerate the observed fibers and drops the exact-syndrome consistency filter. Corruption in the normalized outer record produces the separate term $\boldsymbol{D}_z$. This decomposition upgrades the exact-record list of Section IV to the residual cover in Theorem 1.

For one sector, write the normalized record as $(\widetilde{\boldsymbol{y}}, \widetilde{\boldsymbol{z}})$. Relative to a hidden error $\boldsymbol{x}$, the outer part has the form $\widetilde{\boldsymbol{z}} = \mathbf{H}\bar{\pi}^L\boldsymbol{x} + \boldsymbol{D}_z$. Here $\boldsymbol{D}_z$ records transformed measurement error. The local difference $\boldsymbol{\mu}_y = \widetilde{\boldsymbol{y}} - \mathbf{J}^L\boldsymbol{x}$ identifies the corrupted affine fibers, and α_{f} denotes their fraction. The quantities $\boldsymbol{D}_z$ and $\boldsymbol{\mu}_y$ are analysis variables defined relative to the hidden error; they are not separately available to the decoder. Sparse normalization preserves the discrepancy needed to bound their grouped supports.

The same affine translation is now applied to the observed, possibly wrong fibers. A corrupted outcome has two effects. It can remove the true local error from the enumerated list, and it can invalidate the inner-distance test on assignments selected from that fiber. The robust affine-fiber lemma in Appendix F charges both effects to $O(\alpha_{\mathrm{f}} n)$ outer positions. Its regularity event depends only on the observed zero-agreement graphs. One event therefore supplies a fixed finite assignment family for every hidden error compatible with the record.

Each corrupted local outcome touches only a bounded number of blocks, and the outer normalization has bounded degree. Hence both α_{f} and the grouped support of $\boldsymbol{D}_z$ are bounded by fixed multiples of the raw record discrepancy. This is a deterministic incidence bound and uses no noise independence.

For the assignment matched to a hidden error, let $\boldsymbol{a} = \iota^L \boldsymbol{\mu}_y$ be the lifted local-outcome error. Robust regularity produces the outer data residual $\boldsymbol{v}$ and the residual syndrome relation

$$\boldsymbol{v} = \pi^L(\boldsymbol{x} + \boldsymbol{a} - \boldsymbol{e}'), \qquad \widetilde{\boldsymbol{z}} - \mathbf{H}\bar{\pi}^L \boldsymbol{e}' = \mathbf{H}\boldsymbol{v} + \boldsymbol{D}_z, \quad (7)$$

where $|\operatorname{supp}(\boldsymbol{v})| \leq (q_0 + c_{\mathrm{f}}\alpha_{\mathrm{f}})n$ for fixed constants q_0 and c_{f}. The regularity loss q_0 is chosen inside the published outer decoding radius. Measurement faults add only the term proportional to α_{f}. The reduction has removed the original high-weight data error from the decoder input.

The final step invokes the single-shot quantum-Tanner theorem of Gu et al. [11]. We use that theorem as published. It accepts a small data residual together with a faulty syndrome and returns a stabilizer-reduced residual proportional to the syndrome corruption. The coordinate audit in Appendix F verifies that $\mathbf{H}\boldsymbol{v} + \boldsymbol{D}_z$ uses the same unreduced generator rows and support norms required by that theorem. The exact quotient decoder in Section IV is its zero-measurement-noise specialization on the same outer family.

Run the two CSS-sector regularity routines with failure budgets summing to η_{alg}. On their joint success event, the finite assignment families simultaneously cover every hidden error compatible with the observed record. For sufficiently small fixed raw discrepancy, Eq. (7) lies inside the single-shot regime in both sectors. Lifting the two repaired outer residuals and taking the Cartesian product of the sector lists gives Eq. (3). The final quotient residual is bounded by $\kappa\tau_{\widetilde{\boldsymbol{\sigma}}}(\boldsymbol{e})n$, while the total runtime remains randomized near-linear. Appendix F contains the parameter ledger and two-sector support calculation.

Combining the two CSS sectors now gives Theorem 1. Each candidate presented to the published repair decoder lies in its small-residual regime, and the repair radius itself is unchanged. The proof uses a folded alphabet and returns a list; the two converse bounds below show why both features are structurally relevant.

VI. NECESSARY TRADE-OFFS AND SCOPE

Two questions remain after the decoding theorem. Could a binary code retain a constant list equally close to the Singleton radius, and could the decoder select one correct candidate from the same record? The answers explain why folded blocks and a list-valued output are structural features of the result.

The first obstruction concerns alphabet size. For a Q-ary CSS sector of length N, quotient dimension K_Q, data radius t, residual radius r, and list bound L, double counting directly in the stabilizer quotient gives the following bound.

Theorem 4 (Quotient residual sphere bound)**.** *With $V_Q(N,u)$ denoting the volume of a Q-ary Hamming ball of radius u, every exact-syndrome residual cover satisfies*

$$Q^{K_Q} V_Q(N,t) \leq L Q^N V_Q(N,r). \quad (8)$$

Degeneracy cannot remove this bound because the counting is already performed on logical classes.

Corollary 3 (Binary near-Singleton obstruction)**.** *For every asymptotic binary CSS rate $R \in (0,1)$, define $\gamma_*(R) = (1-R)/2 - H_2^{-1}(1-R) > 0$, where H_2 is binary entropy on $[0,1/2]$. An exact binary logical list at radius $((1-R)/2-\gamma)N$ is exponential in N whenever $0 < \gamma < \gamma_*(R)$.*

For each fixed γ, our construction uses an alphabet fixed independently of blocklength. The obstruction shows that a binary constant-size list cannot persist in the indicated closer-to-Singleton interval. This motivates the larger alphabet held fixed for each chosen γ and packaged into fixed-size folded blocks. Determining the optimal alphabet remains open. Stronger alphabet lower bounds in classical generalized Singleton list decoding address a related but different setting [19].

The second obstruction concerns selection. Suppose several compatible errors produce the same observation and are farther apart than twice the desired residual radius. Independent randomness cannot reveal which hidden error occurred.

Theorem 5 (Same-observation selector bound)**.** *If $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_M$ are compatible with one record and have pairwise quotient distance greater than $2r$, every selector whose conditional output law depends only on that record succeeds uniformly with probability at most $1/M$.*

A minimum-weight logical Pauli already supplies a two-point witness once the allowed data radius reaches half the code distance. Extra trusted information can change the observation model, but it must distinguish the packed alternatives. Appendix G gives the quotient sphere proof, the minimum-logical witness, a Fano lower bound, and the corresponding limits of dense syndrome sections and direct linear tags.

The two obstructions explain the shape of the theorem. A larger fixed alphabet keeps the logical list bounded in the stated near-Singleton regime. The output remains a cover because the record alone cannot generally identify one logical alternative. Thus the alphabet and list output reflect the information available in the record.

Several boundaries remain. The constants, alphabet, and degree may deteriorate as γ tends to zero, and the outer component choice is finite but nonuniform. The noise model covers one adversarial raw record. Repeated extraction, hook propagation, and local-stochastic closure lie outside this model. The theorem provides a folded-block residual cover; a qubit-normalized threshold or a recovery channel would require additional results. Lowering the alphabet, making the outer choice uniform, embedding the interface in a concrete extraction circuit, and adding trusted candidate-selection information remain separate theoretical problems. Within these boundaries, the two obstructions place the list and alphabet on an information-theoretic footing.

VII. CONCLUSION

We have shown that one sparse QLDPC syndrome record can support a constant-size logical list near the folded-block Singleton radius. The list is generated in randomized near-linear time. Below a fixed corruption threshold, faulty outcomes leave a stabilizer-reduced residual proportional to the record corruption. Measurement faults therefore reach the final decoder only through a small residual.

The construction succeeds by handling high-weight logical ambiguity before repair. This division keeps high-radius list geometry and small-residual single-shot decoding in the regimes for which each is suited. The resulting interface can be paired with future schemes that add trusted candidate-selection information or a concrete extraction circuit. Unique selection and circuit-level fault tolerance remain separate problems.

Appendix A: Notation and CSS quotient identities

Let $\mathcal{A}$ be a finite-dimensional vector space over a finite field. For $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{A}^n$, write

$$\Delta(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{n}\big|\{i : x_i \neq y_i\}\big|.$$

If $\bar{\boldsymbol{g}} \in (\mathcal{A} \cup \{\perp\})^n$ has an erased set of relative size s, $\Delta(\bar{\boldsymbol{g}}, \boldsymbol{y})$ ignores erased positions but remains normalized by n. For a CSS pair $\mathcal{C}_{\mathrm{Z}}^{\perp} \subseteq \mathcal{C}_{\mathrm{X}} \subseteq \mathcal{A}^n$ and an X-logical class $\Gamma = \boldsymbol{h} + \mathcal{C}_{\mathrm{Z}}^{\perp}$, set

$$\Delta_{\mathrm{X}}(\bar{\boldsymbol{g}}, \Gamma) = \min_{\boldsymbol{z} \in \mathcal{C}_{\mathrm{Z}}^{\perp}} \Delta(\bar{\boldsymbol{g}}, \boldsymbol{h} + \boldsymbol{z}),$$

and define Δ_{Z} symmetrically. These conventions are also used after folding, with one right vertex counted as one block.

For the AEL construction, let $\mathcal{G} = (\mathcal{L}, \mathcal{R}, \mathcal{E}_{\mathcal{G}})$ be a d-regular balanced bipartite graph with $|\mathcal{L}| = |\mathcal{R}| = n$. Let $(\mathcal{D}_{\mathrm{X}}, \mathcal{D}_{\mathrm{Z}})$ be the outer CSS pair and $(\mathcal{C}_{\mathrm{X}}, \mathcal{C}_{\mathrm{Z}})$ the length-d inner CSS pair. Fix complements

$$\mathcal{C}_{\mathrm{X}} = \mathcal{W}_{\mathrm{X}} \oplus \mathcal{C}_{\mathrm{Z}}^{\perp}, \qquad \mathcal{C}_{\mathrm{Z}} = \mathcal{W}_{\mathrm{Z}} \oplus \mathcal{C}_{\mathrm{X}}^{\perp}, \tag{A1}$$

and duality-preserving quotient sections ϕ_X and ϕ_Z into $\mathcal{W}_{\mathrm{X}}$ and $\mathcal{W}_{\mathrm{Z}}$. On the edge coordinates define

$$\begin{aligned}\mathcal{E}_{\mathrm{X}} &= \{\boldsymbol{h} : \boldsymbol{h}_\ell = \phi_X(\boldsymbol{u}_\ell) + \boldsymbol{z}_\ell,\ \boldsymbol{u} \in \mathcal{D}_{\mathrm{X}},\ \boldsymbol{z}_\ell \in \mathcal{C}_{\mathrm{Z}}^{\perp}\},\\ \mathcal{E}_{\mathrm{Z}} &= \{\boldsymbol{h} : \boldsymbol{h}_\ell = \phi_Z(\boldsymbol{v}_\ell) + \boldsymbol{x}_\ell,\ \boldsymbol{v} \in \mathcal{D}_{\mathrm{Z}},\ \boldsymbol{x}_\ell \in \mathcal{C}_{\mathrm{X}}^{\perp}\}.\end{aligned} \tag{A2}$$

Let $\mathrm{Fold}_{\mathrm{R}} : \mathcal{A}^{\mathcal{E}_{\mathcal{G}}} \to (\mathcal{A}^d)^{\mathcal{R}}$ be the fixed coordinate regrouping at right vertices, and set

$$\mathcal{F}_{\mathrm{X}} = \mathrm{Fold}_{\mathrm{R}}(\mathcal{E}_{\mathrm{X}}), \qquad \mathcal{F}_{\mathrm{Z}} = \mathrm{Fold}_{\mathrm{R}}(\mathcal{E}_{\mathrm{Z}}).$$

Folding is a coordinate permutation, so it preserves orthogonal complements. For edge assignments $\boldsymbol{h}, \boldsymbol{h}'$, write

$$\begin{aligned}\Delta_{\mathrm{L}}(\boldsymbol{h}, \boldsymbol{h}') &= \frac{1}{n}|\{\ell : \boldsymbol{h}_\ell \neq \boldsymbol{h}'_\ell\}|,\\ \Delta_{\mathrm{R}}(\boldsymbol{h}, \boldsymbol{h}') &= \frac{1}{n}|\{b : \boldsymbol{h}_b \neq \boldsymbol{h}'_b\}|.\end{aligned}$$

The identities proved below are

$$\mathcal{E}_{\mathrm{Z}}^{\perp} = \phi_X(\mathcal{D}_{\mathrm{Z}}^{\perp}) + (\mathcal{C}_{\mathrm{Z}}^{\perp})^L, \tag{A3}$$

$$\mathcal{F}_{\mathrm{X}}/\mathcal{F}_{\mathrm{Z}}^{\perp} \cong \mathcal{E}_{\mathrm{X}}/\mathcal{E}_{\mathrm{Z}}^{\perp} \cong \mathcal{D}_{\mathrm{X}}/\mathcal{D}_{\mathrm{Z}}^{\perp}. \tag{A4}$$

Lemma 1 (Average radius gives a short logical list). *If one CSS sector has* $\mathrm{QAR}(\delta, k, \varepsilon)$*, then every open quotient ball of radius*

$$\tau_k = \frac{k-1}{k}(\delta - \varepsilon)$$

contains at most $k-1$ distinct logical classes in that sector.

Proof of Lemma 1. If k distinct classes all had distance strictly smaller than τ_k from one center, their distance sum would be strictly smaller than $k\tau_k = (k-1)(\delta - \varepsilon)$, contradicting the definition with no erasures. □

Lemma 2 (CSS quotient, rate, and locality of quantum AEL). *The folded pair $(\mathcal{F}_{\mathrm{X}}, \mathcal{F}_{\mathrm{Z}})$ defined from Eq. (A2) is CSS and satisfies Eqs. (A3)–(A4) in the X sector, together with the symmetric identities in the Z sector. If the outer and inner quantum rates are r_{out} and r_{in}, the folded quantum rate is $r_{\mathrm{out}} r_{\mathrm{in}}$. For fixed d and fixed alphabet, bounded outer check weight and coordinate degree imply bounded folded check weight and coordinate degree.*

Proof of Lemma 2. Orthogonality to every local summand $(\mathcal{C}_{\mathrm{X}}^{\perp})^L$ first forces a vector in $\mathcal{E}_{\mathrm{Z}}^{\perp}$ to lie in $\mathcal{C}_{\mathrm{X}}^L$. The unique splitting in Eq. (A1) writes it as $\phi_X(u) + z$ with $z \in (\mathcal{C}_{\mathrm{Z}}^{\perp})^L$. Duality of ϕ_X and ϕ_Z then says that it is orthogonal to the global part $\phi_Z(\mathcal{D}_{\mathrm{Z}})$ exactly when $u \in \mathcal{D}_{\mathrm{Z}}^{\perp}$, proving Eq. (A3). Since $\mathcal{D}_{\mathrm{Z}}^{\perp} \subseteq \mathcal{D}_{\mathrm{X}}$, this space lies in $\mathcal{E}_{\mathrm{X}}$, and folding preserves the CSS inclusion.

Taking the quotient of Eq. (A2) by Eq. (A3) cancels the independent local stabilizer summands and gives Eq. (A4). If one outer symbol has dimension $b = \dim(\mathcal{C}_{\mathrm{X}}/\mathcal{C}_{\mathrm{Z}}^{\perp})$ over the base field, the logical dimension is $r_{out}nb$ while the final physical dimension is $nd\dim(A)$; their ratio is $r_{out}b/(d\dim A) = r_{out}r_{\mathrm{in}}$. Finally, a check of the edge code is either one fixed inner check or one outer check lifted through a length-d local map. Fixed d, bounded outer row weight, and bounded outer column weight therefore bound both the folded check weight and the number of checks incident on one folded symbol. □

Appendix B: Explicit folded Reed–Solomon CSS inner pair

A classical code $C \subseteq A^d$ has $\mathrm{AR}(\delta, k, \xi)$ with erasures when every erased center $\bar{g}$, with erasure fraction s, and every set $H \subseteq C$ of at most k distinct codewords satisfy

$$\sum_{h \in H} \Delta(\bar{g}, h) \geq (|H| - 1)(\delta - s - \xi).$$

Chen and Zhang prove the folded Reed–Solomon average-radius theorem [7]. Corollary 4.6 of Jeronimo et al. [6] gives exactly the erased-center form used here. Their reduction punctures whole folded blocks, observes that the remaining evaluation points are still distinct, and rescales the punctured rate. We invoke that corollary directly.

Lemma 3 (Explicit dual-containing folded Reed–Solomon inner pair)**.** *Fix an inner quantum rate $r_{\mathrm{in}} \in (0,1)$, a list parameter $k \geq 2$, and positive list-geometry slack ξ and rate slack ν. There is an explicit vector-space CSS pair of rate $r'_{\mathrm{in}} \in [r_{\mathrm{in}}, r_{\mathrm{in}} + \nu)$ whose two components have the folded Reed–Solomon average-radius guarantee with erasures $\mathrm{AR}((1 - r'_{\mathrm{in}})/2, k, \xi)$. The quotient maps and binary presentation are explicit. The construction admits unbounded admissible choices of the constant inner length d while preserving these properties.*

Proof of Lemma 3. Evaluation set. Choose an odd folding parameter B large enough for the published folded Reed–Solomon erased-center guarantee. For a sufficiently large integer t, set

$$q = 2^{t\varphi(B)}, \qquad N = q - 1, \qquad d = N/B,$$

where φ is Euler's totient. Euler's theorem gives $B \mid N$. Let ω generate $\mathbb{F}_q^*$ and put

$$\alpha_j = \omega^{Bj}, \qquad 0 \leq j < d.$$

Block j contains the B consecutive powers of ω starting at ω^{Bj}. Across all blocks these points are exactly $\mathbb{F}_q^*$.

Let

$$K = \left\lceil \frac{1 + r_{\mathrm{in}}}{2} N \right\rceil$$

and let $E_K \subseteq \mathbb{F}_q^N$ be the evaluation code of polynomials of degree less than K on $\mathbb{F}_q^*$. Fold its coordinates into the d blocks above and call the resulting code $C \subseteq (\mathbb{F}_q^B)^d$. Increase t until $K < N$ and $2/N < \nu$. The folded rate is $\rho = K/N$. Any further increase of t preserves these inequalities, so the admissible blocklengths $d = (2^{t\varphi(B)} - 1)/B$ are unbounded.

Dual containment. We compute the dual directly. For $1 \leq e \leq N - 1$,

$$\sum_{a \in \mathbb{F}_q^*} a^e = 0.$$

Consequently every evaluation of $xg(x)$ with $\deg g < N - K$ is orthogonal to every word in E_K, because all monomial products have degrees in $[1, N-1]$. The two spaces have complementary dimensions, so

$$E_K^{\perp} = \{(ag(a))_{a \in \mathbb{F}_q^*} : \deg g < N - K\}.$$

Because $K > N/2$, the right-hand side is contained in E_K. Folding preserves the coordinate dot product, hence $C^{\perp} \subseteq C$. Taking $\mathcal{C}_{\mathrm{X}} = \mathcal{C}_{\mathrm{Z}} = C$ therefore defines a CSS code.

Its normalized quantum rate is

$$r'_{\mathrm{in}} = \frac{2K - N}{N} = 2\rho - 1,$$

and the choice of K gives $r_{\mathrm{in}} \leq r'_{\mathrm{in}} < r_{\mathrm{in}} + 2/N < r_{\mathrm{in}} + \nu$. Since the folded evaluation tuple is appropriate, the published folded Reed–Solomon theorem and its puncturing corollary give both components

$$\mathrm{AR}(1 - \rho, k, \xi) = \mathrm{AR}\left(\frac{1 - r'_{\mathrm{in}}}{2}, k, \xi\right).$$

Quotient coordinates. For completeness, the quotient maps are explicit as well. Before folding, a complement of $E_K^{\perp}$ in E_K is generated by $w_a = \mathrm{ev}(x^a)$ for

$$J = \{0\} \cup \{N - K + 1, \ldots, K - 1\}.$$

The set J has $2K - N$ elements. The power-sum identity gives $\langle w_0, w_0 \rangle = -1$ and, for nonzero $a, b \in J$, a nonzero pairing only when $a + b = N$, again with value -1. Thus $z_0 = -w_0$ and $z_a = -w_{N-a}$ form the basis dual to $(w_a)_{a \in J}$. Folding these two bases supplies the required duality-preserving local maps.

To match a binary outer QLDPC family, let $m = [\mathbb{F}_q : \mathbb{F}_2]$ and choose a basis $(e_i)_{i=1}^m$ of $\mathbb{F}_q$ together with its trace-dual basis $(e_i^*)_{i=1}^m$. Expand X coordinates in the first basis and Z coordinates in the second. Then the binary dot product of the two expansions equals $\mathrm{Tr}_{\mathbb{F}_q/\mathbb{F}_2}(xy)$ on every field coordinate. Thus the two expanded components form a binary vector-space CSS pair, their block Hamming geometry is unchanged, and their quotient symbol has binary dimension $b = m(2K - N)$. This fixed-dimensional binary presentation is the one used in the outer-inner interface below. □

Appendix C: Quotient compatibility with the published AEL transfer

The erased-center local-to-global theorem is Theorem 3.2 of Jeronimo et al. [6]. We do not repeat its partition induction. The only point to verify is that a finite set of center-minimizing quotient representatives satisfies the classical theorem's outer-distance hypothesis.

Proof of Corollary 1. It suffices to consider the X sector. Fix an erased right-block center $\bar{\boldsymbol{g}}$ and distinct classes $\Gamma_1, \ldots, \Gamma_t \in \mathcal{E}_X/\mathcal{E}_Z^{\perp}$, where $t \leq k$. Choose $\boldsymbol{h}_i \in \mathcal{E}_X$ attaining $\Delta_X(\bar{\boldsymbol{g}}, \Gamma_i)$, and let $\boldsymbol{u}_i \in \mathcal{D}_X$ be its outer word.

If $i \neq j$, then $\boldsymbol{u}_i - \boldsymbol{u}_j \notin \mathcal{D}_Z^{\perp}$ by Eq. (A4). The outer quotient distance therefore makes $\boldsymbol{u}_i$ and $\boldsymbol{u}_j$ differ on at least $\delta_{\mathrm{out}} n$ left vertices. The direct-sum splitting in Eq. (A1) implies that the corresponding local inner words $\boldsymbol{h}_{i,\ell}$ and $\boldsymbol{h}_{j,\ell}$ differ at every such vertex.

Consequently, $\{(\boldsymbol{h}_{i,\ell})_{\ell \in \mathcal{L}} : i \in [t]\} \subseteq \mathcal{C}_X^{\mathcal{L}}$ is a possibly nonlinear outer code of minimum distance at least δ_{out}. Its AEL image under the same expander and the identity inner encoding is exactly the selected set of folded representatives. Theorem 3.2 of [6], with the assumed inner erased-center average-radius bound and $\lambda \leq \delta_{\mathrm{out}} \zeta/(6k^k)$, gives

$$\sum_{i=1}^{t} \Delta_R(\bar{\boldsymbol{g}}, \boldsymbol{h}_i) \geq (t-1)(\delta_0 - s - \zeta),$$

where s is the erased fraction. The left side is $\sum_i \Delta_X(\bar{\boldsymbol{g}}, \Gamma_i)$ by the choice of the representatives. This proves the required quotient statement. The Z sector is identical after exchanging X and Z. $\square$

Appendix D: Near-Singleton parameter selection

All choices in this section depend only on the fixed pair (R, γ) and are made before the asymptotic blocklength grows. In particular, the outer component-code choice may be nonuniform, but the selected alphabet, inner length, graph degree, decoder radii, and spectral gap remain constant along the resulting family.

Lemma 4 (Compatible quantum-Tanner outer family)**.** *For every fixed $r_{\mathrm{out}} < 1$, there is an infinite binary CSS quantum-Tanner family with rate at least r_{out}, positive relative distance in both CSS sectors, bounded check weight, and bounded coordinate degree. The same family has, in both sectors,*

1. *a deterministic linear-time quotient syndrome decoder correcting a positive fraction δ_{dec} of outer data errors; and*
2. *a deterministic linear-time single-shot decoder in the standard vertex-grouped generator coordinates used by the quantum-Tanner decoder, without row reduction and possibly with global dependencies, with fixed positive constants $A_P, B_P, C_{\mathrm{ss},P}, \beta_{\mathrm{ss},P}$ satisfying Eqs. (F11)–(F12).*

The component-code pair is selected from a finite good set and may depend nonuniformly on r_{out}.

Proof of Lemma 4. Choose $\rho \in (0, 1/2)$ with $(1-2\rho)^2 > r_{\mathrm{out}}$. The component-code existence argument for quantum-Tanner codes, together with the two-sided product-expansion theorem, supplies a finite binary component pair of complementary rates near ρ whose two codes and two duals have positive relative distance and whose primal and dual tensor sums are product-expanding [14, 20]. Fix one such pair. The published quantum-Tanner parameter theorem on the corresponding Ramanujan-complex family then gives rate above r_{out}, positive relative distance in both sectors, and bounded locality [11, 14].

It remains only to align the decoder coordinates. Let $\mathcal{Q}_0$ and $\mathcal{Q}_1$ be the two Tanner kernels in the standard construction and set

$$\mathcal{D}_{X,0} = \mathcal{Q}_1, \quad \mathcal{T}_{X,0} = \mathcal{Q}_0^{\perp}, \qquad \mathcal{D}_{Z,0} = \mathcal{Q}_0, \quad \mathcal{T}_{Z,0} = \mathcal{Q}_1^{\perp}.$$

Use the ordered vertex-local generator matrices $\mathbf{H}_{X,0}$ and $\mathbf{H}_{Z,0}$ without row reduction. These are exactly the possibly dependent, vertex-grouped coordinates read by the single-shot decoder, and $\ker \mathbf{H}_{P,0} = \mathcal{D}_{P,0}$ with $\mathcal{T}_{P,0} \subseteq \mathcal{D}_{P,0}$.

The sequential theorem of Gu et al. [11], in sector P, provides fixed positive constants $A_P, B_P, C_{\delta,P}, \beta_{\mathrm{ss},P}$ and a deterministic linear-time map. Given $\mathbf{H}_{P,0}\boldsymbol{v} + \boldsymbol{D}$, its hypothesis is

$$A_P |\boldsymbol{v}|_{R,\mathrm{Gu}} + B_P |\boldsymbol{D}|_V \leq \min\{C_{\delta,P} N, d_{\mathrm{code}}/\Delta\}, \tag{D1}$$

and its conclusion is

$$\min_{\boldsymbol{t} \in \mathcal{T}_{P,0}} \mathrm{wt}(\boldsymbol{v} - \widehat{\boldsymbol{v}} - \boldsymbol{t}) \leq \beta_{\mathrm{ss},P} |\boldsymbol{D}|_V. \tag{D2}$$

Here $|\boldsymbol{v}|_{R,\mathrm{Gu}}$ is stabilizer-reduced Hamming weight and $|\boldsymbol{D}|_V$ counts the literal vertex groups. Positive relative distance lets us choose one fixed $\delta_{\mathrm{dec}} > 0$ small enough that $\mathrm{wt}(\boldsymbol{v}) \leq \delta_{\mathrm{dec}} N$ and $\boldsymbol{D} = 0$ satisfy the hypothesis in both sectors. The conclusion then gives $\boldsymbol{v} - \widehat{\boldsymbol{v}} \in \mathcal{T}_{P,0}$. Thus exact quotient correction is the zero-measurement-noise specialization of the same published single-shot map, on the same matrix and the same component pair. This is the required common contract. $\square$

Proof of Theorem 2. Parameter and graph choice. Set

$$\eta = \gamma/16, \qquad r_{\mathrm{out}} = 1 - \eta, \qquad r_{\mathrm{in}} = R/r_{\mathrm{out}},$$
$$k = \lceil 4/\gamma \rceil, \qquad \zeta = \gamma/8.$$

The restriction on γ ensures $r_{\mathrm{in}} < 1$. Choose the outer family from Lemma 4, and denote its minimum sector distance by δ_{out} and its quotient decoding radius by $\delta_{\mathrm{dec}} > 0$.

Set

$$\theta = \gamma/2, \qquad \epsilon = \gamma/4,$$
$$h_* = 13\gamma/16, \qquad K_* = k-1. \tag{D3}$$

Choose a fixed $0 < g \le 1$ so that

$$A_P(K_*+1)g(1+h_*^{-1}) < C_{\mathrm{ss},P}/4 \qquad (P \in \{\mathrm{X},\mathrm{Z}\}). \tag{D4}$$

Now set

$$\xi_{\mathrm{reg}} = \frac{\gamma\delta_{\mathrm{dec}}}{16k},$$
$$\lambda_* = \min\left\{\begin{matrix} \dfrac{\delta_{\mathrm{out}}\zeta}{6k^k}, \dfrac{\gamma\xi_{\mathrm{reg}}}{4}, \xi_{\mathrm{reg}}, \dfrac{\xi_{\mathrm{reg}}^2}{2500}, \\ g\epsilon,\ g,\ \dfrac{g^2}{2500}, \min\limits_P \dfrac{C_{\mathrm{ss},P}h_*}{4A_P} \end{matrix}\right\}. \tag{D5}$$

The first four terms are the transfer and exact-record regularity thresholds; the remaining terms reserve the robust and single-shot margins. This is the single spectral choice used by both the exact and faulty-record arguments.

By Lemma 3, with $\nu = \gamma/8$, choose an explicit constant inner CSS pair of rate r_* with $r_{\mathrm{in}} \le r_* < r_{\mathrm{in}} + \gamma/8$ whose two classical components have

$$\mathrm{AR}(\delta_0, k, \zeta/2), \qquad \delta_0 = \frac{1-r_*}{2}.$$

Use the unbounded choice of the inner construction parameter t to ensure that its blocklength d satisfies $2\sqrt{d-1}/d \le \lambda_*$. For every outer quantum-Tanner blocklength N in the chosen family, select a bipartite d-regular Ramanujan graph with N vertices on each side. The all-sizes theorem of [21] provides this exact length alignment, and the normalized second singular value is at most λ_*. Thus the AEL graph degree and the inner blocklength are the same parameter, rather than independently chosen constants.

Alphabet matching. Let b be the binary dimension of its quotient symbol from Lemma 3. Replace the outer code by the coordinatewise direct sum of b copies. Each outer position is now a symbol in $\mathbb{F}_2^b$. Rate and sector relative distance do not decrease, check weight stays constant, and quotient decoding consists of b independent calls to the original decoder. If the block support is at most $\delta_{\mathrm{dec}}n$, every binary component has support at most $\delta_{\mathrm{dec}}n$, so the component outputs concatenate to a decoder for the direct-sum quotient with the same radius. The outer symbol is therefore exactly compatible with the inner quotient, without a field-extension assumption on the published outer decoder.

Quantum AEL amplification then has asymptotic rate at least $r_{\mathrm{out}}r_* \ge R$, remains LDPC, and has $\mathrm{QAR}(\delta_0, k, \zeta)$ by Corollary 1.

Radius calculation. Let $A = (1-R)/2$. Since $\eta \le 1/16$,

$$A - \delta_0 \le \frac{R\eta}{2(1-\eta)} + \frac{\gamma}{16} \le \frac{\gamma}{8}.$$

The terms in λ_* beyond the transfer threshold certify the later exact-record, robust, and single-shot interfaces. Indeed, the inner average-radius property with two words gives $\delta(\mathcal{C}_\mathrm{X}), \delta(\mathcal{C}_\mathrm{Z}) \ge \delta_0 - \zeta/2$, so for $\beta = A - \gamma$ one may take $\theta = \gamma/2$ in Theorem 3. The same one-line list-size consequence of the classical inner AR inequality bounds its local lists by $k-1$ through radius

$$\frac{k-1}{k}(\delta_0 - \zeta/2) \ge A - 5\gamma/16 > \beta + \theta/2.$$

Thus $K \le k-1$ in Lemma 5, and ξ_{reg} is a valid choice for its regularity accuracy. The same distance estimate gives

$$\delta(C) - \beta \ge h_* = 13\gamma/16, \qquad \beta + 2\epsilon \le \delta(C).$$

These are precisely the margin and local-list hypotheses needed in the faulty-record composition.

The choices also give $1/k \le \gamma/4$ and $\delta_0 - \zeta \le 1/2$. Hence

$$\begin{aligned} \frac{k-1}{k}(\delta_0 - \zeta) &\ge A - \gamma/8 - \gamma/8 - \frac{1}{2k} \\ &\ge A - 3\gamma/8 > A - \gamma = \tau. \end{aligned}$$

Lemma 1 therefore bounds the list at radius τ by $k-1$. Integer rate and blocklength roundings can be absorbed into the strict slack between $A - 3\gamma/8$ and $A - \gamma$. $\square$

Appendix E: Syndrome-native regularity and algorithm

For the regularity interface, fix the local-list bound K, local margin $\varepsilon_{\mathrm{loc}}$, and regularity edge accuracy γ_{reg}. The spectral conditions used below are

$$\lambda \le \min\left\{\gamma_{\mathrm{reg}}\varepsilon_{\mathrm{loc}}, \gamma_{\mathrm{reg}}, \frac{\gamma_{\mathrm{reg}}^2}{2500}\right\}. \tag{E1}$$

The family of assignments is required to cover every h obeying

$$\Delta_{\mathrm{R}}(g,h) \le \beta. \tag{E2}$$

For one CSS sector, write $\mathcal{S}_{\mathrm{in}} \subseteq \mathcal{C}_{\mathrm{in}}$ and $\mathcal{T}_{\mathrm{out}} \subseteq \mathcal{D}_{\mathrm{out}}$ for the inner and outer nested pairs, and choose $\mathcal{C}_{\mathrm{in}} = \phi(\mathcal{U}_{\mathrm{out}}) \oplus \mathcal{S}_{\mathrm{in}}$. Write $\pi : \mathcal{C}_{\mathrm{in}} \to \mathcal{U}_{\mathrm{out}}$ for the quotient coordinate map $\pi(\phi(\boldsymbol{u}) + \boldsymbol{s}) = \boldsymbol{u}$. Choose a surjective local syndrome map $\mathbf{J} : \mathcal{A}_{\mathrm{in}}^d \to \mathcal{Y}$ with $\ker \mathbf{J} = \mathcal{C}_{\mathrm{in}}$ and a fixed linear right inverse $\iota : \mathcal{Y} \to \mathcal{A}_{\mathrm{in}}^d$. The extension

$$\bar{\pi} = \pi \circ (I - \iota\mathbf{J}) : \mathcal{A}_{\mathrm{in}}^d \longrightarrow \mathcal{U}_{\mathrm{out}}$$

agrees with π on $\mathcal{C}_{\mathrm{in}}$. For a sparse outer syndrome map $\mathbf{H} : \mathcal{U}_{\mathrm{out}}^L \to \mathcal{Z}$ with $\ker \mathbf{H} = \mathcal{D}_{\mathrm{out}}$, define

$$\Sigma(\boldsymbol{x}) = \left((\mathbf{J}\boldsymbol{x}_\ell)_{\ell \in \mathcal{L}},\, \mathbf{H}\bar{\pi}^L\boldsymbol{x}\right). \tag{E3}$$

For $\boldsymbol{y} \in \mathcal{Y}$ and $\rho \in [0,1]$, the local affine list is

$$\mathcal{E}(\boldsymbol{y}, \rho) = \{\boldsymbol{e} \in \mathcal{A}_{\mathrm{in}}^d : \mathbf{J}\boldsymbol{e} = \boldsymbol{y},\ \Delta(\boldsymbol{e}, 0) \le \rho\}.$$

Finally, the exact compatible logical fiber used in the completeness proof is

$$\{\boldsymbol{x}+\mathcal{E}_{\mathrm{stab}} : \Sigma(\boldsymbol{x})=\boldsymbol{\sigma},\ \Delta_{\mathrm{R}}(\boldsymbol{x},0)\le\beta\}. \tag{E4}$$

For the syndrome-native specialization, assume $\beta+\theta\le\delta(\mathcal{C}_{\mathrm{in}})$, that every radius- $(\beta+\theta/2)$ inner ball has size at most K with $1\le K\le|\mathcal{C}_{\mathrm{in}}|$, and set

$$\xi_{\mathrm{reg}}=\frac{\theta\delta_{\mathrm{dec}}}{8K}, \qquad \lambda\le\min\{\xi_{\mathrm{reg}}\theta/2,\xi_{\mathrm{reg}},\xi_{\mathrm{reg}}^2/2500\}. \tag{E5}$$

Lemma 5 (Outer-code-free regularity predecoder)**.** *Fix a linear code $C_{\mathrm{in}}\subseteq A^d$ of relative distance δ, a radius $\beta\ge0$, a local margin $\varepsilon_{\mathrm{loc}}>0$, a regularity edge accuracy $0<\gamma_{\mathrm{reg}}\le1$, and a failure probability $\eta_{\mathrm{reg}}\in(0,1)$. Assume $\beta+\varepsilon_{\mathrm{loc}}\le\delta$, every radius- $(\beta+\varepsilon_{\mathrm{loc}})$ ball in A^d contains at most K codewords, and $1\le K\le|C_{\mathrm{in}}|$. Let $G=(L,R,E)$ be a balanced d-regular bipartite graph with normalized second singular value λ satisfying Eq. (E1).*

Given any edge word $g\in A^E$, form at each $\ell\in L$ the ordered list

$$\mathcal{L}_\ell=\{c\in C_{\mathrm{in}}:\Delta(c,g_\ell)\le\beta+\varepsilon_{\mathrm{loc}}\},$$

padded by distinct codewords to size K. There is a randomized $\widetilde{O}(n)$-time procedure which, using only these lists and their agreement graphs with g, enumerates at most

$$M=M(K,\gamma_{\mathrm{reg}},\eta_{\mathrm{reg}})$$

assignments $h'\in C_{\mathrm{in}}^L$. With probability at least $1-\eta_{\mathrm{reg}}$, every locally inner-coded *word $h\in C_{\mathrm{in}}^L$ satisfying $\Delta_{\mathrm{R}}(g,h)\le\beta$ has an enumerated assignment satisfying*

$$\Delta_{\mathrm{L}}(h,h')\le\frac{\lambda}{\varepsilon_{\mathrm{loc}}}+K\gamma_{\mathrm{reg}}+\frac{\lambda+K\gamma_{\mathrm{reg}}+\lambda/\varepsilon_{\mathrm{loc}}}{\delta-\beta}. \tag{E6}$$

No outer-code membership is assumed. The same success event covers all such h simultaneously, and M is independent of n. A deterministic $O(n^{3.5})$ version has the same coverage.

Proof of Lemma 5. We invoke, rather than reproduce, the regularity and signature-net machinery of Srivastava and Tulsiani [8]. Lemma 3.9 of that work constructs a common regular family for the K agreement graphs, and Claim 2.17 enumerates a constant-size net of realizable disjoint-set signatures. The only issue is whether their candidate-generation stage still applies before outer-code membership is known.

Fix any $h\in C_{\mathrm{in}}^L$ with $\Delta_{\mathrm{R}}(g,h)\le\beta$. The proof of their Claim 4.3 uses only this right-block distance and the fact that each h_ℓ belongs to C_{in}. With local-list margin $\varepsilon_{\mathrm{loc}}$ it gives a missed-left-vertex fraction $a\le\lambda/\varepsilon_{\mathrm{loc}}$. Order and pad the local lists, and let A_i contain the vertices where the ith entry equals h_ℓ. On the right-block agreement set $B=\{b:g_b=h_b\}$, the identity

$$E_{H_i}(A_i,B)=E_G(A_i,B)$$

is exact. The proof of their rigidity Lemma 4.4 therefore applies without change: it uses only this identity, inner distance, and expander mixing, not membership in the outer code. For the disjoint sets S_i supplied by the signature net, its unsimplified estimate is

$$\frac{1}{n}\sum_i|S_i\setminus A_i|\le\frac{\lambda+K\gamma_{\mathrm{reg}}+a}{\delta-\beta}.$$

The uncovered true vertices contribute at most $a+K\gamma_{\mathrm{reg}}$ after choosing the regularity tolerance no larger than γ_{reg}. Substituting $a\le\lambda/\varepsilon_{\mathrm{loc}}$ gives exactly Eq. (E6).

The cited regularity construction and signature net enumerate only a parameter-dependent number of assignments. Their success event depends on the observed agreement graphs, so once it holds the preceding argument applies simultaneously to every eligible h. Repeating the complete construction a parameter-dependent number of times gives failure probability η_{reg} without changing the asymptotic runtime. For the finitely many blocklengths below the regularity threshold, brute-force enumeration has parameter-only size. The deterministic regularity routine from the same work gives the stated $O(n^{3.5})$ alternative. No outer decoder is invoked in this candidate-generation stage. □

Lemma 6 (Canonical nested-pair syndrome)**.** *The sparse map Σ in Eq. (E3) satisfies* $\ker\Sigma=\mathcal{E}_{\mathrm{code}}$. *If d and the local maps are fixed and $\mathbf{H}$ has bounded row and column weight, then Σ has bounded check weight and bounded coordinate degree and can be evaluated in $O(n)$ time.*

Proof of Lemma 6. If $\Sigma(\boldsymbol{x})=0$, then $\boldsymbol{x}\in\mathcal{C}_{\mathrm{in}}^L$ and $\bar{\pi}^L\boldsymbol{x}=\pi^L\boldsymbol{x}\in\mathcal{D}_{\mathrm{out}}$. Hence $\boldsymbol{x}=\phi^L(\pi^L\boldsymbol{x})+\boldsymbol{s}$ for some $\boldsymbol{s}\in\mathcal{S}_{\mathrm{in}}^L$, so $\boldsymbol{x}\in\mathcal{E}_{\mathrm{code}}$. The converse follows directly from the definition of $\mathcal{E}_{\mathrm{code}}$. Each local map has constant size, while every nonzero outer coefficient expands through one fixed length-d map, proving the sparsity and runtime claims. □

Definition 3 (Outer quotient syndrome decoder)**.** A radius-δ_{dec} syndrome decoder for $\mathcal{D}_{\mathrm{out}}/\mathcal{T}_{\mathrm{out}}$ is a map $\mathsf{Dec}^{\mathrm{syn}}_{\mathcal{D}_{\mathrm{out}}/\mathcal{T}_{\mathrm{out}}}:\mathcal{Z}\to\mathcal{U}_{\mathrm{out}}^L$ such that every $\boldsymbol{v}\in\mathcal{U}_{\mathrm{out}}^L$ with $\mathrm{wt}(\boldsymbol{v})\le\delta_{\mathrm{dec}}n$ satisfies

$$\boldsymbol{v}-\mathsf{Dec}^{\mathrm{syn}}_{\mathcal{D}_{\mathrm{out}}/\mathcal{T}_{\mathrm{out}}}(\mathbf{H}\boldsymbol{v})\in\mathcal{T}_{\mathrm{out}}.$$

The interface is required only for errors inside the promised radius; its values outside that set may be arbitrary.

Proposition 1 (Parameterized syndrome-native interface)**.** *Fix $\beta\ge0$, $\theta>0$, $\delta_{\mathrm{dec}}\in(0,1]$, and $\eta_{\mathrm{reg}}\in(0,1)$ with $\beta+\theta\le\delta(\mathcal{C}_{\mathrm{in}})$. Suppose every inner ball of radius $\beta+\theta/2$ around an arbitrary ambient center has size at most K, where $1\le K\le|\mathcal{C}_{\mathrm{in}}|$. Set $\xi_{\mathrm{reg}}=\theta\delta_{\mathrm{dec}}/(8K)$ and assume $\lambda\le\min\{\xi_{\mathrm{reg}}\theta/2,\xi_{\mathrm{reg}},\xi_{\mathrm{reg}}^2/2500\}$. If the outer quotient syndrome decoder runs in time $T_{\mathrm{syn}}(n)$, then the conclusion of Theorem 3 holds with $M=M(K,\theta,\delta_{\mathrm{dec}},\eta_{\mathrm{reg}})$ and runtime $\widetilde{O}(n)+MT_{\mathrm{syn}}(n)$. The deterministic alternative runs in $O(n^{3.5})+MT_{\mathrm{syn}}(n)$.*

Proof of Proposition 1 and Theorem 3. Write $\boldsymbol{\sigma} = (\boldsymbol{y}, \boldsymbol{z})$ and choose the canonical affine origin

$$\boldsymbol{r}_\ell = \iota(\boldsymbol{y}_\ell), \qquad \boldsymbol{g}_\ell = -\boldsymbol{r}_\ell. \tag{E7}$$

Every element of the observed fiber has a unique form $\boldsymbol{e}_{\ell,i} = \boldsymbol{r}_\ell + \boldsymbol{c}_{\ell,i}$ with $\boldsymbol{c}_{\ell,i} \in \mathcal{C}_{\rm in}$. Moreover,

$$\begin{aligned} \Delta(\boldsymbol{e}_{\ell,i}, 0) &= \Delta(\boldsymbol{c}_{\ell,i}, \boldsymbol{g}_\ell), \\ (\boldsymbol{e}_{\ell,i})_a = 0 &\iff (\boldsymbol{c}_{\ell,i})_a = (\boldsymbol{g}_\ell)_a. \end{aligned} \tag{E8}$$

Thus translation by $\boldsymbol{r}_\ell$ maps $\mathcal{E}(\boldsymbol{y}_\ell, \beta + \theta/2)$ bijectively to the ordinary $\mathcal{C}_{\rm in}$-list of radius $\beta + \theta/2$ around the synthetic received block $\boldsymbol{g}_\ell$, and its zero-coordinate graph is exactly the corresponding agreement graph. Pad the translated list to size K with distinct words of $\mathcal{C}_{\rm in}$, which is possible because $K \le |\mathcal{C}_{\rm in}|$, and run Lemma 5 with local margin $\theta/2$ and regularity edge error ξ_{reg}. For every enumerated $\mathcal{C}_{\rm in}$-valued assignment $\boldsymbol{c}'$, set

$$\boldsymbol{e}' = \boldsymbol{r} + \boldsymbol{c}'. \tag{E9}$$

Then $\mathbf{J}^L \boldsymbol{e}' = \boldsymbol{y}$. This gives exactly the zero-agreement enumeration described in the main text, but now as an equality of inputs to the regularity predecoder rather than an analogy with received-word decoding.

For each assignment compute

$$\boldsymbol{z}' = \boldsymbol{z} - \mathbf{H}\bar{\pi}^L \boldsymbol{e}', \qquad \widehat{\boldsymbol{v}} = \mathsf{Dec}^{\rm syn}_{\mathcal{D}_{\rm out}/\mathcal{T}_{\rm out}}(\boldsymbol{z}').$$

Discard this assignment unless $\mathbf{H}\widehat{\boldsymbol{v}} = \boldsymbol{z}'$. Otherwise output

$$\widehat{\boldsymbol{x}} = \boldsymbol{e}' + \phi^L(\widehat{\boldsymbol{v}}).$$

The filter is only a syndrome-consistency test; it invokes neither distance pruning nor an all-syndrome section. Every retained output has syndrome $\boldsymbol{\sigma}$ by Lemma 6.

To prove coverage, fix $\boldsymbol{x}$ occurring in Eq. (E4) and put $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{r}$. Since $\mathbf{J}^L\boldsymbol{x} = \boldsymbol{y} = \mathbf{J}^L\boldsymbol{r}$, one has $\boldsymbol{h} \in \mathcal{C}^L_{\rm in}$. Equations (E7)– (E8) also give the exact right-block identity

$$\Delta_{\rm R}(\boldsymbol{g}, \boldsymbol{h}) = \Delta_{\rm R}(0, \boldsymbol{x}) \le \beta. \tag{E10}$$

The target $\boldsymbol{h}$ need not obey the outer code; this is precisely why Lemma 5 was proved without that assumption. Its explicit bound, with $a \le 2\lambda/\theta \le \xi_{\rm reg}$, $\eta \le \xi_{\rm reg}$, $\delta(\mathcal{C}_{\rm in}) - \beta \ge \theta$, and $\lambda \le \xi_{\rm reg}$, is

$$\begin{aligned} \Delta_{\rm L}(\boldsymbol{e}', \boldsymbol{x}) &= \Delta_{\rm L}(\boldsymbol{c}', \boldsymbol{h}) \\ &\le \frac{(2K+3)\xi_{\rm reg}}{\theta} \le \frac{5}{8}\delta_{\rm dec} < \delta_{\rm dec}. \end{aligned} \tag{E11}$$

This proves, with strict margin, that the affine zero test supplies the assignment needed by the outer quotient decoder.

Now $\boldsymbol{c} = \boldsymbol{x} - \boldsymbol{e}' \in \mathcal{C}^L_{\rm in}$ and $\boldsymbol{v} = \pi^L\boldsymbol{c}$ has support contained in the left-block mismatch set in Eq. (E11). Hence $\mathrm{wt}(\boldsymbol{v}) \le \delta_{\rm dec} n$. Since $\bar{\pi}|_{\mathcal{C}_{\rm in}} = \pi$,

$$\boldsymbol{z}' = \mathbf{H}\bar{\pi}^L(\boldsymbol{x} - \boldsymbol{e}') = \mathbf{H}\boldsymbol{v}.$$

The outer decoder returns $\widehat{\boldsymbol{v}}$ with $\boldsymbol{v} - \widehat{\boldsymbol{v}} \in \mathcal{T}_{\rm out}$, and this promised branch necessarily passes the filter. Writing $\boldsymbol{c} = \phi^L(\boldsymbol{v}) + \boldsymbol{s}$ for $\boldsymbol{s} \in \mathcal{S}^L_{\rm in}$ gives

$$\boldsymbol{x} - \widehat{\boldsymbol{x}} = \phi^L(\boldsymbol{v} - \widehat{\boldsymbol{v}}) + \boldsymbol{s} \in \phi^L(\mathcal{T}_{\rm out}) + \mathcal{S}^L_{\rm in} = \mathcal{E}_{\rm stab}.$$

Thus the output contains the class of $\boldsymbol{x}$. Once the regular families are valid, the same graphs and enumeration prove this simultaneously for every close $\boldsymbol{x}$ with syndrome $\boldsymbol{\sigma}$.

The inner data are fixed, so all affine lists and zero-agreement graphs cost $O(n)$ to construct. The cited randomized regularity stages cost $\widetilde{O}(n)$ and enumerate only M assignments. Each residual sparse matrix-vector product, outer decoder call, filter, embedding, and folding costs $O(n) + T_{\rm syn}(n)$. The cited deterministic regularity alternative costs $O(n^{3.5})$. □

Proof of Corollary 2. Apply Proposition 1, with failure budget $\eta_{\rm reg}/2$ in each sector, to $(\mathcal{S}_{\rm in}, \mathcal{C}_{\rm in}, \mathcal{T}_{\rm out}, \mathcal{D}_{\rm out}) = (\mathcal{C}_{\rm Z}^\perp, \mathcal{C}_{\rm X}, \mathcal{D}_{\rm Z}^\perp, \mathcal{D}_{\rm X})$ and to the symmetric Z pair, with $\beta = (1-R)/2 - \gamma$ and $\theta = \gamma/2$. The inner list bound and spectral inequalities were verified in the proof of Theorem 2. Lemma 4 and the direct-sum construction in that proof show that the same literal matrix $\mathbf{H}$ has a linear-time quotient decoder: it is the zero-measurement-noise specialization of the published quantum-Tanner single-shot decoder [10, 11]. This is exactly Definition 3. Lemma 6 keeps the measured checks LDPC. Each Pauli component has support contained in the full Pauli support, so the Cartesian product of the two constant output sequences covers every allowed Pauli class and remains bounded independently of n. A union bound over the two regularity events gives total failure probability at most $\eta_{\rm reg}$. □

Appendix F: Robust noisy-record composition

Work first in one nested-pair sector. A physically measured sparse record may use a local extension Ψ of π different from $\bar{\pi}$. Since both maps agree on $\mathcal{C}_{\rm in}$, there is a fixed local map $\mathbf{L}_0$ with $\Psi - \bar{\pi} = \mathbf{L}_0\mathbf{J}$. In this sector, the literal raw record map is

$$\Sigma_{raw}(\boldsymbol{x}) = \big(\mathbf{J}^L\boldsymbol{x}, \mathbf{H}\Psi^L\boldsymbol{x}\big).$$

The full CSS raw record in the main text is the concatenation of the two sector maps. Write a corrupted physical record as

$$\begin{aligned} \widetilde{\boldsymbol{y}} &= \mathbf{J}^L\boldsymbol{x} + \boldsymbol{\mu}_y, \\ \widetilde{\boldsymbol{z}}^{\rm phys} &= \mathbf{H}\Psi^L\boldsymbol{x} + \boldsymbol{\mu}_z^{\rm phys}. \end{aligned} \tag{F1}$$

The fixed sparse normalization $\mathcal{N}(\widetilde{\boldsymbol{y}}, \widetilde{\boldsymbol{z}}^{\rm phys}) = (\widetilde{\boldsymbol{y}}, \widetilde{\boldsymbol{z}})$ gives

$$\begin{aligned} \widetilde{\boldsymbol{z}} &= \widetilde{\boldsymbol{z}}^{\rm phys} - \mathbf{H}\mathbf{L}_0^L\widetilde{\boldsymbol{y}} = \mathbf{H}\bar{\pi}^L\boldsymbol{x} + \boldsymbol{D}_z, \\ \boldsymbol{D}_z &= \boldsymbol{\mu}_z^{\rm phys} - \mathbf{H}\mathbf{L}_0^L\boldsymbol{\mu}_y. \end{aligned} \tag{F2}$$

In particular, $\mathcal{N}(\Sigma_{raw}(\boldsymbol{x})) = \Sigma(\boldsymbol{x})$ when the record is exact. This is the bridge between the raw-record definition in Section II, the exact canonical fiber in Section IV, and the noisy residual below. Let $|\boldsymbol{D}_z|_{V,\oplus}$ denote the summed vertex-support weight of $\boldsymbol{D}_z$ in the standard vertex-grouped outer coordinates. The two terms of Eq. (F2) are bounded separately. The term $\boldsymbol{\mu}_z^{\mathrm{p}hys}$ already lives in those coordinates, and each of its scalar entries belongs to at most a fixed number c_1 of vertex groups, since the unreduced outer generator family has bounded coordinate degree; hence its vertex-support weight is at most $c_1|\boldsymbol{\mu}_z^{\mathrm{p}hys}|$. For the term $\mathbf{H}\mathbf{L}_0^L\boldsymbol{\mu}_y$, the fixed local map $\mathbf{L}_0$ has a support-expansion constant c_{L_0} in the direct-sum coordinates, so $|\mathbf{L}_0^L\boldsymbol{\mu}_y| \le c_{L_0}|\boldsymbol{\mu}_y|$. Every such block feeds into at most w outer checks, where w is the bounded column weight of $\mathbf{H}$, so $\mathbf{H}\mathbf{L}_0^L\boldsymbol{\mu}_y$ is nonzero on at most $wc_{L_0}|\boldsymbol{\mu}_y|$ scalar outer coordinates, each again lying in at most c_1 vertex groups. Adding the two contributions gives

$$\begin{aligned} |\boldsymbol{D}_z|_{V,\oplus} &\le c_1|\boldsymbol{\mu}_z^{\mathrm{p}hys}| + c_1 w c_{L_0}|\boldsymbol{\mu}_y| \\ &\le c_N\big(|\boldsymbol{\mu}_y| + |\boldsymbol{\mu}_z^{\mathrm{p}hys}|\big), \\ c_N &= c_1 \max\{1, wc_{L_0}\}. \end{aligned} \tag{F3}$$

All three constants c_1, w, c_{L_0} are fixed with the construction and independent of n.

Let $\mathcal{B}_y = \{\ell : (\boldsymbol{\mu}_y)_\ell \neq 0\}$, $\alpha_{\mathrm{f}} = |\mathcal{B}_y|/n$, and $h = \delta(\mathcal{C}_{\mathrm{in}}) - \beta > 0$. Since every nonzero local outcome contributes at least one scalar discrepancy, $|\mathcal{B}_y| \le |\boldsymbol{\mu}_y|$. Fix $\epsilon > 0$ with $\beta + 2\epsilon \le \delta(\mathcal{C}_{\mathrm{in}})$, and suppose every affine-fiber list of radius $\beta+\epsilon$ has size at most $K_{\mathrm{loc}} \le |\mathcal{C}_{\mathrm{in}}|$. For regularity accuracy $g > 0$ define

$$q(g, \alpha_{\mathrm{f}}) = (K_{\mathrm{loc}}+1)g + 2\alpha_{\mathrm{f}} + \frac{\lambda + (K_{\mathrm{loc}}+1)g + \alpha_{\mathrm{f}}}{h}. \tag{F4}$$

The robust regularity step uses

$$\lambda \le g\epsilon, \qquad \lambda \le g, \qquad \lambda \le g^2/2500, \tag{F5}$$

and returns an assignment satisfying

$$\mathbf{J}^L\boldsymbol{e}' = \widetilde{\boldsymbol{y}}, \qquad \Delta_{\mathrm{L}}(\boldsymbol{e}', \boldsymbol{x}+\boldsymbol{a}) \le q(g,\alpha_{\mathrm{f}}), \qquad \boldsymbol{a} = \iota^L\boldsymbol{\mu}_y. \tag{F6}$$

Because $\bar{\pi} = \pi(I - \iota\mathbf{J})$ and $\mathbf{J}\iota = I$, one has $\bar{\pi}^L\boldsymbol{a} = 0$. For this assignment,

$$\begin{aligned} \boldsymbol{c} &= \boldsymbol{x} + \boldsymbol{a} - \boldsymbol{e}' \in \mathcal{C}_{\mathrm{in}}^L, \\ \boldsymbol{v} &= \pi^L\boldsymbol{c}, \\ |\operatorname{supp}(\boldsymbol{v})| &\le q(g,\alpha_{\mathrm{f}})n, \qquad &(\text{F7}) \\ \widetilde{\boldsymbol{z}} - \mathbf{H}\bar{\pi}^L\boldsymbol{e}' &= \mathbf{H}\boldsymbol{v} + \boldsymbol{D}_z. \qquad &(\text{F8}) \end{aligned}$$

For each sector $P \in \{\mathrm{X}, \mathrm{Z}\}$, let b_P be the binary dimension of the inner quotient symbol. The actual outer interface is the literal direct sum

$$\mathcal{D}_P = \bigoplus_{j=1}^{b_P}\mathcal{D}_{P,0}, \qquad \mathbf{H}_P = \bigoplus_{j=1}^{b_P}\mathbf{H}_{P,0}, \qquad \mathcal{T}_P = \bigoplus_{j=1}^{b_P}\mathcal{T}_{P,0}.$$

For component j, the generic syndrome-error vector $\boldsymbol{D}$ in Eq. (D1) is set to $\boldsymbol{D}_{z,P}^{(j)}$; no additional data-error variable is introduced. To avoid confusing the paper's folded-block support with the reduced Hamming weight denoted $|\cdot|_R$ in [11], define

$$\begin{aligned} \mathrm{wt}_{\mathrm{b}lk}(\boldsymbol{v}) &= \big|\{\ell : \boldsymbol{v}_\ell^{(j)} \neq 0 \text{ for some } j \in [b_P]\}\big|, \\ |\boldsymbol{D}_{z,P}|_{V,\oplus} &= \sum_{j=1}^{b_P}|\boldsymbol{D}_{z,P}^{(j)}|_V. \end{aligned} \tag{F9}$$

For every component, its Gu stabilizer-reduced weight is bounded by

$$|\boldsymbol{v}^{(j)}|_{R,\mathrm{Gu}} \le \mathrm{wt}(\boldsymbol{v}^{(j)}) \le \mathrm{wt}_{\mathrm{b}lk}(\boldsymbol{v}). \tag{F10}$$

Thus the following aggregate condition is a sufficient condition for every componentwise invocation of Eq. (D1):

$$A_P\,\mathrm{wt}_{\mathrm{b}lk}(\boldsymbol{v}) + B_P|\boldsymbol{D}_{z,P}|_{V,\oplus} \le C_{\mathrm{ss},P}n. \tag{F11}$$

Let $\widehat{\boldsymbol{v}}^{(j)}$ be the component outputs. Choosing a minimizing $t_j \in \mathcal{T}_{P,0}$ in each component and using that union support is at most the sum of component supports gives

$$\begin{aligned} \min_{t\in\mathcal{T}_P}\mathrm{wt}_{\mathrm{b}lk}(\boldsymbol{v} - \widehat{\boldsymbol{v}} - t) &\le \sum_{j=1}^{b_P}\min_{t_j\in\mathcal{T}_{P,0}}\mathrm{wt}(\boldsymbol{v}^{(j)} - \widehat{\boldsymbol{v}}^{(j)} - t_j) \\ &\le \beta_{\mathrm{ss},P}\sum_{j=1}^{b_P}|\boldsymbol{D}_{z,P}^{(j)}|_V \\ &= \beta_{\mathrm{ss},P}|\boldsymbol{D}_{z,P}|_{V,\oplus}. \end{aligned} \tag{F12}$$

There is no missing factor b_P: the output union is bounded by a sum, and the syndrome-noise norm was defined as the same sum. No row-equivalent replacement of the vertex-local checks is used.

For the final two-sector composition, the unified parameter choice above has already fixed $\theta, \epsilon, h_*, K_*, g$, and λ. The inner length and the expander family were then chosen once with $\lambda \le \lambda_*$, so the robust and single-shot inequalities in Eqs. (D4) and (D5) require no second construction choice. Put

$$q_*(g,\alpha_{\mathrm{f}}) = (K_*+1)g + 2\alpha_{\mathrm{f}} + \frac{\lambda + (K_*+1)g + \alpha_{\mathrm{f}}}{h_*}. \tag{F13}$$

The parameter calculation above gives $K_{\mathrm{loc}} \le K_*$ and $h \ge h_*$, so $q(g,\alpha_{\mathrm{f}}) \le q_*(g,\alpha_{\mathrm{f}})$. For sector P, write $c_{N,P}$ for the fixed normalization constant c_N obtained from Eq. (F3) in that sector.

If $j_P = \operatorname{rank} J_P$ and $m_{P,0}(n)$ is the number of literal outcomes in one binary outer component, with $m_{P,0}(n) \le c_{H,P}n$, then

$$\begin{aligned} m_n &= n(j_X + j_Z) + b_X m_{X,0}(n) + b_Z m_{Z,0}(n) \le c_m n, \\ c_m &= j_X + j_Z + b_X c_{H,X} + b_Z c_{H,Z}. \end{aligned} \tag{F14}$$

The raw discrepancy for a hidden error is $\tau(\boldsymbol{e}) = d_H(\widetilde{\boldsymbol{\sigma}}, \Sigma_{raw}(\boldsymbol{e}))/m_n$. One valid threshold and final residual constant are

$$\tau_0 = \min_{P\in\{\mathrm{X},\mathrm{Z}\}} \frac{C_{\mathrm{ss},P} - A_P q_*(g,0)}{c_m\left[A_P(2+h_*^{-1}) + B_P c_{N,P}\right]}, \tag{F15}$$

$$\kappa = c_m \sum_{P\in\{\mathrm{X},\mathrm{Z}\}} \left(c_{\iota,P} + c_{\phi,P}\beta_{\mathrm{ss},P}c_{N,P}\right). \tag{F16}$$

Here $c_{\iota,P}, c_{\phi,P} \leq d$ are fixed folded-support Lipschitz constants. Equations (D4)–(D5) make every numerator in Eq. (F15) positive.

Lemma 7 (Robust affine-fiber regularity)**.** *Assume $\beta + 2\epsilon \leq \delta(\mathcal{C}_{\mathrm{in}})$ and $\lambda \leq \min\{g\epsilon, g, g^2/2500\}$. On a valid regularity event, the zero-agreement enumeration contains, simultaneously for every compatible hidden error $\boldsymbol{x}$ with $\Delta_{\mathrm{R}}(\boldsymbol{x},0) \leq \beta$, an assignment $\boldsymbol{e}'$ satisfying*

$$\mathbf{J}^L\boldsymbol{e}' = \widetilde{\boldsymbol{y}}, \qquad \Delta_{\mathrm{L}}(\boldsymbol{e}', \boldsymbol{x}+\boldsymbol{a}) \leq q_0 + (2+h^{-1})\alpha_{\mathrm{f}}. \tag{F17}$$

Proof of Lemma 7. Let $B \subseteq R$ be the right blocks on which $\boldsymbol{x}$ vanishes, so $|B| \geq (1-\beta)n$. Outside B_y, the true local error lies in the observed affine fiber. As in Eq. (E7), set $\widetilde{\boldsymbol{r}}_\ell = \iota(\widetilde{\boldsymbol{y}}_\ell)$ and $\widetilde{\boldsymbol{g}}_\ell = -\widetilde{\boldsymbol{r}}_\ell$. Write each padded observed affine candidate uniquely as $\boldsymbol{e}_{\ell,i} = \widetilde{\boldsymbol{r}}_\ell + \boldsymbol{c}_{\ell,i}$ with $\boldsymbol{c}_{\ell,i} \in C$. The translated words $\boldsymbol{c}_{\ell,i}$ form ordinary C-lists around $\widetilde{\boldsymbol{g}}$, and the zero graphs are their exact agreement graphs. Run the exhaustive factor/signature-net construction from Lemma 5, including $\mathbf{1}_L$ in the common function family and using edge accuracy g. This construction depends only on the observed lists and graphs. Every translated assignment $\boldsymbol{c}'$ returned by it is mapped back as $\boldsymbol{e}' = \widetilde{\boldsymbol{r}} + \boldsymbol{c}'$, which immediately gives $\mathbf{J}^L\boldsymbol{e}' = \widetilde{\boldsymbol{y}}$.

Let $L_{\mathrm{miss}} \subseteq L \setminus B_y$ be the vertices where its local weight exceeds $\beta+\epsilon$. Every nonzero edge of $\boldsymbol{x}$ enters $R\setminus B$, so expander mixing gives

$$(\beta+\epsilon)d|L_{\mathrm{miss}}| \leq |E_G(L_{\mathrm{miss}}, R\setminus B)| \leq \beta d|L_{\mathrm{miss}}| + \lambda dn.$$

Hence $|L_{\mathrm{miss}}| \leq \lambda n/\epsilon \leq gn$. After padding every list to K_{loc} elements, define the disjoint true-index sets

$$A_i = \{\ell \in L : \boldsymbol{e}_{\ell,i} = \boldsymbol{x}_\ell\}.$$

Equality of the syndromes shows $A_i \cap B_y = \varnothing$, and these sets cover at least $(1-g-\alpha_{\mathrm{f}})n$ vertices. On each A_i, the identity

$$E_{\mathcal{G}}(A_i, B) = E_{\mathcal{H}_i}(A_i, B) \tag{F18}$$

remains exact because $\boldsymbol{x}$ is zero on the edges incident to B.

If $n < 4\cdot 2^P/\eta$, the algorithm uses the brute-force branch of Lemma 5. Among all K_{loc}^n assignments it includes one that chooses the true index on every $\bigcup_i A_i$; that assignment obeys

$$\Delta_{\mathrm{L}}(\boldsymbol{e}', \boldsymbol{x}+\boldsymbol{a}) \leq g + \alpha_{\mathrm{f}} \leq q(g,\alpha_{\mathrm{f}}).$$

Its enumeration size is a fixed-parameter constant and it covers every hidden error simultaneously. We may therefore assume below that $n \geq 4\cdot 2^P/\eta$, where the signature-net branch applies.

For the signature of these A_i, the exhaustive net contains disjoint approximants S_i with function-family distance at most its tolerance $\eta \leq g$. Regularity loses at most gdn edges for each graph, and uncovered true vertices lose at most $(g+\alpha_{\mathrm{f}})dn$ more. Hence

$$\sum_i |E_{\mathcal{H}_i}(S_i, B)| \geq d|B| - (K_{\mathrm{loc}}+1)gdn - \alpha_{\mathrm{f}}dn. \tag{F19}$$

The same-fiber distance argument is unavailable on B_y. Define the three disjoint sets

$$T = \bigcup_i (S_i \cap A_i),$$
$$Q = \bigcup_i (S_i \cap B_y),$$
$$W = \bigcup_i \big(S_i \setminus (A_i \cup B_y)\big),$$

and write $w = |W|/n$. Since the S_i are disjoint, $|T\cup Q| \leq n - |W|$. On T, zero agreement with $\boldsymbol{x}$ is exact. On Q we use only $H_i \subseteq G$, making no same-fiber claim. On W, a selected candidate and $\boldsymbol{x}$ lie in the same uncorrupted fiber but are distinct, so their difference is a nonzero word of C and has weight at least $\delta(C)d$. Consequently

$$\begin{aligned}\sum_i |E_{\mathcal{H}_i}(S_i, B)| &\leq |E_G(T\cup Q, B)| + (1-\delta(C))d|W| \\ &\leq (1-w)d|B| + \lambda dn + (1-\delta(C))wdn.\end{aligned}$$

Using $|B|/n \geq 1-\beta$ gives

$$\sum_i |E_{\mathcal{H}_i}(S_i, B)| \leq d|B| - wdn\big(\delta(C)-\beta\big) + \lambda dn. \tag{F20}$$

Thus the corrupted selected vertices are absorbed into the single mixing term on $T\cup Q$ rather than being assigned a fictitious code-distance bound. Combining Eqs. (F19) and (F20) gives

$$w\big(\delta(C)-\beta\big) \leq \lambda + (K_{\mathrm{loc}}+1)g + \alpha_{\mathrm{f}}. \tag{F21}$$

Because $\mathbf{1}_L$ belongs to the signature family, $|S_i| \geq |A_i| - gn$. The final mismatch consists of the at most $g+\alpha_{\mathrm{f}}$ uncovered true vertices, $K_{\mathrm{loc}}g$ of size discrepancy, the fraction w, and one additional charge for selected vertices in B_y; therefore

$$\Delta_{\mathrm{L}}(\boldsymbol{e}', \boldsymbol{x}+\boldsymbol{a}) \leq (K_{\mathrm{loc}}+1)g + 2\alpha_{\mathrm{f}} + w.$$

Substituting Eq. (F21) gives exactly Eq. (F4). On a correctly selected uncorrupted block, $\boldsymbol{e}'_\ell = \boldsymbol{x}_\ell$ and $\boldsymbol{a}_\ell = 0$; on a corrupted block the comparison is to $\boldsymbol{x}_\ell + \boldsymbol{a}_\ell$. This proves Eq. (F6). The random regularity event depends only on the observed zero-agreement graphs, which establishes simultaneous coverage. $\square$

Proof of Theorem 1. Fix one observed pair of raw CSS records and run the regularity-family construction in each sector with failure probability at most $\eta_{\rm alg}/2$. Condition on the two success events. They depend only on the observed zero-agreement graphs and therefore hold simultaneously for all compatible hidden errors.

For a compatible error $\boldsymbol{e} = (\boldsymbol{x}_X, \boldsymbol{x}_Z)$, define the sector noise variables as in Eq. (F1), and set

$$\alpha_{P,\mathrm{f}} = |\{\ell : (\boldsymbol{\mu}_{y,P})_\ell \neq 0\}|/n.$$

Lemma 7 supplies an enumerated $\boldsymbol{e}'_P$ with

$$\Delta_{\mathrm{L}}(\boldsymbol{e}'_P, \boldsymbol{x}_P + \iota^L_P \boldsymbol{\mu}_{y,P}) \leq q_*(g, \alpha_{P,\mathrm{f}}). \tag{F22}$$

For $\boldsymbol{v}_P = \pi^L_P(\boldsymbol{x}_P + \iota^L_P \boldsymbol{\mu}_{y,P} - \boldsymbol{e}'_P)$, Eqs. (F7) and (F8) give

$$\begin{aligned} \mathrm{wt}_{\mathrm{b}lk}(\boldsymbol{v}_P) &\leq q_*(g, \alpha_{P,\mathrm{f}})n, \\ \widetilde{\boldsymbol{z}}_P - \mathbf{H}_P \bar{\pi}^L_P \boldsymbol{e}'_P &= \mathbf{H}_P \boldsymbol{v}_P + \boldsymbol{D}_{z,P}. \end{aligned} \tag{F23}$$

If $\tau(\boldsymbol{e}) \leq \tau_0$, then Eq. (F14) and the definition of the raw Hamming discrepancy imply

$$\alpha_{P,\mathrm{f}} \leq c_m \tau(\boldsymbol{e}), \qquad |\boldsymbol{D}_{z,P}|_{V,\oplus}/n \leq c_{N,P} c_m \tau(\boldsymbol{e}).$$

Consequently

$$q_*(g, \alpha_{P,\mathrm{f}}) \leq q_*(g, 0) + c_m \tau(\boldsymbol{e})(2 + h_*^{-1}),$$

and Eq. (F15) makes the joint small-set condition (F11) valid in both sectors. Apply the componentwise outer single-shot decoder to Eq. (F23) and output

$$\widehat{\boldsymbol{x}}_P = \boldsymbol{e}'_P + \phi^L_P(\widehat{\boldsymbol{v}}_P) \tag{F24}$$

for every enumerated assignment. There is no exact-syndrome consistency filter, since such a filter would reject the promised branch when the record is corrupted.

Writing $\boldsymbol{x}_P + \iota^L_P \boldsymbol{\mu}_{y,P} - \boldsymbol{e}'_P = \phi^L_P(\boldsymbol{v}_P) + \boldsymbol{s}_{0,P}$ and using Eq. (F12) gives

$$\begin{aligned} \min_{\boldsymbol{s}_P \in E_{S,P}} \mathrm{wt}_{\mathrm{R}}(\boldsymbol{x}_P - \widehat{\boldsymbol{x}}_P - \boldsymbol{s}_P) \leq c_{\iota,P}|\boldsymbol{\mu}_{y,P}| \\ + c_{\phi,P}\beta_{\mathrm{ss},P}|\boldsymbol{D}_{z,P}|_{V,\oplus}. \end{aligned} \tag{F25}$$

Take the Cartesian product of the two constant sector lists. The support of a full Pauli residual is the union of its sector supports and is at most their sum. Equation (F3), the definition of $\tau(\boldsymbol{e})$, and Eq. (F16) then give $d_{\mathcal{S}}(\boldsymbol{e}, \widehat{\boldsymbol{e}}_j) \leq \kappa \tau(\boldsymbol{e}) n$.

For fixed parameters, affine-list construction and all sparse transforms cost $O(n)$. Regularity costs $\widetilde{O}(n)$ and returns only a constant number of assignments. Each direct-sum outer call is linear, so the total running time is $\widetilde{O}(n)$ and the product list has constant size L. A union bound over the two regularity events gives success probability at least $1 - \eta_{\rm alg}$. Rate, locality, and the data radius are those of the single capacity family fixed in Appendix D. This proves the theorem. □

Appendix G: Converse and selection barriers

For a prime power Q, write

$$V_Q(N,t) = \sum_{j=0}^{t} \binom{N}{j} (Q-1)^j$$

for the Q-ary Hamming-ball volume. Let

$$\begin{aligned} h_2(x) &= -x\log_2 x - (1-x)\log_2(1-x), \\ H_Q(x) &= \frac{h_2(x) + x\log_2(Q-1)}{\log_2 Q}, \end{aligned}$$

with the usual continuous endpoint convention.

Proof of Theorem 4. Choose any transversal $\mathcal{T}$ of $\mathcal{C}_{\mathrm{X}}/\mathcal{C}_{\mathrm{Z}}^{\perp}$. For $y \in \mathbb{F}_Q^N$ and $u \in \mathcal{T} \cap B(y,t)$, the error $e_u = u - y$ has weight at most t. All such errors have the same syndrome, and distinct transversal elements give distinct stabilizer classes. One residual center contains at most $V_Q(N,r)$ such classes, because every covered class has a representative in a radius-r physical ball. Hence $|\mathcal{T} \cap B(y,t)| \leq LV_Q(N,r)$ for every y. Double counting pairs (u,y) gives

$$\sum_y |\mathcal{T} \cap B(y,t)| = |\mathcal{T}|V_Q(N,t) = Q^{K_Q}V_Q(N,t),$$

Thus $Q^{K_Q}V_Q(N,t) \leq LQ^N V_Q(N,r)$, as claimed. □

Proof of Corollary 3. Strict concavity places $H_2(p)$ above the chord $2p$ on $0 < p < 1/2$, proving $\gamma_*(R) > 0$. Set $\beta = (1-R)/2 - \gamma$. The chosen γ gives $R + H_2(\beta) > 1$. Apply Theorem 4 with $r = 0$ and the standard binary type estimate $V_2(N, \beta N) = 2^{N(H_2(\beta) - o(1))}$ to obtain

$$L_N \geq 2^{N[R + H_2(\beta) - 1 - o(1)]}.$$

Setting measurement noise to zero and restricting to one CSS sector shows that randomization, redundant checks, and full-Pauli output do not change the count. □

Proof of Theorem 5. The quotient balls $B_r(e_i;\mathcal{S})$ are disjoint. Any deterministic output lies in at most one; for a random output C,

$$\sum_i \Pr\{C \in B_r(e_i;\mathcal{S})\} = \mathbb{E}\sum_i \mathbf{1}\{C \in B_r(e_i;\mathcal{S})\} \leq 1.$$

At least one success probability is no larger than $1/M$. □

Proposition 2 (Minimum-logical witness)**.** *Let ℓ be a logical Pauli representative attaining the code distance d, the minimum quotient weight over all nontrivial logical classes. There are errors e_0, e_1 with the same exact syndrome, each of weight at most $\lceil d/2 \rceil$, and $d_{\mathcal{S}}(e_0, e_1) = d$. Consequently, whenever the allowed data radius contains $\lceil d/2 \rceil$ and $r < d/2$, every syndrome-only selector has worst-case failure probability at least $1/2$.*

Proof. Partition the support of ℓ into A and B of size at most $\lceil d/2\rceil$, and set $e_0 = \ell|_A$ and $e_1 = -\ell|_B$. Then $e_0 - e_1 = \ell$, so the syndromes agree and their quotient distance is d. Apply Theorem 5 with $M = 2$. □

Proposition 3 (Trusted-information lower bound)**.** *Let J be uniform on the M packed errors in Theorem 5. Let K collect decoder coins and secret randomness independent of J at fixed $\widetilde{\boldsymbol{\sigma}}$, and let T be an additional transcript whose distribution may depend on J. If a decoder using $(\widetilde{\boldsymbol{\sigma}}, K, T)$ selects the correct residual ball with error at most ε, then*

$$I(J;T \mid K,\widetilde{\boldsymbol{\sigma}}) \geq \log_2 M - h_2(\varepsilon) - \varepsilon \log_2(M-1). \quad \text{(G1)}$$

Proof. A successful center identifies at most one packed hypothesis and therefore induces an estimator of J. Conditional Fano's inequality [22], together with uniform J and the independence of K, gives Eq. (G1) directly. □

1. Why a global syndrome section is not the algorithm

Proposition 4 (AEL syndrome-section equivalence)**.** *Assume the outer symbol dimension is fixed and choose a full-row-rank outer check matrix $\mathbf{H}_{\mathrm{X}}^{\mathrm{out}}$ with $O(n)$ nonzero entries. Use the canonical check basis consisting of its rows lifted through ϕ_{Z}, together with a fixed basis of $\mathcal{C}_{\mathrm{X}}^{\perp}$ on every left block. Then*

$$\mathcal{E}_{\mathrm{X}}^{\perp} = \phi_{\mathrm{Z}}(\mathcal{D}_{\mathrm{X}}^{\perp}) \oplus (\mathcal{C}_{\mathrm{X}}^{\perp})^{L}$$

and the same identity holds after folding. Let $T_{F,X}(n)$ be the worst-case time for an algorithm mapping every consistent final X syndrome to one ambient representative, and define $T_{D,X}(n)$ analogously for the outer code. Then

$$\begin{aligned} T_{F,X}(n) &\leq T_{D,X}(n) + O(n), \\ T_{D,X}(n) &\leq T_{F,X}(n) + O(n). \end{aligned} \quad \text{(G2)}$$

The same equivalence holds in the Z sector.

Proof. The displayed sum is direct because $\operatorname{im}\phi_{\mathrm{Z}} = \mathcal{W}_{\mathrm{Z}}$ and $\mathcal{W}_{\mathrm{Z}} \cap \mathcal{C}_{\mathrm{X}}^{\perp} = \{0\}$. Removing dependent rows from an LDPC check matrix preserves sparsity, so the stated canonical basis is available.

Let $\mathbf{J}_{\mathrm{X}}$ be a full-row-rank check matrix for the fixed inner code $\mathcal{C}_{\mathrm{X}}$. Define Ψ_{X} by

$$\langle \boldsymbol{q}, \Psi_{\mathrm{X}}(\boldsymbol{v})\rangle = \langle \phi_{\mathrm{Z}}(\boldsymbol{q}), \boldsymbol{v}\rangle$$

for every outer dual symbol $\boldsymbol{q}$. Duality preservation gives $\Psi_{\mathrm{X}}\phi_{\mathrm{X}} = I$, while on $\mathcal{C}_{\mathrm{X}}$ the map Ψ_{X} is the projection along $\mathcal{C}_{\mathrm{X}} = \phi_{\mathrm{X}}(\mathcal{U}) \oplus \mathcal{C}_{\mathrm{Z}}^{\perp}$. The canonical syndrome is therefore

$$\Sigma_{\mathrm{X}}(\boldsymbol{v}) = \left((\mathbf{J}_{\mathrm{X}}\boldsymbol{v}_\ell)_{\ell\in\mathcal{L}}, \mathbf{H}_{\mathrm{X}}^{\mathrm{out}}(\Psi_{\mathrm{X}}(\boldsymbol{v}_\ell))_{\ell\in\mathcal{L}}\right). \quad \text{(G3)}$$

Given a final syndrome, apply a fixed right inverse of $\mathbf{J}_{\mathrm{X}}$ independently on every block to obtain $\boldsymbol{a}$ with the prescribed local components. Subtract the outer component of $\Sigma_{\mathrm{X}}(\boldsymbol{a})$, call an outer syndrome section on the resulting residual outer syndrome, and embed its output $\boldsymbol{u}$ through ϕ_{X}. Equation (G3) and $\Psi_{\mathrm{X}}\phi_{\mathrm{X}} = I$ show that $\boldsymbol{a} + \phi_{\mathrm{X}}(\boldsymbol{u})$ has the requested final syndrome. Every operation except the outer section costs $O(n)$.

Conversely, call a final section on a syndrome with zero local components and an arbitrary consistent outer component. Its output lies blockwise in $\mathcal{C}_{\mathrm{X}}$. Projecting through $\mathcal{C}_{\mathrm{X}} = \phi_{\mathrm{X}}(\mathcal{U}) \oplus \mathcal{C}_{\mathrm{Z}}^{\perp}$ gives an outer word with the requested syndrome by Eq. (G3). Folding and unfolding cost $O(n)$. This proves Eq. (G2); the Z sector is symmetric. □

Proposition 5 (Dense linear-section barrier)**.** *Let $\mathbf{H} \in \mathbb{F}_q^{m\times n}$ have full row rank and row weight at most w, let $\mathcal{C} = \ker \mathbf{H}$, and let $\mathcal{S} \subseteq \mathcal{C}$. Put $\ell = \dim(\mathcal{C}/\mathcal{S})$ and $d_Q = \min_{\boldsymbol{c}\in\mathcal{C}\setminus\mathcal{S}} \mathrm{wt}(\boldsymbol{c})$. Every linear right inverse $\mathbf{R} \in \mathbb{F}_q^{n\times m}$ with $\mathbf{HR} = I_m$ satisfies*

$$\mathrm{nnz}(\mathbf{R}) \geq \frac{\ell(d_Q - 1)}{w}. \quad \text{(G4)}$$

In particular, a positive-rate, linear-sector-distance LDPC family has no linear syndrome right inverse with $O(n)$ nonzero entries.

Proof. Set $\mathbf{Q} = I - \mathbf{RH}$. Then $\mathbf{HQ} = 0$ and $\mathbf{Q}$ restricts to the identity on $\mathcal{C}$. Consequently the composite $\mathbb{F}_q^n \xrightarrow{\mathbf{Q}} \mathcal{C} \to \mathcal{C}/\mathcal{S}$ has rank ℓ, so at least ℓ columns $\mathbf{Q}\boldsymbol{e}_j$ lie in $\mathcal{C} \setminus \mathcal{S}$ and have weight at least d_Q. Write $\boldsymbol{r}_i = \mathbf{R}\boldsymbol{f}_i$ for the ith standard syndrome vector and $a_i = \mathrm{wt}(\boldsymbol{r}_i)$. For each selected column,

$$\mathbf{Q}\boldsymbol{e}_j = \boldsymbol{e}_j - \sum_{i:\mathbf{H}_{ij}\neq 0} \mathbf{H}_{ij}\boldsymbol{r}_i \quad \Longrightarrow \quad d_Q - 1 \leq \sum_{i:H_{ij}\neq 0} a_i.$$

Summing over ℓ selected columns counts each a_i at most the weight of row i, hence at most w times. Therefore $\ell(d_Q - 1) \leq w\sum_i a_i = w\,\mathrm{nnz}(\mathbf{R})$, which proves Eq. (G4). □

2. Why a direct linear tag does not select a residual center

Lemma 8 (Low-weight preimages of a linear tag)**.** *Let $\mathbf{H} : \mathbb{F}_q^{2n} \to \mathbb{F}_q^{\ell}$ be a linear stabilizer-tag syndrome map of rank a. Every syndrome in $\operatorname{im}\mathbf{H}$ has a Pauli preimage of physical support at most a. Consequently,*

$$\mathbf{H}(B_r) = \operatorname{im}\mathbf{H} \qquad \textit{whenever } r \geq a. \quad \text{(G5)}$$

Proof. The $2n$ columns of $\mathbf{H}$ are the syndromes of the weight-one X- and Z-basis Paulis and span $\operatorname{im}\mathbf{H}$. Choose a independent columns. Every image syndrome is a linear combination of these columns and is therefore realized by a Pauli using at most a physical coordinates. If an X column and a Z column share a qudit, the physical support only decreases. This proves Eq. (G5). □

Proposition 6 (Direct logical-tag locality barrier)**.** *Let a base stabilizer code have stabilizer S_{base} and distance d. Suppose a direct logical-layer tag enlarges it to $S_{\text{base}} \subseteq S_{\text{comp}} \subseteq N(S_{\text{base}})$ with $\dim(S_{\text{comp}}/S_{\text{base}}) = \ell$. Every generating set of S_{comp} contains at least ℓ elements outside S_{base}, and each such element has physical weight at least d. Thus this architecture has no bounded-weight stabilizer generating set when $d = \Omega(n)$ and $\ell > 0$.*

Proof. The images of any generating set in the quotient $S_{\text{comp}}/S_{\text{base}}$ must span its ℓ dimensions, so at least ℓ generators have nonzero quotient image. Each such generator belongs to $N(S_{\text{base}}) \setminus S_{\text{base}}$ and therefore represents a nontrivial logical Pauli of the base code. By the definition of code distance, every physical representative of that class has weight at least d. $\square$

The last two results concern only direct linear syndrome tags and direct logical-layer stabilizer composition. They do not exclude hidden-key average-case tests, blockwise traps, subsystem redesigns, interactive transcripts, or a new residual-cleaning stage.

DATA AVAILABILITY

No data were created or analyzed in this theoretical study.

References

[1] D. Leung and G. Smith, Communicating over adversarial quantum channels using quantum list codes, IEEE Transactions on Information Theory **54**, 883 (2008), arXiv:quant-ph/0605086.

[2] R. Arvind, N. Bansal, D. E. Koh, T. Haug, and K. Bharti, Quantum error correction in adversarial regimes (2025), arXiv:2509.08943 [quant-ph].

[3] T. Bergamaschi, L. Golowich, and S. Gunn, Approaching the quantum singleton bound with approximate error correction, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (2024) pp. 1507–1516, arXiv:2212.09935.

[4] G. L. Matthews and J. Shapiro, List-decodable folded quantum hermitian codes (2026), arXiv:2605.10534 [cs.IT].

[5] T. Bergamaschi, F. G. Jeronimo, T. Mittal, S. Srivastava, and M. Tulsiani, List decodable quantum ldpc codes (2024), arXiv:2411.04306 [cs.IT].

[6] F. G. Jeronimo, T. Mittal, S. Srivastava, and M. Tulsiani, Explicit codes approaching generalized singleton bound using expanders, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (2025) pp. 843–854, arXiv:2502.07308.

[7] Y. Chen and Z. Zhang, Explicit folded reed–solomon and multiplicity codes achieve relaxed generalized singleton bounds, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (2025) pp. 1–12, arXiv:2408.15925.

[8] S. Srivastava and M. Tulsiani, List decoding expander-based codes up to capacity in near-linear time, in *2025 IEEE 66th Annual Symposium on Foundations of Computer Science* (2025) pp. 1465–1487, arXiv:2504.20333.

[9] H. Bombin, Single-shot fault-tolerant quantum error correction, Physical Review X **5**, 031043 (2015), arXiv:1404.5504 [quant-ph].

[10] A. Leverrier and G. Zémor, Efficient decoding up to a constant fraction of the code length for asymptotically good quantum codes, in *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms* (2023) pp. 1216–1244, arXiv:2206.07571.

[11] S. Gu, E. Tang, L. Caha, S. H. Choe, Z. He, and A. Kubica, Single-shot decoding of good quantum LDPC codes, Communications in Mathematical Physics **405**, 85 (2024), arXiv:2306.12470 [quant-ph].

[12] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, PRX Quantum **2**, 040101 (2021), arXiv:2103.06309 [quant-ph].

[13] P. Panteleev and G. Kalachev, Asymptotically good quantum and locally testable classical LDPC codes, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (2022) pp. 375–388, arXiv:2111.03654.

[14] A. Leverrier and G. Zémor, Quantum tanner codes, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science* (2022) pp. 872–883, arXiv:2202.13641.

[15] M. B. Hastings, J. Haah, and R. O'Donnell, Fiber bundle codes: Breaking the $n^{1/2}$ polylog(n) barrier for quantum LDPC codes, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (2021) pp. 1276–1288, arXiv:2009.03921 [quant-ph].

[16] 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 (2024), arXiv:2308.07915 [quant-ph].

[17] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, Physical Review A **54**, 1098 (1996), arXiv:quant-ph/9512032.

[18] A. M. Steane, Multiple-particle interference and quantum error correction, Proceedings of the Royal Society of London A **452**, 2551 (1996), arXiv:quant-ph/9601029.

[19] O. Alrabiah, V. Guruswami, and R. Li, AG codes have no list-decoding friends: Approaching the Generalized Singleton Bound requires exponentially large alphabets, IEEE Transactions on Information Theory **71**, 2443 (2025), arXiv:2308.13424 [cs.IT].

[20] G. Kalachev and P. Panteleev, Two-sided robustly testable codes (2022), arXiv:2206.09973 [cs.IT].

[21] A. W. Marcus, D. A. Spielman, and N. Srivastava, Interlacing families iv: Bipartite ramanujan graphs of all sizes, SIAM Journal on Computing **47**, 2488 (2018), arXiv:1505.08010.

[22] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley, Hoboken, NJ, 2006).