hexagon:2610.00530v2

A MATCHING THRESHOLD OF 2×10^{-3} FOR THE ROTATED SURFACE CODE UNDER CIRCUIT-LEVEL NOISE

TEJASVI SINGH TOMAR

Abstract. We consider the rotated surface code, decoded by minimum-weight perfect matching, with its standard eight-step syndrome-extraction circuit, under circuit-level depolarizing noise in which every gate, idle step, preparation and measurement fails independently with probability p. For the two usual log-likelihood weights, computed from the first-order or from the exact edge probabilities, we prove a threshold of at least 2×10^{-3}, and for all log-likelihood weights computed from a probability between $q(1-q)$ and the first-order edge probability q, a threshold of at least 1.6×10^{-3}. For p up to these values, the logical error rate per round of a memory experiment at odd distance d is at most an explicit bound that decays exponentially in d. The proof bounds the weighted paths that a decoding failure leaves in the decoding graph by a local criterion with memory for non-backtracking walks, after excluding short stretches that are incompatible with a minimum-weight correction chosen canonically in the proof; its computations are checked in exact arithmetic, and the results are also proved in the proof assistant Lean 4, with all computations checked by its kernel.

1. Setting and result

The rotated surface code decoded by minimum-weight perfect matching is a standard model of a fault-tolerant quantum memory. Under circuit-level noise, simulations place its threshold near 0.5–1% [9, 2], and for the circuit considered here the logical failure curves cross near 5×10^{-3} (Remark 11.6). Dennis, Kitaev, Landahl and Preskill reduced matching failure to self-avoiding walks and proved a threshold of 1.14% for independent qubit and measurement errors [1]. Rigorous thresholds for circuit-level noise are much smaller: they proved 1.7×10^{-4} per qubit and time step for the toric code with a CNOT syndrome-extraction circuit and worst-case Pauli faults, assuming that classical decoding takes one time step, and Fowler proved 7.4×10^{-4} for the planar code with nearest-neighbour circuits and unit weights [3], using the bound $14p/5$ on line probabilities of the circuits of [9]; for the circuit considered here his argument, with unit weights, gives at most 4.31×10^{-4}. General threshold theorems for local stochastic noise cover concatenated codes [6] and low-density parity-check codes [7, 8], and finite thresholds have recently been proved for the union-find and greedy decoders of the surface code under circuit-level local stochastic noise [10].

For the standard syndrome-extraction circuit we prove that the threshold is at least 1.04×10^{-3} (Theorem 1.1) and at least 1.6×10^{-3} (Theorem 1.2(a)) for all weights considered here, and at least 2×10^{-3} for profile weights, a class that contains both common choices of the weights (Theorem 1.2(b) and Section 1.3). The thresholds proved in Theorems 1.1 and 1.2(b) exceed 4.31×10^{-4} by factors of at least 2.4 and 4.6 (Remark 11.1). The proof of Theorem 1.1 checks one inequality for each of the 29 types of detectors in each basis, where two detectors have the same type if their edges have the same directions and probabilities; Theorem 1.2 rests on certificates checked by computer.

Date: 28 October 2026.

2020 *Mathematics Subject Classification.* 81P73; 94B35, 05C38, 68V05.

Key words and phrases. Quantum error correction, surface code, fault tolerance, threshold, minimum-weight perfect matching, circuit-level noise, computer-assisted proof.

All proofs start from the reduction of [1]: if the decoder fails, the symmetric difference of the error and the correction contains a self-avoiding path in the decoding graph from one boundary to the other, and the error edges carry at least half of its weight. For Theorem 1.1 we bound the sum over such paths, with each edge entering through its own probability rather than through the largest one, by a local criterion for weighted non-backtracking walks, which is sharp on trees. For Theorem 1.2 we label the edges of the path as error and correction edges and exclude the labellings of short stretches that are incompatible with a minimum-weight correction, chosen canonically for part (b) (Section 6). We count the remaining labelled walks with a positive potential on states of three consecutive labelled edges. The inequalities for this potential depend only on a bounded neighbourhood of each detector, so the instances with $d \le 25$ and $T \le 13$ contain all of them, and each certificate, checked in exact integer arithmetic, covers every odd distance and every number of rounds.

1.1. **Code and circuit.** Fix an odd distance $d \ge 3$. The data qubits sit at the points of

$$\mathcal{D} = \{(x, y) : x, y \in \{1, 3, \dots, 2d-1\}\}.$$

A point $m = (x, y)$ with $x, y \in \{0, 2, \dots, 2d\}$ has the diagonal neighbourhood $N(m) = \{m + (\pm 1, \pm 1)\} \cap \mathcal{D}$. It has *X type* if $x/2 + y/2$ is odd and *Z type* otherwise. The measure qubits are the points m with $|N(m)| \ge 2$, excluding X-type points with $x \in \{0, 2d\}$ and Z-type points with $y \in \{0, 2d\}$. There are $d^2 - 1$ of them. An X-type (Z-type) measure qubit m measures the stabilizer $\prod_{q \in N(m)} X_q$ ($\prod_{q \in N(m)} Z_q$). Write

$$\delta^X = \big((1,1), (-1,1), (1,-1), (-1,-1)\big), \qquad \delta^Z = \big((1,1), (1,-1), (-1,1), (-1,-1)\big).$$

One round of syndrome extraction consists of eight steps.

(S1) Reset every measure qubit to $|0\rangle$.
(S2) Apply H to every X-type measure qubit.
(S3–S6) In step $2 + k$, $k = 1, \dots, 4$, apply $\mathrm{CX}(m \to m + \delta^X_k)$ for each X-type m with $m + \delta^X_k \in \mathcal{D}$, and $\mathrm{CX}(m + \delta^Z_k \to m)$ for each Z-type m with $m + \delta^Z_k \in \mathcal{D}$.
(S7) Apply H to every X-type measure qubit.
(S8) Measure every measure qubit in the Z basis.

Each qubit takes part in at most one operation per step. The layout and the four CX layers coincide with those of the circuits `surface_code:rotated_memory_x/z` generated by Stim [4]. The orders δ^X and δ^Z are the standard ones that orient hook errors perpendicular to the logical operator they could shorten [5, 2].

In the *memory experiment in basis Z*, every data qubit is prepared in $|0\rangle$ in step (S1) of round 0. Rounds $0, \dots, T-1$ follow, for some $T \ge 1$, and then all data qubits are measured in the Z basis. Let $s_{m,r}$ be the outcome of m in round r and f_q the final outcome of q. For every Z-type m the *detectors* are

$$D_{m,0} = s_{m,0}, \qquad D_{m,r} = s_{m,r} \oplus s_{m,r-1} \ (1 \le r \le T-1), \qquad D_{m,T} = s_{m,T-1} \oplus \bigoplus_{q \in N(m)} f_q.$$

We call r the *layer* of $D_{m,r}$. The logical outcome is $\bigoplus_{q \in \mathcal{D},\, y(q)=1} f_q$, the value of Z_L on the bottom row. The experiment in basis X is the same with $|+\rangle$, X-basis final measurements, X-type detectors and X_L on the left column $x(q) = 1$. Without noise all detectors and the logical outcome equal 0 (Table 3, V1).

1.2. **Noise.** A *location* is one of the following. Each location fails independently with probability p.

- A reset or preparation. A failed reset to $|0\rangle$ (preparation of $|+\rangle$) is followed by X (by Z).
- A measurement. A failed measurement reports the flipped outcome.

- An H gate, or a qubit that is idle during a step. Failure is followed by X, Y or Z, each with probability 1/3.
- A CX gate. Failure is followed by one of the 15 non-identity two-qubit Paulis, each with probability 1/15.

A qubit is idle in a step if no reset, gate or measurement acts on it. Thus data qubits are idle in steps (S1) (except in round 0), (S2), (S7), (S8) and in each of (S3)–(S6) in which they have no CX partner. Z-type measure qubits are idle in (S2) and (S7) and in the CX steps without a partner, while X-type measure qubits are idle in the CX steps without a partner. A *fault outcome* is a location together with the Pauli or flip that follows its failure. Its probability is $p\,\pi(f)$ with $\pi(f) \in \{1, \frac{1}{3}, \frac{1}{15}\}$.

1.3. **Decoding graph and decoder.** We describe basis Z. Basis X is the same with the roles of the two types exchanged. Let V be the set of detectors $D_{m,r}$. The *signature* $\sigma(f) \subseteq V$ of a fault outcome f is the set of detectors that change when f is the only failure, and $\ell(f) \in \{0,1\}$ records whether the logical outcome changes. By Lemmas 2.1 and 2.3, $|\sigma(f)| \le 2$, and $\ell(f)$ depends only on $\sigma(f)$. The *decoding graph* G has vertex set $V \cup \{b_1, b_2\}$ and one edge $e(f)$ for each nonempty signature: $e(f) = \sigma(f)$ if $|\sigma(f)| = 2$, and $e(f) = \{u, b_1\}$ or $\{u, b_2\}$ if $\sigma(f) = \{u\}$ and $\ell(f) = 1$ or $\ell(f) = 0$, respectively. For an edge e put

$$c_e = \sum_{f\colon e(f)=e} \pi(f), \qquad q_e = c_e\, p.$$

The weights are $w_e = \log\bigl((1-r_e)/r_e\bigr)$ for numbers r_e with $q_e(1-q_e) \le r_e \le q_e$. Two choices are common: $r_e = q_e$, and r_e equal to the probability that an odd number of outcomes with $e(f) = e$ occur. The second lies in this range by the Bonferroni inequality. For a location ℓ let $c_{\ell,e}$ be the sum of $\pi(f)$ over the outcomes f of ℓ with $e(f) = e$. We call the multiset of the nonzero numbers $c_{\ell,e}$, over all locations ℓ, the *profile* of e. Because the outcomes of a location exclude each other, the second choice equals $\bigl(1 - \prod_\ell (1 - 2c_{\ell,e}p)\bigr)/2$. Thus both choices depend only on p and on the profile of e. We call weights as above *profile weights* if $r_e = \psi(p, \text{profile of } e)$ for one function ψ. For $p \le 2 \times 10^{-3}$ each q_e is at most $\frac{154}{15} \cdot 2 \times 10^{-3} < 0.021$ (Lemmas 2.1 and 2.3), so all weights are positive.

For a set C of edges let $\partial C \subseteq V$ be the set of detectors of odd degree in C. Given the set $\mathcal{A} \subseteq V$ of detectors that equal 1, the decoder returns a set C of minimum weight with $\partial C = \mathcal{A}$, and it predicts that the logical outcome was flipped if and only if $|C \cap \delta(b_1)|$ is odd, where $\delta(b_1)$ is the set of edges at b_1. A minimum-weight perfect matching of $\mathcal{A}$, with boundary copies and shortest-path distances, yields such a set as the symmetric difference of its paths. We prove the results below for every minimizer C. The decoder *fails* if its prediction differs from the actual flip of the logical outcome.

1.4. **Result.**

Theorem 1.1. *Let $d \ge 3$ be odd, $T \ge 1$, $B \in \{X, Z\}$ and $0 < p \le p_* = 1.04 \times 10^{-3}$. Put $\lambda = 31\sqrt{p}$. The memory experiment in basis B with T rounds, decoded as in Section 1.3, fails with probability*

$$P_{\text{fail}} \le \frac{(T+1)(d+1)}{2} \cdot 2\sqrt{\frac{c_{\max} p}{1 - c_{\max} p}} \cdot \frac{\lambda^{d-1}}{1-\lambda} \le 0.104\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}, \qquad c_{\max} = \tfrac{154}{15}.$$

We call P_{fail}/T the logical error rate per round of the memory experiment in basis B. Since $T + 1 \le 2T$, it obeys $P_{\text{fail}}/T \le 0.208\,(d+1)\lambda^d/(1-\lambda)$ uniformly in T. For $p \le p_*$ we have $\lambda \le 0.99972$, so this decays exponentially in d. At $p = 10^{-4}$ the bound reads $P_{\text{fail}}/T \le$

$0.31\,(d+1)\,(0.31)^d$; it is below 1 for all $d \geq 3$ when $p \leq 5 \times 10^{-4}$, but only for $d \geq 63063$ when $p = p_*$.

Theorem 1.2. *Let $d \geq 3$ be odd, $T \geq 1$ and $B \in \{X, Z\}$, and consider the memory experiment in basis B with T rounds, decoded as in Section 1.3.*

(a) *If $0 < p \leq 1.6 \times 10^{-3}$ and $\lambda = 0.997\left(p/(1.6 \times 10^{-3})\right)^{21/50}$, then*

$$P_{\text{fail}} \leq 0.032\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}.$$

(b) *If the weights are profile weights, $0 < p \leq p_{**} = 2 \times 10^{-3}$ and $\lambda = 0.999\,(p/p_{**})^{9/25}$, then*

$$P_{\text{fail}} \leq 0.061\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}.$$

Part (b) implies $P_{\text{fail}}/T \leq 0.122\,(d+1)\lambda^d/(1-\lambda)$, and $\lambda \leq 0.999$ for $p \leq p_{**}$. For $p < 1.92 \times 10^{-4}$, however, the factor $\lambda = 31\sqrt{p}$ of Theorem 1.1 is smaller than that of part (b). The exponents $21/50$ and $9/25$ are parameters of the two certificates, the tilts s of Section 6.

The circuit of Section 1.1 has eight steps per round. A common leaner schedule measures and resets the measure qubits in a single step and prepares and measures the X-type ones in the X basis directly (Section 7). It has fewer idle steps, and the same two arguments yield larger thresholds.

Theorem 1.3. *For the lean schedule of Section 7, with the noise model, memory experiments and decoder of Sections 1.2 and 1.3, let $d \geq 3$ be odd, $T \geq 1$ and $B \in \{X, Z\}$.*

(a) *If $0 < p \leq 1.53 \times 10^{-3}$ and $\lambda = 25.55\sqrt{p}$, then*

$$P_{\text{fail}} \leq 0.0985\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}.$$

(b) *If $0 < p \leq 2.4 \times 10^{-3}$ and $\lambda = 0.997\left(p/(2.4 \times 10^{-3})\right)^{21/50}$, then*

$$P_{\text{fail}} \leq 0.028\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}.$$

(c) *If the weights are profile weights, $0 < p \leq 3 \times 10^{-3}$ and $\lambda = 0.997\left(p/(3 \times 10^{-3})\right)^{9/25}$, then*

$$P_{\text{fail}} \leq 0.052\,(T+1)(d+1)\,\frac{\lambda^d}{1-\lambda}.$$

At $p = 10^{-3}$ the bounds on P_{fail}/T compare as follows:

	λ at $p = 10^{-3}$	bound on P_{fail}/T at $p = 10^{-3}$
Theorem 1.1	0.98031	$10.6\,(d+1)\,\lambda^d$, below 1 for $d \geq 423$
Theorem 1.2(a)	< 0.8185	$0.353\,(d+1)\,(0.8185)^d$
Theorem 1.2(b)	< 0.7784	$0.551\,(d+1)\,(0.7784)^d$, below 1 for $d \geq 5$ and below 10^{-6} for $d \geq 71$
Theorem 1.3(c)	< 0.6714	below 0.39 for all $d \geq 3$ and below 10^{-6} for $d \geq 43$

Section 8 extends both arguments to error rates that depend on the type of operation (Theorem 8.1) and applies them to six error models built from the published component error rates of six devices (Appendix B). With certificates for these models, Theorem 8.1 proves exponential decay for the models of the Quantinuum Helios and H2 data sheets and for one with the error rates of IBM's Heron r3 processor ibm_boston placed on a square grid. The models of Google's and USTC's processors are not certified: for the tilts we tried, floating-point estimates indicate that no positive potential satisfies the inequalities of Lemma 6.5 with $\lambda < 1$.

Outline. Section 2 describes the decoding graph and its node types: each fault outcome produces at most one edge, and the bulk has twelve edges per node. In Section 3 we reduce failure to weighted paths from b_1 to b_2. We prove the local criterion in Section 4 and combine the pieces into Theorem 1.1 in Section 5. Section 6 proves Theorem 1.2, Section 7 treats the lean schedule, and Section 8 treats non-uniform error rates. Computer-verified statements carry the mark †, and Section 9 lists the computations behind them. Section 10 describes the formal proofs in Lean 4, which cover Theorems 1.1–1.3 in full and Theorem 8.1 for the models it certifies. Section 11 compares our argument with Fowler's and discusses the limits of the method, other schedules, the weights, simulations and Google's Willow schedule. Appendix A lists the node types, and Appendix B gives the sources of the device error models and the details of the Willow comparison.

2. The decoding graph

We write $G_{d,T}$ for the decoding graph of the experiment with distance d and T rounds, in basis Z or X. The *height* of a detector $D_{m,r}$ is the second coordinate y of m in basis Z and the first coordinate x in basis X.

Lemma 2.1 (†)**.** *For $B \in \{X, Z\}$, odd $d \le 13$ and $T \le 7$ the following hold.*

(a) *Every fault outcome flips at most two detectors.*
(b) *Outcomes with the same signature have the same value of ℓ, and $\ell(f) = 1$ only if $|\sigma(f)| = 1$.*
(c) *The edges at b_1 join it to detectors of height 2, and the edges at b_2 join it to detectors of height $2d - 2$. No detector is adjacent to both.*
(d) *The heights of the two ends of an edge in V differ by at most 2.*
(e) *b_1 has exactly $(T + 1)(d + 1)/2$ neighbours, and $c_e \le 154/15$ for every edge e.*

Part (a) says that the circuit-level noise decomposes into a graph in each sector. Here each Y component splits into its X part, seen by Z-type detectors, and its Z part, seen by X-type detectors. Lemma 2.3 below will show that Lemma 2.1 holds for all odd $d \ge 3$ and all $T \ge 1$.

2.1. **The bulk.** Let m be a Z-type measure qubit far from the boundary and $1 \le r \le T - 1$ far from 0 and T. Then the data qubit $m + \delta_k^Z$ meets m in step $2 + k$. There are two kinds of data qubits. A data qubit q of kind A has Z-type neighbours $q \mp (1, 1)$, which it meets in steps (S3) and (S6), and X-type neighbours $q + (1, -1)$, $q + (-1, 1)$, which it meets in (S4) and (S5). A data qubit of kind B has Z-type neighbours $q + (-1, 1)$, $q + (1, -1)$, met in (S4) and (S5), and X-type neighbours $q \mp (1, 1)$, met in (S3) and (S6).

Lemma 2.2. *In basis Z, a bulk detector $v = D_{m,r}$ has exactly twelve edges. Their other ends, written as offsets $(\Delta x, \Delta y, \Delta r)$, and their coefficients are*

offset	$\pm(2,-2,0)$	$\pm(0,0,1)$	$\pm(2,2,0)$	$\pm(2,2,-1)$	$\pm(2,-2,-1)$	$\pm(4,0,-1)$
$15\,c_e$	72	66	48	16	8	8

In basis X the coefficients are the same, with offsets $\pm(2,-2,0)$, $\pm(0,0,1)$, $\pm(2,2,0)$, $\pm(2,2,-1)$, $\pm(2,-2,1)$ and $\pm(0,4,-1)$.

Proof. We treat basis Z, where only X components matter, and list the outcomes that produce each edge in Table 1, in units of $p/15$. Three propagation rules determine the list. First, a CX copies an X on its control to its target. Second, an X on the target of a CX is not copied. Third, an X on an X-type measure qubit before its first CX acts trivially, because that qubit is in $|+\rangle$ at that point. By the first rule, an X on such a qubit between its k-th and $(k+1)$-st CX

is copied onto its remaining data qubits. For $k = 1$ this equals an X on its first data qubit, for $k = 2$ it is the hook pair, and for $k = 3$ it is an X on its last data qubit.

An X on a data qubit q changes the next outcome of each Z-type neighbour that has not yet met q in the current round, and the outcome in the next round of each neighbour that has. Hence an X on a kind-A qubit between (S6) of round $r-1$ and (S3) of round r flips its two neighbours in layer r, and it produces the edge with offset $(2, 2, 0)$. An X between (S3) and (S6) of round r flips $q + (1, 1)$ in layer r and $q - (1, 1)$ in layer $r + 1$, and it gives the offset $(-2, -2, 1)$. Similarly, for kind B the windows are (S5) of round $r-1$ to (S4) of round r, with offset $(2, -2, 0)$, and (S4) to (S5), with offset $(-2, 2, 1)$. The hook pair of an X-type measure qubit m' is $\{m' + (1, -1), m' + (-1, -1)\}$. It flips $m' + (2, 0)$ in layer r and $m' + (-2, 0)$ in layer $r + 1$, while the flips of $m' + (0, -2)$ cancel, which leads to the offset $(-4, 0, 1)$. Finally, an X on the Z-type qubit m after its reset changes $s_{m,r}$ only, which gives the offset $(0, 0, 1)$. Applying these rules to all locations in the neighbourhood, we obtain the entries of Table 1. For a CX we record which of the 15 outcomes contribute, as pairs (control part, target part).

The same derivation with Z components yields the offsets and coefficients stated for basis X. □

edge	contributing outcomes (units of $p/15$)	$15c_e$
$(0, 0, \pm 1)$	reset flip 15; idle in (S2) 10; each of the 4 CX: $(I/Z, X/Y)$ on (data, m) 4; idle in (S7) 10; measurement flip 15	66
$\pm(2, 2, 0)$	kind A, q: CX in (S6) of round $r-1$, $(X/Y, I/Z)$ 4; idles in (S7), (S8), (S1), (S2) 4×10; CX in (S3), $(X/Y, X/Y)$ 4	48
$\pm(2, -2, 0)$	kind B, q: CX in (S5) of round $r-1$, $(X/Y, I/Z)$ 4; X-type $q + (1, 1)$ in (S5) of round $r-1$, $(X/Y, I/Z)$ 4; its CX onto q in (S6), X/Y on q 8; idles 4×10; X-type $q - (1, 1)$: CX onto q in (S3), $(I/Z, X/Y)$ and $(X/Y, I/Z)$ 8, its CX in (S4), $(X/Y, X/Y)$ 4; CX $q \to q + (-1, 1)$ in (S4), $(X/Y, X/Y)$ 4	72
$\pm(2, 2, -1)$	kind A, q between (S3) and (S6): CX in (S3) $(X/Y, I/Z)$ 4; CX onto q in (S4) and (S5), $(I/Z, X/Y)$ $4 + 4$; CX in (S6) $(X/Y, X/Y)$ 4	16
$\pm(2, -2, -1)$	kind B, q between (S4) and (S5): CX in (S4) $(X/Y, I/Z)$ 4; CX in (S5) $(X/Y, X/Y)$ 4	8
$\pm(4, 0, -1)$	X-type m' after (S4): its CX in (S4) $(X/Y, I/Z)$ 4; its CX in (S5) $(X/Y, X/Y)$ 4	8

TABLE 1. Outcomes producing the twelve edges at a bulk detector (basis Z, Lemma 2.2). The kind-A and kind-B qubits are the data qubits shared by the two ends of the edge.

2.2. **Locality and node types.** The *type* of a detector $v \in V$ is the multiset of pairs (offset of the other end, c_e) over the edges e at v, where an edge to b_1 or b_2 is recorded with the label b_1 or b_2 in place of the offset.

Lemma 2.3. *Let $v = D_{m,r}$.*

(a) *If $v \in \sigma(f)$, then f acts on a qubit at ℓ^∞-distance at most 2 from m, in round $r-1$ or r (in round $T-1$ or in the final measurement if $r = T$).*

(b) *The type of v is determined by the positions of the qubits at ℓ^∞-distance at most 6 from m and the X or Z types of the measure qubits among them, by $\min(r, 2)$ and by $\min(T-r, 2)$.*

(c) *For all odd $d \geq 3$ and all $T \geq 1$, every detector of $G_{d,T}$ has the type of a detector of $G_{d',T'}$ for some $d' \in \{3,5,7,9\}$ and $T' \in \{1,\dots,5\}$. The same reduction applies to the statements of Lemma 2.1.*

Proof. (a) We use the propagation rules from the proof of Lemma 2.2. In basis Z only X components matter, and basis X is the same with Z components and the roles of the two types of measure qubits exchanged. An X on a data qubit q is copied by the gate $\mathrm{CX}(q \to m')$ onto each Z-type measure qubit m' with $q \in N(m')$, where it flips the outcome of m' and is removed by the next reset of m'. It is copied nowhere else, because q is the target of its CX gates with X-type measure qubits. An X on a Z-type measure qubit flips its outcome and is removed by the next reset, whereas an X on an X-type measure qubit m' is copied onto the data qubits of $N(m')$ that m' meets later in the same round. The latter does not change the outcome of m', and it is removed by the next reset. Hence the propagated X part of a single-qubit outcome lies, on data qubits, within distance 1 of the faulty qubit, and the detectors it flips lie within distance 2 of it. A two-qubit outcome is the product of its two factors, and detector values are linear in the Pauli frame. In time, a measure qubit is reset at the start of each round, and a data error that is present from round ρ on changes $s_{m',\rho'}$ for all $\rho' \geq \rho$ or for all $\rho' \geq \rho+1$, so it flips layers ρ and $\rho+1$ only. The final layer compares the last round with the data measurement, which sees the same error.

(b) By (a), the outcomes with v in their signature act within distance 2 of m, so their CX partners lie within distance 3. By the rules in (a), the signature and the logical effect of such an outcome in round ρ are determined by the steps at which CX gates act, in rounds ρ and $\rho+1$, on its qubits, on the data qubits within distance 1 of them and on the measure qubits adjacent to those. Hence all of these qubits lie within distance 5 of m. Which gates act on a qubit, and in which step, depends only on which of its diagonal neighbours exist, and these lie within distance 6 of m. Whether round $r-1$ is round 0, and whether layer r or $r+1$ is the final one, is recorded by $\min(r,2)$ and $\min(T-r,2)$.

(c) The bulk pattern of qubit positions and measure-qubit types is invariant under the lattice generated by $(4,0)$, $(0,4)$ and $(2,2)$, and the measure qubits of each of the two types form a single coset of it. Along the side $y=0$ the X-type measure qubits sit at $x \equiv 2 \pmod 4$, along $y=2d$ at $x \equiv 0$, along $x=0$ the Z-type ones sit at $y \equiv 0$ and along $x=2d$ at $y \equiv 2 \pmod 4$. For odd d this pattern, measured from the nearest corner, does not depend on d. Hence, for a detector in a given basis, the configuration within distance 6 of m is determined by the distances from m to the four sides, truncated at 7. If $d \geq 9$, each quadruple of truncated distances that occurs in $G_{d,T}$ also occurs in $G_{9,T}$. Indeed, the range $[7,11]$ of coordinates at distance at least 7 from both opposite sides of G_9 contains the even coordinates 8 and 10, one of each residue modulo 4, so the parity condition that makes m X-type or Z-type can be met there. Likewise, for $T \geq 5$ the layer data $(\min(r,2), \min(T-r,2))$ of all layers occur for $T'=5$. Part (b) then implies the claim for the types of detectors. The statements of Lemma 2.1 concern single outcomes, edges and the neighbours of b_1, and they reduce in the same way. For the logical effect, after round $\rho+1$ the propagated frame of an outcome in round ρ no longer changes on the data qubits, so $\ell(f)$ is the parity of its X part on the bottom row at that time, which is again determined locally. An outcome with empty signature is handled in the same way, with the configuration around the outcome in place of the configuration around a detector. Finally, the number of neighbours of b_1 is a sum over detectors. By the types, a detector is adjacent to b_1 exactly when it has height 2, and each layer has $(d+1)/2$ detectors of height 2. □

Corollary 2.4 (†)**.** *In each basis the detectors of all graphs $G_{d,T}$, $d \geq 3$ odd and $T \geq 1$, have exactly 29 types. They are listed in Appendix A.*

3. From matching failure to weighted paths

Fix an experiment and let $\mathcal{F}$ be the set of fault outcomes that occur. At most one outcome occurs per location. Let E be the set of edges e such that the number of $f \in \mathcal{F}$ with $e(f) = e$ is odd. Since detector values are linear in the Pauli frame, the set of detectors equal to 1 is ∂E. The logical outcome is linear in the Pauli frame as well, and by Lemma 2.1(b) an outcome flips it if and only if its edge ends at b_1. Therefore the logical outcome flips if and only if $|E \cap \delta(b_1)|$ is odd.

Lemma 3.1. *If the decoder fails with the minimizer C, then $E \triangle C$ contains a simple path from b_1 to b_2.*

Proof. Let $R = E \triangle C$. Since $\partial C = \partial E$, each detector has even degree in R. Failure means that $|C \cap \delta(b_1)|$ and $|E \cap \delta(b_1)|$ have different parity, so b_1 has odd degree in R. As the degrees in R sum to $2|R|$, the vertex b_2 has odd degree as well. In each connected component of $(V \cup \{b_1, b_2\}, R)$ the number of odd-degree vertices is even. Hence b_1 and b_2 lie in the same component, and that component contains a simple path between them. □

Lemma 3.2. *Every simple path from b_1 to b_2 in $G_{d,T}$ has at least d edges.*

Proof. By Lemma 2.1(c,d), the path enters V at height 2, leaves it from height $2d-2$, and each of its edges inside V changes the height by at most 2. Thus it has at least $(2d-4)/2 = d-2$ edges inside V, and two edges at the boundary vertices. □

Lemma 3.3. *Let $\gamma \subseteq E \triangle C$ be a simple path from b_1 to b_2. Then $w(\gamma \cap E) \geq w(\gamma)/2$.*

Proof. We follow the argument of [1, Sec. 4]. Because the detector boundary of γ is empty, $C \triangle \gamma$ is also a valid output, and $w(C) \leq w(C \triangle \gamma) = w(C) - w(C \cap \gamma) + w(\gamma \setminus C)$. Since $\gamma \subseteq E \triangle C$, we have $\gamma \setminus C = \gamma \cap E$ and $w(C \cap \gamma) = w(\gamma) - w(\gamma \cap E)$. □

Lemma 3.4. *For every set S of edges, $\Pr(S \subseteq E) \leq \prod_{e \in S} q_e$.*

Proof. If $S \subseteq E$, then for each $e \in S$ some $f_e \in \mathcal{F}$ has $e(f_e) = e$. The f_e are distinct, because each outcome has one edge. Two distinct outcomes at the same location never occur together, while outcomes at distinct locations are independent. Hence

$$\Pr(S \subseteq E) \leq \sum_{(f_e)_{e \in S}} \Pr(f_e \in \mathcal{F} \ \forall e) \leq \sum_{(f_e)_{e \in S}} \prod_{e \in S} p\, \pi(f_e) = \prod_{e \in S} q_e. \qquad \square$$

Lemma 3.5. *Put $a_e = 2\sqrt{q_e/(1-q_e)}$. For every path γ,*

$$\sum_{S \subseteq \gamma,\ w(S) \geq w(\gamma)/2} \prod_{e \in S} q_e \ \leq \ \prod_{e \in \gamma} a_e.$$

Proof. Each term on the left satisfies $w(S) - w(\gamma \setminus S) \geq 0$, so it is at most $\prod_{e \in S} q_e\, e^{(w(S) - w(\gamma \setminus S))/2}$. Summing this over all $S \subseteq \gamma$, we obtain $\prod_{e \in \gamma} \big(q_e e^{w_e/2} + e^{-w_e/2}\big)$. With $e^{w_e} = (1-r_e)/r_e$ and $q_e(1-q_e) \leq r_e \leq q_e$ we get

$$q_e \sqrt{\frac{1-r_e}{r_e}} \leq \frac{q_e}{\sqrt{q_e(1-q_e)}} = \sqrt{\frac{q_e}{1-q_e}}, \qquad \sqrt{\frac{r_e}{1-r_e}} \leq \sqrt{\frac{q_e}{1-q_e}},$$

so each factor is at most a_e. □

Proposition 3.6. *Let Γ be the set of simple paths from b_1 to b_2 in $G_{d,T}$. Then $P_{\text{fail}} \leq \sum_{\gamma \in \Gamma} \prod_{e \in \gamma} a_e$.*

Proof. By Lemmas 3.1 and 3.3, failure implies that some $\gamma \in \Gamma$ and $S = \gamma \cap E$ satisfy $S \subseteq E$ and $w(S) \geq w(\gamma)/2$. We apply the union bound over γ and S, then Lemmas 3.4 and 3.5. □

4. A local criterion for non-backtracking walks

Lemma 4.1. *Let G be a finite graph with vertex set $V \cup \{b_1, b_2\}$, without an edge between b_1 and b_2, and let $a_e > 0$ for each edge e. Let $\lambda > 0$ and suppose that*

$$\Sigma(v) := \sum_{e \ni v} \frac{a_e}{\lambda + a_e} \le 1 \qquad \text{for every } v \in V,$$

where the sum includes edges to b_1 and b_2. Put $\kappa = \lambda + \max_{e \ni b_2} a_e$ and $\Sigma(b_1) = \sum_{e \ni b_1} a_e/(\lambda + a_e)$. Then for every $L \ge 2$ the simple paths from b_1 to b_2 with L edges satisfy

$$\sum_{\gamma} \prod_{e \in \gamma} a_e \;\le\; \kappa\, \Sigma(b_1)\, \lambda^{L-1}.$$

Proof. For an oriented edge $\varepsilon = (u \to v)$ put $h(\varepsilon) = 1/(\lambda + a_\varepsilon)$. For $m \ge 1$ and an oriented edge $(x_0 \to x_1)$ with $x_0 \in V \cup \{b_1\}$ and $x_1 \in V$, let $W_m(x_0 \to x_1)$ be the sum of $\prod_{i=1}^{m} a_{x_i x_{i+1}}$ over all walks $x_1, x_2, \dots, x_{m+1}$ with $x_{m+1} = b_2$, $x_2, \dots, x_m \in V$, consecutive vertices adjacent and $x_{i+1} \ne x_{i-1}$ for $1 \le i \le m$. We show by induction on m that

$$W_m(x_0 \to x_1) \le \kappa\, \lambda^m\, h(x_0 \to x_1). \tag{4.1}$$

Because the edge $\{x_0, x_1\}$ contributes $a_{x_0x_1}/(\lambda + a_{x_0x_1})$ to $\Sigma(x_1)$, the hypothesis implies

$$\sum_{e \ni x_1,\ e \ne \{x_0, x_1\}} \frac{a_e}{\lambda + a_e} \le 1 - \frac{a_{x_0x_1}}{\lambda + a_{x_0x_1}} = \lambda\, h(x_0 \to x_1). \tag{4.2}$$

For $m = 1$, $W_1(x_0 \to x_1) = a_{x_1b_2}$ if x_1 is adjacent to b_2 and 0 otherwise. Since $a_{x_1b_2} = (\lambda + a_{x_1b_2}) \cdot a_{x_1b_2}/(\lambda + a_{x_1b_2}) \le \kappa\, a_{x_1b_2}/(\lambda + a_{x_1b_2})$, (4.1) follows from (4.2). For $m \ge 2$,

$$W_m(x_0 \to x_1) = \sum_{\substack{x_2 \in V \setminus \{x_0\} \\ x_2 \sim x_1}} a_{x_1x_2} W_{m-1}(x_1 \to x_2) \le \kappa \lambda^{m-1} \sum_{x_2} \frac{a_{x_1x_2}}{\lambda + a_{x_1x_2}} \le \kappa \lambda^m h(x_0 \to x_1),$$

by induction and (4.2). A simple path $b_1, x_1, \dots, x_{L-1}, b_2$ is a walk of this kind, so the sum in the lemma is at most $\sum_{x_1 \sim b_1} a_{b_1x_1} W_{L-1}(b_1 \to x_1) \le \kappa \lambda^{L-1} \Sigma(b_1)$. □

On the k-regular tree with all $a_e = a$ the criterion reads $ka/(\lambda + a) \le 1$, that is $\lambda \ge (k-1)a$. Since this is the exact growth rate of non-backtracking walks there, the criterion is sharp on trees. On a lattice with unequal edges it gives the growth rate of weighted non-backtracking walks exactly: if all nodes have the same multiset $\{a_e\}$ and $\Sigma \equiv 1$, then h is a positive eigenvector of the weighted non-backtracking operator with eigenvalue λ.

5. Proof of Theorem 1.1

5.1. **The criterion holds at every detector.** Let $\lambda = 31\sqrt{p}$ and write $a_e = \sqrt{p}\, \alpha_e(p)$ with $\alpha_e(p) = 2\sqrt{c_e/(1 - c_e p)}$. Then

$$\Sigma(v) = \sum_{e \ni v} \frac{\alpha_e(p)}{31 + \alpha_e(p)},$$

which increases with p. Thus it suffices to check $\Sigma(v) \le 1$ at $p = p_*$. There $31\sqrt{p_*} = 0.999719\ldots \ge 0.9997$, and $\Sigma(v)$ decreases in λ, so it suffices to check

$$\sum_{e \ni v} \frac{a_e}{0.9997 + a_e} \le 1 \qquad \text{at } p = p_*. \tag{5.1}$$

By Corollary 2.4 we need to check only 29 types per basis. For the bulk type of Lemma 2.2 the computation is short:

$15\,c_e$	72	66	48	16	8
multiplicity	2	2	2	2	4
a_e	0.141662	0.135603	0.115570	0.066650	0.047116
$a_e/(0.9997+a_e)$	0.124117	0.119442	0.103625	0.062503	0.045009

and the weighted sum is 0.999409. The other 28 types satisfy (5.1) with left side at most $0.9629^{\dagger}$. We list the values for all types in Appendix A. The check is done in exact rational arithmetic, with rational upper bounds for the square roots (Table 3, V6).

5.2. **Assembly.** We first check that the graph $G_{d,T}$ satisfies the hypotheses of Lemma 4.1. Each edge comes from a nonempty signature, which contains a detector, so there is no edge between b_1 and b_2. Since there is one edge per signature, there are no parallel edges, and the number of edges at b_1 equals the number of its neighbours. Combining Proposition 3.6, Lemma 3.2 and Lemma 4.1, we obtain

$$P_{\mathrm{fail}} \le \sum_{L \ge d} \kappa\, \Sigma(b_1)\, \lambda^{L-1} = \kappa\, \Sigma(b_1) \frac{\lambda^{d-1}}{1-\lambda}.$$

Let $a_{\max} = 2\sqrt{c_{\max} p/(1-c_{\max} p)}$. By Lemma 2.1(e), $\kappa \le \lambda + a_{\max}$ and each term of $\Sigma(b_1)$ is at most $a_{\max}/(\lambda + a_{\max})$, so $\kappa\, \Sigma(b_1) \le \frac{(T+1)(d+1)}{2}\, a_{\max}$. This is the first bound of the theorem. For $p \le p_*$ we have $1 - c_{\max} p \ge 0.98932$ and $a_{\max} \le 6.4429\sqrt{p}$, so

$$\begin{aligned} P_{\mathrm{fail}} &\le \frac{(T+1)(d+1)}{2} \cdot 6.4429\sqrt{p} \cdot \frac{\lambda^{d-1}}{1-\lambda} \\ &= \frac{6.4429}{62}\,(T+1)(d+1)\frac{\lambda^{d}}{1-\lambda} \le 0.104\,(T+1)(d+1)\frac{\lambda^{d}}{1-\lambda}. \end{aligned}$$

□

6. Proof of Theorem 1.2

In the proof of Theorem 1.1 we sum over all ways of labelling the edges of a path as errors or corrections, although most of these labellings cannot occur. Indeed, the correction edges on a stretch of the failure path belong to a minimum-weight correction. Their weight is therefore at most that of the error edges on the stretch plus any detour between its ends, and at most that of any set of edges with the same boundary. We turn these two inequalities into local exclusion rules (Definition 6.2) and count only the labelled paths that obey them. For profile weights a third rule uses the freedom in the choice of the minimizer: among the minimizers for which the decoder fails we take one that minimizes a fixed linear functional, and such a minimizer cannot turn a corner around which an equally heavy route of smaller value exists. These rules are the main new ingredient of the proof. For part (a) we fix $p_{**} = 1/625$ and the tilt $s = 21/50$, and for part (b) $p_{**} = 1/500$ and $s = 9/25$.

6.1. **Windows.** For a path or walk P we write $c(P) = \prod_{e \in P} c_e$, with multiplicity, and G/B for the graph obtained from G by identifying b_1 and b_2 to a single vertex B. Let C be a minimizer and $\gamma \subseteq E \triangle C$ the path of Lemma 3.1. A *window* is a sub-path σ of γ. We write $\sigma_C = \sigma \cap C$ and $\sigma_E = \sigma \cap E$, so that σ is the disjoint union of σ_C and σ_E.

Lemma 6.1. (a) *For every window σ and every set P' of edges of G/B with the same detector boundary as σ, $w(\sigma_C) - w(\sigma_E) \le w(P')$.*

(b) *For every $X \subseteq C$ and every set P' of edges of G/B with the same detector boundary as X, $w(X) \le w(P')$.*

Proof. (a) The set $C' = C \triangle \sigma \triangle P'$ satisfies $\partial C' = \partial C$. Since $C \triangle \sigma = (C \setminus \sigma_C) \cup \sigma_E$, we have $w(C') \le w(C) - w(\sigma_C) + w(\sigma_E) + w(P')$, and $w(C) \le w(C')$.

(b) Similarly, the set $C' = C \triangle X \triangle P'$ satisfies $\partial C' = \partial C$ and $w(C') \le w(C) - w(X) + w(P')$, and $w(C) \le w(C')$. □

For a detector $D_{m,r}$ with $m = (x, y)$ put $\phi(D_{m,r}) = x + 9y + 81r$, and for a set C of edges let $\Phi(C)$ be the sum of $\phi(u) + \phi(v)$ over the edges $\{u, v\}$ of C with $u, v \in V$.

A *labelled window* is a path $\sigma = (v_0, \dots, v_m)$ in G with $2 \le m \le 4$ and a label E or C on each edge. We write n_C and n_E for the numbers of its C- and E-edges, σ_C and σ_E for the sets of these edges, and $T(\sigma)$ for the set of detectors of odd degree in σ_C. The *head* of σ is v_{m-1}. The *ball* of a detector v is the set of detectors joined to v by a chain of at most four edges between detectors. Since the detectors of σ lie within three such edges of its head, they belong to the ball of the head. An *admissible walk* for σ is a walk in G/B, with b_1 and b_2 read as B, that visits B at most once and whose detectors all lie in the ball of the head.

Definition 6.2. A labelled window σ is *excluded* if it satisfies one of the following conditions.

(M) There is an admissible walk P' from v_0 to v_m with h' edges, empty if $v_0, v_m \in \{b_1, b_2\}$, such that $k = n_C - n_E - h' \ge 0$ and

$$\Big(\frac{1}{p_{**}}\Big)^k \frac{c(\sigma_E)\, c(P')}{c(\sigma_C)} > \prod_{e \in \sigma} \frac{1}{1 - c_e\, p_{**}} \cdot \Big(\frac{1}{1 - c_{\max}\, p_{**}}\Big)^{h'}. \tag{6.1}$$

(J) There are admissible walks, each from a vertex of $T(\sigma)$ to another vertex of $T(\sigma)$ or to B, such that each vertex of $T(\sigma)$ is an end of exactly one of them, their union P' has h' edges counted with multiplicity, $k = n_C - h' \ge 0$ and

$$\Big(\frac{1}{p_{**}}\Big)^k \frac{c(P')}{c(\sigma_C)} > \prod_{e \in \sigma_C} \frac{1}{1 - c_e\, p_{**}} \cdot \Big(\frac{1}{1 - c_{\max}\, p_{**}}\Big)^{h'}. \tag{6.2}$$

If $T(\sigma)$ is empty, P' is empty.

(S) (Only for Theorems 1.2(b) and 1.3(c).) The window has $m = 2$, its vertices $u = v_0$, $b = v_1$ and $w = v_2$ are detectors, both of its edges carry the label C, and some detector $x \notin \{u, b, w\}$ is adjacent to u and to w, with $\phi(x) < \phi(b)$, such that the profiles of ux and xw are those of ub and bw, in some order.

Rule (M) compares the window with an alternative route between its ends. Rule (J) compares its correction edges with a re-pairing of their endpoints. Rule (S) compares the corner u, b, w with the opposite corner u, x, w of a square. Restricting the walks to the ball of the head keeps (M) and (J) local, and the detector x of (S) lies within two edges of the head b (Lemma 6.6). Because dropping walks only drops exclusions, any restriction of this kind would do.

Lemma 6.3. (a) *Let $p \le p_{**}$, let the weights be as in Section 1.3, and let C be a minimizer. Label the edges of $\gamma \cap E$ by E and those of $\gamma \cap C$ by C. Then no window of γ, with these labels, is excluded by (M) or (J).*

(b) *Let $p \le p_{**}$, let the weights be profile weights and suppose that the decoder fails. Let C^* be a minimizer for which the decoder fails, with the smallest value of Φ among these minimizers, and label the path $\gamma \subseteq E \triangle C^*$ as in (a). Then no window of γ is excluded by (S).*

Proof. (a) Suppose that a window σ of γ is excluded, and let P' be the walk or the union of walks of Definition 6.2. Let P'' be the set of edges that P' uses an odd number of times, so that $w(P'') \le w(P')$. Put $\varepsilon_e = -\ln(1 - q_e)$. Since $q_e(1 - q_e) \le r_e \le q_e$, each weight satisfies $|w_e + \ln q_e| \le \varepsilon_e$, and each ε_e with $e \in P'$ is at most $-\ln(1 - c_{\max} p)$.

Under (M), P'' has the same detector boundary as σ, because b_1, b_2 and B are not detectors, and we have

$$w(\sigma_C) - w(\sigma_E) - w(P') \geq k \ln \frac{1}{p} + \ln \frac{c(\sigma_E)\, c(P')}{c(\sigma_C)} - \sum_{e \in \sigma} \varepsilon_e - \sum_{e \in P'} \varepsilon_e.$$

At $p = p_{**}$ the right side is positive by (6.1). As p decreases, the first term does not decrease, because $k \geq 0$, and each ε_e decreases, so the right side stays positive. Hence $w(\sigma_C) - w(\sigma_E) > w(P'')$, which contradicts Lemma 6.1(a).

Under (J), each vertex of $T(\sigma)$ is an end of exactly one walk, so P'' has the detector boundary $T(\sigma)$ of σ_C. The same estimate without the terms of σ_E, and with (6.2) in place of (6.1), shows that $w(\sigma_C) > w(P'')$, which contradicts Lemma 6.1(b) for $X = \sigma_C$.

(b) Suppose that the window (u, b, w) of γ is excluded by (S), with the detector x, and let Q be the cycle with the edges ub, bw, wx and xu. The edges ub and bw belong to C^*. The set $C' = C^* \triangle Q$ satisfies $\partial C' = \partial C^*$, and since Q has no edge at b_1, the decoder predicts the same logical flip with C' as with C^*. Since profile weights agree on edges with the same profile, $w(ux) + w(xw) = w(ub) + w(bw)$. If neither ux nor xw belongs to C^*, then $w(C') = w(C^*)$, so C' is a minimizer for which the decoder fails, and $\Phi(C') - \Phi(C^*) = 2\big(\phi(x) - \phi(b)\big) < 0$, which contradicts the choice of C^*. If exactly one of them, say xw, belongs to C^*, then $w(C') = w(C^*) - w(ub) - w(bw) - w(xw) + w(ux) = w(C^*) - 2w(xw)$, and if both do, then $w(C') = w(C^*) - w(Q)$. In both cases $w(C') < w(C^*)$, which contradicts the minimality of C^*. □

For example, for a window made of C-edges only, $k \geq 0$ says that the alternative has at most as many edges as the window. For mixed windows, rule (M) also uses the E-edges, through the term $w(\sigma_E)$ of Lemma 6.1(a). Rule (J) uses only the correction edges, but it can re-pair the ends of two separate runs of them.

6.2. **Labelled paths.** A *labelled path* is a simple path $b_1 = v_0, v_1, \ldots, v_L = b_2$ in G with a label E or C on each edge, none of whose sub-paths with 2, 3 or 4 edges is, with its labels, an excluded window. Its weight is

$$\omega(\gamma) = \prod_{\text{E-edges}} q_e^{1-s}(1-q_e)^{-s} \prod_{\text{C-edges}} \Big(\frac{q_e}{1-q_e}\Big)^s,$$

and we write $\omega(e, l)$ for the factor of an edge e with label l.

Proposition 6.4. *For every $p \leq p_{**}$, P_{fail} is at most the total weight of all labelled paths, provided that for part (b) the weights are profile weights.*

Proof. Suppose the decoder fails. For part (a) let C be the minimizer that it returns, and for part (b) let $C = C^*$ be the minimizer of Lemma 6.3(b). In both cases the decoder fails with C. Let $\gamma \subseteq E \triangle C$ be the path of Lemma 3.1, and label the edges of $\gamma \cap E$ by E and those of $\gamma \cap C$ by C. By Lemma 6.3 no sub-path of γ with at most four edges is, with these labels, an excluded window, so γ with these labels is a labelled path. Its set S of E-edges satisfies $S \subseteq E$ and, by Lemma 3.3, $w(S) \geq w(\gamma \setminus S)$. By the union bound and Lemma 3.4, P_{fail} is at most the sum of $\prod_{e \in S} q_e$ over labelled paths. As in the proof of Lemma 3.5, we have $\prod_{e \in S} q_e \leq \prod_{e \in S} q_e e^{s w_e} \prod_{e \in \gamma \setminus S} e^{-s w_e}$. Because $q_e(1-q_e) \leq r_e \leq q_e$, we have $q_e e^{s w_e} \leq q_e^{1-s}(1-q_e)^{-s}$ and $e^{-s w_e} \leq (q_e/(1-q_e))^s$. □

6.3. **A criterion with memory.** A *state* is a path v_0, v_1, v_2, v_3 of distinct vertices with $v_1, v_2, v_3 \in V$ and $v_0 \in V \cup \{b_1\}$, together with labels $l_1 l_2 l_3$ of its three edges. Its *successors* are the states $(v_1, v_2, v_3, v_4; l_2 l_3 l_4)$ with $v_4 \in V \setminus \{v_0, v_1, v_2\}$, and its *terminal moves* are the edges $\{v_3, b_2\}$ with a label l_4. In both cases we require that none of the windows (v_2, v_3, v_4),

$(v_1,\dots,v_4)$ and $(v_0,\dots,v_4)$, with the labels l_3l_4, $l_2l_3l_4$ and $l_1\cdots l_4$, is excluded, where $v_4 = b_2$ for terminal moves. These three windows have the head v_3.

Lemma 6.5. *Let $\lambda > 0$, and let H assign a positive number to every state. Suppose that every state $\sigma = (v_0, v_1, v_2, v_3; l_1l_2l_3)$ satisfies*

$$\sum_{\text{successors}} \omega(v_3v_4, l_4)\, H(v_1, v_2, v_3, v_4; l_2l_3l_4) + \sum_{\text{terminal moves}} \omega(v_3b_2, l_4) \le \lambda\, H(\sigma). \tag{6.3}$$

Then for every $L \ge 4$ the labelled paths with L edges have total weight at most

$$\lambda^{L-3} \sum \omega(b_1v_1, l_1)\, \omega(v_1v_2, l_2)\, \omega(v_2v_3, l_3)\, H(b_1, v_1, v_2, v_3; l_1l_2l_3),$$

where the sum runs over the states with $v_0 = b_1$.

Proof. For a state σ and $m \ge 1$ let $W_m(\sigma)$ be the total weight of the sequences of m edges that continue σ by successors and end with a terminal move, the weight being the product of the factors of the new edges. Then $W_1(\sigma)$ is the terminal sum in (6.3), so $W_1(\sigma) \le \lambda H(\sigma)$, and for $m \ge 2$, by induction,

$$W_m(\sigma) = \sum_{\text{successors}} \omega(v_3v_4, l_4)\, W_{m-1}(v_1, v_2, v_3, v_4; l_2l_3l_4) \le \lambda^m H(\sigma).$$

A labelled path $b_1, v_1, \dots, v_L$ begins with the state $(b_1, v_1, v_2, v_3; l_1l_2l_3)$, and each quadruple of consecutive vertices is a successor of the previous one. Indeed $v_{i+1} \notin \{v_{i-2}, v_{i-1}, v_i\}$ because the path is simple, and the windows ending at v_{i+1} are not excluded because the path is labelled. Hence the labelled paths with L edges have total weight at most $\sum \omega(b_1v_1, l_1)\, \omega(v_1v_2, l_2)\, \omega(v_2v_3, l_3)\, W_{L-3}(b_1, v_1, v_2, v_3; l_1l_2l_3)$. □

6.4. **Classes and locality.** The certificate assigns the same potential to states with the same *class* and the same labels. Here the class of $(v_0, v_1, v_2, v_3; l_1l_2l_3)$ records the types of v_0 (or the label b_1), v_1, v_2 and v_3 and the offsets $v_1 - v_0$, $v_2 - v_1$ and $v_3 - v_2$. We call the inequality (6.3), written in terms of classes, labels and edge coefficients, the *row* of the state. The *growth rate* of a set of rows is the infimum of the λ for which some positive potential satisfies all of them. The *typed ball* of a detector v is the set of pairs (offset $u - v$, type of u) over the detectors u in the ball of v.

Lemma 6.6. *The rows of the states $(v_0, v_1, v_2, v_3; l_1l_2l_3)$ with a given last vertex v_3 are determined by the typed ball of v_3. The sum of $\omega(b_1v_1, l_1)\, \omega(v_1v_2, l_2)\, \omega(v_2v_3, l_3)\, H(b_1, v_1, v_2, v_3; l_1l_2l_3)$ over the states with a given neighbour v_1 of b_1 is determined by the typed ball of v_1. Every typed ball of every graph $G_{d,T}$ occurs in some $G_{d',T'}$ with $d' \le 25$ and $T' \le 13$.*

Proof. The vertices of a state ending at v_3 and of its successors lie within three edges of v_3, and their classes and the coefficients of the edges involved are read from types and offsets. Since the windows tested in these rows have the head v_3, the exclusion tests use only walks whose detectors lie in the ball of v_3, the edges among these detectors and their boundary edges, all of which the types record. Rule (S) uses a detector within two edges of the head and the profiles of four edges between detectors of the ball. The profiles of the edges at a detector v are determined by the same local data as its type in Lemma 2.3(b), because the outcomes with v in their signature and their locations are, and they are in fact determined by the type of $v^\dagger$. So the types record these profiles as well. The start sums involve $v_2 \sim v_1$, $v_3 \sim v_2$ and no exclusions.

A chain of at most four edges between detectors changes each coordinate by at most 16 and the layer by at most 4. By Lemma 2.3(b), the typed ball of v_3 is therefore determined by the positions of the qubits within ℓ^∞-distance 22 of v_3 and the X or Z types of the measure qubits among them, by $\min(r, 6)$ and by $\min(T - r, 6)$, where r is the layer of v_3. The last statement

follows as in Lemma 2.3(c), with 6 replaced by 22 and 2 by 6. Indeed, the even coordinates of G_{25} at distance at least 23 from both opposite sides are 24 and 26, one of each residue modulo 4, and the layer data of every layer occur for $T' = 13$. □

6.5. **Proof of Theorem 1.2.** Put $\lambda_0 = 997/1000$ for part (a) and $\lambda_0 = 999/1000$ for part (b). The instances with $d \le 25$ and $T \le 13$ have 2,546,882 distinct rows for part (a) and 2,545,552 for part (b)†, and by Lemma 6.6 these are all the rows of all graphs $G_{d,T}$. Each of the two certificates assigns a positive integer h to each of the 953,120 pairs of a class and a label sequence $l_1 l_2 l_3$ that occur, and we put $H = h/10^{15}$. All rows hold in exact integer arithmetic at $p = p_{**}$ with $\lambda = \lambda_0$, when each factor $\omega(e, l)$ is replaced by a rational upper bound computed with integer roots of order 50 for part (a) and 25 for part (b)† (Table 4). Let $p \le p_{**}$ and $\rho = (p/p_{**})^s$. Then each factor satisfies $\omega(e, l)(p) \le \rho\, \omega(e, l)(p_{**})$, because $(p/p_{**})^{1-s} \le \rho$ and $(1 - q_e)^{-1}$ decreases as p decreases. Windows excluded by (M) or (J) stay excluded by Lemma 6.3(a), and rule (S) does not depend on p. Hence all rows hold at p with $\lambda_p = \lambda_0 \rho$. Let K be an upper bound, computed with the same rational bounds of the factors, for the start sums of Lemma 6.6 at p_{**} over all neighbours v_1 of b_1†. Let K_3 be an upper bound of the same kind for the sums of $\omega(b_1 v_1, l_1)\, \omega(v_1 v_2, l_2)\, \omega(v_2 b_2, l_3)$ over the paths b_1, v_1, v_2, b_2 with a given v_1 and all labels†; such paths exist only for $d = 3$. For part (a) $K = 0.031062\ldots$ and $K_3 = 0.032106\ldots$, and for part (b) $K = 0.056229\ldots$ and $K_3 = 0.064013\ldots$. From Proposition 6.4, Lemma 6.5 and Lemma 3.2 we obtain

$$P_{\text{fail}} \le |N(b_1)| \Big(\rho^3 K_3\, [d = 3] + \sum_{L \ge \max(d,4)} \lambda_p^{L-3} \rho^3 K \Big) \le \rho^3\, |N(b_1)|\, (K + K_3)\, \frac{\lambda_p^{d-3}}{1 - \lambda_p}.$$

Recall that $|N(b_1)| = (T + 1)(d + 1)/2$ by Lemma 2.1(e). With $\rho = \lambda_p/\lambda_0$ the right side equals $(K + K_3)(T + 1)(d + 1)\lambda_p^d/(2\lambda_0^3(1 - \lambda_p))$, and $(K + K_3)/(2\lambda_0^3)$ is at most 0.032 for part (a) and 0.061 for part (b). □

7. The lean schedule

In the lean schedule all measure qubits and data qubits are prepared in a first step, the Z-type measure qubits in $|0\rangle$ and the X-type ones in $|+\rangle$. Each round then consists of the steps (S3)–(S6) of Section 1.1 and one more step. In it each Z-type measure qubit is measured in the Z basis and reset to $|0\rangle$, each X-type measure qubit is measured in the X basis and reset to $|+\rangle$, and the data qubits are idle. The noise model of Section 1.2 applies to these locations, the measurement and the reset in this step failing independently, each with probability p. The detectors, the logical outcomes and the decoder are as before.

Section 2 carries over with the following changes†. Each fault outcome again flips at most two detectors, and Lemma 2.1 holds with $c_{\max} = 94/15$. In the bulk of basis Z, the coefficients $15c_e$ of the edges with offsets $\pm(0, 0, 1)$, $\pm(2, -2, 0)$, $\pm(2, 2, 0)$, $\pm(2, 2, -1)$, $\pm(2, -2, -1)$ and $\pm(4, 0, -1)$ are 46, 42, 18, 16, 8 and 8, against 66, 72, 48, 16, 8 and 8 in Lemma 2.2. The difference comes from the three idle steps per round of the data qubits and the two of the Z-type measure qubits that the lean schedule does not have. There are again 29 node types per basis. We count the reset in the measure-and-reset step, and the first step, as the start of the next round. With this convention a measure qubit is again reset at the start of each round, and the proofs of Lemmas 2.3 and 6.6, which use only the layout, the CX orders and this structure of the rounds, apply verbatim. All edges again change each coordinate by at most 4 and the layer by at most 1.

Proof of Theorem 1.3. (a) As in Section 5, with $\lambda = 25.55\sqrt{p}$. At $p = 1.53 \times 10^{-3}$ we have $25.55\sqrt{p} \ge 0.9993$, and the criterion (5.1), with 0.9997 replaced by 0.9993, holds for all types,

with largest left side 0.99794 at the bulk type and 0.95318 for the others†. With $c_{\max} = 94/15$, the assembly of Section 5 then leads to the prefactor $\sqrt{c_{\max}/(1 - c_{\max}p)}/25.55 \le 0.0985$.

(b) As in Section 6, with $p_{**} = 3/1250$, $s = 21/50$ and $c_{\max} = 94/15$ in Definition 6.2. The instances with $d \le 25$ and $T \le 13$ have 2,548,765 distinct rows, all of which already occur for $d \le 15$ and $T \le 7$. A certificate with positive integers on the 953,120 pairs of a class and a label sequence satisfies all rows in exact arithmetic with $\lambda = 997/1000$, and the start constants are $K = 0.026464\ldots$ and $K_3 = 0.028384\ldots$, with $(K + K_3)/(2 \cdot 0.997^3) \le 0.028$†.

(c) As in part (b) of Theorem 1.2, with $p_{**} = 3/1000$, $s = 9/25$, $\lambda_0 = 997/1000$ and $c_{\max} = 94/15$. The instances with $d \le 25$ and $T \le 13$ have 2,547,432 distinct rows. Again, a certificate on the same 953,120 pairs satisfies all rows in exact arithmetic with $\lambda = 997/1000$, and the start constants are $K = 0.045261\ldots$ and $K_3 = 0.057416\ldots$, with $(K + K_3)/(2 \cdot 0.997^3) \le 0.052$†. □

8. Non-uniform error rates

In a real device the gates, measurements and idle steps fail at different rates. To model this, we sort the locations of the circuit of Section 1.1 into seven categories: resets and preparations; measurements; H gates; CX gates; qubits idle in steps (S2)–(S7); data qubits idle in step (S1); data qubits idle in step (S8). A *rate vector* θ assigns a failure probability θ_k to the locations of category k, with the failure outcomes of Section 1.2. The probability of an edge becomes

$$q_e(\theta) = \sum_k c_{e,k}\,\theta_k,$$

where $c_{e,k}$ is the sum of $\pi(f)$ over the outcomes f with $e(f) = e$ at locations of category k, so that $\sum_k c_{e,k} = c_e$. The graph, the signatures and the logical effects do not depend on θ, and Sections 2 and 3 hold with $q_e(\theta)$ in place of $c_e p$. The types, which now record the coefficients $c_{e,k}$, are again 29 per basis, and Lemma 2.3 holds with the same proof†.

Theorem 8.1. *Let θ be a rate vector with $\max_e q_e(\theta) < 1/2$, and consider the memory experiment with rates $t\theta$ for some $0 < t \le 1$.*

(a) *Suppose that $\sum_{e \ni v} a_e/(\lambda + a_e) \le 1$ at every detector v, where $a_e = 2\sqrt{q_e(\theta)/(1 - q_e(\theta))}$ and $\lambda < 1$. Then*

$$P_{\text{fail}} \le \frac{(T+1)(d+1)}{2} \cdot \frac{a_{\max}}{\lambda} \cdot \frac{(\lambda\sqrt{t})^d}{1 - \lambda\sqrt{t}},$$

where $a_{\max}$ is the largest a_e over the edges at b_1 and b_2.

(b) *Let $q(P) = \prod_{e \in P} q_e(\theta)$, with multiplicity, and $q_{\max} = \max_e q_e(\theta)$. In Definition 6.2, replace* (6.1) *by*

$$\frac{q(\sigma_E)\,q(P')}{q(\sigma_C)} > \prod_{e \in \sigma} \frac{1}{1 - q_e(\theta)} \cdot \Big(\frac{1}{1 - q_{\max}}\Big)^{h'} \tag{8.1}$$

and (6.2) *by*

$$\frac{q(P')}{q(\sigma_C)} > \prod_{e \in \sigma_C} \frac{1}{1 - q_e(\theta)} \cdot \Big(\frac{1}{1 - q_{\max}}\Big)^{h'}, \tag{8.2}$$

keeping the conditions $k \ge 0$, and define labelled paths, states and rows as in Section 6 with the windows excluded by these two rules, with the factors $\omega(e, l)$ evaluated at $q_e(\theta)$ and with a tilt $0 < s \le 1/2$. Suppose that a positive potential satisfies all rows with some $\lambda < 1$, and let K and K_3 be the largest start sums, as in Section 6. Then

$$P_{\text{fail}} \le \frac{(T+1)(d+1)}{2} \cdot \frac{K + K_3}{\lambda^3} \cdot \frac{(\lambda t^s)^d}{1 - \lambda t^s}.$$

(c) *For a location ℓ of category k let $c_{\ell,e}(\theta)$ be the sum of $\theta_k\pi(f)$ over the outcomes f of ℓ with $e(f)=e$, and call the multiset of the nonzero numbers $c_{\ell,e}(\theta)$ the profile of e at θ. Suppose that r_e is a function of t and of this profile, and add rule (S) of Definition 6.2, with these profiles, to the two rules of part (b). If a positive potential satisfies all rows with some $\lambda<1$, then the bound of part (b) holds with the start sums of these rows.*

Proof. (a) For the rates $t\theta$ each a_e is at most $\sqrt{t}\,a_e$, because $1-tq_e(\theta)\ge 1-q_e(\theta)$. Since $x/(\lambda\sqrt{t}+x)$ increases in x, the criterion holds with $\lambda\sqrt{t}$ in place of λ. The bound then follows from the assembly of Section 5, with $\sqrt{t}\,a_{\max}$ in place of $a_{\max}$.

(b) Lemma 6.3(a) holds for the rates $t\theta$, with (8.1) and (8.2) in place of (6.1) and (6.2). In its proof the first two terms of the lower bound become $k\ln(1/t)+\ln(q(\sigma_E)q(P')/q(\sigma_C))$ under (M) and $k\ln(1/t)+\ln(q(P')/q(\sigma_C))$ under (J), and $k\ge 0$. Similarly, Lemma 6.6 holds with the same proof. Each factor satisfies $\omega(e,l)(t\theta)\le t^s\,\omega(e,l)(\theta)$, because $t^{1-s}\le t^s$ for $s\le 1/2$. Hence all rows hold for the rates $t\theta$ with λt^s in place of λ, and every start sum is at most t^{3s} times its value at θ. With $\rho=t^s$, the assembly of Section 6 yields the bound.

(c) The profiles at $t\theta$ are t times those at θ, so rule (S) does not depend on t, and Lemma 6.3(b) holds with the same proof: profile weights agree on edges with the same profile. The proof of part (b) then applies with the rows of part (c), and Proposition 6.4 uses the minimizer C^* of Lemma 6.3(b). □

Table 2 lists six error models, built from the published component error rates of six devices and mapped onto the seven categories; Appendix B gives their sources and the conversions. The models idealize the devices, so Theorem 8.1 applies to the models, not to the devices. Part (b), with $s=43/100$, certifies three models:

Quantinuum Helios:	$P_{\text{fail}}\le 0.033\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.75\,t^{43/100}$,
Quantinuum H2:	$P_{\text{fail}}\le 0.022\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.715\,t^{43/100}$,
IBM Heron r3, square grid:	$P_{\text{fail}}\le 0.023\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.97\,t^{43/100}$,

for their rates scaled by any $t\le 1$. For profile weights, part (c) with $s=19/50$ improves these to

Quantinuum Helios:	$P_{\text{fail}}\le 0.058\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.675\,t^{19/50}$,
Quantinuum H2:	$P_{\text{fail}}\le 0.039\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.643\,t^{19/50}$,
IBM Heron r3, square grid:	$P_{\text{fail}}\le 0.039\,(T+1)(d+1)\,\lambda^d/(1-\lambda)$,	$\lambda=0.873\,t^{19/50}$.

Part (a) also certifies Helios and H2, but it fails at $\lambda=1$ for the other four models†. For the models of Google's and USTC's processors, the floating-point growth rates of the rows in Table 2 exceed 1.

9. Computer verification

Table 3 lists the computations behind the statements marked †. We compute the fault table by forward propagation of Pauli frames, independently of Stim's detector-error-model code, and keep each Pauli outcome of each location with its exact probability. All comparisons in V6 use exact rational arithmetic.

Table 4 lists the checks behind Theorem 1.2. Two further scripts, one for each part, regenerate the rows of all instances with $d\le 25$ and $T\le 13$ and check them against the two certificates. By Lemma 6.6 the rows at a detector depend only on its typed ball, so each script computes the rows at one detector for each distinct typed ball of the range, which we call the heads; balls are compared as full tuples. The rows are generated with exact integer arithmetic. The supplement

model	CX	H	meas.	reset	idle	idle (S1)	idle (S8)	(a)	(b)	(c)
Helios [13]	1e-3	4.5e-5	5e-4	5e-4	9e-4	9e-4	9e-4	$0.9325^\dagger$	$0.75^\dagger$	$0.675^\dagger$
H2 [14]	1.25e-3	4.5e-5	1e-3	1e-3	3e-4	3e-4	3e-4	$0.8957^\dagger$	$0.715^\dagger$	$0.643^\dagger$
Heron r3 [17]	1.9375e-3	2.76e-4	3.48e-3	2e-3	1e-4	1e-3	2e-3	(1.212)	$0.97^\dagger$	$0.873^\dagger$
Google 72q [11]	3.55e-3	6.2e-4	8e-3	1.5e-3	3.7e-4	3e-3	6.25e-3	(1.741)	(1.393)	(1.249)
Willow [15]	4.125e-3	5.25e-4	7.7e-3	1.5e-3	3.7e-4	3e-3	6.25e-3	(1.813)	(1.445)	(1.295)
Zuchongzhi 3.2 [12]	5.5e-3	8.5e-4	9.5e-3	1.5e-3	3e-4	1.5e-3	3.5e-3	(1.915)	(1.516)	(1.360)

TABLE 2. Error models (Pauli error probabilities per location) and the rates λ of parts (a), (b) and (c) of Theorem 8.1 at $t = 1$. Daggered values are certified exactly, those of (b) with $s = 43/100$ and those of (c) with $s = 19/50$. Values in parentheses are not certificates: in column (a) they are floating-point roots of the criterion, and in columns (b) and (c) floating-point growth rates of the rows, computed from the instances with $d \le 15$ and $T \le 7$, the smallest over $s \in \{0.40, 0.43, 0.45\}$ for (b) and over a grid of s between 0.30 and 0.44 for (c).

	check	range
V1	noiseless circuit: all detectors and the logical outcome are deterministic, with value 0	$d \le 13$, $T \le 7$
V2	fault table: each outcome, injected alone into Stim's Pauli-frame simulator, gives the signature and logical effect computed here; the set of signatures equals that of Stim's detector error model; the outcome probabilities of each location sum to 1; the coefficients c_e agree with Stim's edge probabilities at $p = 10^{-7}$ to 2×10^{-6} relative	$d \le 13$, $T \le 7$
V3	Lemma 2.1(a,b)	$d \le 13$, $T \le 7$
V4	Lemma 2.1(c,d,e)	$d \le 13$, $T \le 7$
V5	exactly 29 types per basis; all occur already for $d \le 9$, $T \le 5$; the set of types is the same for all $d \ge 9$, $T \ge 5$ in range (Lemma 2.3)	$d \le 13$, $T \le 7$
V6	criterion (5.1) for all types at $p = p_*$, $\lambda = 0.9997$, and at $p = 10^{-3}$, $\lambda = 0.98$; the unique maximum is the bulk type; the second largest value at p_* is at most 0.9629	all $29 + 29$ types

TABLE 3. Computer verification for Theorem 1.1. Lemma 2.3(c) needs only $d \le 9$ and $T \le 5$.

also contains the iteration that found the potentials and tests of the row generator and of the exclusions, none of which the proof uses.

The checks behind Theorem 1.3 (Table 5) are those of Tables 3 and 4, run on the lean schedule with its own certificates.

For Section 8, a further script runs the checks H1–H4 of Table 6 on the error models of Table 2. For H4 it runs the script of rows S1–S5 with the edge probabilities $q_e(\theta)$, written as integers over a common denominator, in place of the coefficients, and check H5 runs the script of row S6 in the same way.

10. Formal verification in Lean 4

We prove Theorems 1.1–1.3, and Theorem 8.1 for the models it certifies, in the proof assistant Lean 4, using the Mathlib library; Lean's *kernel* checks every step of a proof against a small set of logical rules. The formal development defines the circuit, the noise, the memory experiment

	check	range
S1	parameters of the certificate: $p_{**} = 1/625$, $\lambda = 997/1000$, $s = 21/50$	
S2	rows of all instances: 2,546,882 distinct rows, all of which already occur for $d \le 15$, $T \le 7$; 13,658 distinct typed balls; 29 node types per basis, with injective identifiers; largest coefficient $c_{\max} = 154/15$	$d \le 25$, $T \le 13$
S3	every class in a row, as head or successor, has a positive potential in the certificate	all rows
S4	every row holds in exact integer arithmetic at p_{**} and $\lambda = 997/1000$, with rational upper bounds of the factors ω; largest ratio of the two sides 0.99900	all rows
S5	start constants $K = 0.031062\ldots$ and $K_3 = 0.032106\ldots$, with $(K + K_3)/(2 \cdot 0.997^3) \le 0.032$	all neighbours of b_1
S6	part (b), as S1–S5: $p_{**} = 1/500$, $\lambda = 999/1000$, $s = 9/25$, rules (M), (J) and (S); the profile of each edge at a detector is determined by the type of the detector; 2,545,552 distinct rows; positive potentials; all rows exact, largest ratio 0.99900; $K = 0.056229\ldots$, $K_3 = 0.064013\ldots$, $(K + K_3)/(2 \cdot 0.999^3) \le 0.061$	$d \le 25$, $T \le 13$

TABLE 4. Computer verification for Theorem 1.2. Rows S1–S5 concern part (a). Lemma 6.6 needs $d \le 25$ and $T \le 13$.

	check	range
L1	as V1–V5: deterministic noiseless values 0; fault table equal to Stim's outcome by outcome and in first-order coefficients; Lemma 2.1 with $c_{\max} = 94/15$; 29 types per basis, stable for $d \ge 9$, $T \ge 5$	$d \le 13$, $T \le 7$
L2	as V6: criterion at $p = 1.53 \times 10^{-3}$, $\lambda = 0.9993 \le 25.55\sqrt{p}$; largest left side 0.99794 (bulk), second 0.95318	all $29 + 29$ types
L3	as S1–S5: $p_{**} = 3/1250$, $s = 21/50$; 2,548,765 rows, all present for $d \le 15$, $T \le 7$; positive potentials; all rows exact at $\lambda = 997/1000$, largest ratio 0.99799; $K = 0.026464\ldots$, $K_3 = 0.028384\ldots$, $(K + K_3)/(2 \cdot 0.997^3) \le 0.028$	$d \le 25$, $T \le 13$
L4	as S6: $p_{**} = 3/1000$, $s = 9/25$, rules (M), (J) and (S); 2,547,432 rows; positive potentials; all rows exact at $\lambda = 997/1000$, largest ratio 0.99900; $K = 0.045261\ldots$, $K_3 = 0.057416\ldots$, $(K + K_3)/(2 \cdot 0.997^3) \le 0.052$	$d \le 25$, $T \le 13$

TABLE 5. Computer verification for Theorem 1.3.

and the decoder of Section 1, and derives the bounds of the theorems from these definitions. Besides the definitions and proofs written by hand, it contains generated data files (node types, certificates) and generated check files. Each formal theorem is proved in two versions: in the *kernel* version every computation is checked by the kernel, and in the *native* version the large computations are evaluated outside the kernel, which takes far less time. In this section we state what is modelled, what is proved and what has to be trusted. The formal development is part of the supplement. Its file `FORMALIZATION.md` lists the formal theorems with their Lean names

	check	range
H1	every location has one of the seven categories; with equal rates the coefficients $c_{e,k}$ add up to c_e	$d \le 9$, $T \le 5$
H2	category-resolved types: 29 per basis, all present for $d = 9$, $T = 5$, and the same set for all $d \ge 9$, $T \ge 5$ in range	$d \le 13$, $T \le 7$
H3	part (a) holds exactly for Helios at $\lambda = 0.9325$ and for H2 at $\lambda = 0.8957$; for the other four models its left side exceeds 1 at $\lambda = 1$ for some type	all $29 + 29$ types
H4	part (b) for Helios, H2 and Heron r3, with $s = 43/100$, as S1–S5: $q_{\max}$ recomputed, and every edge probability of the range at most $q_{\max}$; 2,546,246, 2,546,404 and 2,546,625 distinct rows; positive potentials; all rows exact at $\lambda = 3/4$, $143/200$ and $97/100$; $(K + K_3)/(2\lambda^3) \le 0.033$, 0.022 and 0.023	$d \le 25$, $T \le 13$
H5	part (c) for Helios, H2 and Heron r3, with $s = 19/50$, as S6, with the profiles at the model rates: 2,544,853, 2,545,240 and 2,545,179 distinct rows; positive potentials; all rows exact at $\lambda = 27/40$, $643/1000$ and $873/1000$; $(K + K_3)/(2\lambda^3) \le 0.058$, 0.039 and 0.039	$d \le 25$, $T \le 13$

TABLE 6. Computer verification for Theorem 8.1 and the error models of Table 2.

and files, the versions of Lean and Mathlib, the commands and the measured time and memory of the checks, and its file `AXIOMS.md` lists the axioms of every formal theorem.

10.1. **What is modelled.** *The circuit and the experiment.* The layout of the rotated surface code, the eight steps of a round, the noise locations with their outcomes and the probabilities $p\,\pi(f)$ are those of Sections 1.1 and 1.2. We define the effect of the faults by Pauli-frame propagation: a CX gate copies an X component from its control to its target and a Z component from its target to its control, an H gate exchanges the two components, and a reset clears both. For any set of faults, at most one outcome per location, this determines the recorded measurement outcomes, the detectors $D_{m,r}$ of Section 1.1 and the flip of the logical outcome. The formal development proves that these are the parities of the single-fault values, so they agree with the signatures $\sigma(f)$ and the logical effects $\ell(f)$ of Section 1.3.

The decoder and the failure probability. Here the decoding graph and the weights are those of Section 1.3, for any numbers r_e with $q_e(1 - q_e) \le r_e \le q_e$; this range contains both common choices. The failure probability P_{fail} is the probability, under the product distribution of the faults, that some set C of minimum weight with ∂C equal to the observed detectors predicts the logical flip wrongly. Hence a bound on this probability holds for every rule that selects a minimizer.

Profile weights. The profile of an edge and profile weights are those of Section 1.3, and the theorems that use rule (S) are stated for profile weights.

The lean schedule. We model the lean schedule of Section 7 by its own Pauli-frame semantics: a preparation step, then per round four CX steps and a fused measure-and-reset step, in which the X-type measure qubits are measured in the X basis. Its noise locations are those of Section 7, including the reset failures after the last measurement. The development proves that every lean location changes the measurement outcomes and the final data values exactly as the corresponding location of the eight-step circuit restricted to the lean locations, and that the reset failures after the last measurement change nothing. Hence the lean experiment has the failure probability of the restricted eight-step circuit, and the formal theorems are about the lean schedule itself.

Non-uniform rates. A rate vector θ assigns to each of the seven categories of Section 8 the failure probability θ_k, and the outcomes of a failed location are uniform. The formal theorems are stated for the three certified models, whose rates are those of the accompanying files, scaled by any $0 < t \leq 1$.

10.2. **What is proved.** The formal theorems are Theorems 1.1, 1.2 and 1.3, and the instances of Theorem 8.1 for the models it certifies: part (a) for Helios and H2, and parts (b) and (c) for Heron r3, Helios and H2. Each holds for all odd $d \geq 3$, all $T \geq 1$ and both bases, and states the inequalities of the corresponding theorem of this paper, with its constants, for the failure probability defined above. Each is proved in a kernel version and a native version.

10.3. **What has to be trusted.** A Lean proof can rely on axioms, which the kernel accepts without proof. Lean reports, for every theorem, the axioms it depends on. Every formal theorem depends on three standard axioms of Lean's foundations (propositional extensionality, the axiom of choice and the soundness of quotients). The kernel versions depend on nothing else.

The native versions also depend on axioms of a second kind, one for each large computation. For such a computation Lean evaluates the Boolean test outside the kernel, with compiled code or with its interpreter, and records the outcome "true" as an axiom naming the test. Trusting these axioms means trusting, in addition to the kernel, Lean's compiler, interpreter and runtime, the C compiler that builds the precompiled modules, and the assumption that the loaded libraries were built from the checked sources. The kernel itself already relies on the GMP library for arithmetic on large natural numbers. The large computations are formal versions of checks of Tables 3–6: the signatures and node types near representative detectors, which by Lemma 2.3(c) lie in instances with $d \leq 9$ and $T \leq 5$; the criterion for the node types; the type tables, with the edge profiles for rule (S); the coverage of all typed balls by the heads; and the rows, start sums, factor bounds and numerator bounds of each certificate.

In the kernel versions the kernel checks the same facts, in reformulated forms whose soundness is proved, split into generated files that are checked one per process. A further file compares the statement of each kernel theorem of Theorems 1.2 to 8.1 with the statement of its native version, and the two are identical; the two versions of Theorem 1.1 are stated by the same source text.

The certificates are read from the files of the supplement; the formal data files record the SHA-256 hash of each certificate file. Programs outside Lean generate the node types, the tables, the heads and the other auxiliary data from which the checks start. Because every item they produce is checked inside Lean, an error in them can only make a check fail. The search for exclusion witnesses is likewise untrusted, in both versions: the rows accept a window as excluded only after the witness found by the search passes an exact check whose soundness is proved.

Two cross-checks support the model without being part of the proof. The node types computed in Lean coincide, as strings including all coefficients, with the node types of the accompanying programs for all five models (the eight-step circuit, the lean schedule built by those programs from its own circuit description through Stim, and the three non-uniform models). All five models have the same 13,658 typed balls on the range of Lemma 6.6.

11. Remarks

Remark 11.1 (Comparison with [3]). Fowler's argument uses unit weights, bounds the number of non-backtracking paths of length m by $3n^2\,11^{m-1}$ and bounds each line probability by the largest one, for which it takes the value $14p/5$ of the circuits of [9]. In the circuit considered here the largest bulk coefficient is $72/15$, so that argument gives at most $1/(484 \cdot \frac{72}{15})$, the bound quoted in Section 1, before the heavier boundary edges are taken into account. Two changes account for the improvement to 1.04×10^{-3}. First, log-likelihood weights let each edge contribute

$2\sqrt{q_e}$ instead of $2\sqrt{q_{\max}}$. Second, the criterion of Lemma 4.1 replaces the factor 11 by an exact count of the twelve unequal edges at each node. At the bulk node with $\lambda = 1$ the left side of the criterion is below 1 at $p = 1.0419 \times 10^{-3}$ but above 1 at $p = 1.0420 \times 10^{-3}$, so the method cannot reach beyond 1.042×10^{-3}. As $p \to 0$ the bulk root of $\Sigma = 1$ is $\lambda = \Lambda\sqrt{p}$ with $\Lambda = 30.929$, the solution of $\sum_e 2\sqrt{c_e}/(\Lambda + 2\sqrt{c_e}) = 1$ over the bulk coefficients. With unit weights the same method gives 8.7×10^{-4} in the bulk.

Remark 11.2 (Limits of the certificate). Consider the rows of Section 6 computed from the instances with $d \le 15$ and $T \le 7$, with the exclusions at each p. Their growth rate is 0.99407 at $p = 1.6 \times 10^{-3}$ for the rules of part (a) with $s = 0.42$, and 0.99744 to 0.99751 at $p = 2 \times 10^{-3}$ for the rules of part (b) with $s = 0.36$. But for every tilt s it is above 1 at $p = 1.62 \times 10^{-3}$ for the rules of part (a) and at $p = 2.05 \times 10^{-3}$ for those of part (b); for the lean schedule the same holds at 2.47×10^{-3} and at 3.05×10^{-3}. These statements rest on lower and upper Collatz–Wielandt bounds computed in floating point on a grid of tilts, together with the log-convexity of the growth rate in s [18]. Since the full range of d and T only adds rows, certificates of these kinds, with any tilt, cannot reach these values. Longer windows raise these values only slowly: on the bulk lattice, without boundaries, the rules of part (a) reach growth rate 1 at 1.6185×10^{-3} with windows of four edges and at 1.6936×10^{-3} with windows of five edges. Rule (S) extends to runs of three or four correction edges: such a run can be excluded if an edge-disjoint route between its ends uses only edges between detectors, has the same multiset of profiles and has a smaller value of Φ. This excludes more windows but leaves the growth rates at the values of p above unchanged to five digits.

Remark 11.3 (Limits of arguments by increasing events). Call a set of outcomes *failing* if, when exactly these outcomes occur, some minimizer makes the decoder fail. An event is *increasing* if it is preserved when outcomes are added. All bounds in this paper bound P_{fail} by the probability of an increasing event that contains the failing sets, and so does any union bound over paths, clusters or other witnesses that only require certain outcomes to occur. Rule (S) does not change this, because it too depends only on the labels and the graph. Every such event contains the event U that some subset of the occurring outcomes is failing, and U is itself increasing. Lemma 6.1(b) describes U: every subset X of a minimizer has minimum weight among the sets with boundary ∂X. Hence if the decoder fails and γ is the path of Lemma 3.1, then $\gamma \cap C$ is a minimizer for the error $\gamma \cap E$, which is produced by a subset of the occurring outcomes, and with this minimizer the decoder fails on that error alone. So U occurs if and only if some simple path γ from b_1 to b_2 contains a set E' of edges, each produced by an occurring outcome, such that $\gamma \setminus E'$ has minimum weight among the sets with boundary $\partial E'$.

Monte Carlo estimates of $\Pr(U)$ for the memory experiment with $T = d$, with U decided in each sample by integer programming, are not rigorous; they are lower bounds, because the samples that the program did not resolve, below 1% for $d \le 9$ at $p \le 4.25 \times 10^{-3}$ and about 3% at $d = 11$ and $p = 4 \times 10^{-3}$, count as not in U. For the circuit of Section 1.1 and $d = 5, 7, 9, 11$, they decrease with d at $p = 3.5 \times 10^{-3}$, agree within their statistical errors at $p = 4 \times 10^{-3}$, at about 0.044, and increase with d at $p = 5 \times 10^{-3}$. For the lean schedule they are roughly constant in d at 6×10^{-3} and increase from $d = 5$ to $d = 7$ at 7×10^{-3}. The threshold of U for the 8-step circuit therefore lies between 3.5×10^{-3} and 5×10^{-3}, probably near 4×10^{-3}, and no argument of this kind can prove a threshold above it. Where $\Pr(U)$ no longer decays but matching still succeeds, failing sets are present but are undone by the other outcomes, and a proof there would need an argument that is not monotone in the faults.

Remark 11.4 (Other schedules). Each of the following modifications produces a circuit whose fault outcomes have the signatures and probabilities of a subset of our outcomes: deleting noise locations; merging (S8) with (S1) of the next round into a single measure-and-reset step; and

preparing and measuring the X-type measure qubits in the X basis directly, without (S2) and (S7). For such a circuit each coefficient is at most the corresponding c_e, so its edge weights a_e are smaller, and Theorem 1.1 holds for it with the same constants. In particular, the noisy circuits `surface_code:rotated_memory_x/z` generated by Stim with all noise parameters equal to p are of this kind. Theorem 1.2, however, does not transfer in this way, because its exclusions compare coefficients of alternative paths, and these change when noise locations are removed. For the lean schedule, Theorem 1.3 provides sharper bounds.

Remark 11.5 (Weights). Theorems 1.1–1.3 assume $q_e(1-q_e) \le r_e \le q_e$, and Theorems 1.2(b) and 1.3(c) also assume profile weights. Stim's detector error model replaces each depolarizing location by independent mechanisms. If these are projected to a sector and the mechanisms that produce the same edge are combined as independent events, then $1-2r_e = \prod_\ell (1-x_\ell)^{k_\ell/m_\ell}$, where k_ℓ outcomes of ℓ produce e, and (x_ℓ, m_ℓ) is $(16p/15, 8)$, $(4p/3, 2)$ or $(2p, 1)$ for two-qubit, single-qubit and flip locations. In our circuits $c_{\ell,e}$ determines k_ℓ and the kind of ℓ, so these are profile weights, and for $d \le 13$, $T \le 7$ and $p \le 3\times 10^{-3}$ they lie in this range for both schedules. The heuristic decomposition of Stim's detector error model sometimes splits an outcome whose signature in the sector is $\{u, v\}$ into the two boundary edges at u and v. This changes some edge probabilities by factors up to 2, and a matching decoder built on that decomposition is not covered. But the canonical decomposition used here sends the X part and the Z part of each outcome to its unique edge in each sector.

Remark 11.6 (Numerical threshold). In simulations of this circuit with d rounds, decoded by matching on the decoding graph of Section 1.3 with $r_e = q_e$, the logical failure curves for $d = 3, 5, 7, 9$ cross near $p \approx 5\times 10^{-3}$, about 2.5 times $p_{**} = 2\times 10^{-3}$ and 5 times p_*. For the lean schedule they cross near $p \approx 8\times 10^{-3}$, about 2.7 times the value 3×10^{-3} of Theorem 1.3(c).

Remark 11.7 (Google's Willow schedule). Google's surface-code experiments on the Willow processor [11] use a different circuit: the ZXXZ variant of the code, CZ gates conjugated by Hadamard gates, an echo pulse on all qubits in the middle of each round, and a fused measure-and-reset step. We modelled it with the noise of Section 1.2 at the Willow rates of Table 2, all multiplied by a factor t (CZ gates get the CX rate, Hadamard gates and echo pulses the H rate, and the data qubits, idle during measurement and reset, the rates of steps (S8) and (S1)). With these rules, the circuits published with [16] that we compared have the same detector error model after three changes that do not alter it: an explicit initial reset, no random Pauli layer at the start, and two consecutive resets merged into one. In this model $\Pr(U)$ stops decreasing with d between $t = 0.9$ and $t = 1$, and the detection events measured on the device, decoded with the decoder of Section 1.3, match the model at $t \approx 1.1$; Appendix B.2 gives these statistical estimates. For this decoder and in this model, the device thus operates at or above the threshold of U, where no argument by increasing events applies, while matching still suppresses errors.

12. Conclusion

Two questions remain open. Between 2×10^{-3} and the limit of Remark 11.3 there is room for arguments by increasing events with new exclusion rules, since Remark 11.2 shows that the present rules and longer windows gain little. Beyond that limit, where matching still succeeds because failing sets of faults are undone by the others, a proof would have to control this cancellation, which no argument by increasing events can do.

Code and data

The programs, the certificates, the formal proofs and the outputs of all checks in Tables 3–6 are archived at `https://doi.org/10.5281/zenodo.23498518`.

Appendix A. Node types

Each line gives one type: an example detector (x, y, r) in $G_{9,5}$, the number of detectors of this type in $G_{9,5}$, the degree, the coefficients $15c_e$ of its edges (a superscript $\flat$ marks the edge to b_1 and $\sharp$ the edge to b_2), and the left side of (5.1), that is Σ at $p = p_*$ and $\lambda = 0.9997$, rounded up to four decimals. The first row is the bulk type.

(x, y, r)	#	deg	$\{15c_e\}$	Σ
$(4, 4, 1)$	72	12	$\{72, 72, 66, 66, 48, 48, 16, 16, 8, 8, 8, 8\}$	0.9995
$(2, 6, 1)$	12	11	$\{72, 70, 66, 66, 48, 48, 22, 16, 8, 8, 8\}$	0.9629
$(16, 4, 1)$	12	11	$\{72, 70, 66, 66, 48, 48, 22, 16, 8, 8, 8\}$	0.9629
$(4, 16, 1)$	12	9	$\{148^\sharp, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
$(6, 2, 1)$	12	9	$\{148^\flat, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
$(2, 2, 1)$	4	8	$\{154^\flat, 70, 66, 66, 48, 16, 8, 8\}$	0.7897
$(16, 16, 1)$	4	8	$\{154^\sharp, 70, 66, 66, 48, 16, 8, 8\}$	0.7897
$(4, 4, 5)$	21	8	$\{66, 51, 51, 39, 39, 16, 8, 8\}$	0.6737
$(0, 4, 1)$	12	7	$\{78, 78, 70, 48, 22, 8, 8\}$	0.6459
$(18, 6, 1)$	12	7	$\{78, 78, 70, 48, 22, 8, 8\}$	0.6459
$(16, 4, 5)$	3	7	$\{66, 51, 49, 39, 39, 22, 8\}$	0.6368
$(4, 4, 0)$	21	8	$\{66, 41, 41, 29, 29, 16, 8, 8\}$	0.6298
$(0, 16, 1)$	4	6	$\{78, 78, 72^\sharp, 70, 8, 8\}$	0.5939
$(18, 2, 1)$	4	6	$\{78, 78, 72^\flat, 70, 8, 8\}$	0.5939
$(2, 6, 0)$	3	7	$\{66, 41, 39, 29, 29, 22, 8\}$	0.5927
$(2, 2, 5)$	4	6	$\{100^\flat, 66, 51, 39, 16, 8\}$	0.5711
$(4, 16, 5)$	3	6	$\{108^\sharp, 66, 51, 39, 8, 8\}$	0.5584
$(4, 16, 0)$	4	6	$\{80^\sharp, 66, 41, 29, 16, 8\}$	0.5359
$(6, 2, 0)$	3	6	$\{88^\flat, 66, 41, 29, 8, 8\}$	0.5239
$(16, 16, 5)$	1	5	$\{114^\sharp, 66, 49, 39, 8\}$	0.5149
$(0, 4, 5)$	3	6	$\{78, 51, 39, 22, 8, 8\}$	0.4920
$(2, 2, 0)$	1	5	$\{94^\flat, 66, 39, 29, 8\}$	0.4807
$(18, 6, 0)$	3	6	$\{78, 41, 29, 22, 8, 8\}$	0.4701
$(0, 16, 5)$	1	5	$\{78, 63^\sharp, 51, 8, 8\}$	0.4421
$(18, 2, 0)$	1	5	$\{78, 53^\flat, 41, 8, 8\}$	0.4235
$(18, 2, 5)$	1	3	$\{78, 49, 39^\flat\}$	0.3275
$(18, 6, 5)$	3	3	$\{78, 49, 39\}$	0.3275
$(0, 4, 0)$	3	3	$\{78, 39, 29\}$	0.3053
$(0, 16, 0)$	1	3	$\{78, 39, 29^\sharp\}$	0.3053

Table 7. Node types in basis Z.

(x, y, r)	#	deg	$\{15c_e\}$	Σ
$(6, 4, 1)$	72	12	$\{72, 72, 66, 66, 48, 48, 16, 16, 8, 8, 8, 8\}$	0.9995
$(4, 2, 1)$	12	11	$\{72, 70, 66, 66, 48, 48, 22, 16, 8, 8, 8\}$	0.9629
$(6, 16, 1)$	12	11	$\{72, 70, 66, 66, 48, 48, 22, 16, 8, 8, 8\}$	0.9629
$(2, 4, 1)$	12	9	$\{148^\flat, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
$(16, 6, 1)$	12	9	$\{148^\sharp, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
$(2, 16, 1)$	4	8	$\{146^\flat, 72, 66, 66, 48, 22, 8, 8\}$	0.7975
$(16, 2, 1)$	4	8	$\{146^\sharp, 72, 66, 66, 48, 22, 8, 8\}$	0.7975
$(4, 2, 5)$	21	8	$\{66, 51, 51, 39, 39, 16, 8, 8\}$	0.6737
$(4, 18, 1)$	12	7	$\{78, 78, 70, 48, 22, 8, 8\}$	0.6459
$(6, 0, 1)$	12	7	$\{78, 78, 70, 48, 22, 8, 8\}$	0.6459
$(6, 16, 5)$	3	7	$\{66, 51, 49, 39, 39, 22, 8\}$	0.6368
$(6, 4, 0)$	21	8	$\{66, 41, 41, 29, 29, 16, 8, 8\}$	0.6298
$(2, 0, 1)$	4	6	$\{80^\flat, 78, 78, 48, 22, 8\}$	0.6083
$(16, 18, 1)$	4	6	$\{80^\sharp, 78, 78, 48, 22, 8\}$	0.6083
$(4, 2, 0)$	3	7	$\{66, 41, 39, 29, 29, 22, 8\}$	0.5927
$(2, 4, 5)$	3	6	$\{100^\flat, 66, 51, 39, 16, 8\}$	0.5711
$(16, 2, 5)$	4	6	$\{108^\sharp, 66, 51, 39, 8, 8\}$	0.5584
$(16, 6, 0)$	3	6	$\{80^\sharp, 66, 41, 29, 16, 8\}$	0.5359
$(2, 16, 5)$	1	5	$\{98^\flat, 66, 51, 39, 22\}$	0.5348
$(2, 4, 0)$	4	6	$\{88^\flat, 66, 41, 29, 8, 8\}$	0.5239
$(16, 2, 0)$	1	5	$\{78^\sharp, 66, 41, 29, 22\}$	0.4995
$(6, 0, 5)$	3	6	$\{78, 51, 39, 22, 8, 8\}$	0.4920
$(4, 18, 0)$	3	6	$\{78, 41, 29, 22, 8, 8\}$	0.4701
$(2, 0, 5)$	1	5	$\{78, 61^\flat, 39, 22, 8\}$	0.4558
$(16, 18, 0)$	1	5	$\{78, 51^\sharp, 29, 22, 8\}$	0.4350
$(4, 18, 5)$	3	3	$\{78, 49, 39\}$	0.3275
$(16, 18, 5)$	1	3	$\{78, 49^\sharp, 39\}$	0.3275
$(2, 0, 0)$	1	3	$\{78, 39^\flat, 29\}$	0.3053
$(6, 0, 0)$	3	3	$\{78, 39, 29\}$	0.3053

TABLE 8. Node types in basis X.

Appendix B. Device error models

B.1. **Sources and conversions.** The error models of Table 2 use the typical values of the Quantinuum Helios and H2 data sheets [13, 14]; the medians of IBM's Heron r3 processor ibm_boston on 2026-10-18 [17], placed on a square grid, which the device does not have; the error budget of Google's 72-qubit processor [11, Table S4], with the CZ crosstalk and leakage rows added to the CX rate, and the leakage-heating row added to the data-idle error during measurement and reset, which is split over steps (S1) and (S8); the mean simultaneous randomized-benchmarking errors of the Willow spec sheet [15]; and the cross-entropy-benchmarking errors of the 97-qubit region of Zuchongzhi 3.2 [12], used as reported. Errors reported as average infidelities r, including the memory errors of the Quantinuum data sheets, are converted to Pauli error probabilities $3r/2$ (one qubit) and $5r/4$ (two qubits), and errors reported as Pauli errors are used as they are. Readout and SPAM errors are used as measurement flip probabilities, and the SPAM errors of Helios and H2 also as reset failure probabilities. The other rates are assumptions: the rates in the column "idle" for the IBM, Google and USTC models (for the Google models estimated from coherence times); the reset rates and the (S1) and (S8) idle rates of the IBM and USTC models; for Willow, the reset and data-idle rates of the 72-qubit budget; and for both Google models, the split of the data-idle error over (S1) and (S8). The models assume independent stochastic Pauli noise on the circuit of Section 1.1, whereas the devices use other native gates and schedules and also suffer leakage and correlated or coherent errors.

B.2. **Google's Willow schedule.** The circuits published with [16] that we compared with the model of Remark 11.7 are those for $d = 3, 5, 7$ in basis X and $d = 7$ in basis Z, with ten rounds. For $d \leq 15$ the decoding graph of the model is graph-like in each sector, with twelve edges at each bulk node, and each bulk coefficient is at least the corresponding one of Lemma 2.2. In Monte Carlo simulations of the model, the logical failure curves cross at $t \approx 1.39$. The estimates of $\Pr(U)$ behind Remark 11.7 are lower bounds as in Remark 11.3, from 1,900 to 8,000 samples per point, and at $t = 0.9$ they still decrease with d. Decoded separately in each sector, the detection events measured on the device [16] give logical error rates per cycle, geometric means over patches and bases, of 9.9×10^{-3}, 6.6×10^{-3} and 4.3×10^{-3} for $d = 3, 5, 7$.

References

[1] E. Dennis, A. Kitaev, A. Landahl and J. Preskill, *Topological quantum memory*, J. Math. Phys. **43** (2002), 4452–4505.

[2] A. G. Fowler, M. Mariantoni, J. M. Martinis and A. N. Cleland, *Surface codes: Towards practical large-scale quantum computation*, Phys. Rev. A **86** (2012), 032324.

[3] A. G. Fowler, *Proof of finite surface code threshold for matching*, Phys. Rev. Lett. **109** (2012), 180502.

[4] C. Gidney, *Stim: a fast stabilizer circuit simulator*, Quantum **5** (2021), 497.

[5] Y. Tomita and K. M. Svore, *Low-distance surface codes under realistic quantum noise*, Phys. Rev. A **90** (2014), 062320.

[6] P. Aliferis, D. Gottesman and J. Preskill, *Quantum accuracy threshold for concatenated distance-3 codes*, Quantum Inf. Comput. **6** (2006), 97–165.

[7] A. A. Kovalev and L. P. Pryadko, *Fault tolerance of quantum low-density parity check codes with sublinear distance scaling*, Phys. Rev. A **87** (2013), 020304(R).

[8] D. Gottesman, *Fault-tolerant quantum computation with constant overhead*, arXiv:1310.2984.

[9] D. S. Wang, A. G. Fowler and L. C. L. Hollenberg, *Surface code quantum computing with error rates over 1%*, Phys. Rev. A **83** (2011), 020302(R).

[10] S. Yoshida, E. Lake and H. Yamasaki, *Proof of a finite threshold for the union-find decoder*, arXiv:2602.20238.

[11] Google Quantum AI and Collaborators, *Quantum error correction below the surface code threshold*, arXiv:2408.13687.

[12] K. Zhang, Z. Yi, S. Guo et al., *Learning to decode in parallel: self-coordinating neural network for real-time quantum error correction*, arXiv:2601.09921, Table 2.

[13] Quantinuum, *Quantinuum Helios product data sheet*, version 1.2, June 9, 2026, Table 1.

[14] Quantinuum, *System Model H2 product data sheet*, version 4.00, September 17, 2025, Table 1.

[15] Google Quantum AI, *Willow spec sheet*, December 9, 2024.

[16] Google Quantum AI, *Data for "Quantum error correction below the surface code threshold"*, Zenodo (2024), `https://doi.org/10.5281/zenodo.13273331`.

[17] IBM Quantum Platform, calibration data of the processor ibm_boston (Heron r3), read on October 18, 2026.

[18] J. F. C. Kingman, *A convexity property of positive matrices*, Quart. J. Math. Oxford **12** (1961), 283–284.

ORCID: `https://orcid.org/0000-0003-4668-3639`