

A MATCHING THRESHOLD OF 1.6×10^{-3} FOR THE ROTATED SURFACE CODE UNDER CIRCUIT-LEVEL NOISE

TEJASVI SINGH TOMAR

ABSTRACT. We consider the rotated surface code of odd distance d with its standard syndrome-extraction circuit, under circuit-level depolarizing noise in which every gate, idle step, preparation and measurement fails independently with probability p , and decoded by minimum-weight perfect matching with the log-likelihood weights of single faults. We prove that a memory experiment with T rounds fails with probability at most $0.032(T+1)(d+1)\lambda^d/(1-\lambda)$, where $\lambda = 0.997(p/p_{**})^{21/50}$, for every $p \leq p_{**} = 1.6 \times 10^{-3}$. Hence the logical error rate per round decays exponentially in d for all $p \leq p_{**}$. Earlier rigorous bounds for matching on surface-code circuits are 1.7×10^{-4} , due to Dennis, Kitaev, Landahl and Preskill, and 7.4×10^{-4} , due to Fowler. We keep the reduction of Dennis, Kitaev, Landahl and Preskill from a matching failure to a weighted self-avoiding path in the circuit-level decoding graph, and bound these paths with a local criterion for non-backtracking walks, which already gives 1.04×10^{-3} in closed form. The sharper value comes from the minimality of the correction along the path: we exclude the labellings of short stretches of the path into error and correction edges that contradict it, and count the remaining labelled walks with a potential function, which is checked in exact arithmetic on all local configurations. For a leaner schedule with fused measure-and-reset steps the same arguments give 1.53×10^{-3} and 2.4×10^{-3} . Both arguments extend to error rates that depend on the type of operation. They certify exponential decay for error models built from the published component error rates of Quantinuum’s Helios and H2 and of IBM’s Heron r3 processor `ibm_boston`, mapped onto this circuit. Any argument that bounds the failure probability by the probability of an increasing event, ours included, is limited by the smallest such event that contains all failures. Monte Carlo estimates place its threshold for this circuit near 4×10^{-3} . The computations that the proofs use are exact, and the fault table is cross-checked against Stim.

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 10.6). Rigorous thresholds are much smaller. The path argument of Dennis, Kitaev, Landahl and Preskill [1] gives 1.7×10^{-4} , and Fowler’s refinement [3] gives 7.4×10^{-4} , each for its own circuit and noise model. For the standard syndrome-extraction circuit we prove 1.04×10^{-3} in closed form and 1.6×10^{-3} with a computer-verified certificate, by a local criterion for weighted non-backtracking walks and by the minimality of the correction along the failure path.

1.1. Code and circuit. Fix an odd distance $d \geq 3$. 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

Date: 19 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.

are the points m with $|N(m)| \geq 2$, 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

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

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 $CX(m \rightarrow m + \delta_k^X)$ for each X-type m with $m + \delta_k^X \in \mathcal{D}$, and $CX(m + \delta_k^Z \rightarrow m)$ for each Z-type m with $m + \delta_k^Z \in \mathcal{D}$.
- (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 $T \geq 1$, 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

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

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. 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 Lemma 2.1, $|\sigma(f)| \leq 2$, 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

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

The weights are $w_e = \log((1 - r_e)/r_e)$ for numbers r_e with $q_e(1 - q_e) \leq r_e \leq q_e$. 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 $p \leq p_{**} = 1.6 \times 10^{-3}$ each q_e is at most $\frac{154}{15} p_{**} < 0.017$ (Lemma 2.1), 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. The results below hold 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 $d \geq 3$ be odd, $T \geq 1$, $B \in \{X, Z\}$ and $0 < p \leq p_* = 1.04 \times 10^{-3}$. Put $\lambda = 31\sqrt{p}$. The memory experiment in basis B with T rounds, decoded as in Section 1.3, fails with probability*

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

We call P_{fail}/T the logical error rate per round of the memory experiment in basis B . Since $T+1 \leq 2T$, it obeys $P_{\text{fail}}/T \leq 0.208 (d+1) \lambda^d / (1 - \lambda)$ uniformly in T . For $p \leq p_*$ we have $\lambda \leq 0.99972$, so this decays exponentially in d . At $p = 10^{-3}$, where $\lambda = 0.98031$, the bound reads $P_{\text{fail}}/T \leq 10.6 (d+1) \lambda^d$, and at $p = 10^{-4}$ it reads $0.31 (d+1) (0.31)^d$. The bound on P_{fail}/T is below 1 for all $d \geq 3$ when $p \leq 5 \times 10^{-4}$, for $d \geq 423$ when $p = 10^{-3}$, and only for $d \geq 63063$ when $p = p_*$.

Theorem 1.2. *In the setting of Theorem 1.1, let $0 < p \leq p_{**} = 1.6 \times 10^{-3}$ and put $\lambda = 0.997 (p/p_{**})^{21/50}$. Then*

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

Theorem 1.2 gives $P_{\text{fail}}/T \leq 0.064 (d+1) \lambda^d / (1 - \lambda)$, and $\lambda \leq 0.997$ for $p \leq p_{**}$. At $p = 10^{-3}$ we have $\lambda < 0.8185$, and the bound on P_{fail}/T is at most $0.353 (d+1) (0.8185)^d$, which is below 0.78 for all $d \geq 3$ and below 10^{-6} for $d \geq 87$. For $p < 1.047 \times 10^{-4}$ the rate $31\sqrt{p}$ of Theorem 1.1 is the smaller one.

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 give 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 $\lambda = 0.997 (p/(2.4 \times 10^{-3}))^{21/50}$, then

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

At $p = 10^{-3}$, part (b) gives $\lambda < 0.6903$, and the bound on P_{fail}/T is below 0.24 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 error models built from the published component error rates of six devices. Three models are certified: those of the Quantinuum Helios and H2 data sheets, and one with the error rates of IBM's Heron r3 processor `ibm.boston` on a square grid. The models of Google's and USTC's processors are not. For them, floating-point estimates show that no positive potential satisfies the inequalities of Lemma 6.5 with $\lambda < 1$.

Relation to earlier work. 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, and of 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 [1]. 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]. 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]. Our argument keeps the path reduction of [1] and replaces the path count by the local criterion of Section 4, applied to the exact edge probabilities of the circuit. For Theorem 1.2 it also uses the minimality of the correction along the path, beyond the half-weight inequality of [1]. For the circuit considered here, Fowler's argument gives 4.3×10^{-4} (Remark 10.1), so Theorems 1.1 and 1.2 improve the rigorous value for this circuit by factors of about 2.4 and 3.7.

Outline. Section 2 describes the decoding graph: each fault outcome produces one edge, the bulk has twelve edges per node, and the graph has 29 node types in each sector. Section 3 reduces failure to weighted paths from b_1 to b_2 , following [1]. Section 4 proves the local criterion, and Section 5 combines the pieces into Theorem 1.1. 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 are collected in Tables 3–6.

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 $d \leq 13$ and $T \leq 7$ 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 $c_e \leq 154/15$ for every edge e .*

Part (a) says that the circuit-level noise decomposes into a graph in each sector. 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 $d \geq 3$ and all $T \geq 1$.

2.1. The bulk. Let m be a Z -type measure qubit far from the boundary and $1 \leq r \leq T - 1$ far from 0 and T . 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. A CX copies an X on its control to its target. An X on the target of a CX is not copied. 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 gives 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 $(-2, -2, 1)$. 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 gives the offset $(-4, 0, 1)$. 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 gives 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 gives the offsets and coefficients stated for basis X . $\square$

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 $\min(T - r, 2)$.*

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 $q \rightarrow q + (-1, 1)$ 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.

- (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 $CX(q \rightarrow m')$ 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. 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. It 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. 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 gives 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. 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. $\square$

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. Detector values are linear in the Pauli frame, so 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, 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 . The degrees in R sum to $2|R|$, so 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. $\square$

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. It therefore has at least $(2d - 4)/2 = d - 2$ edges inside V , and two edges at the boundary vertices. $\square$

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. This is the argument of [1, Sec. 4]. The detector boundary of γ is empty, so $C \triangle \gamma$ 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 $w(C \cap \gamma) = w(\gamma) - w(\gamma \cap E)$. $\square$

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,

and outcomes at distinct locations are independent. Hence

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

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 $S \subseteq \gamma$ gives $\prod_{e \in \gamma} (q_e e^{w_e/2} + e^{-w_e/2})$. With $e^{w_e} = (1 - r_e)/r_e$ and $q_e(1 - q_e) \leq r_e \leq q_e$ we get

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

so each factor is at most a_e . $\square$

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 $w(S) \geq w(\gamma)/2$. Apply the union bound over γ and S , then Lemmas 3.4 and 3.5. $\square$

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

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

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 $L \geq 2$ the simple paths from b_1 to b_2 with L edges satisfy

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

Proof. For an oriented edge $\varepsilon = (u \rightarrow v)$ put $h(\varepsilon) = 1/(\lambda + a_\varepsilon)$. For $m \geq 1$ and an oriented edge $(x_0 \rightarrow x_1)$ with $x_0 \in V \cup \{b_1\}$ and $x_1 \in V$, let $W_m(x_0 \rightarrow x_1)$ 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 $x_{i+1} \neq x_{i-1}$ for $1 \leq i \leq m$. We show by induction on m that

$$(4.1) \quad W_m(x_0 \rightarrow x_1) \leq \kappa \lambda^m h(x_0 \rightarrow x_1).$$

The edge $\{x_0, x_1\}$ contributes $a_{x_0 x_1}/(\lambda + a_{x_0 x_1})$ to $\Sigma(x_1)$, so the hypothesis gives

$$(4.2) \quad \sum_{e \ni x_1, e \neq \{x_0, x_1\}} \frac{a_e}{\lambda + a_e} \leq 1 - \frac{a_{x_0 x_1}}{\lambda + a_{x_0 x_1}} = \lambda h(x_0 \rightarrow x_1).$$

For $m = 1$, $W_1(x_0 \rightarrow x_1) = a_{x_1 b_2}$ if x_1 is adjacent to b_2 and 0 otherwise. Since $a_{x_1 b_2} = (\lambda + a_{x_1 b_2}) \cdot a_{x_1 b_2} / (\lambda + a_{x_1 b_2}) \leq \kappa a_{x_1 b_2} / (\lambda + a_{x_1 b_2})$, (4.2) gives (4.1). For $m \geq 2$,

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

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 $\sum_{x_1 \sim b_1} a_{b_1 x_1} W_{L-1}(b_1 \rightarrow x_1) \leq \kappa \lambda^{L-1} \Sigma(b_1)$. $\square$

On the k -regular tree with all $a_e = a$ the criterion reads $ka/(\lambda + a) \leq 1$, that is $\lambda \geq (k-1)a$. This is the exact growth rate of non-backtracking walks there, so 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 . It therefore suffices to check $\Sigma(v) \leq 1$ at $p = p_*$. There $31\sqrt{p_*} = 0.999719\dots \geq 0.9997$, and $\Sigma(v)$ decreases in λ , so it suffices to check

$$(5.1) \quad \sum_{e \ni v} \frac{a_e}{0.9997 + a_e} \leq 1 \quad \text{at } p = p_*.$$

By Corollary 2.4 only 29 types per basis need to be checked. 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[†]. The values for all types are listed 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. 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 . Each signature gives one edge, so there are no parallel edges, and the number of edges at b_1 equals the number of its neighbours. By Proposition 3.6, Lemma 3.2 and Lemma 4.1,

$$P_{\text{fail}} \leq \sum_{L \geq d} \kappa \Sigma(b_1) \lambda^{L-1} = \kappa \Sigma(b_1) \frac{\lambda^{d-1}}{1-\lambda}.$$

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

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

6. PROOF OF THEOREM 1.2

Theorem 1.1 sums over all ways of labelling the edges of a path as errors or corrections, although most of these labellings cannot occur. 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. These rules are the main new ingredient of the proof. Throughout this section we fix $p_{**} = 1/625$ and the tilt $s = 21/50$.

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 , $w(\sigma_C) - w(\sigma_E) \leq w(P')$.*

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

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 $w(C') \leq w(C) - w(\sigma_C) + w(\sigma_E) + w(P')$, and $w(C) \leq w(C')$.

(b) The set $C' = C \triangle X \triangle P'$ satisfies $\partial C' = \partial C$ and $w(C') \leq w(C) - w(X) + w(P')$, and $w(C) \leq w(C')$. $\square$

A *labelled window* is a path $\sigma = (v_0, \dots, v_m)$ in G with $2 \leq m \leq 4$ 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. The detectors of σ lie within three such edges of its head, so 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 $k = n_C - n_E - h' \geq 0$ and

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

(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, $k = n_C - h' \geq 0$ and

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

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

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. Restricting the walks to the ball of the head keeps both rules local (Lemma 6.6). Dropping walks only drops exclusions, so any restriction of this kind would do.

Lemma 6.3. *Let $p \leq p_{**}$, 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 an excluded labelled window.*

Proof. 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 $w(P'') \leq w(P')$. Put $\varepsilon_e = -\ln(1 - q_e)$. Since $q_e(1 - q_e) \leq r_e \leq q_e$, each weight satisfies $|w_e + \ln q_e| \leq \varepsilon_e$, 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

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

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), gives $w(\sigma_C) > w(P'')$, which contradicts Lemma 6.1(b) for $X = \sigma_C$. $\square$

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) uses the E-edges as well: the correction edges of a window weigh at most as much as its error edges together with any alternative. 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 $b_1 = v_0, v_1, \dots, v_L = b_2$ 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

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

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.*

Proof. Suppose the decoder fails, and let $\gamma \subseteq E \triangle C$ be the path of Lemma 3.1. 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, $\prod_{e \in S} q_e \leq \prod_{e \in S} q_e e^{s w_e} \prod_{e \in \gamma \setminus S} e^{-s w_e}$. 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 $e^{-s w_e} \leq (q_e / (1 - q_e))^s$. $\square$

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 $l_3 l_4$, $l_2 l_3 l_4$ and $l_1 \dots l_4$, 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 $\sigma = (v_0, v_1, v_2, v_3; l_1 l_2 l_3)$ satisfies*

$$(6.3) \quad \sum_{\text{successors}} \omega(v_3 v_4, l_4) H(v_1, v_2, v_3, v_4; l_2 l_3 l_4) + \sum_{\text{terminal moves}} \omega(v_3 b_2, l_4) \leq \lambda H(\sigma).$$

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

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

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

Proof. For a state σ and $m \geq 1$ 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 $W_1(\sigma) \leq \lambda H(\sigma)$, and for $m \geq 2$ induction gives

$$W_m(\sigma) = \sum_{\text{successors}} \omega(v_3 v_4, l_4) W_{m-1}(v_1, v_2, v_3, v_4; l_2 l_3 l_4) \leq \lambda^m H(\sigma).$$

A labelled path $b_1, v_1, \dots, v_L$ begins with the state $(b_1, v_1, v_2, v_3; l_1 l_2 l_3)$, 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 $\sum \omega(b_1 v_1, l_1) \omega(v_1 v_2, l_2) \omega(v_2 v_3, l_3) W_{L-3}(b_1, v_1, v_2, v_3; l_1 l_2 l_3)$. $\square$

6.4. Classes and locality. The certificate gives the same potential to states with the same *class* and the same labels. The class of $(v_0, v_1, v_2, v_3; l_1 l_2 l_3)$ 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 $v_3 - v_2$. The *row* of a state is the inequality (6.3) written in terms of classes, labels and edge coefficients. 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 $(v_0, v_1, v_2, v_3; l_1 l_2 l_3)$ with a given last vertex v_3 are determined by the typed ball of v_3 . The sum of $\omega(b_1 v_1, l_1) \omega(v_1 v_2, l_2) \omega(v_2 v_3, l_3) H(b_1, v_1, v_2, v_3; l_1 l_2 l_3)$ 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 $d' \leq 25$ and $T' \leq 13$.*

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. The windows tested in these rows have the head v_3 , so 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. 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. 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 $T' = 13$. $\square$

6.5. Proof of Theorem 1.2. The instances with $d \leq 25$ and $T \leq 13$ have 2,546,882 distinct rows[†], and by Lemma 6.6 these are all the rows of all graphs $G_{d,T}$. The certificate 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 $\lambda = 997/1000$, when each factor $\omega(e, l)$ is replaced by a rational upper bound computed with integer 50-th roots[†] (Table 4). Let $p \leq p_{**}$ and $\rho = (p/p_{**})^s$. Each factor satisfies $\omega(e, l)(p) \leq \rho \omega(e, l)(p_{**})$, because $(p/p_{**})^{1-s} \leq \rho$ and $(1 - q_e)^{-1}$ decreases as p decreases. The excluded windows stay excluded by Lemma 6.3. Hence all rows hold at p with $\lambda_p = 0.997 \rho$. Let $K = 0.031062 \dots$ 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 $K_3 = 0.032106 \dots$ be an upper bound of the same kind for the sums of $\omega(b_1 v_1, l_1) \omega(v_1 v_2, l_2) \omega(v_2 v_3, l_3)$ over the paths b_1, v_1, v_2, v_3 with a given v_1 and

all labels[†]; such paths exist only for $d = 3$. By Proposition 6.4, Lemma 6.5 and Lemma 3.2,

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

Recall that $|N(b_1)| = (T + 1)(d + 1)/2$ by Lemma 2.1(e). With $\rho = \lambda_p/0.997$ the right side equals $(K + K_3)(T + 1)(d + 1)\lambda_p^d/(2 \cdot 0.997^3(1 - \lambda_p))$, and $(K + K_3)/(2 \cdot 0.997^3) \leq 0.032$. $\square$

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 $25.55\sqrt{p} \geq 0.9993$, 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[†]. The assembly of Section 5 with $c_{\max} = 94/15$ gives the prefactor $\sqrt{c_{\max}/(1 - c_{\max}p)}/25.55 \leq 0.0985$.

(b) As in Section 6, with $p_{**} = 3/1250$, $s = 21/50$ and $c_{\max} = 94/15$ in Definition 6.2. The instances with $d \leq 25$ and $T \leq 13$ have 2,548,765 distinct rows, all of which already occur for $d \leq 15$ and $T \leq 7$. 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 $K = 0.026464\dots$ and $K_3 = 0.028384\dots$, with $(K + K_3)/(2 \cdot 0.997^3) \leq 0.028$ [†]. $\square$

8. NON-UNIFORM ERROR RATES

The gates, measurements and idle steps of a real device 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 $0 < t \leq 1$.*

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

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

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*

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

and (6.2) by

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

keeping the conditions $k \geq 0$, and define labelled paths, states and rows as in Section 6 with these excluded windows, with the factors $\omega(e, l)$ evaluated at $q_e(\theta)$ and with a tilt $0 < s \leq 1/2$. 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

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

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

(b) Lemma 6.3 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 $k \geq 0$. Lemma 6.6 holds with the same proof. Each factor satisfies $\omega(e, l)(t\theta) \leq t^s \omega(e, l)(\theta)$, because $t^{1-s} \leq t^s$ for $s \leq 1/2$. 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 . The assembly of Section 6, with $\rho = t^s$, gives the bound. $\square$

Table 2 lists error models built from published component error rates of six devices, mapped onto the seven categories. 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). Errors reported as Pauli errors, and the cross-entropy-benchmarking errors of Zuchongzhi 3.2, 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. Rates that a source does not report are assumptions, listed with the sources in the accompanying files. The models are idealizations: they 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. Theorem 8.1 applies to the stated models, not to the devices.

Part (b), with $s = 43/100$, certifies three models:

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

for their rates scaled by any $t \leq 1$. Part (a) also certifies Helios and H2. For the other three models part (a) fails at $\lambda = 1^\dagger$, and floating-point estimates of the growth rate of the rows of part (b) are 1.39, 1.45 and 1.52.

model	CX	H	meas.	reset	idle	idle (S1)	idle (S8)	(a)	(b)
Helios [13]	1e-3	4.5e-5	5e-4	5e-4	9e-4	9e-4	9e-4	0.9325 [†]	0.75 [†]
H2 [14]	1.25e-3	4.5e-5	1e-3	1e-3	3e-4	3e-4	3e-4	0.8957 [†]	0.715 [†]
Heron r3 [16]	1.9375e-3	2.76e-4	3.48e-3	2e-3	1e-4	1e-3	2e-3	(1.212)	0.97 [†]
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)
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)
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)

TABLE 2. Error models (Pauli error probabilities per location) and the smallest rate λ of parts (a) and (b) of Theorem 8.1 at $t = 1$. Daggered values are exact certificates, those of (b) with $s = 43/100$. Values in parentheses are not certificates: in column (a) they are floating-point roots of the criterion, and in column (b) floating-point growth rates of the rows, computed from the instances with $d \leq 15$ and $T \leq 7$, the smallest over $s \in \{0.40, 0.43, 0.45\}$. Helios and H2: typical values of the data sheets. Heron r3: medians of `ibm_boston` on 2026-10-18, placed on a square grid, which the device does not have. Google 72q: the error budget of the 72-qubit processor in [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). Willow: mean simultaneous randomized-benchmarking errors of the spec sheet. Zuchongzhi 3.2: cross-entropy-benchmarking errors of the 97-qubit region, used as reported. Assumptions: the rates in the column “idle” for the IBM, Google and USTC 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).

9. COMPUTER VERIFICATION

Table 3 lists the computations behind the statements marked [†]. The fault table is computed by forward propagation of Pauli frames, independently of Stim’s detector-error-model code, and each Pauli outcome of each location is kept with its exact probability. Over the full range there are 1,924,076 outcomes. The circuit file also declares the detectors of the other type, and the decoding graph of each sector ignores them. All comparisons in V6 use exact rational arithmetic. One run of the accompanying script performs all checks in the table and passes only if all of them hold.

Table 4 lists the checks behind Theorem 1.2. A second script regenerates the rows of all instances with $d \leq 25$ and $T \leq 13$ and checks them against the certificate, which lists the 953,120 integers h . By Lemma 6.6 the rows at a detector depend only on its typed ball, so the script computes the rows at one detector for each distinct typed ball of the range; balls are compared as

	check	range
V1	noiseless circuit: all detectors and the logical outcome are deterministic, with value 0	$d \leq 13, T \leq 7$
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	$d \leq 13, T \leq 7$
V3	Lemma 2.1(a,b)	$d \leq 13, T \leq 7$
V4	Lemma 2.1(c,d,e)	$d \leq 13, T \leq 7$
V5	exactly 29 types per basis; all occur already for $d \leq 9, T \leq 5$; the set of types is the same for all $d \geq 9, T \geq 5$ in range (Lemma 2.3)	$d \leq 13, T \leq 7$
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 $d \leq 9$ and $T \leq 5$.

full tuples. The rows are generated by a C++ program with exact integer arithmetic. A Python implementation of the same rules gives the same rows at all detectors of five test instances with $d \leq 9$ and $T \leq 5$, and of two instances with the rates of Section 8. Classes are named by 64-bit hashes of the types and offsets they record. A collision would merge two classes, which can only make the system of rows harder to satisfy. The potential was found by iteration on the system of rows, a step that the proof does not use. A further test checks the exclusions and is not used in the proof. Let $w_e^+ = -\ln q_e - \ln(1 - q_e)$ and $w_e^- = -\ln q_e + \ln(1 - q_e)$ be the upper and lower bounds of the weights at p_{**} . For each excluded window, Dijkstra's algorithm on G/B with the weights w^+ finds an edge set with the same detector boundary as σ that is lighter than $w^-(\sigma_C) - w^+(\sigma_E)$, or an edge set with the same detector boundary as σ_C that is lighter than $w^-(\sigma_C)$. The test covers all 23.7 million excluded windows of the instances with $d \leq 9$ and $T \leq 5$, with no failure.

Table 5 lists the checks behind Theorem 1.3. They are the checks of Tables 3 and 4, run on the lean schedule with its own certificates. The Dijkstra test of the exclusions finds no failure among the 23.6 million excluded windows of the 40 instances with $d \leq 9$ and $T \leq 5$.

Table 6 lists the checks behind Section 8. A third script runs them on the error models of Table 2. For part (b) it reuses the code of Table 4, with the edge probabilities $q_e(\theta)$, written as integers over a common denominator, in place of the coefficients. As a test, with uniform rates it reproduces the rows of the code of Table 4 exactly for $(d, T) = (5, 3), (7, 4)$ and $(9, 5)$.

10. REMARKS

Remark 10.1 (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 $1/(484 \cdot \frac{72}{15}) = 4.3 \times 10^{-4}$, before the heavier boundary edges are taken into account. Two changes account for the improvement to 1.04×10^{-3} . Log-likelihood weights let each edge contribute $2\sqrt{q_e}$ instead of $2\sqrt{q_{\max}}$. The criterion of Lemma 4.1 replaces the factor 11 by an exact count of the twelve

	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 $d \leq 15$, $T \leq 7$; 13,658 distinct typed balls; 29 node types per basis, with injective identifiers; largest coefficient $c_{\max} = 154/15$	$d \leq 25$, $T \leq 13$
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 $K = 0.031062\dots$ and $K_3 = 0.032106\dots$, with $(K + K_3)/(2 \cdot 0.997^3) \leq 0.032$	all neighbours of b_1

TABLE 4. Computer verification for Theorem 1.2. Lemma 6.6 needs $d \leq 25$ and $T \leq 13$.

	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 $d \geq 9$, $T \geq 5$	$d \leq 13$, $T \leq 7$
L2	as V6: criterion at $p = 1.53 \times 10^{-3}$, $\lambda = 0.9993 \leq 25.55\sqrt{p}$; 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 $d \leq 15$, $T \leq 7$; positive potentials; all rows exact at $\lambda = 997/1000$, largest ratio 0.99799; $K = 0.026464\dots$, $K_3 = 0.028384\dots$, $(K + K_3)/(2 \cdot 0.997^3) \leq 0.028$	$d \leq 25$, $T \leq 13$

TABLE 5. Computer verification for Theorem 1.3.

	check	range
H1	every location has one of the seven categories; with equal rates the coefficients $c_{e,k}$ add up to c_e	$d \leq 9$, $T \leq 5$
H2	category-resolved types: 29 per basis, all present for $d = 9$, $T = 5$, and the same set for all $d \geq 9$, $T \geq 5$ in range	$d \leq 13$, $T \leq 7$
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 $97/100$; $(K + K_3)/(2\lambda^3) \leq 0.033$, 0.022 and 0.023	$d \leq 25$, $T \leq 13$

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

unequal edges at each node. At the bulk node with $\lambda = 1$ the left side of the criterion is below 1 at $p = 1.0419 \times 10^{-3}$ and above 1 at $p = 1.0420 \times 10^{-3}$, so the method cannot reach beyond

1.042×10^{-3} . As $p \rightarrow 0$ 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 10.2 (Limits of the certificate). For the rows of Section 6, computed from the instances with $d \leq 15$ and $T \leq 7$ with the exclusions at each p , the growth rate of the rows in floating point at $s = 0.42$ is 0.9941 at $p = 1.6 \times 10^{-3}$ and 1.0003 at $p = 1.62 \times 10^{-3}$. So a certificate of this kind cannot reach much beyond 1.6×10^{-3} . For the lean schedule the corresponding values are 0.9894 at 2.4×10^{-3} and 0.9997 at 2.45×10^{-3} . On the bulk lattice, without boundaries, the growth rate of the same scheme reaches 1 at 1.6185×10^{-3} for the 8-step circuit and at 2.451×10^{-3} for the lean schedule. Longer windows raise this bulk value slowly: with windows of five edges it reaches 1 at 1.6936×10^{-3} and 2.565×10^{-3} , respectively. Counting labelled paths edge by edge thus levels off near 1.7×10^{-3} and 2.6×10^{-3} . Counting each correction run once, by its two endpoints, goes further. A correction run from y to x is a minimum-weight path, so its weight is at least the distance between y and x ; and, by Lemma 6.1(b), the union of two correction runs has minimum weight among the sets with the same four endpoints as boundary. We counted error edges one by one and correction runs by their endpoints, with weight $e^{-sd(y,x)}$ for a run from y to x , and excluded the pairs of correction runs, joined by one error edge, that violate this minimality or the inequality of Lemma 6.1(a). On the bulk lattice the growth rate of this count reaches 1 at 1.91×10^{-3} for the 8-step circuit and at 2.90×10^{-3} for the lean schedule, with $s = 0.34$ and correction runs of up to six and five edges; the longer runs carry less than 1% of the weight. A certificate of this kind needs transitions over up to six edges and states that remember the previous correction run. Near the boundaries this gives far more rows than the certificate of Section 6, and we have not attempted one.

Remark 10.3 (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 bound obtained by counting paths, clusters or witnesses with a union bound. 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 $\partial E'$.

We estimated $\Pr(U)$ by Monte Carlo sampling for the memory experiment with $T = d$, deciding U in each sample by integer programming. Samples in which the program reached its round limit count as not in U , so the estimates are lower bounds; the share of such samples is below 1% for $d \leq 9$ at $p \leq 4.25 \times 10^{-3}$ and about 3% at $d = 11$ and $p = 4 \times 10^{-3}$. These estimates are not rigorous. For the circuit of Section 1.1 and $d = 5, 7, 9, 11$, with 800 to 2,550 samples per point, $\Pr(U)$ decreases with d at $p = 3.5 \times 10^{-3}$ and increases with d at $p = 5 \times 10^{-3}$. At $p = 4 \times 10^{-3}$ the four estimates agree within their statistical errors, at about 0.044, and at $p = 4.25 \times 10^{-3}$ they increase from $d = 7$ to $d = 11$. For the lean schedule the estimates 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 10.4 (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. 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 gives sharper bounds.

Remark 10.5 (Weights). Theorems 1.1 and 1.2 assume $q_e(1 - q_e) \leq r_e \leq q_e$. 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. The canonical decomposition used here sends the X part and the Z part of each outcome to its unique edge in each sector.

Remark 10.6 (Numerical threshold). We simulated this circuit with d rounds and decoded it by matching on the decoding graph of Section 1.3 with $r_e = q_e$, using PyMatching and 2×10^5 samples per point. The logical failure curves for $d = 3, 5, 7, 9$ cross near $p \approx 5 \times 10^{-3}$. The rigorous values p_{**} and p_* are smaller by factors of about 3.1 and 5. For the lean schedule the curves cross near $p \approx 8 \times 10^{-3}$, about 3.3 times the value 2.4×10^{-3} of Theorem 1.3.

CODE AND DATA

The programs, the certificates and the outputs of all checks in Tables 3–6 are archived at <https://doi.org/10.5281/zenodo.23472228>.

USE OF AI TOOLS

Generative AI tools were used to find and prove the results, to write the programs, and to draft the text.

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 b 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	$\{148^b, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
(2, 2, 1)	4	8	$\{154^b, 70, 66, 66, 48, 16, 8, 8\}$	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	$\{78, 78, 72^b, 70, 8, 8\}$	0.5939
(2, 6, 0)	3	7	$\{66, 41, 39, 29, 29, 22, 8\}$	0.5927
(2, 2, 5)	4	6	$\{100^b, 66, 51, 39, 16, 8\}$	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	$\{88^b, 66, 41, 29, 8, 8\}$	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	$\{94^b, 66, 39, 29, 8\}$	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	$\{78, 53^b, 41, 8, 8\}$	0.4235
(18, 2, 5)	1	3	$\{78, 49, 39^b\}$	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	$\{148^b, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
(16, 6, 1)	12	9	$\{148^\sharp, 72, 66, 66, 48, 16, 8, 8, 8\}$	0.8334
(2, 16, 1)	4	8	$\{146^b, 72, 66, 66, 48, 22, 8, 8\}$	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	$\{80^b, 78, 78, 48, 22, 8\}$	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	$\{100^b, 66, 51, 39, 16, 8\}$	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	$\{98^b, 66, 51, 39, 22\}$	0.5348
(2, 4, 0)	4	6	$\{88^b, 66, 41, 29, 8, 8\}$	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	$\{78, 61^b, 39, 22, 8\}$	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	$\{78, 39^b, 29\}$	0.3053
(6, 0, 0)	3	3	$\{78, 39, 29\}$	0.3053

TABLE 8. Node types in basis X .

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

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