Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Sharing flag circuits across code blocks: an exact local criterion for joint logical measurements

Annan Chen

annan.chen24@imperial.ac.uk

Imperial College London https://orcid.org/0009-0000-1308-0466

Research Article

Keywords: quantum error correction, fault-tolerant quantum computation, flag circuits, joint logical measurements, ancilla sharing

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11759691/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Sharing flag circuits across code blocks: an exact local criterion for joint logical measurements

Annan Chen[1, *]

[1] *Department of Physics, Imperial College London, London, United Kingdom*

Joint logical measurements can reuse a single flag circuit across code blocks while retaining its auxiliary qubits, flag gates and correction table. We establish an exact local criterion for this reuse. For layered CNOT circuits under full Pauli faults, we prove that a flag circuit can be shared among any number of blocks, with any gate order inside each data layer, if and only if on a single block it bounds the full Pauli weight of the residual error by the number of faults; every block of the shared circuit then obeys the same bound. Two blocks suffice to witness any violation. We also determine the exact fault order guaranteed by a certificate that bounds the X weight alone. The criterion extends the Chao–Reichardt construction to joint measurements at arbitrary distance and certifies the single-round mixed-Pauli circuit of a five-qubit-code teleportation experiment. In a complete two-block Steane-code protocol with noisy error correction and standard decoding, a shared one-flag circuit removes every first-order failure, whereas the bare circuit fails with probability $8p/5 + O(p^2)$ at physical error rate p. Local certification thus provides a reusable design rule for joint logical measurements with auxiliary overhead independent of the number of blocks.

I. INTRODUCTION

A joint measurement of logical Pauli operators can entangle encoded qubits while keeping them encoded. It is the basic operation of encoded Bell-state preparation, teleportation and measurement-based computation [1, 2] and underlies lattice surgery [3]. Flag circuits perform measurements of moderate weight with few auxiliary qubits: a measurement qubit couples to the data, and flag qubits detect the faults on it that would spread to several data qubits [4, 5]. They have also been used to measure logical operators [1, 6]. Trapped-ion experiments have demonstrated flag-based fault tolerance and logical teleportation [7–9] and, recently, a flagged joint measurement between two five-qubit-code blocks [10].

The joint-measurement circuit of that experiment has the structure of an economical construction, which we call *sharing* [Fig. 1(a)]. Each data gate of a single-block flag circuit becomes a layer of m copies, one to the corresponding qubit of each of the m blocks; the auxiliary qubits, flag gates and correction table are kept. The auxiliary overhead is then independent of the number of blocks. The one-flag circuit of Fig. 1(c), for example, measures a product of weight-three operators on m blocks with two auxiliary qubits and $3m + 2$ CNOTs. The measurement qubit visits the data in sequence, so depth grows with m. The question is whether protection certified once, on one block, carries over: when can a *given* auxiliary circuit and its correction table be kept unchanged as blocks are added, whatever the order of the copied gates?

Sharing can fail because only the product is measured. In a single-block measurement of an operator M, the errors E and EM act identically on the projected state, and a fault-tolerance certificate bounds the lighter of the two. A joint measurement provides this freedom only for the product $M_1 \cdots M_m$: all blocks share one choice. A fault inside a layer can make two blocks need opposite choices. In the *bare* (unflagged) weight-three circuit, an X fault on the measurement qubit between the two copies of the second data gate leaves X_3 on one block and X_2X_3 on the other, and either common choice leaves weight two on one block [Fig. 1(b)]. On a single block the same circuit bounds each Pauli component by one: a residual such as Z_2X_3 has X weight one and Z weight one, which separate X and Z recovery corrects in a distance-three Calderbank–Shor–Steane (CSS) code [11, 12], while its full Pauli weight is two.

This distinction decides the question. We prove that, up to a fault order t, the shared circuit leaves at most s errors on every block after any $s \le t$ faults, for every number of blocks and every layer order, *if and only if the single-block circuit bounds the full Pauli weight of its residual error by the number of faults* (Theorem 2). Two blocks suffice to witness every violation. One local certification therefore covers the size and ordering of the shared measurement.

Full-Pauli guarantees carry over to the shared circuit: the zero-flag condition of Chamberland and Beverland [5, Def. 6] and the flag-selected corrections of Chao and Reichardt [13, Def. 2] give their corresponding shared guarantees. For a certificate that bounds the X weight alone, as used in CSS criteria [14], the implied Pauli-weight order is exactly $q = \min\{t, \lfloor w/2 \rfloor - 1\}$, where w is the weight of the measured operator on one block (Corollary 9, Fig. 3). Thus $q = t$ for $w \ge 2t + 2$, and $q = t - 1$ at $w = 2t + 1$, the weight of a minimum-weight logical operator at distance $2t + 1$.

We apply the criterion at three scales.

(i) *Arbitrary distance.* The Chao–Reichardt construction [13] is shared with its decoder unchanged and at most $d+1$ auxiliary qubits at distance d, for any number of blocks (Section V). At weights three and

* annan.chen24@imperial.ac.uk; ORCID: 0009-0000-1308-0466

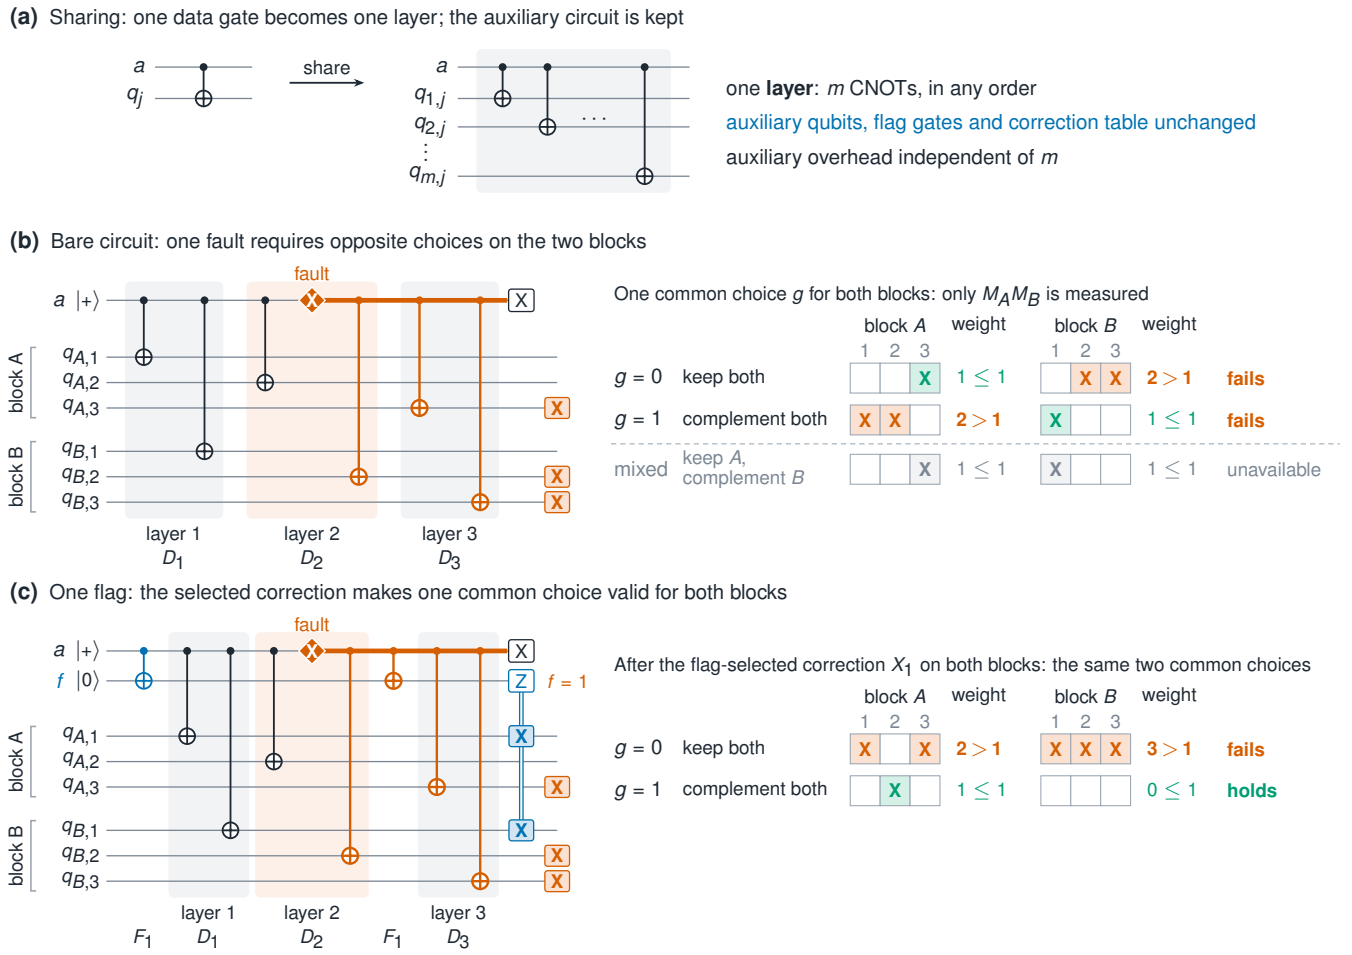


FIG. 1. Sharing a flag circuit. (a) Every data gate becomes a layer of m gates; the auxiliary qubits, flag gates and correction table are kept. (b) The bare weight-three circuit $D_1D_2D_3$ on two blocks. An X fault on the measurement qubit a inside layer 2 is copied to every later target (orange) and leaves X_3 on A and X_2X_3 on B. Keeping both residuals ($g = 0$) or multiplying both by their measured operators ($g = 1$) leaves weight two on one block. The mixed choice, which would satisfy both blocks, needs M_B by itself, whereas the circuit measures M_AM_B. (c) The one-flag circuit $F_1D_1D_2F_1D_3$ with the same fault. The raised flag selects the correction X_1 on both blocks (blue). Of the same two common choices, $g = 1$ now leaves weights one and zero. The gate sequence is printed below each circuit: D_j is the layer of CNOTs from a to data qubit j of every block and F_1 a CNOT from a to the flag.

five, circuits with one and four flags give the one- and two-fault guarantees with far fewer gates, and four flags are the minimum at weight five within a classified circuit family (Section VI).

(ii) *An experimental primitive.* The single-round one-flag circuit of Freund *et al.* [10, Fig. 12] satisfies the single-block condition in its mixed-Pauli form. Its propagation bound therefore holds for every number of blocks and every layer order (Section VII).

(iii) *A complete protocol.* In a two-block Steane-code measurement with repeated readout, noisy quantum error correction (QEC) and standard blockwise decoding, the shared one-flag circuit tolerates every single fault, whereas the bare circuit fails with probability $8p/5$ to first order in the physical error rate p. At $p = 3 \times 10^{-5}$ without idle errors the flag lowers the failure probability sevenfold (Section VIII).

The criterion is a code-independent propagation guarantee. Because it bounds the error on each block separately, the shared measurement composes with standard blockwise error correction. Two neighboring lines of work take other routes. Protocols that decode the joint measurement together with later syndrome measurements [1, 10] use more information and can employ weaker primitives; the Steane example compares the two routes. Constructions that design new flag circuits for replicated interactions and general Clifford circuits [15, 16], or that share flags among the stabilizer measurements of one code [17], build the protection into the circuit. Here the auxiliary circuit and its correction table are given, and the criterion decides when they can

TABLE I. Summary of notation.

S, M	Support $\{1,\dots,w\}$; operator $X_1\cdots X_w$
w, m	Weight of M; number of blocks
a, D_j, F_i	Measurement qubit; its CNOTs to data qubit j and to flag i
ω, $s(\omega)$	Fault event; its cost, Eq. (1)
$E(\omega)$	Residual data Pauli X^xZ^z
f, σ	Record of Z-basis outcomes; ideal branch
$r = R(f)$	Correction X^r selected by the record
$g \in \{0,1\}$	Common choice between E_b and E_bM_b
t	Order of a certificate
X_t, L_t	Guarantees of Eqs. (2) and (3)

be reused.

Section II defines the model, Section III proves the criterion, and Section IV gives the effective order. Sections V–VIII develop the applications. Appendices A and B contain the full proofs and the arbitrary-distance construction. The Supplemental Material [18] provides the detailed protocols, finite certificates, Shor-code noise calculations and comparisons with earlier constructions.

II. CIRCUIT MODEL AND PROPAGATION GUARANTEES

A flag circuit measures the Pauli operator $M = X_1\cdots X_w$ on the support $S = \{1,\dots,w\}$ of one code block; the other data qubits of the block are called spectators. A *measurement qubit* a, prepared in $|+\rangle$ and measured in the X basis, controls one CNOT to each data qubit of S, and its outcome is the measured eigenvalue. *Flag qubits*, prepared in $|0\rangle$ and measured in the Z basis, are targets of further CNOTs from a; their outcomes signal X errors on a that spread to several data qubits. We write such a circuit as its gate sequence, with D_j the CNOT from a to data qubit j and F_i a CNOT from a to flag i; the one-flag circuit of Fig. 1(c) is $F_1D_1D_2F_1D_3$. Other Pauli operators are reached by single-qubit basis changes (Corollary 7).

The results hold for a wider class of circuits, defined by Assumptions A1–A4; Table I collects the notation.

(A1) **Data-target circuit.** The auxiliary qubits undergo $|0\rangle$ or $|+\rangle$ preparations, CNOTs, X- or Z-basis measurements and resets, which may be chosen adaptively. Each coordinate $j \in S$ is the target of exactly one CNOT, controlled by an auxiliary qubit called its *port*, and the data qubits take part in no other gate.

(A2) **Sharing.** Sharing such a circuit, the *base*, among m blocks gives its complete *lift* (m,π): the data CNOT at each j becomes a *layer* of m consecutive copies from the same port, one to qubit j of every block, in an order π_j. The auxiliary operations are kept, and operations on different auxiliary qubits keep their relative order. Statements about the base are called *local*.

(A3) **Record and correction.** The *record* f is the ordered list of Z-basis outcomes of the auxiliary qubits. It determines the control path, whether the run is accepted, and the correction $X^{R(f)}$, an X Pauli on S that is applied to every block; here $X^r = \prod_j X_j^{r_j}$ for a binary vector r. X-basis outcomes are reported as measurement results.

(A4) **Full Pauli faults.** A faulty CNOT is followed by any of the fifteen nonidentity Paulis on its two qubits. Single-qubit fault locations, called *cuts*, admit X, Y and Z; they lie on every qubit before and after each of its operations, in particular after a preparation, before a measurement, between the copies of a layer, and on the data qubits. A spectator fault acts on one spectator.

For a fixed control path, let Λ be its set of fault locations. A *fault event* ω assigns a Pauli, modulo phase, to each location, and its cost is the number of faulty locations,

$$s(\omega) = \#\{\ell \in \Lambda : \omega(\ell) \neq I\}. \tag{1}$$

Propagating the faults to the end of the circuit gives the record and the residual data Pauli $E(\omega) = X^xZ^z$, up to phase, with binary vectors x and z.

The record of a fault-free run can itself be random. With the operations selected by a record f held fixed, an *ideal branch* σ is a list of Z outcomes that the circuit can produce without faults, and an event is *compatible* with f when its faults change some ideal branch into f. The guarantees below quantify over all compatible events, whereas the correction depends on f alone. This convention includes adaptive control and unrecorded resets (Appendix A 1).

Vectors add modulo two, $|v|_A$ counts the nonzero entries of v on a set A, and $\mathbf{1}$ is the indicator vector of S; the Pauli weight $\mathrm{wt}(E)$ counts all data qubits, spectators included. Fix an order $t \geq 0$, a record f and a correction value r. The X-weight and Pauli-weight guarantees are

$$\mathsf{X}_t(f,r): \quad \min_{g\in\{0,1\}} |x + r + g\mathbf{1}|_S \leq s(\omega), \tag{2}$$

$$\mathsf{L}_t(f,r): \quad \min_{g\in\{0,1\}} \mathrm{wt}(X^rE(\omega)M^g) \leq s(\omega), \tag{3}$$

for every event ω compatible with f of cost $s(\omega) \leq t$. The choice g expresses the freedom of the measurement: E and EM act identically on the projected state, and the guarantee bounds the lighter one. Both bounds are *graded*: s faults leave weight at most s, for every $s \leq t$. A circuit with correction table R has an *X-weight* or a *Pauli-weight certificate* of order t, and is said to satisfy X_t or L_t, when the guarantee holds with $r = R(f)$ for every accepted record f. The table is *deterministic* when it accepts every reachable record, and a *zero-flag guarantee* is the case $f = 0$, $r = 0$. The Pauli-weight guarantee of the base is the *local condition*.

In a lift, block b has support S_b, operator M_b and residual $E_b(\omega) = X^{x_b} Z^{z_b}$, and the circuit measures the product $\prod_b M_b$. Its Pauli-weight guarantee is

$$\min_{g\in\{0,1\}} \max_b \mathrm{wt}\left[X^{R_b(f)} E_b(\omega) M_b^g\right] \le s(\omega), \qquad (4)$$

with one *common choice* g for all blocks, because only the product is measured. Sharing copies the correction, $R_b = R$; Lemma 3 also treats block-dependent corrections. The X-weight version replaces each Pauli weight by $|x_b + R_b(f) + g\mathbf{1}_{S_b}|_{S_b}$. The choice g may differ from event to event, whereas the applied correction is a function of f. Each block then carries at most s errors and the whole register at most ms.

Both guarantees are worst-case statements over all Pauli faults at all locations. They therefore apply to any noise model whose faults are of this kind, including biased noise and the circuit-level noise of Section VIII.

These guarantees bound the data error; a measurement needs two further properties. The ideal-measurement condition *G0* requires the Kraus operators of the fault-free circuit to be proportional to the projectors of the measured product (Lemma 1). The reported eigenvalue is protected by repetition in a complete protocol (Section VIII), since a single Z fault on the measurement qubit flips the X readout and leaves the flags unchanged.

Lemma 1 (Ideal measurement for cat-state flag circuits)**.** *Grow a Greenberger–Horne–Zeilinger (GHZ) state from one $|+\rangle$ root and $|0\rangle$ targets. Let flags be targets only, coupled an even number of times to ports that already carry the root computational bit. Then the ideal flags are zero. Let c be the number of GHZ ports, all finally measured in the X basis. For their outcomes $y \in \{0,1\}^c$, define $\Pi_\mu = [I + (-1)^\mu \prod_b M_b]/2$, $\mu \in \{0,1\}$, with identity I. The Kraus operator is*

$$\begin{aligned} K_y &= 2^{-(c+1)/2}\left[I + (-1)^{|y|} \textstyle\prod_b M_b\right] \\ &= 2^{(1-c)/2} \Pi_{|y| \bmod 2}. \end{aligned}$$

The parity $|y| \bmod 2$ is the reported eigenvalue bit of the product, so G0 holds for this class, which contains the flag circuits with one measurement qubit ($c = 1$); for other circuits of Assumption A1 it is verified separately.

Proof. On each root branch every active port carries the same bit. Even flag couplings cancel their controlled flips. After all data gates the two branches are $|0^c\rangle|\psi\rangle$ and $|1^c\rangle \prod_b M_b|\psi\rangle$ with equal amplitude. Taking the X-basis overlap gives the formula. □

III. THE LOCAL SHARING CRITERION

Theorem 2 (Local sharing criterion)**.** *Under Assumptions A1–A4, fix a record f, a correction r and an order t. Then*

$$\begin{aligned} \mathsf{L}_t^{\mathrm{base}}(f, r) &\iff \left[\forall \pi\ \mathsf{X}_t^{(2,\pi)}(f; r, r)\right] \\ &\iff \left[\forall m \ge 2, \pi\ \mathsf{L}_t^{(m,\pi)}(f; r, \dots, r)\right]. \end{aligned}$$

Here $\mathsf{L}_t^{\mathrm{base}}$ is the local condition, Eq. (3) for the base. The superscript (m, π) denotes the guarantee of Eq. (4), or its X-weight version, for the complete lift with m blocks and layer orders $\pi = (\pi_1, \dots, \pi_w)$, and the arguments after the semicolon are the corrections of the blocks.

In words, the local condition holds if and only if every two-block lift obeys the X-weight bound, and if and only if every lift obeys the Pauli-weight bound. The weakest shared requirement thus implies the strongest.

a. Mechanism. Take a shared event and set aside its spectator faults. Let $J \subseteq S$ be the set of coordinates j with a faulty data CNOT, a port fault inside the layer, or a fault on data qubit j of some block. At every other coordinate all blocks receive the same X error, so they share one X pattern x_0 outside J, and their Z errors lie within J. A single-block event, the *compression* of the shared one, reproduces the X errors on the auxiliary qubits, hence the record, and places a Z error, a *mark*, on each data qubit of J. Distinct marks are paid for by distinct faults of the shared event, and the full Pauli alphabet of A4 allows each replacement. If the marked event has cost s', its local certificate gives one g with

$$|r + x_0 + g\mathbf{1}|_{S\setminus J} + |J| \le s'. \qquad (5)$$

For this g the marked support contains the corrected support of every block, and restoring the spectator faults proves sufficiency.

Conversely, take a local event and keep its faults with an X component on an auxiliary qubit, u in number. Let J' be the set of coordinates at which such a fault occurs at a data CNOT, and let x_0 be the data X pattern after their data factors are deleted. In a two-block lift, placing the fault of $j \in J'$ between the two copies of layer j gives the block visited second an extra X at j. The layer orders can therefore assign the complementary sets

$$P_0 = J' \setminus \mathrm{supp}(r + x_0), \qquad P_1 = J' \setminus P_0$$

to the two blocks as extra X support. One block then fills J' for $g = 0$ and the other fills it for $g = 1$, so the two-block X-weight bound gives $|r + x_0 + g\mathbf{1}|_{S\setminus J'} + |J'| \le u$, and the remaining faults add at most their number to the Pauli weight. Every local failure thus becomes a two-block witness. Appendix A gives the proofs, in which every step maps an event to one of equal or lower cost; Fig. 2 shows the supports.

Lemma 3 (Projection from a fixed complete lift)**.** *Fix a complete lift (m, π) and a record f. If corrections $r_1, \dots, r_m$ satisfy its Pauli-weight guarantee, then $\mathsf{L}_t^{\mathrm{base}}(f, r_b)$ holds for each block b.*

Embed any base event into the chosen block at the corresponding gates and cuts. The other blocks acquire X errors only, so the residual on the chosen block, the errors on the auxiliary qubits and the record are those of the base event, at equal or lower cost. Restricting the shared bound to that block proves the claim (Appendix A 4).

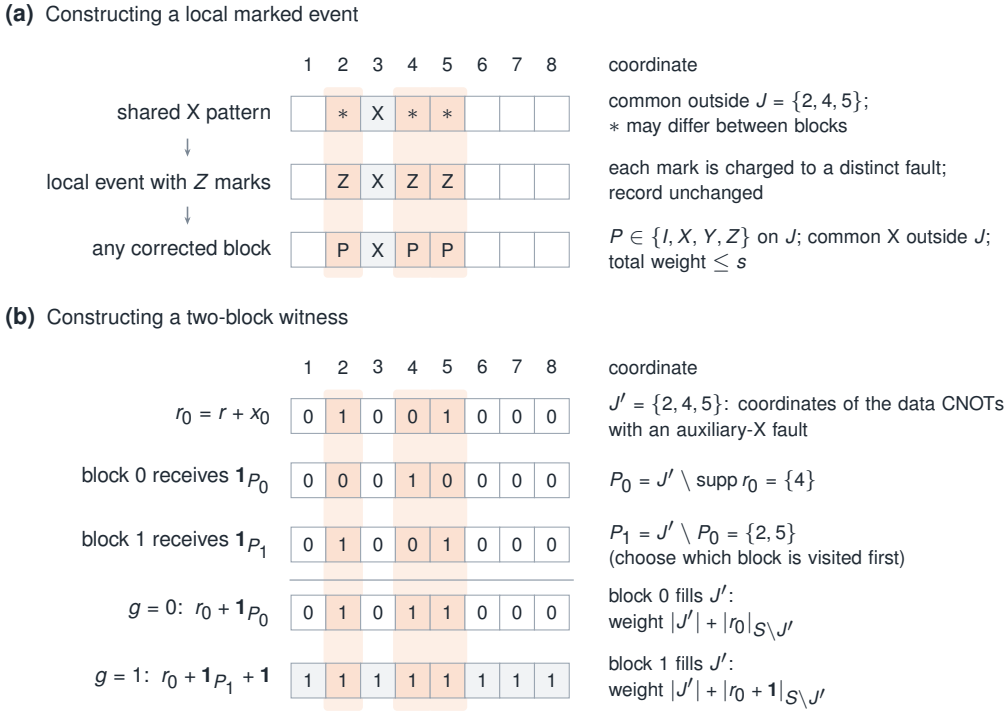


FIG. 2. Support mechanism of Theorem 2, drawn for eight coordinates. (a) From a shared event to a local one, for the identity correction and $g = 0$; entries are Paulis and blank cells are identities. Outside the set J of coordinates with a faulty data CNOT, a port fault inside the layer or a data-qubit fault, all blocks carry the same X pattern. Replacing the data factors on J by Z marks gives a single-block event of equal or lower cost with the same record. Its certificate bounds the support of the corrected marked event, which contains the corrected support of every block on S. (b) From two blocks to the local condition; entries are bits. The coordinates J' of the data CNOTs with an auxiliary-X fault split into complementary subsets P_0, P_1, realized on two blocks by choosing which block is visited first. For either common choice one of the two blocks fills all of J', so the two-block X-weight bound returns the marked single-block bound. Orange marks J or J'; gray marks X support outside it.

Corollary 4 (Admissible corrections)**.** *At a fixed order t, let $C_1^L(f)$ be the set of corrections r with $\mathsf{L}_t^{\mathrm{base}}(f, r)$, and $C_{m,\pi}^L(f)$ the set of tuples $(r_1, \dots, r_m)$ that satisfy the Pauli-weight guarantee of a fixed complete lift. Then*

$$\operatorname{diag} C_1^L(f) \subseteq C_{m,\pi}^L(f) \subseteq \left(C_1^L(f)\right)^m,$$

where $\operatorname{diag} C = \{(r, \dots, r) : r \in C\}$ uses m entries. Thus that lift admits a tuple precisely when the base admits a correction. A common correction works for any fixed complete lift precisely when it works locally.

The right inclusion can be strict, because individually valid block corrections may require different choices g. Lemma 3 makes the local condition necessary for the Pauli-weight guarantee of any fixed lift, and the two-block X-weight witness of Theorem 2 ranges over the layer orders.

Corollary 5 (Reachable records)**.** *For a fixed complete lift and every $s \geq 0$, the records reachable with at most s faults are exactly those of the base. The equality also holds for each fixed ideal branch.*

Compression gives one inclusion and embedding gives the other; both maps preserve the record at equal or lower cost.

Corollary 6 (Zero flags and restricted supports)**.** *Theorem 2 applies to the zero-flag guarantee. Its forward implication, from the local condition to the shared guarantee, also holds if block b receives only the gates of a subset $T_b \subseteq S$, with correction $r|_{T_b}$ and operator $M_b = \prod_{j \in T_b} X_{b,j}$. All blocks still use one common choice g.*

With ideal flags deterministically zero, the zero-flag part of the standard t-flag definition of Chamberland and Beverland [5, Def. 6] therefore gives this blockwise guarantee for every number of blocks. A complete flag-selected correction table, as in Chao and Reichardt [13, Def. 2], extends it to every record.

Corollary 7 (Other Pauli axes)**.** *For each block choose a product of single-qubit Clifford operators U_b. Conjugate its data interactions, Pauli faults and copied correction by U_b, so that its measured operator is $P_b = U_b M_b U_b^\dagger$ and its correction is $U_b X^r U_b^\dagger$. The same sharing guarantees hold, at the same fault order, modulo the one product $\prod_b P_b$.*

Conjugation preserves Pauli support and maps the full fault alphabet onto itself. The corollary covers XX, XZ or YY logical products and non-CSS codes such as the

five-qubit code. The sign of a Hermitian observable enters the reported eigenvalue (Section VII), and physical basis-change gates, where used, count as locations of the circuit.

IV. WHEN X-WEIGHT CERTIFICATES SUFFICE

Proposition 8 (Equivalence for $w \geq 2t+2$)**.** *For every fixed record and correction value, $w \geq 2t+2$ implies $\mathsf{X}_t(f,r) \iff \mathsf{L}_t(f,r)$ in the full Pauli model of Section II.*

A circuit with an X-weight certificate at $w \geq 2t+2$ can therefore be shared with its flags and correction table unchanged.

Proof. Temporarily exclude spectator faults. For an s-fault event let J collect the coordinates at which a fault has a factor on the data qubit, and put $n = |J| \leq s$. Delete these data factors, retaining the auxiliary history and record, and denote by x_0 the remaining X pattern on S. Let $\alpha_g = |r + x_0 + g\mathbf{1}|_{S\setminus J}$ for $g = 0, 1$, so that $\alpha_0 + \alpha_1 + n = w$. Failure of the Pauli-weight bound for both values of g requires $\alpha_0 + n > s$ and $\alpha_1 + n > s$. Full Pauli closure permits every choice of X factors on J at cost at most s, with unchanged record. Consequently, for every $k = 0, \ldots, n$,

$$\alpha_0 + k \leq s \quad \text{or} \quad \alpha_1 + n - k \leq s.$$

The endpoints $k = 0$ and $k = n$ imply $\alpha_0, \alpha_1 \leq s$. Taking $k = s - \alpha_0 + 1$, which lies in $[1, n]$, forces $\alpha_1 + n - k \leq s$, hence $w \leq 2s + 1$, a contradiction. If the original event contains s_{sp} spectator faults, remove them and apply this argument with budget $s - s_{\mathrm{sp}}$; restoring them adds weight at most s_{sp}. The reverse implication is immediate. □

Corollary 9 (Effective order)**.** *For $w \geq 2$ and $t \geq 0$, put $q = \min\{t, \lfloor w/2 \rfloor - 1\}$. For every fixed record and correction, $\mathsf{X}_t(f,r) \iff \mathsf{L}_q(f,r)$.*

The X-weight test thus certifies full-Pauli order q. A circuit with a higher-order Pauli-weight certificate shares at that higher order.

Proof. Every X pattern on S, or its complement, has weight at most $\lfloor w/2 \rfloor$. The X-weight bounds at orders $s \geq \lfloor w/2 \rfloor$ therefore hold automatically, so $\mathsf{X}_t(f,r)$ is exactly $\mathsf{X}_q(f,r)$. Since $w \geq 2q+2$, Proposition 8 gives $\mathsf{X}_q(f,r) \iff \mathsf{L}_q(f,r)$. For $q = 0$ both sides concern zero-fault events alone. At a record unreachable through order q, both guarantees hold vacuously at that order. Spectator weight enters the second step through the proof of Proposition 8. □

Corollary 10 (The minimum-weight boundary)**.** *For $t \geq 1$ and $w = 2t+1$, $\mathsf{X}_t(f,r) \iff \mathsf{L}_{t-1}(f,r)$.*

At the boundary the X-weight certificate guarantees sharing through order $t-1$; order t is provided by a Pauli-weight certificate such as those of Section VI.

Proof. Every X pattern on S, or its complement, has weight at most t, so the order-t X-weight bound holds automatically. At the remaining orders, Proposition 8 applies since $2t + 1 \geq 2(t-1) + 2$. □

The regimes $w \geq 2t+1$ are those of a logical representative in a code of distance $2t+1$. The boundary $w = 2t+1$ is reached by a minimum-weight representative, such as the weight-three logical operators of the Steane and five-qubit codes at $t = 1$. Heavier representatives lie in the regime $w \geq 2t+2$, where the X-weight certificate suffices.

Figure 3 shows the rule; the marked boundary points are the circuits of Section VI and the Supplemental Material.

V. SHARING AT ARBITRARY DISTANCE

Corollary 11 (Sharing the padded Chao–Reichardt construction)**.** *Let $t \geq 1$, $d = 2t+1$, and take the weight-w construction of Chao–Reichardt [13, Lemma 4], with $h = (d+1)^2/4 = (t+1)^2$ initial flag rounds, as described in Appendix B. Replace each data CNOT by a complete m-block layer and copy its flag-selected X correction $R_{\mathrm{CR}}(f)$ to every block. For every $m \geq 2$, all layer orders, and every event of $s \leq t$ physical faults,*

$$\min_{g \in \{0,1\}} \max_{b} \mathrm{wt}\left[X^{R_{\mathrm{CR}}(f)} E_b M_b^g\right] \leq s.$$

The ideal circuit measures only $\prod_b M_b$. With fast reset its peak auxiliary count is at most $d+1 = 2t+2$.

Proof. Definition 2 and the proof of Lemma 4 in Ref. [13] supply the local Pauli-weight certificate: at a faulty data CNOT the factor $P_\ell X_{\ell+1} \cdots X_w$, $P_\ell \in \{I, X, Y, Z\}$, is corrected to P_ℓ or $P_\ell X_\ell$. The proof uses equivalence under the measured operator M alone. Each data coordinate is a CNOT target exactly once. Z-basis flag records select the X correction. These data X corrections commute with subsequent data-target gates and can be deferred to a Pauli frame. Initial padding adds only auxiliary operations. Hence Theorem 2 applies. Between its preparation and its readout every flag is coupled to the measurement qubit twice, so ideal flags vanish and the measurement readout implements the product projector as in Lemma 1. The peak count follows from the original fast-reset schedule. □

The construction, padding and decoder are those of Chao and Reichardt. Sharing copies the data gates and keeps the auxiliary circuit and correction table. If $N_{\mathrm{CX}}^{\mathrm{base}}$ counts the base CNOTs, sharing gives

$$N_{\mathrm{CX}}^{\mathrm{shared}} = N_{\mathrm{CX}}^{\mathrm{base}} + (m-1)w. \tag{6}$$

For the endpoint schedule of Fig. 2 in Ref. [13], with this padding, there are $(d-1)(w+h)$ flag lifetimes, each running from a preparation through two couplings to a readout. Between rounds the new flag opens just before the

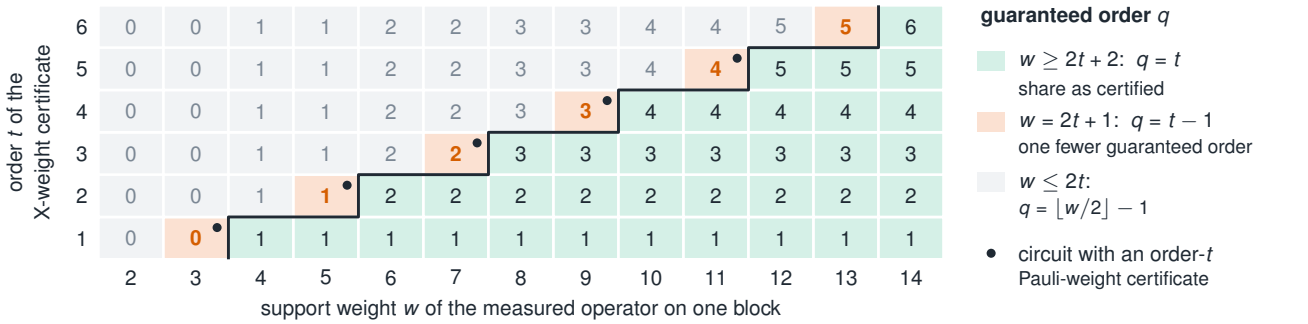


FIG. 3. What an X-weight certificate guarantees for sharing. Each cell gives the order $q = \min\{t, \lfloor w/2 \rfloor - 1\}$ of the Pauli-weight certificate implied by an X-weight certificate of order t at support weight w (Corollary 9). Right of the staircase line, $w \geq 2t+2$ and $q = t$: a circuit with an X-weight certificate can be shared unchanged. On the boundary $w = 2t + 1$ the certificate guarantees one fewer order. Dots mark boundary points for which a circuit with an order-t Pauli-weight certificate is given (Section VI and the Supplemental Material).

corresponding old flag closes. Counting both endpoints gives

$$N_{\mathrm{CX}}^{\mathrm{shared}} = mw + 2(d-1)(w+h),$$
$$N_{\mathrm{CX}}^{\mathrm{direct}} = mw + 2(d-1)(mw+h).$$

The second expression applies the same construction directly to the joint product of weight mw. Sharing saves $2(d-1)(m-1)w$ CNOTs with the same peak auxiliary bound $d+1$. The direct construction bounds the total residual weight by s, and the shared one bounds each block by s.

At distance d, the construction uses at most $d+1$ auxiliary qubits for every number of blocks. Its flag-gate count, $2(d-1)(w+h)$, is 28 at $d = w = 3$ and 112 at $d = w = 5$ (24 and 104 with the padding of the published version, Appendix B); the circuits of Section VI use 2 and 8. Appendix B gives the circuit, decoder and padding.

VI. SMALL CIRCUITS AT THE BOUNDARY

Two circuits reach the boundary $w = 2t + 1$ with few gates (Fig. 4). The sequence $F_1D_1D_2F_1D_3$, with corrections I and X_1 for flag values 0 and 1, satisfies L_1 and G0; shared by m blocks it uses two auxiliary qubits and $3m + 2$ CNOTs. At weight five, the four-flag sequence of Fig. 4(b) satisfies L_2 with the sixteen corrections of Table II; shared, it uses five auxiliary qubits and $5m + 8$ CNOTs. Exhaustive classification shows that four flags are the minimum at $w = 5$, $t = 2$ among circuits with one measurement qubit whose flags are each coupled twice, around at least one data gate, and whose correction table is deterministic. A two-flag circuit of this family has an X-weight certificate of order two and fails on two blocks with two faults. The Supplemental Material [18] gives the classification and verified circuits through distance eleven.

TABLE II. Corrections $X^{R(f)}$ of the four-flag circuit of Fig. 4(b), modulo M, for every record $f = (f_1, f_2, f_3, f_4)$; $f_i = 1$ denotes outcome -1 of flag i.

f	$X^{R(f)}$	f	$X^{R(f)}$
0000	I	1000	I
0001	I	1001	I
0010	X_4	1010	X_2
0011	$X_1X_2X_3$	1011	X_2X_3
0100	X_1	1100	X_1
0101	X_3	1101	X_1
0110	X_1X_2	1110	X_2
0111	X_1X_2	1111	X_3

VII. A SHARED MEASUREMENT IN A FIVE-QUBIT-CODE EXPERIMENT

Freund *et al.* teleport a logical qubit between two blocks of the five-qubit code [19] with a flagged joint measurement [10]. The circuit of their Fig. 12 is a shared one-flag circuit for a mixed-Pauli logical operator. Label the two blocks A and B and let $L_b = Z_{b,1}X_{b,2}Z_{b,3}$ be the operator measured on block b; coordinates $4, 5$ are spectators. Read from left to right, the circuit is

$$\begin{aligned} &\mathrm{CZ}_{a,B1},\ \mathrm{CZ}_{a,A1},\ \mathrm{CX}_{a,f},\\ &\mathrm{CX}_{a,A2},\ \mathrm{CX}_{a,B2},\ \mathrm{CX}_{a,f},\\ &\mathrm{CZ}_{a,A3},\ \mathrm{CZ}_{a,B3}, \end{aligned} \tag{7}$$

where $\mathrm{CZ}_{a,Aj}$ and $\mathrm{CX}_{a,Aj}$ are a controlled-Z gate and a CNOT from the measurement qubit a to qubit j of block A. The measurement qubit is prepared in $|+\rangle$ and measured in X; the flag f is prepared in $|0\rangle$ and measured in Z. Conjugating coordinates $1, 3$ by Hadamards gives the sequence $D_1F_1D_2F_1D_3$ with layer orders (BA, AB, AB).

This circuit satisfies the local condition with the flag-selected correction $\mathcal{R}(0) = I$, $\mathcal{R}(1) = Z_1$. On a single block, write $L = Z_1X_2Z_3$ and $\mathcal{P} = \{I, X, Y, Z\}$, let $\mathcal{P}_{\leq 1}^{(5)}$ be the set of five-qubit Paulis of weight at most one, and let $\mathcal{E}_f$ be the set of residuals, up to phase, of all events

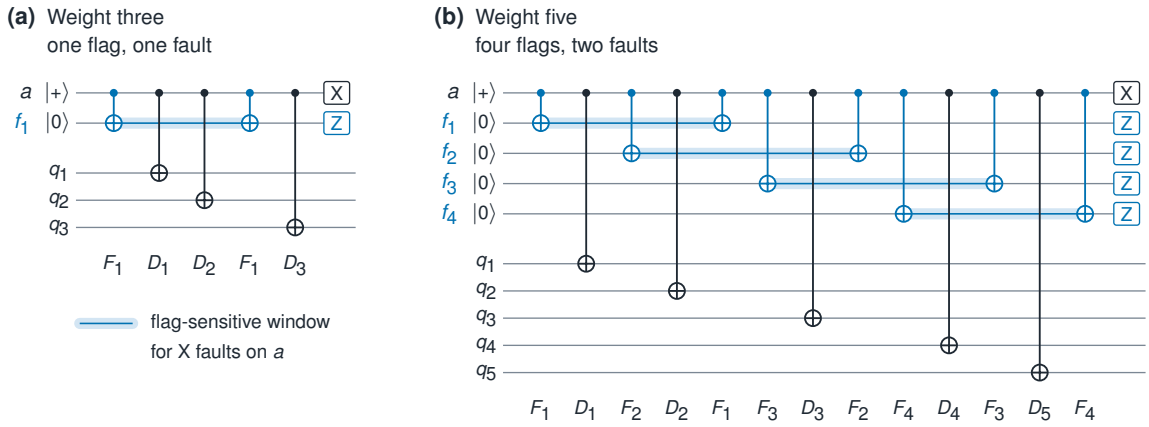


FIG. 4. Circuits with Pauli-weight certificates on the boundary $w = 2t + 1$. All CNOTs are controlled by the measurement qubit a. A shaded segment marks the window between the two couplings of a flag, in which an X error on a is recorded by that flag. Consecutive windows overlap; where the windows of flags i and $i+2$ meet, the earlier one closes before the later one opens. Gate sequences are printed below each circuit. The corrections are I and X_1 in (a) and those of Table II in (b). To share either circuit, replace each D_j by a data layer as in Fig. 1(a).

with at most one fault and flag outcome f. Then

$$\begin{aligned} \mathcal{E}_0 &= \mathcal{P}^{(5)}_{\le 1} \cup \{LP_1 : P \in \mathcal{P}\}, \\ \mathcal{E}_1 &= \{I\} \cup \{P_2 Z_3 : P \in \mathcal{P}\}, \end{aligned} \tag{8}$$

where P_j is the Pauli P on coordinate j. Every residual of $\mathcal{E}_0$ has weight at most one modulo L. For those of $\mathcal{E}_1$, $\mathcal{R}(1)P_2Z_3L = X_2P_2$ up to phase and $\mathcal{R}(1)I = Z_1$. Corollary 7 therefore extends the one-fault bound to any number of blocks and all layer orders, with the flag bit selecting the correction $\mathcal{R}(f)$ on every block. The ideal instrument has Kraus operators

$$K_{y,f} = \delta_{f,0}\,\frac{I + (-1)^y \prod_b L_b}{2}. \tag{9}$$

This is a nondemolition measurement of the product that preserves all other encoded information. With the stabilizer signs and $X_L = X^{\otimes 5}$ in Eq. (1) of Ref. [10], $L_b = -X_{L,b}$ on the code. Hence the logical product bit is $y \oplus (m \bmod 2)$; the sign cancels for two blocks.

The certificate covers the single-round circuit with its own flag readout. The experiment measures the operator twice, reads one flag after both rounds and uses syndrome information [20]; the Supplemental Material [18] describes that procedure. An earlier trapped-ion experiment implemented a logical CNOT on the five-qubit code with pieceable fault tolerance [21].

VIII. A COMPLETE STEANE-CODE MEASUREMENT

We now use the shared one-flag circuit in a nondestructive measurement of M_AM_B, with $M_b = X_{b,1}X_{b,2}X_{b,3}$, on two blocks of the Steane $[[7, 1, 3]]$ code [12]. Each factor M_b is a logical X representative. The ideal instrument is $\rho \mapsto \Pi_y \rho \Pi_y$ with $\Pi_y = [I + (-1)^y M_A M_B]/2$; both blocks are retained, and the input may be entangled with a reference.

Because a single fault can flip one joint outcome (Section II), the protocol performs three rounds, each consisting of the joint measurement, QEC on A and QEC on B, and reports the majority of the three outcomes. Each joint measurement uses $F_1D_1D_2F_1D_3$ with A before B in every layer and a freshly prepared flag; a raised flag selects $X_{A,1}X_{B,1}$. Each QEC module is the two-auxiliary flagged syndrome extraction of Chao and Reichardt [4, Sec. III.A and Fig. 3(b), arXiv v1]: it scans the six Steane stabilizers and, at the first nontrivial syndrome or flag, remeasures all six without flags and applies the corresponding correction. All corrections are Pauli-frame updates, and each joint outcome is interpreted using the Pauli frame at the time of measurement. We call this rule, in which every module is decoded from its own outcomes, *standard decoding*. The protocol accepts every record and always runs to completion; the Supplemental Material lists the gates.

Proposition 12 (One-fault tolerance of the complete protocol)**.** *Consider this protocol with an arbitrary encoded input. For an error-free input and at most one internal Pauli fault, the reported outcome and the conditional logical output are correct, and at most one Pauli error remains on each block, modulo the code stabilizers and M_AM_B. For one input Pauli error on each block and no internal fault, the outcome and output are correct and every input error is removed.*

Proof. A fault in a joint measurement leaves at most one error per block, for one common choice g, by its local certificate and Theorem 2. The subsequent fault-free QEC modules remove these errors. Only that round's joint readout can be wrong. A fault in QEC leaves at most one correctable error; it can affect the next joint readout, after which fault-free QEC removes it. An input

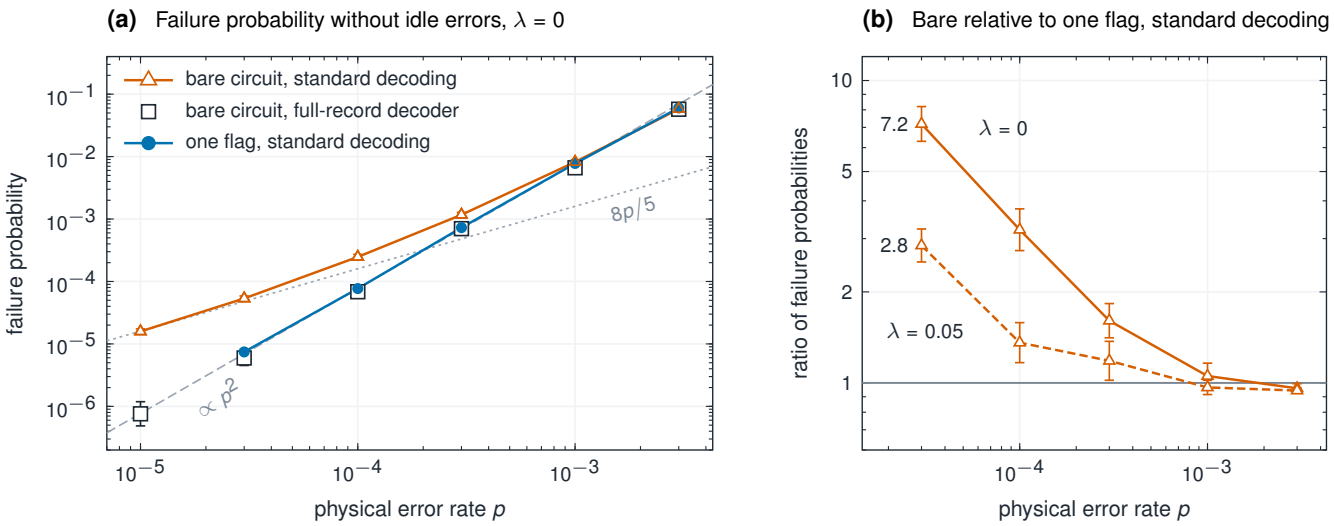


FIG. 5. Complete two-block Steane-code measurement under circuit noise. (a) Failure probability against the physical error rate p without idle errors. With standard decoding the bare circuit follows its first-order term $8p/5$ (dotted) at low p, and the shared one-flag circuit, sampled for $p \geq 3 \times 10^{-5}$, scales as p^2 (dashed guide through its estimate at $p = 10^{-4}$). A full-record decoder built for this protocol gives the bare circuit the same scaling; for $p \geq 10^{-4}$ its estimates lie within 15% of the one-flag ones and the symbols coincide. (b) Ratio of the bare to the one-flag failure probability under standard decoding, for idle-to-gate error ratios $\lambda = 0$ and 0.05; it is 7.2 and 2.8 at $p = 3 \times 10^{-5}$. Bars are 95% intervals: Wilson intervals in (a), most of them smaller than the symbols, and log-ratio intervals in (b).

error can similarly affect only the first round. Thus at most one of three interpreted votes is wrong. The QEC circuits have the two properties used here: they remove a single input error when fault-free and leave at most one error from a single internal fault on a clean block. Both were checked at every physical location and for every Pauli fault. All ideal operations measure only code stabilizers and $M_A M_B$, so the argument holds for the conditional quantum instrument, including its action on a reference. □

Exhaustive single-fault enumeration and independent stabilizer simulations [22] of representative runs confirm the proposition; the Supplemental Material gives their coverage.

a. Noise model. Each gate is followed by a depolarizing error of total probability p, each preparation and readout flips with probability p, and each idle qubit is depolarized with probability λp per time step, where λ is the idle-to-gate error ratio; every gate, reset, readout and feedback decision takes one time step. A run fails when the reported outcome or the conditional logical output is incorrect, as assessed by an ideal decoder after the final noisy QEC.

b. Standard decoding. Replace the joint measurement by the bare sequence $D_1 D_2 D_3$ and keep everything else: the protocol, the QEC circuits, the layer order and standard decoding. A single-fault assignment places one nonidentity Pauli at one location. Of the 26,154 single-fault assignments of this bare protocol, 24 produce a logical failure: four two-qubit Pauli faults on each of the CNOTs to the data qubits $q_{A,2}$ and $q_{B,2}$, in each of the three rounds. Each has probability $p/15$, so

$$P_{\text{fail}}^{\text{bare,std}}(p, \lambda) = \frac{8}{5} p + O(p^2), \tag{10}$$

independently of λ. All 26,784 single-fault assignments of the flagged protocol give the correct outcome and output (Proposition 12). This is the fault of Fig. 1(b), and its repair in Fig. 1(c), in a complete protocol.

c. Decoding with the full record. The first-order failures of the bare circuit can also be removed by a decoder that reads the complete record of the protocol: the three joint outcomes and all syndromes and flags of the six QEC modules. Such a decoder is built for this protocol, block count and layer order, and it corrects any single input error or internal fault. The flagged protocol reaches the same result with standard decoding and, in a fault-free run, also removes one input error on each block (Proposition 12). For the bare circuit these two tasks conflict: an X fault on the measurement qubit after the first-round CNOT to $q_{A,2}$ and the input error $X_{A,3} X_{B,1}$ produce the same complete record and require corrections that differ by a logical X. A correction that depends on the flag bit alone composes with standard blockwise QEC, and this is what the local criterion certifies.

d. Finite noise. Both protocols use two auxiliary and fourteen data qubits. An error-free run of the flagged protocol has 240 CNOTs in 513 time steps, six CNOTs and twelve steps more than the bare one.

Figure 5 shows fixed-sample Monte Carlo estimates.

With standard decoding the bare protocol approaches the linear term of Eq. (10) as p decreases, and the flagged protocol scales as p^2 [Fig. 5(a)]. Their ratio therefore grows as p decreases [Fig. 5(b)]. At $p = 3\times10^{-5}$ and $\lambda = 0$ the failure probabilities are 5.3×10^{-5} and 7.4×10^{-6}, a ratio of 7.2; with idle errors at $\lambda = 0.05$ the ratio is 2.8. Near $p = 10^{-3}$, where higher-order events dominate both protocols, the ratio reaches one; at $p = 3\times10^{-3}$ it is 0.96 and 0.94 for the two idle ratios. With the full-record decoder the bare protocol also scales as p^2. At $p = 10^{-4}$ and $\lambda = 0$ its estimate, 6.85×10^{-5}, and the flagged one, 7.70×10^{-5}, have overlapping intervals. At the two largest sampled error rates, $p = 10^{-3}$ and 3×10^{-3}, its estimated failure probability is about 5–15% lower than the flagged one. The shared flag circuit obtains first-order protection with standard blockwise decoding.

IX. DISCUSSION

A single-block Pauli-weight certificate exactly determines when a flag circuit and its correction table can be reused across code blocks. The guarantee is uniform in block count and in the order of the copied gates; two blocks suffice to witness failure. The effective-order formula gives the corresponding guarantee from an X-weight certificate. These results turn local certification into a reusable step in logical-measurement design.

Sharing suits processors in which one measurement qubit reaches every block, trading auxiliary space for serial depth. In a leading-order comparison of two-block measurements on the Shor $[[25, 1, 5]]$ code, the shared four-flag circuit has a smaller failure coefficient than cat-state measurement for idle-to-gate error ratios below 0.022 with reused auxiliaries and below 0.158 with parallel ones (Supplemental Material). The Chao–Reichardt application keeps the peak auxiliary count at most $d+1$ for any number of blocks; small circuits reduce the flag-gate count at low weight. The circuit of Freund *et al.* connects the criterion to an experimental logical measurement, and the Steane protocol demonstrates composition with standard QEC and suppression of first-order failures.

The criterion suggests extensions to code-dependent correctable sets and repeated data interactions with aggregated flags, as in Ref. [20]. Whether the two-block X-weight guarantee for a single fixed layer order already implies the local condition remains open.

ACKNOWLEDGMENTS

Funding. The author received no specific funding for this work.

Competing interests. The author declares no competing interests.

AI-assisted method. During the preparation of this work, the author used OpenAI ChatGPT-6 and Claude Fable 5.2 in order to improve language and expression. After using the tool, the author reviewed and edited the content as needed and took full responsibility for the content of the publication.

DATA AVAILABILITY

The data that support the findings of this article are openly available [23].

Appendix A: Proofs of the sharing criterion

All statements use the locations and event cost of Section II. We write $X(\omega)$ for the auxiliary X history: the X component on every auxiliary qubit immediately before and after each of its operations, together with its changes inside data layers. Only Pauli errors are taken modulo phase; a measured Hermitian operator's sign is retained. Table III lists the uses of the assumptions.

1. Conventions, records and compatible branches

With adaptive control, every terminating control path is a circuit of the form of A1. A *lifetime* of an auxiliary qubit runs from a preparation or reset to its measurement or next reset, and a lift keeps all lifetimes. A preparation fault is placed just after the preparation, and a readout flip is an anticommuting Pauli just before the measurement. Factors placed at the same location are multiplied, and a location whose product is the identity is not counted. The component of a fault on a data qubit is its *data factor*, or *target factor* at a data CNOT, which A4 allows to be replaced by any Pauli or deleted. A fault at a cut of a data qubit is a *direct-data fault*; a *port fault* lies at a data CNOT or at a cut of the port. An event is *legal* when every location carries a Pauli admitted by A4.

Freeze the operations selected by a record f. Let σ and τ respectively label all Z and X outcomes of this frozen ideal circuit, and let $K^f_{\sigma,\tau}$ be its data Kraus operator. Any unrecorded reset environment has a branch index η, suppressed in $K^f_{\sigma,\tau}$; all identities below hold for each η, which is unavailable to the decoder. The compatible branches are those with a nonzero Kraus operator. Commuting an event's Pauli factors forward gives its data residual $E(\omega)$ and outcome-flip strings $\delta_Z(\omega)$ and $\delta_X(\omega)$. The record is

$$f = \sigma + \delta_Z(\omega),$$

and the X outcomes are $\tau + \delta_X(\omega)$. The frozen controls follow f, even if $\sigma \neq f$. This description therefore includes faults that change a later control decision. A certificate quantifies over every σ, τ, ω with a nonzero branch and the given f, and it selects one $r = R(f)$ for all of them.

Lemma 13 (Invariance of the record)**.** *At a fixed compatible ideal branch, two error histories with the same auxiliary X history have the same Z-outcome flips. Changing direct-data factors or auxiliary Z factors while keeping that history preserves the record and its control path.*

Proof. Under A1, a CNOT copies an X error from its control to its target and a Z error from its target to its control. An X error on a data target therefore stays on the data, and a Z error on it adds only Z to the auxiliary qubits. A Z measurement is flipped exactly by the measured qubit's X component. Induction over the Z measurements gives identical outcomes and hence, by A3, identical later operations. Preparation and reset discard the incoming error and start the next lifetime. Auxiliary Z can change X outcomes only, which A3 reports as measurement results. □

a. Why compatible branches are preserved. For fixed f, σ, τ, expand each ideal data gate as $|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X_j$. All other gates and projections are on auxiliaries. Consequently

$$K^f_{\sigma,\tau} = \sum_{x \in \{0,1\}^w} c_x X^x.$$

The coefficients contain the auxiliary circuit and its measurement outcomes. Sharing makes the substitution

$$X_j \longmapsto \prod_{b=1}^{m} X_{b,j}$$

without changing those coefficients: every controlled copy uses the same port bit inside its layer, by A2. Distinct monomials on the right are distinct Paulis, since a complete block retains all coordinates. They are linearly independent, as are the base monomials. Thus the base and a complete lift have exactly the same nonzero ideal branches on a frozen control path. For restricted supports the substitution may identify monomials, and a vanishing base operator still gives a vanishing lifted one; this direction suffices for the forward implication. After an error-history modification, the X-outcome labels may change by δ_X; they remain unrestricted. Lemma 13 preserves δ_Z and hence f. This supplies the branch compatibility used by the event maps below.

For a concrete random-branch example, append an independent auxiliary prepared in $|+\rangle$ and measured in Z to a deterministic even-flag circuit. Write $f = (f_{\text{flag}}, e)$ and choose $R(f_{\text{flag}}, e) = R_0(f_{\text{flag}})$ for $e = 0, 1$. The ideal value of e is a fair bit, and an X fault can exchange its two values. The same correction covers both ideal branches for each record f. Theorem 2 covers random and deterministic ideal branches alike.

2. Compression and the marked event

Lemma 14 (Layer compression)**.** *Remove the spectator faults, s_{sp} in number, from a shared event of cost s. Let $J \subseteq S$ contain the coordinate of every faulty data CNOT, every port fault strictly inside its data layer, and every remaining direct-data fault. One can choose a base auxiliary X history and a vector x_0 such that the record is unchanged and*

$$x_b|_{S\setminus J} = x_0|_{S\setminus J}, \qquad \operatorname{supp}(z_b) \subseteq J$$

for each block, with spectators temporarily omitted.

Proof. At $j \notin J$ the port's X component is constant throughout the layer, so every data copy receives the same X component. Data Z can originate only from a data factor of a fault, hence is supported in J. For a touched layer combine the port's X changes at the base gate to retain its X value at the layer exit. Boundary cuts stay on their original side of the layer. A2 permits this compression because the copies of a layer are consecutive on that port. Disjoint auxiliary operations retain their order and all their X factors; all other auxiliary X histories are unchanged. The resulting base propagation defines x_0. The values of x_0 on J do not enter this equality. Lemma 13 preserves the record and control path. □

Lemma 15 (Legal marking and distinct charges)**.** *The construction of Lemma 14 admits a legal base event of cost $s' \le s - s_{\text{sp}}$ with the same record, the same data X component off J, and Z component one at every coordinate of J.*

Proof. If a layer contains a data-port fault or an internal port cut fault, put its combined auxiliary X factor and one Z_j at the base data CNOT. The full two-qubit alphabet of A4 permits this even when that auxiliary factor is identity. Delete the other data factors at j. If only independent direct-data faults touch j, replace one by Z_j and delete the others. A CNOT leaves its target's Z component in place and copies it to the control; this auxiliary Z does not enter Lemma 13.

Choose one original touching fault to pay for each mark. A data CNOT touches one data coordinate by A1, an interior port cut belongs to one layer by A2, and a direct-data fault touches one coordinate. Each fault therefore pays for at most one mark. A combined auxiliary-X factor at a touched port is included in the same charged location. Retain other auxiliary X faults at their original locations. The constructed event thus has at most as many faulty locations as the original event has nonspectator faults. Multiplication at coincident locations and deletion of identities can only reduce their number. A4 makes every replacement and merger legal. The record and compatible branch are preserved as above. □

Forward implication of Theorem 2. Apply $\mathsf{L}^{\text{base}}_t(f, r)$ to the marked event. Its cost is $s' \le s - s_{\text{sp}} \le t$, so the graded guarantee supplies a g with

$$|r + x_0 + g\mathbf{1}|_{S\setminus J} + |J| \le s' \le s - s_{\text{sp}}.$$

The X correction and multiplication by M act on the X component alone, so every Z mark remains. With this same g, every block's corrected support on S is contained in $J \cup \mathrm{supp}((r + x_0 + g\mathbf{1})|_{S\setminus J})$. Restoring the spectator faults adds at most s_{sp} to the weight of any block. This gives Eq. (4) with bound s. The construction is pointwise in f, r and every compatible ideal branch, so it preserves any acceptance rule permitted by A3. □

3. The complementary two-block witness

Lemma 16 (Two-block bridge)**.** *If* $\mathsf{X}_t^{(2,\pi)}(f; r, r)$ *holds for every allowed layer order* π*, then* $\mathsf{L}_t^{\mathrm{base}}(f, r)$ *holds.*

Proof. Take a base event of cost $s \le t$ and record f. First multiply same-location factors and delete identities. Keep in fault group $\mathcal{U}$ every location with a nonzero auxiliary X component, at cost u. Delete its auxiliary Z component and, at a data port, its target factor. Put the remaining data-affecting faults in group $\mathcal{V}$, at cost v; discard pure auxiliary Z faults. Then $u + v \le s$. Every fault of $\mathcal{V}$ touches at most one data coordinate, possibly a spectator. These changes retain the auxiliary X history and record.

Let J' be the set of data-CNOT coordinates of the faults in $\mathcal{U}$. By A1 and same-location merging, at most one such port fault occurs at each coordinate. Auxiliary cuts, including cuts adjacent to a data gate, are retained at their corresponding boundaries and lie outside J'. Let x_0 be the base data X pattern after the target factors have been deleted, and put

$$r_0 = r + x_0, \qquad P_0 = J' \setminus \mathrm{supp}(r_0), \qquad P_1 = J' \setminus P_0.$$

Use blocks $b = 0, 1$. At each $j \in P_b$, visit block $1 - b$ first and insert the retained port-X fault immediately after that earlier data CNOT. Visit block b second. Keep all other faults of $\mathcal{U}$, with their auxiliary X factors and lifetimes unchanged. This is a legal event under A2 and A4, of cost at most u. The exit X history and record are unchanged. Its two data patterns are

$$x_b = x_0 + \mathbf{1}_{P_b}, \qquad b = 0, 1.$$

If the common choice is $g = 0$, block 0 is nonzero at all of J'; if it is $g = 1$, block 1 is nonzero there. Outside J' both blocks agree. The X-weight guarantee, applied at the cost of this event, therefore gives

$$|r + x_0 + g\mathbf{1}|_{S\setminus J'} + |J'| \le u.$$

All deleted target factors of $\mathcal{U}$ lie in J', so restoring them keeps this bound. Restoring $\mathcal{V}$ adds at most v to the local Pauli weight. The original local event consequently has weight at most $u + v \le s$ for this choice of g. Auxiliary Z has no data support, so the conclusion stands. This proves the local certificate for every compatible event. □

The Pauli-weight guarantee of all lifts immediately implies the two-block X-weight guarantee; together with the forward implication, the bridge completes Theorem 2. The layer orders of the witness depend on the event and on the correction r, which is why the two-block statement of the theorem quantifies over all layer orders.

4. Projection and reachable records

Proof of Lemma 3. Choose one complete block b in a fixed lift and any base event of cost $s \le t$. Put each faulty base data CNOT at its copy on b, retaining both its port and target factors. Place direct-data and spectator faults on the corresponding block-b wires. Auxiliary operations and lifetimes are unchanged. A port cut before or after a base gate is put before or after that block's copy. Such cuts are admitted by A4; conjugation through the adjacent CNOT alternatively represents the same factor within the full two-qubit gate alphabet, at the same cost.

The added targets in other blocks carry no inserted faults. They acquire X errors only, which stay on the data. By induction over gates, deleting these other blocks leaves exactly the base auxiliary error history and the selected residual E. Record flips and the control path agree, and the branch argument above ensures compatibility. Merging any coincident factors only lowers the embedded event cost to $s' \le s$. Its shared certificate gives one g such that

$$\mathrm{wt}(X^{r_b} E M^g) \le s' \le s.$$

Since the base event was arbitrary, this is its local certificate. □

Forward lifting gives the diagonal inclusion in Corollary 4, and projection gives its product inclusion. These statements also hold for record-to-correction tables because every step is pointwise in the record. Compression and marking map lifted events to base events with the same record at equal or lower cost; projection embeds base events in the opposite direction. This proves Corollary 5, including its branchwise statement.

5. Restricted supports and the fault alphabet

For restricted supports T_b, perform the marking construction on the full base support. Treat direct faults on a coordinate inactive in its block as spectator faults. The remaining residual of block b is contained in the restriction of the marked support to T_b. Restriction keeps or lowers the weight, so the same common choice g proves the forward part of Corollary 6. A compatible restricted ideal branch also gives a compatible base branch by the substitution argument.

The reverse implication requires complete blocks, even when the retained supports have full union. For the bare sequence $D_1 D_2 D_3$ at $t = 1$, a fault $X_a Z_2$ after D_2 leaves

TABLE III. Where the circuit assumptions enter the proof. The bounds are graded by the event cost throughout.

Step	Assumptions	Mathematical use
Record invariance	A1, A3	CNOTs separate X/Z histories; data are targets; the record alone controls the path.
Layer compression	A1, A2	Untouched copies receive the same X; the copies of a layer are consecutive on the port.
Marked-event construction	A1, A4	Each fault touches at most one support coordinate; full Pauli replacement pays for each Z mark at no extra cost.
Complementary two-block bridge	A1–A4	One data port per coordinate; selectable layer orders; unchanged record; legal target deletion.
Fixed-lift projection	A1–A4	A complete selected block retains every coordinate and admits the base cuts; other targets acquire X only.
Branch compatibility	A1–A3	Free commuting data-X monomials have an injective complete-lift substitution; unrecorded branches are quantified uniformly.
Effective-order reduction	A1, A3, A4	Arbitrary target-X replacements preserve the record at equal or lower cost.

Z_2X_3. The two equivalent residuals, Z_2X_3 and X_1Y_2, both have weight two. Thus the base fails $\mathsf{L}_1(0,0)$. Retain only $T_A = \{1\}$ and $T_B = \{2,3\}$ in two blocks. A nonspectator fault introduces Z support at at most one of block B's two coordinates. Choose g to remove the X component on its other coordinate; if there is no Z, choose the lighter of the two-bit X pattern and its complement. Block B then has weight at most one, and block A has at most one for either g. A sole spectator fault is handled with $g = 0$. The zero-fault event uses $g = 0$ as well. This incomplete lift satisfies $\mathsf{L}_1(0,0)$ although its base does not.

The guarantees are combinatorial: they hold for every assignment of probabilities to the Pauli faults, so a certificate in the full model covers every noise distribution supported on it, including biased noise. The full alphabet is needed for the certificate itself. With pure X faults alone, the bare weight-three sequence satisfies the restricted $\mathsf{L}_1(0,0)$, since a three-bit X string or its complement has weight at most one. In its two-block A-then-B lift, an X fault on a between the gates to $q_{A,2}$ and $q_{B,2}$ leaves $X_{A,3}$ and $X_{B,2}X_{B,3}$, with weights $(1,2)$ and $(2,1)$ for the two common choices. The transfer to shared circuits therefore starts from a certificate in the full alphabet.

6. Data-factor replacement and effective order

Lemma 17 (Independent target replacements)**.** *For a base event, let J be the coordinates carrying directly inserted data factors. After deleting those factors, any X pattern can be placed on J with the same auxiliary X history and record, at equal or lower cost.*

Proof. For each $j \in J$, choose one original fault with a target factor on j. Replace that factor by the desired X or identity and delete the other data factors at j, keeping auxiliary X factors fixed. Different coordinates select different faults, by A1. Data X stays on its qubit, and the auxiliary Z generated by an old data Z does not enter the record, by Lemma 13. A4 permits each replacement; locations that become identity are deleted. Compatible ideal branches are retained because the frozen circuit and Z flips are unchanged. □

Lemma 17 justifies the target-factor substitutions in the proof of Proposition 8. That proof gives the effective-order relation in Corollary 9, including zero faults and records unreachable at the reduced order.

Appendix B: The padded Chao–Reichardt construction and its decoder

The construction in Corollary 11 uses the following record-to-correction map. Chao and Reichardt's Definition 2 requires that every set of $s \le t$ faults leave at most s data errors after correction from the flag pattern alone [13]. Their Lemma 4 supplies this full-Pauli guarantee. The conservative padding $h = (t+1)^2$ used here is the statement in arXiv:1912.09549v1; the published Lemma 4 reduces it to $h = t(t+2)$. Both use the padded construction, whereas their Conjecture 3 concerns the unpadded circuit.

For a weight-w product of X operators, prepare one measurement qubit in $|+\rangle$. Each of the w CNOTs from the measurement qubit to the data is enclosed by a fresh group of $2t$ flag intervals. Between adjacent data CNOTs, turn on flag k of the next group immediately before turning off flag k of the previous group, for $k = 1, \ldots, 2t$. Each flag is prepared in $|0\rangle$, coupled to the measurement qubit twice, and read in Z. Prepend h such groups without their data CNOTs. These dummy groups remove the decoder's initial boundary ambiguity. The terminal measurement qubit readout is in X. There are $4t(w+h)$ flag CNOTs, and each data coordinate is targeted exactly once. With fast reset, $2t+1$ physical flag qubits can be reused, giving $2t+2$ auxiliaries including the measurement qubit.

The decoder of Ref. [13] uses the flag record alone. A group is a binary string of length $2t$, in the physical order of its flag intervals. Find t consecutive all-zero groups among the initial dummy groups; the padding guarantees such a run under at most t faults. Discard the groups before and within this run, and index the remaining groups as $f_1, f_2, \ldots$. If every remaining bit is zero, return the identity correction. Otherwise let Ω be the index of the final nonzero group and let ϵ be the unit string marking its rightmost nonzero flag. Initialize $u_0 = 0^{2t}$ and an empty set of suffix corrections. For $i = 1, \ldots, \Omega$:

1. Set $u_i = u_{i-1} \oplus f_i$.
2. If $|u_i| > t$, or if $i = \Omega$, $|u_i| = t$ and $|u_i \oplus \epsilon| > t$, replace u_i by $u_i \oplus 1^{2t}$ and multiply the correction by X on every data target from group i onward.

A suffix beginning in a dummy group acts on all later data targets. Products of suffix corrections are taken modulo two.

The cited proof counts full Pauli support. After fixing the final nonzero group and its rightmost bit, the proof adds measurement-qubit faults from left to right. Adding one fault between data positions ℓ and $\ell + 1$ changes the suffix correction in exactly one of the adjacent groups. A physical fault with data component P_ℓ therefore leaves either P_ℓ or $P_\ell X_\ell$ at that coordinate after correction. The induction counts the support of arbitrary Paulis. Flag-readout faults are included in the same fault budget. This is the local certificate used by Theorem 2.

The fault-free measurement qubit implements the product projector. Sharing replaces each data CNOT by its layer and copies each suffix correction to every block, which gives the counts of Section V.

[1] T. Tansuwannont, C. Chamberland, and D. Leung, Phys. Rev. A **101**, 012342 (2020), arXiv:1803.09758v3.
[2] D. Litinski, Quantum **3**, 128 (2019), arXiv:1808.02892.
[3] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, New J. Phys. **14**, 123011 (2012), arXiv:1111.4022.
[4] R. Chao and B. W. Reichardt, Phys. Rev. Lett. **121**, 050502 (2018), arXiv:1705.02329v1.
[5] C. Chamberland and M. E. Beverland, Quantum **2**, 53 (2018), arXiv:1708.02246v3.
[6] C. Chamberland and A. W. Cross, Quantum **3**, 143 (2019), arXiv:1811.00566.
[7] C. Ryan-Anderson *et al.*, Phys. Rev. X **11**, 041058 (2021), arXiv:2107.07505.
[8] L. Postler *et al.*, Nature **605**, 675 (2022), arXiv:2111.12654.
[9] C. Ryan-Anderson *et al.*, Science **385**, 1327 (2024), arXiv:2404.16728.
[10] R. Freund, F. Butt, C. Benito, I. Pogorelov, M. Meyer, A. Steiner, A. Bermudez, M. Müller, and T. Monz, Modular fault-tolerant quantum computing on a non-CSS code (2026), arXiv:2609.22572v1 [quant-ph].
[11] A. R. Calderbank and P. W. Shor, Phys. Rev. A **54**, 1098 (1996).
[12] A. M. Steane, Phys. Rev. Lett. **77**, 793 (1996).
[13] R. Chao and B. W. Reichardt, PRX Quantum **1**, 010302 (2020), arXiv:1912.09549v1 [quant-ph].
[14] P. Prabhu and B. W. Reichardt, Quantum **7**, 1154 (2023), arXiv:2108.02184v2.
[15] B. Anker and M. Marvian, PRX Quantum **5**, 040340 (2024), arXiv:2212.10738v2.
[16] B. Anker and M. Marvian, Universal fault tolerance with non-transversal Clifford gates (2025), arXiv:2510.08402v1.
[17] P.-H. Liou and C.-Y. Lai, IEEE Trans. Quantum Eng. **6**, 3101524 (2025), arXiv:2407.00607v1.
[18] See Supplemental Material at [URL will be inserted by publisher] for circuit certificates, protocol details, Shor-code noise calculations, verification methods and comparisons with earlier constructions, including Refs. [24–28].
[19] R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, Phys. Rev. Lett. **77**, 198 (1996).
[20] R. Freund *et al.*, Modular fault-tolerant quantum computing on a non-CSS code, Zenodo (2026), https://doi.org/10.5281/zenodo.22796458.
[21] C. Ryan-Anderson, N. C. Brown, M. S. Allman, *et al.*, Implementing fault-tolerant entangling gates on the five-qubit code and the color code (2022), arXiv:2208.01863v1 [quant-ph].
[22] C. Gidney, Quantum **5**, 497 (2021).
[23] A. Chen, Code and data for "Sharing flag circuits across code blocks: an exact local criterion for joint logical measurements", Zenodo (2026), https://doi.org/10.5281/zenodo.24203932.
[24] A. Benhemou and N. Berthusen, Automated logical Clifford gadgets for heterogeneous architectures via chain maps (2026), arXiv:2607.02482v1.
[25] D. P. DiVincenzo and P. Aliferis, Phys. Rev. Lett. **98**, 020501 (2007), arXiv:quant-ph/0607047v2.
[26] S. Ghosh, H. Boche, and A. Tanggara, Fault tolerance of quantum circuits with tensor networks and symplectic geometry (2026), arXiv:2610.03429v3.
[27] P. Prabhu, *Lower overhead fault-tolerant building blocks for noisy quantum computers*, Ph.D. thesis, University of Southern California (2024), arXiv:2605.12385v1.
[28] P. W. Shor, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science* (IEEE, 1996) pp. 56–65, arXiv:quant-ph/9605011v2.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplement.pdf