Exact and Approximate Bulk Recovery in Self-Similar Tensor-Network Codes

Seth Douglas Nidhal Mghirbi

22 November 2026

Abstract

We study whether every bulk subsystem of a recursively generated tensor-network code can be recovered from its prescribed boundary footprint after one of five descendant intervals is erased, at every finite cutoff. Exact recovery is decidable for fixed-dimensional, seam-free local isometric circuits of bounded depth per scale, with finite-state placement and effective exact coefficients. A doubled purity calculation gives a finite endpoint transfer system. On explicit templates, exact recovery becomes Π_1^0-complete with unbounded per-scale depth, growing radial dimension or a sequential seam. A construction on $\{3, 32\}$ gives completeness on an explicit template, with local dimension finite per instance but unbounded across instances, for cyclic qudit tensors satisfying specified local identities and normalized disk isometry, without isometric annuli. Approximate recovery depends on input control. A homogeneous depth-one rational tree template of radial dimension 681408 has a Π_1^0-complete comparison at every rational $0 < \varepsilon < 48/49$, with a robust gap and an infimum not uniformly approximable within $12/49$. A single-excitation template of dimension 575960, with one amplitude matrix per port and a forwarded polarized probe, has a uniformly computable infimum but a Π_1^0-complete comparison at its exact root value $4/5$; all other rational thresholds are decidable. More generally, a single excitation with an inert no-arrival state and no fresh-input control has an exactly computable algebraic recovery value and decidable rational thresholds. In contrast, worst-case fresh bulk inputs can steer a dissipative memory along one ray, giving an undecidable $f_* \geq 1/6$ versus $f_* = 1/49$ promise, with $f_* \leq 9/49$ always. A probability budget bounds nonexceptional disjoint recovery footprints of any fixed source contrast Δ by $2/\Delta$; broadcast to these footprints is not needed. Necessity of dissipation and the case of pooled disconnected receivers remain open. Persistent fidelity gaps also hold for the growing-depth and a reweighted triangular construction.

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

Contents

1 Introduction **3**

1.1 Statements and dividing lines . . . 3
1.2 Established methods and the contribution . . . 5

2 Subsystem recovery and the seam starting point 7
2.1 Finite encodings and effective descriptions . . . 7
2.2 A homogeneous sequential isometry . . . 9

3 Exact recovery at bounded depth 10
3.1 The circuit presentation . . . 10

4 Explicit templates with growing resources 14

5 Approximate recovery at fixed depth and dimension 18
5.1 The performance quantity . . . 18
5.2 Records in a pure isometry . . . 18
5.3 The separated threshold problem . . . 20

6 Single-excitation propagation and the exact root threshold 21
6.1 Rational scattering and a polarized no-arrival state . . . 22
6.2 What the receiver can recover . . . 24
6.3 The role of the no-arrival plane and the known mechanism . . . 26

7 Input control and recovery values 27
7.1 A passive excitation and its finite recovery channel . . . 27
7.2 A polarized readout with unitary continuation . . . 30
7.3 Fresh inputs steering a dissipative memory on one ray . . . 30
7.4 A probability budget for disjoint recovery footprints . . . 34

8 Triangular disk isometries with matching rules 36
8.1 Whole-leg geometry and local identities . . . 36
8.2 An effective history tensor . . . 37

9 Interpretation and remaining boundaries 41

A Physical realization of the triangular matching table 43

B A normalized gap for triangular disks 46
B.1 Weights and boundary defects . . . 46

C Buffered protocols and annular circuits 50
C.1 Finite endpoint control for unequal refinement lengths . . . 50
C.2 Fixed dimensions and numerical variations . . . 52

1 Introduction

An encoding isometry preserves the joint quantum input globally. It need not preserve access to a specified input subsystem from a prescribed part of the output. In a network with repeated local structure, that distinction gives an infinite family of finite questions: for every bulk site and every cutoff containing it, can an observer recover its unknown quantum state after losing a fixed part of its boundary footprint? The observer must use one channel that works for arbitrary states of all other bulk inputs, including entanglement with them and with external references.

We classify this question for several precise presentations of self-similar codes. Our starting example is a fivefold refinement tree with a finite ancilla swept sequentially through each row and across an oriented seam. Its exact recovery problem is Π_1^0-complete (Theorem 4). The seam is an additional resource: its register interacts successively with an entire tape at every scale. Removing it and bounding the local circuit depth changes the answer. Allowing approximation changes the answer again, even when all dimensions and the depth remain fixed.

1.1 Statements and dividing lines

For the fivefold geometry, a site is a word v in $\{0,1,2,3,4\}^*$; its boundary descendants at cutoff n form an interval $I_n(v)$. The tested region is

$$R_n(v,j) = I_n(v) \setminus I_n(vj), \qquad |v| < n, \quad 0 \le j \le 4. \tag{1}$$

Everything outside that region is inaccessible. Exact all-scale recovery, denoted $\mathcal{P}$, means that each such query admits an exact subsystem decoder. Section 2 specifies the encoders and all quantifiers.

The principal exact result is **Theorem 6**: $\mathcal{P}$ is decidable for a single finite local isometric circuit repeated at each fivefold refinement, with fixed radial dimension, fixed depth and range, no seam, and homogeneous, finite-period, or explicitly finite-state address-dependent placement. The coefficient field must be computable, closed under effective conjugation, and have decidable equality. Algebraic numbers suffice; approximation access to computable real numbers alone is not covered. The proof computes a finite failing-cutoff bound from the architecture. No useful complexity bound is asserted.

The complementary results keep their uniformity convention explicit. **Theorem 9** proves Π_1^0-completeness *on an explicit template* of nearest-neighbor circuits with a fixed gate alphabet and fixed local dimension, using $O_w(\log(k+2))$ rounds at radial level k. Here w is the finite input program. **Corollary 12** permits the bound $O_w(1+g(k))$ for every fixed computable nondecreasing unbounded integer function g. The level scheduler evaluates prescribed repetition counts; it does not simulate the program to choose a gate. **Proposition 10** gives a depth-one explicit template when the local dimension may grow with the level. An arbitrary level-dependent gate generator is a broader model: it can put the answer to a bounded halting simulation into even one gate. That observation is not the content of these templates.

Theorem 26 proves Π_1^0-completeness on an explicit seeded tile-history template on the regular triangular graph $\{3,32\}$. Each instance uses one cyclic qudit vertex tensor and a

symmetric unitary edge tensor, and satisfies the one- and two-vertex identities of Steinberg et al., Definition 2 and Fig. 5, in the qudit extension stated in Section 8. Each normalized complete disk is an isometry on all independent bulk inputs. Local dimension is finite per instance and unbounded across instances. These codes are not layer-isometric and are not hyperinvariant in the sense of Evenbly's 2017 state-network construction [18]. Proposition 27 exhibits an annular kernel, also after the logical inputs are frozen. Thus the asserted radial implication of the local criterion fails for $p = 3$ in this qudit setting. Full-rim isometric annuli without bounded circuit presentations remain unclassified here.

For approximation use squared fidelity and optimize one decoder against all input states and references. Write f_* for the infimum over queries. **Theorem 15** proves that, for each fixed rational $0 < \varepsilon < 48/49$, deciding $f_* \geq 1 - \varepsilon$ is Π_1^0-complete, with a separated hardness promise, within homogeneous depth-one rational tree isometries of one fixed radial dimension. Its bulk dimension is seven and its radial dimension is the single constant

$$d = 7 \cdot 39 \cdot 4 \cdot 2 \cdot (8 \cdot 39) = 448 \cdot 39^2 = 681408. \tag{2}$$

Exact recovery in this same class is decidable. There is no uniform value algorithm with absolute error at most $12/49$. The size in (2) is an existence bound, not a practical encoding proposal.

Theorem 18 gives a different approximate boundary. On an explicit single-excitation depth-one tree code, with one amplitude matrix per scattering port, rational entries and radial dimension 575960, deciding whether f_* equals its root value $4/5$ is Π_1^0-complete. Yet its infimum is uniformly computable, and comparison with every other rational threshold is decidable. The no-arrival probe has two polarizations: it is rotated according to the logical sector and forwarded, along with an intact old quantum wire. This theorem supplies no robust gap. With a one-state no-arrival space, the analogous root-value question reduces to decidable exact recovery.

Theorem 20 gives an exactly computable algebraic value and decidable rational thresholds for the rational single-excitation class with an inert no-arrival state; fresh inputs do not steer its excitation. The scalar-unitary continuation extension with a two-sector polarized readout is **Theorem 21**. By contrast, **Theorem 23** realizes an undecidable $f_* \geq 1/6$ versus $f_* = 1/49$ promise when worst-case fresh bulk inputs steer a dissipative memory on one ray. Its radial dimension is $7 \cdot 158 \cdot 7^{80}$, about 5×10^{70}, and its high side lies in $[1/6, 9/49]$. **Proposition 24** bounds disjoint nonexceptional recovery footprints of a fixed input-pair contrast $\Delta > 0$ by $2/\Delta$. It explicitly excludes a bounded initial transient, a footprint with the spine leaf and a footprint containing all children of an off-spine source. Broadcast to the tested footprints is unnecessary. The hardness uses the fresh-input quantifier; no frozen-input reduction is asserted.

There are also late-cutoff gaps. A halting instance in the growing-depth template has $f_n = 1/49$ at every sufficiently late cutoff (Theorem 9). For a reweighted triangular template it has $f_n \leq 16383/16384$ at every sufficiently late cutoff, while nonhalting gives exact recovery throughout (Theorem 32). The latter statement needs reweighting: a nonzero exact witness in the unweighted disk can have vanishing normalized weight.

Presentation and resources	Predicate	Verdict and scope
Fixed local circuit; fixed dimension; finite-state placement; no seam	Exact $\mathcal{P}$	Decidable, Thm. 6
Fixed alphabet and dimension; $O_w(1+g(k))$ rounds, $g \nearrow \infty$ computably	Exact $\mathcal{P}$	Π^0_1-complete on an explicit template, Cor. 12
Depth one; growing radial dimension	Exact $\mathcal{P}$	Π^0_1-complete on an explicit template, Prop. 10
Homogeneous sequential encoder with seam memory	Exact $\mathcal{P}$	Π^0_1-complete, Thm. 4
Cyclic $\{3, 32\}$ disk codes; local identities; dimension varies across instances	Exact sector recovery	Π^0_1-complete on an explicit template, Thm. 26
Full-rim isometric annuli without bounded circuit presentation	Exact recovery	Unclassified here
Homogeneous depth-one rational recorded tree; $d = 681408$	$f_* \geq 1 - \varepsilon$	Π^0_1-complete with a gap for each rational $0 < \varepsilon < 48/49$, Thm. 15
Explicit single-excitation tree; one amplitude matrix per port; $d = 575960$	$f_* = 4/5$	Π^0_1-complete at the exact root threshold, Thm. 18
Same single-excitation template	f_* and $f_* \geq c$, $c \in \mathbb{Q} \setminus \{4/5\}$	Infimum uniformly computable; other thresholds decidable, Thm. 18
Single excitation; inert no-arrival state; no fresh-input control	f_* and all rational thresholds	Exact algebraic value; decidable, Thm. 20
One-ray dissipative memory; worst-case fresh inputs; $d = 7 \cdot 158 \cdot 7^{80}$	$f_* \geq 1/6$ versus $f_* = 1/49$	Undecidable promise; always $f_* \leq 9/49$, Thm. 23

Table 1: The input presentation is part of each decision problem. The exact decider requires effective equality and conjugation; the approximate upper bound uses algebraic entries. Completeness means computable many-one completeness.

1.2 Established methods and the contribution

The exact proof combines established ideas. Subsystem correctability is equivalent to complementary privacy at zero error, as in Kretschmann–Kribs–Spekkens, Theorem 5 [3]. We prove the required finite-dimensional specialization, including all spectator references. Bounded causal cones are standard in MERA [4, 5]; replica swaps leave only bounded endpoint neighborhoods in our macroscopic erasure experiment. Finite-dimensional weighted-automaton zeroness and reachable-span algorithms go back to Schützenberger and Tzeng [6, 7]. Our application appears new: one must compile every marked bulk address, loss and cutoff to a finite family of joint endpoint transfers, without assuming independent endpoints or dropping a different normalization in the two purity terms.

There is substantive overlap with recent circuit verification. Ross–Wesley, Theorem 5.10 and Corollary 5.11 [9], give parameter-sampling cutoffs under their cyclotomic gate hypotheses. Chuklanov, Theorem 1 and Lemma 2 [10], compiles ladder circuit generators into weighted automata of dimension independent of width and reduces equivalences to series zeroness. The shared algebraic step is finite transfer plus zeroness. Here words describe a bulk address, a

marked erasure and subsequent refinements; the geometric and subsystem compiler has to be proved. We use exact field arithmetic, without modular sampling. Neither circuit-verification result is imported as a theorem about recovery. Abdulla et al. [38] verify all-size circuit relations using transducers and synchronized weighted tree automata. Miller–Bakewell [39] gives finite verifying subsets of diagram equations under a connectivity condition, and Huang et al. [40] verify QEC programs with specified error models. Ferris–Poulin [41] and Farrelly et al. [42] develop tensor-contraction decoders, including recursive and holographic code families. These are precedents for represented program verification and finite-code decoding. Our contribution is a new recovery decision theorem via a finite endpoint compiler and standard weighted-automaton zeroness.

Matching-rule simulation and encoding tiles as tensor indices are standard [13, 14, 15, 16]. The triangular contribution is the realization of the constraints inside a nonzero, normalized disk code, with the stated local identities and an observable on actually erased whole legs. Approximate hardness uses standard PFA cutpoint and gap results [21, 22, 23, 24, 25]. The closest operational precedent is Elkouss–Pérez-García [26]: a finite-state channel capacity follows a PFA value and has an undecidable gap. Their finite-state communication channel is information stable, and its capacity has a uniform $1/5$ approximation obstruction. The repeated-isometry experiment here optimizes a worst-case subsystem decoder, has the affine identity (19), and retains a decidable exact endpoint. The one-ray construction likewise uses fresh inputs to program memory, so the distinction between chosen instructions and a source channel distributed to recipients must be retained. Our observable is optimal subsystem fidelity in a pure recursive isometry. Kliesch–Gross–Eisert, Theorem 5 [11], already distinguishes all-length tensor thresholds from linear zeroness. Quantum measurement occurrence [29] and general decidability questions in quantum information [30] provide related precedents, not hardness oracles for our experiment.

The approximate results distinguish passive excitation from input-controlled dissipative memory. When no fresh input can steer a single excitation and its no-arrival state is inert, Theorem 20 computes an algebraic recovery value by standard finite-dimensional methods. The polarized template of Theorem 18 instead has a computable infimum with an exact-root cutpoint obstruction. Theorem 23 gives a robust undecidable gap when worst-case fresh inputs steer a dissipative memory along one ray; Proposition 24 shows that broadcast to the tested footprints is not required. Section 5 supplies a fixed-dimensional recovery-fidelity realization of a known probabilistic-automaton gap, with an infimum that is not uniformly approximable within $12/49$. Whether dissipation is necessary for the one-ray gap remains open. Pooled disconnected receivers also remain open, and these examples do serve such receivers. The linear-amplitude mechanism has precedents in Turakainen [34, 56], Yakaryilmaz–Say (CSR 2009) [35] and Kondacs–Watrous [36].

Strict-cutpoint emptiness for measure-once *unitary* quantum automata with algebraic data remains decidable by Blondel–Jeandel–Koiran–Portier, Theorem 3.2 [27], and the algebraic-group methods of Derksen–Jeandel–Koiran, Theorem 1 [28]. The individual addressed port matrices in our single-excitation construction are contractions, so that unitary-letter hypothesis does not apply.

2 Subsystem recovery and the seam starting point

2.1 Finite encodings and effective descriptions

For a finite encoding $V : B \otimes H \longrightarrow R \otimes E$, B is the target input, H comprises all other bulk inputs, R is accessible and E is inaccessible. All spaces are finite-dimensional. Exact subsystem recovery means existence of a completely positive trace-preserving map $\mathcal{D} : R \to B$ such that

$$\mathcal{D}\Big(\mathrm{Tr}_E(VXV^\dagger)\Big) = \mathrm{Tr}_H X \quad \text{for every } X \in \mathcal{L}(B \otimes H). \tag{3}$$

Equality on operators implies equality after adjoining any reference. A decoder may depend on the finite query and the known encoder. It may not depend on the input state. The all-scale property is the conjunction of (3) over (1) and all cutoffs. The root radial input is a specified pure basis cap. Bulk dimension is $q = 7$ in every fivefold construction below.

A coefficient field is *effective* here if its elements have finite descriptions, field operations and conjugation are computable, and equality is decidable. Gaussian rationals and algebraic numbers are examples. Finite matrix equality and Gaussian elimination are then effective. Exact results do not assume order or an approximation oracle on this field. Approximate finite tests later use real algebraic arithmetic and order.

A set is Π_1^0 if membership is a universal quantifier over a decidable predicate. Blank-tape nonhalting is complete for this class: given $\forall s\, Q(x, s)$ with Q decidable, hardwire x into a program that tests $s = 0, 1, \ldots$ and halts when a test is false. This reduction is effective and is used throughout. Thus a computable map taking nonhalting to recovery and halting to failure proves hardness. A finite failing query gives the complementary enumeration needed for membership.

Lemma 1 (Finite recovery test and purity difference)**.** *Let $V : B \otimes H \to R \otimes E$ be an isometry, $\dim B = q$. For input bases $|a, h\rangle$ and an E basis $|i\rangle$, set $x_{ahi} = (I_R \otimes \langle i|)V|a, h\rangle$. Exact recovery holds if and only if*

$$\langle x_{ahi}, x_{bkt}\rangle = \delta_{ab} G_{hi,kt} \tag{4}$$

for a matrix G independent of a, b. Equivalently, form the normalized pure Choi state of V, naming the target reference A and all spectator references H'. With $p(Z) = \mathrm{Tr}\,\rho_Z^2$,

$$\Delta = p(R) - q^{-1} p(RA) = 0. \tag{5}$$

Consequently every specified finite query is decidable over an effective coefficient field. A nonzero scalar multiple of an isometry can be tested without adjoining normalization square roots.

Proof. Dilate a putative decoder to an isometry $U : R \to B \otimes J$. On $V|a, h\rangle$, the pure target marginal forces $(U \otimes I_E)V|a, h\rangle = |a\rangle|\eta_{ah}\rangle$. For fixed h, apply recovery to $(|a\rangle + |b\rangle)|h\rangle/\sqrt{2}$. Comparing coefficients with the required pure target forces $\eta_{ah} = \eta_{bh}$, including the relative phase. Call the common vector η_h. Projecting its E factor and taking inner products gives (4).

Conversely put $z_{hi} = x_{0hi}$ in a copy of R. Equation (4) says that $x_{ahi} \mapsto |a\rangle z_{hi}$ preserves every Gram entry. A zero linear combination has a zero image, so this defines an isometry on

their span. Enlarge J and complete orthonormal bases to extend it to all of R. On a codeword its action is $|a\rangle \sum_i z_{hi}|i\rangle_E = |a\rangle|\eta_h\rangle$. Isometry of V implies $\langle \eta_h, \eta_k\rangle = \delta_{hk}$. Discarding J, E therefore gives (3) on every matrix unit $|a,h\rangle\langle b,k|$ and hence on every operator.

Set $Y = EH'$. The Choi entries show that (4) is exactly $\rho_{AY} = I_A/q \otimes \rho_Y$. Expanding the square gives

$$\|\rho_{AY} - I_A/q \otimes \rho_Y\|_2^2 = p(AY) - p(Y)/q = p(R) - p(RA)/q,$$

where the last equality uses the complementary Schmidt spectra of the pure Choi state. This proves (5). A finite list of equations (4), taking G from $a = b = 0$, is an effective test. For an unnormalized scalar isometry the Gram equations are homogeneous; alternatively both purities are multiplied by the same positive normalization square. The target and all spectator references have participated in this test. □

This is the zero-error specialization of complementary correctability/privacy [3, Theorem 5]. Its hypotheses hold because V is an isometric dilation of the accessible channel, its complementary output is E, and B, H are the input subsystem split. The direct proof above supplies the precise all-spectator version used here.

Lemma 2 (Two data encoders)**.** *For $a, b \in \mathbb{F}_7$, define*

$$G|a,b\rangle = \frac{1}{7}\sum_{r,s\in\mathbb{F}_7} |f(1), f(2), f(3), f(4), f(5)\rangle, \qquad f(t) = r + st + at^2 + bt^3, \tag{6}$$
$$W|a,b\rangle = |b,a,0,0,0\rangle.$$

Both are rational isometries. Any four shares of G recover both inputs with all reference correlations. Deleting port zero of W removes the fresh input b completely. An all-G tree satisfies $\mathcal{P}$.

Proof. Four distinct evaluations determine a polynomial of degree at most three: a nonzero difference polynomial cannot have four roots. Distinct (r, s, a, b) therefore give orthogonal full evaluation words; for fixed a, b the 49 amplitudes have total squared norm one. The map W is an injective basis map.

If evaluation e is lost, the four retained values can be changed coherently into (r, s, a, b) by the inverse Vandermonde permutation. Replace r by $u = f(e)$ at fixed s, a, b, again a permutation. The resulting joint vector is

$$|a,b\rangle \otimes \left(7^{-1/2}\sum_s |s\rangle\right) \otimes \left(7^{-1/2}\sum_u |u\rangle|u\rangle_{\text{lost}}\right).$$

The junk is normalized and independent of the logical input. Linearity proves the assertion with references. In W the fresh b occurs only in port zero. In a tree, invert the full isometry of each intact child subtree, starting at its leaves, to recover its entering share and discard its later bulk inputs. An inverse on the range extends to a channel by projection and an arbitrary replacement off the range. Apply the four-share decoder at the queried vertex and keep its fresh input. Every operation is within the four intact descendant intervals and works for arbitrary other inputs. □

The evaluation method is the standard quantum polynomial-code method of Cleve–Gottesman–Lo, Eq. (5) [12]; the displayed parameters and erasure decoder are proved above.

2.2 A homogeneous sequential isometry

Sequential ancilla generation has a matrix-product-state description [46], and causal processes admit sequential memory-channel representations [47]. Kari–Salo–Worsch [48] characterize reversible single sweeps without intermediate states. The construction below adds expanding radial storage and an inaccessible terminal carry: it is a seam-assisted recovery realization of the nonhalting problem. Elkouss–Pérez-García [26] give the related finite-state-channel use of PFA memory. In the seam class, one finite rational tensor has ports

$$T : K \otimes D \otimes \mathbb{C}^7 \longrightarrow K \otimes D^{\otimes 5}, \qquad T^\dagger T = I. \tag{7}$$

Vertices are visited breadth by breadth in lexicographic order. The outgoing K of a visit enters the next visit, including the passage from the end of one row to the beginning of the next. This is fixed wiring, with a distinguished root and seam; it is not rotationally invariant. The boundary at cutoff n contains 5^n radial registers and the terminal K. That terminal register is always inaccessible in (1). The root caps are fixed zero basis vectors. Dimensions $\dim D, \dim K$ may vary across finite inputs, and are fixed throughout an instance. Composition makes every initialized V_n an isometry.

Lemma 3 (Finite control on the expanding rows)**.** *From any deterministic blank-tape machine one can effectively construct* (7)*, with entries in* $\{0, 1, 1/7\}$*, such that its data tree is all* G *if the machine never halts. If it halts at time* h*, every vertex at depth greater than* h *uses* W*. At every cutoff the entire control output is a fixed basis product independent of all quantum bulk inputs. The tensor is isometric also on unreachable inputs.*

Proof. Use a two-sided tape with blank β and an absorbing, stationary halting head. Complete missing transitions by halting. Let $\mathcal{A}$ be the finite alphabet of unheaded symbols and symbols carrying a head state. A total rule $F : \mathcal{A}^3 \to \mathcal{A}$ performs one machine step: the center updates its own symbol if headed; the new head arrives from a neighbor moving into it or from a stationary center head. Set conflicting-head triples to β. Then $F(\beta, \beta, \beta) = \beta$ and on valid configurations the parallel rule is exactly a machine step.

A primary radial label is inactive $\perp$ or (x, e), $x \in \mathcal{A}$, $e \in \{0, 1\}$. The bit e marks the terminal active cell of a row. The horizontal state is $(u, v, l) \in \mathcal{A}^2 \times \{0, 1\}$, initialized to $(\beta, \beta, 0)$. The root has the initial headed blank with terminal marker. Inactive visits emit five inactive labels and preserve the horizontal state. An active visit x emits $F(u, v, x)$ as an unmarked active label at child zero. If $e = 0$, other children are inactive and the pair becomes (v, x). If $e = 1$, children one and two also receive $F(v, x, \beta)$ and $F(x, \beta, \beta)$, the latter marked; the pair resets to (β, β). Children three and four are inactive. Update l by OR with the predicate that x has a halting head, before choosing the data encoder.

If the active word is $x_1 \cdots x_m$, its output is the radius-one update on that interval enlarged by one blank on each side. Inactive sites change neither the carry nor the order. The terminal visit resets the pair. Induction from the root gives the tape cells $[-k, k]$ at row k, with its unique head and a unique terminal marker. A halt at time h is read on row h; all later rows see $l = 1$. With no halt l remains zero.

To make this deterministic update isometric on every input, use radial fields (p, z, a): primary label p, archive z, data $a \in \mathbb{F}_7$. Let the archive alphabet contain a zero and an injective copy $c(p, k)$ of every primary/horizontal pair. Put zeros in child archives zero, one

and two, the incoming archive z in child three and $c(p,k)$ in child four. This retains all old control labels regardless of whether their update is injective. Different incoming (p,z,k) have orthogonal outgoing classical words. Within each fixed tuple apply G or W according to the updated bit, an isometry on (a,b). This proves $T^\dagger T = I$ on the entire input basis. All dimensions and entries are obtained from the finite transition table without executing it.

On initialized inputs a visit deterministically updates a fixed control basis tuple and acts on data only with the selected quantum isometry. Induction over the visits proves that controls are independent of arbitrary joint quantum inputs, even with references. No quantum information is carried in K or the archives. Thus later data remain within their descendant subtrees. □

Theorem 4 (Seam class)**.** *For finite rational isometries* (7) *with the stated caps, wiring and inaccessible terminal memory, exact all-scale recovery* $\mathcal{P}$ *is* Π^0_1*-complete. On the reduction image, nonhalting gives fidelity one for every query. A halt at time* h *gives optimal fidelity* $1/49$ *for child-zero loss at every vertex of depth greater than* h*, at every subsequent cutoff. At cutoff* $n \geq h+2$ *more than four fifths of bulk sites have such a failing query.*

Proof. Validate a rational matrix by $T^\dagger T = I$ and reject malformed inputs. Enumerate n, its finitely many vertices and losses, and apply Lemma 1. A failed query enumerates the complement, proving membership. Lemma 3 is a computable reduction from blank nonhalting. In the nonhalting case Lemma 2 supplies all decoders. After a halt choose a site beyond row h, fix all spectators, and erase its child-zero interval. Its unknown fresh datum is a tensor factor routed solely into that interval; control independence and subtree confinement prevent its return. The retained state is independent of it, so an exact decoder is impossible.

For a target maximally entangled with a reference of dimension seven, every decoded output is $\tau \otimes I/7$, with Bell overlap $\operatorname{Tr}\tau/49 = 1/49$. Conversely replacing the target by $I/7$ guarantees at least $1/49$ for arbitrary inputs, by the purification argument in Section 5. This is the optimum. The newest bulk row has 5^{n-1} sites out of $(5^n-1)/4$, a fraction greater than $4/5$. All those sites lie beyond h for $n \geq h+2$. Nonhalting completeness proves the lower bound, including persistence after removal of any finite set of cutoffs. □

3 Exact recovery at bounded depth

3.1 The circuit presentation

Remove K entirely. At row k, $N = 5^k$ sites have addresses of length k in base five, including leading zeros. Children of index x are $5x, \ldots, 5x+4$. One scale step is a finite number of local isometric rounds

$$(D \otimes \mathbb{C}^7)^{\otimes N} \longrightarrow D^{\otimes 5N}.$$

Group five children of a parent in its parent cell when measuring range. Intermediate subregisters, rectangular isometries, clean local ancillas and disentanglers are allowed. Each intermediate wire has an assigned cell. Every nonzero displacement is counted as a routing round; a wire may not circle a row in one uncounted step. Round count, dimensions and integer port offsets are finite input data, independent of k. Gates in a round have disjoint ports.

Placement is homogeneous or specified by a finite automaton reading the most significant digit of a parent address before the remaining digits. Its state indexes a finite table of gate anchors, port offsets and matrices in each round. A product automaton covers all tables. This includes finite period, aligned substitutions of types and finite-state exceptions. For example a residue a modulo t updates as $a \mapsto 5a + d \pmod{t}$, and a level residue updates once per digit. No divisibility of 5^k by t is assumed. Small rings use actual aliases of port locations. Illegal double use of a wire, dangling inputs or conflicting gates do not define instances. The proof below also decides this local legality. The root is a fixed pure cap. These assumptions are the *bounded-depth class*.

Proposition 5 (Finite endpoint compiler)**.** *For each bounded-depth description there is an effective finite directed graph with three phases: address, mark and extension. Its paths enumerate exactly all triples (n, v, j) in (1). At each state s is a finite-dimensional vector space over the coefficient field; edges are linear maps, the root has an initial vector and marked states have terminal linear evaluations. On the path for a query the evaluated scalar is Δ in (5) divided by a strictly positive scalar. All constants, circuit-validation neighborhoods and spaces are computable from the finite description.*

Proof. We give the geometric constants, the doubled contraction, and its finite controller explicitly.

A conservative radius, without an unspecified enlargement. For round i, lift every possible gate's occupied cells and anchor to integer offsets. Routing rounds count both wire endpoints. Let a_i be the maximum absolute occupied offset from its anchor, including zero, and d_i the maximum difference of two occupied offsets on one gate. A gate meeting a set X uses only cells within d_i of X; its anchor is within a_i of a port. Define

$$D_0 = \sum_i d_i, \qquad A_0 = \max_i a_i, \qquad L = D_0 + A_0 + 4, \tag{8}$$
$$w = 10(L+3), \qquad u = w + 2L + 5, \qquad K = 100(w + L + 10), \qquad H_0 = 20K + 100.$$

Empty maxima and sums are zero. If legality or lookup needs a fixed additional table halo, include it in A_0 before these choices.

A fine center is $5c + s$, $0 \le s \le 5$; $s = 5$ covers the right endpoint of a marked cell. A fine offset t, $|t| \le w$, belongs after grouping to parent offset $\lfloor (s+t)/5 \rfloor$. Its absolute value is at most $\lceil (w+5)/5 \rceil$. Two extra cells cushion a cut and its alignment, so take

$$e = \left\lceil \frac{w+5}{5} \right\rceil + 2.$$

A gate straddling a tag cut has its two extreme cells within $d_i + 1$ of that cut; begin with virtual exceptional cells on both sides. Pulling through round i enlarges the exception by at most d_i. Hence its final radius is bounded by

$$e + D_0 \le \frac{w+9}{5} + 2 + D_0 \le w/5 + 4 + L \le w, \tag{9}$$

because $4w/5 = 8(L+3)$. Anchors cost at most A_0 extra cells. A halo u bounds $w+D_0+A_0+2$ even with an extra inspection cone. All fine indices needed to inspect the children of these coarse cells are within $5u+9$ of the linked fine center. Direct subtraction gives

$$K - (5u+9) = 95w + 90L + 966 > 0, \qquad u < K.$$

For updating windows, $\lceil (K+5)/5 \rceil \le K$. Finally $H_0 > 4K+2$, so all coincidences at distance at most $2K$ are visible in capped gaps. These inequalities remain conservative on small rings by computing their actual modular positions.

Pulling uniform swaps through the circuit. Let S_X exchange two copies of X. For every isometry $U : X \to Z$,

$$(U \otimes U)^\dagger I (U \otimes U) = I, \quad (U \otimes U)^\dagger S_Z (U \otimes U) = S_X. \tag{10}$$

The second identity follows from $S_Z(U \otimes U) = (U \otimes U)S_X$. Inserting basis pairs gives $\mathrm{Tr}[(M \otimes N)S] = \mathrm{Tr}(MN)$. Thus a product of physical swaps on R evaluates $p(R)$ in two Choi copies; adding the swap on A evaluates $p(RA)$.

Contracting the two normalized Choi pairs of a fresh bulk with reference tag $t \in \{0,1\}$ applies

$$\mathcal{L}_t(X) = q^{-2} \,\mathrm{Tr}_{B\otimes B}[(I \otimes S_B^t)X]. \tag{11}$$

The reference transpose leaves I, S unchanged. On an isolated full tag S_B^b the scalar is one for $b = t$ and $1/q$ otherwise, since $\mathrm{Tr}\, I = q^2$ and $\mathrm{Tr}\, S = q$. No square roots of Choi normalizations are needed. Every spectator reference has $t = 0$; the selected one is the sole exception.

At each row keep a patch P consisting of all sites within w of the designated centers. Use one arbitrary *joint* operator X on two copies of $D^{\otimes |P|}$. Outside P, insert sitewise I or full S according to a piecewise constant tag background b. Denote this insertion by $J(P, b, X)$. Patches are unions of actual sites; overlaps are identified, never tensorized twice. This is essential when endpoints are close or when the whole ring lies in the patch.

In a reverse scale step group children by parent. Any block meeting a fine patch is within the radius used in (9). Process explicitly all gates touching an exception or bridging a change of tag. Every other gate has uniform output tag, which passes backwards by (10). Inducting over rounds and using (9) places every nonuniform input inside the coarse patch. Pad it to radius w with known tags and use (11) on its fresh inputs. Outside the patch, each tag-one parent contributes $1/q$ and each tag-zero parent contributes one. Fixed pure ancillas contribute one to either uniform tag. Consequently the pullback is

$$J(P_f, b_f, X) \longmapsto \lambda J(P_c, b_c, TX), \qquad \lambda = q^{-m} > 0, \tag{12}$$

where m counts the tag-one parents outside P_c and T is a finite linear map obtained by contracting gates in the halo. Use the same halo and padding convention in both purity terms and for every patch matrix unit. Then λ is independent of X and of the marked reference tag. It may vary with the path, which does not affect a pathwise zero test. Arbitrary correlations among endpoints and input bulk systems have not been discarded.

The three phases and the stored signature. Begin at root circumference one, center $c = 0$, zero background. An address edge $d \in \{0, \dots, 4\}$ sends (N, c) to $(5N, 5c + d)$ and keeps zero background. After any finite digit string, a mark edge j makes four labeled cuts

$$(5c, 5c + j, 5c + j + 1, 5c + 5) \pmod{5N}. \tag{13}$$

Their cyclic gaps are $(j, 1, 4 - j, 5N - 5)$ with tags $(1, 0, 1, 0)$. An arc includes its left endpoint and excludes its right endpoint. Zero gaps contain no sites. This is exactly four retained children. Subsequent extension edges multiply centers and gaps by five, keeping tags. Stopping after h address digits, one mark and t extensions gives v of depth h and $n = h + 1 + t$. Every query has exactly this description.

Store the phase, N and all cyclic gaps capped at H_0, and the placement automaton state at offsets $-K, \dots, K$ from each labeled center. A capped value means at least H_0. The address phase has one center and gap N. A short directed distance is a sum of consecutive nonnegative gaps; any sum smaller than H_0 is exact. Therefore overlap tests and local tag masks are determined. If $N < H_0$ use the exact circumference and its aliases. If $N \geq H_0$, a local window cannot wrap more than once; a long unseen arc has its stored constant tag. Canonically enumerate patch sites by labeled center and signed offset, keeping the earliest occurrence of an actual site. Coincident cuts and multiple zero gaps are harmless with this convention.

To update a state at offset s' about $5c + d$, use the parent state at offset

$$\left\lfloor \frac{d + s'}{5} \right\rfloor \pmod{N} \quad \text{and append digit } (d + s') \bmod 5. \tag{14}$$

Here $|s'| \leq K$ and the parent offset stays inside $[-K, K]$. Euclidean division makes (14) valid across the cyclic origin as well. Mark offsets $d = 0, j, j + 1, 5$ are all covered. On extension use $d = 0$ at each cut. Gaps and circumference multiply by five with saturation; at a mark use (13). A saturated N already makes $5N - 5$ saturated. Root windows contain the state of the empty address. There are finitely many signatures and their deterministic updates are effective. The halo inequalities prove that equal signatures compile identical maps T even when their remote gate layouts differ.

Circuit legality can also be checked. A bad connection, overlapping gate use or dangling port is witnessed in a fixed local neighborhood. Enumerate reachable one-center signatures from the root using all address digits, and inspect the catalog and the actual aliases at every visited signature. This covers every cell at every level. Isometry of every numerical gate is a finite matrix test. This avoids implicit brickwork assumptions on the odd rings 5^k.

Linear spaces and marked subtraction. At a signature the dual patch-operator space has dimension

$$d^{4|P|}, \qquad |P| \leq 4(2w + 1). \tag{15}$$

Initialize its functional by evaluation in two copies of the root cap. Move it outwards by precomposing with each reverse map T. At the mark use

$$T_{j,0} - q^{-1} T_{j,1}, \tag{16}$$

where only the target reference swap differs. Coarse background is zero and the target lies inside the patch, so the omitted scalar in (12) is common. At extensions propagate the already-subtracted functional with the ordinary map. At any marked terminal insert the actual product of I, S on patch sites specified by the tag mask. Reversing a whole path shows that this evaluation times the product of its positive omitted factors is exactly (5). All transfers are computable by their action on matrix units. This proves every assertion. □

Theorem 6 (Bounded-depth decidability)**.** *For every finite bounded-depth circuit description in Section 3, with fixed local dimensions and coefficients in a specified effective conjugation-closed field with decidable equality, exact all-scale subsystem recovery $\mathcal{P}$ is decidable. If it fails, a failing cutoff is effectively bounded by the sum of the finite controller's ambient dimensions. For fixed discrete architecture this is one computable bound for all valid numerical gate assignments.*

Proof. Construct Proposition 5's finite graph. At each state retain a basis of the span of reachable functional vectors. Initialize the root and repeatedly apply every outgoing edge map to every newly admitted independent vector. Admit a destination image exactly when it increases that destination span. Equality and Gaussian elimination are effective in the field. If M is the sum of the dimensions in (15) over all states, at most M admissions occur; the work list terminates.

Every stored basis vector can be retained as the actual path image that was submitted. Every path image lies in the saturated spans, by induction on its length. A terminal linear evaluation is nonzero somewhere on a path if and only if it is nonzero on some stored basis vector: linearity gives one implication and actual-path retention gives the other. By Lemma 1 and the positive scalar in Proposition 5, a nonzero value is precisely a failing physical query. Save predecessors to output it. Otherwise all queries pass.

For a cutoff bound, consider the direct-sum spans of paths of length at most r. They increase with r. A stage with no increase is stable under every edge and hence permanently stable. Each strict increase adds at least one dimension. Starting with the root vector, stabilization occurs by length M; each edge is one scale step, so a failing query exists at cutoff at most M. A finite superset of signatures gives a conservative bound if desired. Signatures, graph edges and ambient dimensions depend only on discrete port dimensions, range, round count and placement table. Changing gate entries changes the reachable span but cannot enlarge that ambient bound. This proves the uniform architectural claim. □

The final span algorithm is the standard zeroness method for a finite weighted automaton [6, 7]; the familiar short-word bound is also stated in Kiefer et al., Proposition 3.1 [8]. Here valid words are restricted by finite control, implemented by the direct sum. Every hypothesis of this method follows from Proposition 5. A finite test at each cutoff alone would have given only membership in Π_1^0, not decidability.

4 Explicit templates with growing resources

The history storage uses Bennett's standard reversible-simulation method [43]. Shiraishi–Matsumoto [44] use reversible halting histories to distinguish asymptotic readouts in a

closed one-dimensional Hamiltonian setting. Here the contribution is an explicit scheduled recovery-hardness template, with an elementary growing-dimension extension.

The input language in this section is a finite program w for one fixed one-tape universal interpreter U. Its finite tape alphabet and radius-one update rule are fixed. It has a stationary absorbing halting state. An effective interpreter can scan a delimited transition table, find the entry for a simulated state and symbol, change the marked cell and move its mark; finite strings can be shifted to make space. These finite scans on one tape give the usual complete halting language. Gate scheduling below depends on k and the literal length and symbols of w, never on whether $U(w)$ halts.

Lemma 7 (Local reversible bounded-time simulation)**.** *There is a fixed alphabet of nearest-neighbor basis permutations, clean-lane preparations and local isometries, with fixed local dimensions, that at row k tests whether $U(w)$ has halted by*

$$m = \left\lfloor \sqrt{\lfloor \log_2(k+2) \rfloor} \right\rfloor$$

steps, whenever the work window fits. It uses $O_w(\log(k+2))$ rounds. Its work is independent of every quantum data input. On rows where it does not fit the test is false.

Proof. Put $\ell = \max(1, |w|)$, $M = m + \ell + 2$, $L_0 = 2M + 1$ and $N = 5^k$. A single deterministic seed bit follows only the middle child at each refinement, hence marks one site c at level k. For $N \le (m+2)L_0$ do no simulation. Otherwise prepare a fixed number of clean work tracks. A shifted seed cursor toggles a window bit on $[c - M, c + M]$, taking $O(M)$ local shifts. Load the literal program starting at c, and its unique interpreter head, in $O(\ell)$ rounds using a shifted seed copy and local controlled XORs. This requires no address-dependent choice of a site.

A right shift of a track H can be paired with a left shift of an equally sized zero track J: swap H_i, J_i on-site and then J_{i-1}, H_i on each edge. It gives $H'_i = H_{i-1}$, $J'_i = J_{i+1}$. If whole-site gates must be disjoint, color cycle edges in three rounds. Since $N = 5^k$, for $k \ge 1$ the repeating edge colors $(1, 2, 1, 2, 3)$ are legal also across the join. The subregister swaps commute and implement the stated shift exactly. All gates have bounded range.

Let F be the fixed three-letter rule for one step of U, with $F(0, 0, 0) = 0$. The operation $Y_i \mapsto Y_i \,\mathrm{XOR}\, F(A_{i-1}, A_i, A_{i+1})$ is a constant-depth permutation, using temporary neighbor copies, paired shifts and their inverses. Pad finite letter encodings to binary words and extend unused entries by a total lookup. Initialize A with the tape, Y blank, and a history track H blank. Repeat m times:

1. Compute $F(A)$ by XOR into Y everywhere.
2. Within the marked window swap A, H, then shift H right by L_0 using paired shifts.
3. Swap A, Y everywhere.

After step i, history windows occupy translates by $L_0, \ldots, iL_0$ of the original window. They are disjoint and avoid it, because $N > (m+2)L_0$. The active head and all nonblank current tape remain within the original window for m steps. Thus depositing the old tape clears A everywhere; the final swap restores Y to blank and makes A the updated tape. This proves the invariant by induction. Counter-moving tracks remain zero on the initialized trajectory. The final halting-head predicate is true at a site exactly when halt has occurred by time m.

Setup costs $O(M+\ell)$ and the loop $O(mL_0)$ rounds, hence $O_w(1+m^2) = O_w(\log(k+2))$. All operations are reversible on their full work spaces, not merely on the initialized trajectory. The fixed interpreter and letter-loading alphabet are independent of w; the word is encoded by a finite sequence of literal loading rounds. Since $m \to \infty$ and 5^k eventually exceeds $(m+2)L_0$, every finite halting time is eventually tested. □

Lemma 8 (Archiving the test and distributing its selector)**.** *The simulation of Lemma 7 can be included in a local isometric fivefold row with fixed radial dimension and only radial outputs. Its data encoder at a parent is W if either an inherited bad bit or its local halting predicate is true, and is G otherwise. Every child inherits that selector. All controls remain basis values independent of all quantum data on initialized inputs.*

Proof. Give each radial site the data, seed and bad bits and an archive large enough for the fixed complete work tuple. Prepare work on separate clean lanes. The reversible work reads the incoming seed without changing it. At the split, forward the seed only to the middle child, set the bad bit of every child to the selector, put the old bad bit in child two's archive, the complete work tuple in child three's archive, and the old archive in child four's archive. Other archive slots are zero. These outputs retain all incoming seed, bad, work and archive labels even when the OR is irreversible. Distinct control inputs therefore have orthogonal output supports; within one control tuple, G or W is isometric on data. Work operations were permutations and clean preparations, so the full circuit is isometric on arbitrary inputs. Intermediate and radial dimensions increase only by fixed constants, since each archive stores one tuple, never an expanding concatenation. Deterministic basis preparation and updates prove control independence by induction, also with arbitrary entangled bulk inputs. □

Theorem 9 (Logarithmic-depth explicit template)**.** *On the explicit language $w \mapsto$ the circuit schedule of Lemmas 7–8, with fixed local dimension and gate alphabet, nearest-neighbor depth $O_w(\log(k+2))$, no seam, and the regions (1), $\mathcal{P}$ is Π_1^0-complete. Nonhalting gives $f_n = 1$ for all n. Halting gives $f_n = 1/49$ at every sufficiently late n, with one fixed target vertex and one fixed loss witnessing the upper bound.*

Proof. If $U(w)$ never halts all selectors are zero, so the data tree is all G and Lemma 2 supplies exact recovery. If it halts, choose a large fitting row whose time budget exceeds its halting time. Some vertex uses W, and its descendants inherit the bad flag. At that vertex erase child zero's entire descendant interval at any later cutoff. Data stay within their own subtrees. Angular work reads only deterministic classical control, so it cannot return the missing qudit. A Bell target and fixed spectators give fidelity at most $1/49$ after any decoder, as in Theorem 4. Maximally mixed replacement gives the matching lower bound for every query. Thus $f_n = 1/49$ thereafter. Bad descendants form a fixed positive fraction of every later row, with fraction depending on the initial bad vertex's depth.

The schedule is effective without testing halt. Every finite encoding is rational and isometric, so Lemma 1 enumerates failures and gives membership in Π_1^0 on this language. Nonhalting reduces with both implications just proved. The unbounded causal neighborhood of the work simulation is precisely the resource absent from Proposition 5. □

Proposition 10 (Growing dimension at depth one)**.** *There is an explicit computable gate schema, with one on-site radial round and no angular gates, whose exact recovery problem is Π_1^0-complete when radial dimension may grow with k.*

Proof. At level k the live classical label is the complete tape word on $[-k,k]$ for an input machine. Its radius-one update, padded by blanks at both ends, is a finite computable table. Copy the next live word to each child and retain the old live word in child four's archive. Pass the incoming archive into child three's archive by an injective padding map. Choose the row-k archive alphabet large enough for the row-k live-word alphabet, with a separate padding symbol. Then both stored old labels fit into the next, wider alphabet. Other archives are zero. Any off-trajectory transition is defined by a total finite extension. Different incoming live/archive labels remain distinguished by the stored copies, so the basis rule is an isometry. Apply G or W according to a halting head in the current word. The root contains the initial headed blank.

Induction gives the k-step configuration at every reachable site of row k. The controls are independent of data. Nonhalting gives all G; halting eventually gives W and a failed query. This is an effective reduction, and finite-query enumeration gives the upper bound. The growing matrix is part of a single specified computable schema; dimension, rather than angular communication, stores the unbounded tape. □

Arbitrarily delayed positive coherent information in quantum channels [45] is a related delayed-witness phenomenon; the following statement concerns the cutoff of subsystem-erasure failure.

Proposition 11 (No architecture-independent horizon)**.** *Within depth-one seam-free exact codes of arbitrarily large finite dimension, the cutoff of a failure can be arbitrarily late, despite Theorem 6.*

Proof. For any $T \geq 1$ attach a classical clock $c \pmod{T+1}$, initialized to zero. At a split copy $c+1$ to all children; apply G for $c \neq T$ and W for $c = T$. Distinct clocks give distinct outgoing words, so the map is isometric. Every cutoff $n \leq T$ is all G; cutoff $T+1$ contains a W row and fails. Each individual instance has fixed finite dimension and is decided by Theorem 6. □

Corollary 12 (Arbitrarily slow unbounded depth)**.** *For every fixed nonnegative integer function g which is computable, nondecreasing and unbounded, $\mathcal{P}$ is Π^0_1-complete on an explicit template with fixed local dimension and gate alphabet and $O_w(1+g(k))$ nearest-neighbor rounds. The late fidelity dichotomy of Theorem 9 holds as well.*

Proof. In Lemma 7 replace m by

$$\left\lfloor \sqrt{\min\{g(k), \lfloor \log_2(k+2) \rfloor\}} \right\rfloor .$$

It tends to infinity. The unchanged window and history proof costs $O_w(1+m^2) \leq O_w(1+g(k))$. The logarithmic cap ensures that sufficiently large rings fit. Each halting time is eventually exceeded, whereas a nonhalting input never sets a selector. Theorem 9's proof applies word for word to these explicitly generated rounds. Evaluating g is an offline scheduling computation with no imposed time bound. The allowance is asymptotic up to a program-dependent constant; no assertion of literally at most $g(k)$ gates on every small row is needed. □

5 Approximate recovery at fixed depth and dimension

5.1 The performance quantity

Use squared Uhlmann fidelity $F(\rho,\sigma)=\|\sqrt{\rho}\sqrt{\sigma}\|_1^2$. For a queried target and region define

$$f(V;v,R)=\max_{\mathcal{D}}\min_{\rho_{BHL}} F\left(((\mathcal{D}\mathcal{N})\otimes \mathrm{id}_L)(\rho),((\mathrm{Tr}_H)\otimes \mathrm{id}_L)(\rho)\right), \qquad \mathcal{N}(X)=\mathrm{Tr}_E(VXV^\dagger). \tag{17}$$

Here $\mathcal{D}$ is a channel from R to B. A reference of dimension $\dim(BH)$ suffices: purify the input, compress its reference by Schmidt decomposition, and use monotonicity of fidelity when discarding the extra purification. The relevant state and channel sets are compact, so the maximum and minimum exist. Put

$$f_n=\min_{v,j} f(V_n;v,R_n(v,j)), \qquad f_*=\inf_n f_n.$$

This is a worst-case subsystem quantity; Bény–Oreshkov [55] give an environment-dual characterization of the optimal worst-case entanglement-fidelity recovery, including subsystems. For the recorded construction below, the formula also holds for standard entanglement fidelity at the maximally mixed target with product spectator inputs. At fidelity one it is exactly (3). A late-cutoff failure means that every sufficiently late cutoff has a failed query, not that every vertex and loss fails.

For any encoding, the channel replacing B by I/q guarantees squared fidelity at least $1/q^2$. To see this, purify any desired target/reference state to $|\psi\rangle_{BLK}$. Its overlap with $I_B/q\otimes\rho_{LK}$ is $\mathrm{Tr}(\rho_B^2)/q\geq 1/q^2$. Discarding K can only increase fidelity. Thus for $q=7$ every instance has $f_*\geq 1/49$, and tolerances $\varepsilon\geq 48/49$ are automatic. Exact tolerance zero is Theorem 6.

5.2 Records in a pure isometry

A rational probabilistic finite automaton has m states, a deterministic start state, row-stochastic matrices M_j for at most five letters and an accepting set. Pad unused letters by duplicates. Let $p(v)$ be its acceptance probability after the child-digit word v, including the empty word at the root. The automaton transition is realized inside a pure homogeneous gate, with no random gate choice.

For every transition probability $a/b\geq 0$ with integers $a\geq 0, b>0$, choose integers $u_1,\dots,u_4$ with $\sum_r u_r^2=ab$. Such a choice exists by Lagrange's four-square theorem, Hardy–Wright, Theorem 369 [31], applied to the positive integer ab; zero uses four zeros. It is effective by finite enumeration. Dividing by b gives rational amplitudes of total squared norm a/b. This is the only arithmetic fact needed to replace square roots of rational probabilities by rational entries. Denote the resulting amplitudes by $c_j(s,t,r)$, $1\leq r\leq 4$; then $\sum_{t,r} c_j(s,t,r)^2=1$ for each s,j.

A radial register has data $a\in\mathbb{F}_7$ and controls (s,r,f,h) of dimensions $(m,4,2,8m)$. Here s is the active state, r is a four-square index, f is a branch bit, and h can store one old tuple (s,r,f). Fix padding labels. On a basis input sum over the five outcome pairs (t_j,k_j) with amplitude $\prod_j c_j(s,t_j,k_j)$. Give child j active state t_j, index k_j and branch bit $f_j=\mathrm{accept}(s)$. Put the entire old tuple (s,r,f) in child zero's h, the old h in child one's h, and padding in

the remaining three. Apply W to (a, b) when s accepts and G otherwise. All fields leave through radial children. Thus the transition choices are stored in the children's state/index pairs, while the old controls are retained in two archive destinations. This specifies one rational gate

$$T : \mathbb{C}^{448m^2} \otimes \mathbb{C}^7 \longrightarrow (\mathbb{C}^{448m^2})^{\otimes 5}, \qquad 448m^2 = 7 \cdot m \cdot 4 \cdot 2 \cdot (8m). \tag{18}$$

Lemma 13 (Isometry and data-independent probabilities)**.** *The gate* (18) *is isometric on its full input space. The probability of acceptance at a vertex v, when its control history is read in the record basis, is $p(v)$ for every normalized joint data input, including entangled inputs. More precisely its effect on earlier bulk inputs is $p(v)I$.*

Proof. Different complete incoming controls have disjoint output supports: child zero's archive retains (s, r, f) and child one's retains h. For one such input, different outcome tuples are distinguished by the five output state/index pairs. Their squared amplitudes sum to $\prod_j \sum_{t,k} c_j(s, t, k)^2 = 1$. Within each summand G or W is an isometry on both data inputs. This proves the full Gram identity, also for arbitrary superpositions of incoming controls. Copying the branch bit is defined on a basis as part of one isometry; it asserts no ability to clone unknown vectors.

After any finite expansion, complete outgoing controls allow the preceding control tuple and all sampled outcomes to be read back recursively. Distinct histories thus remain orthogonal on the full control output. On a fixed history the operation on the complete data input is a composition of isometries. Its squared norm on any normalized input is one, independently of that input and of unprocessed entangled bulk factors. Its probability is consequently the product of its transition weights. Marginalizing sibling outcomes yields exactly $M_j(s, t)$ along child j, so the distribution at v is the PFA distribution after v. For any pure input the acceptance probability is $p(v)$; polarization gives effect $p(v)I$. This proof does not presume product control/data states at intermediate vertices. □

Proposition 14 (Optimal recovery equals an automaton value)**.** *For the homogeneous depth-one template* (18)*, every $n \geq |v| + 1$ satisfies*

$$\begin{aligned} f(V_n; v, R_n(v, 0)) &= 1 - \frac{48}{49} p(v), \\ f(V_n; v, R_n(v, j)) &= 1 \quad (j \neq 0), \qquad f_* = 1 - \frac{48}{49} \sup_{v \in \{0,\ldots,4\}^*} p(v). \end{aligned} \tag{19}$$

Proof. Invert the full isometry on each intact child subtree on its image and discard its later bulk outputs. This gives its entering radial state exactly, with all correlations, and extends to a trace-preserving channel off the image. Read one surviving branch bit, which identifies the selected encoder at v. For G, apply Lemma 2's decoder. For W, retain b if port zero is present; otherwise replace the target by $I/7$.

For a child-zero loss, Lemma 13 makes instrument probabilities data independent. More explicitly, the outcome effect $p(v)I$ on earlier systems, followed by tracing them, leaves any other marginal multiplied by $p(v)$ even when entangled with them. The decoder therefore implements on the target/reference marginal σ_{BL} the channel

$$(1 - p(v))\sigma_{BL} + p(v)(I_B/7) \otimes \sigma_L. \tag{20}$$

The G inverse works for an arbitrary entering radial datum correlated with the fresh input. Purify σ_{BL}. Its pure overlap after replacement is at least $1/49$, as proved above, while the identity part has overlap one. Linearity for this pure ideal and monotonicity on discarding purification give fidelity at least $1-(48/49)p(v)$. For other losses the described decoder is exact.

For the opposite bound take a Bell target/reference pair and fix all spectators in pure product states. On losing port zero, the lost subtree has an orthogonal copy of the branch bit, so the partial trace kills G/W cross terms. This remains true at later cutoffs by invariance of partial trace under that subtree's isometry. In the W branch the reference is independent of everything retained, and after any decoder the target/reference state is $\tau_B \otimes I/7$. Its Bell overlap is $1/49$. In the G branch the overlap is at most one. Their weighted sum is at most $1-(48/49)p(v)$ for every decoder. This proves the optimum. Taking the infimum over cutoffs and vertices gives the last formula, since every finite word is a bulk address. □

5.3 The separated threshold problem

We state precisely the external undecidability input. **Import 1** is Rote, Theorem 1(c) [25]: strict emptiness for binary rational 18-state PFAs at cutpoint $1/18$ is undecidable. Its effective PCP reduction supplies c.e.-hardness of nonemptiness; finite words enumerate it, so emptiness is Π_1^0-complete. The explicit halting-to-PCP source and terminal-marker correction are specified in the proof of Theorem 18. The source start distribution may be rational rather than deterministic. **Import 2** is Rote, Theorem 7 and full-version Appendix A [25], extending Gimbert–Oualhadj [24]: given a d-state PFA and rational thresholds $\lambda_A \in [0,1]$, $\lambda_B \in (0,1)$, it produces a rational $(2d+3)$-state PFA with two extra letters and deterministic start, whose value is at most λ_B if all old probabilities are at most λ_A, and is one otherwise. These imports have no code-recovery conclusion.

Theorem 15 (Every rational $0<\varepsilon<48/49$)**.** *For each fixed rational $0<\varepsilon<48/49$, deciding $f_* \geq 1-\varepsilon$ is Π_1^0-complete on the template (18) with $m=39$ and the fixed radial dimension (2). Hardness holds under a promise separating $f_* \geq 1-\varepsilon+\eta$ from $f_* = 1/49$, for a positive rational η depending only on ε. No uniform algorithm approximates f_* to absolute error at most $12/49$ on this fixed-dimensional template. Exact recovery in the template is decidable.*

Proof. Apply Import 2 to Import 1's automata with $d=18$, $\lambda_A = 1/18$ and $\lambda_B = \lambda$, for rational $0<\lambda<49\varepsilon/48$. The output has 39 states and four letters, a deterministic start and one accepting state. Duplicate a letter for the fifth child. The output satisfies all hypotheses of Lemma 13; rational transition weights give rational gate amplitudes. Its dimensions are

$$d = 7\cdot 39\cdot 4\cdot 2\cdot(8\cdot 39) = 448\cdot 39^2 = 681408,$$

independent of the input automaton and of λ. If old emptiness holds, Proposition 14 gives $f_* \geq 1-(48/49)\lambda$. Otherwise it gives $f_* = 1/49$. Set $\eta = \varepsilon-(48/49)\lambda > 0$. Both cases are separated from the requested threshold. The PCP-based many-one reduction in Import 1 and this effective construction give Π_1^0-hardness, not only undecidability.

For a uniform value obstruction, $\lambda = 1/2$ already separates $f_* \geq 25/49$ from $f_* = 1/49$ and excludes error strictly less than $12/49$. To include the endpoint in the error bound, use

$\lambda = 1/3$ instead: the cases are $f_* \geq 33/49$ and $f_* = 1/49$. Their separation is $32/49 > 24/49$, so error at most $12/49$ would distinguish them by, for example, comparison with $17/49$. This is a uniform impossibility; it makes no assertion that every individual value is a noncomputable real.

For membership, enumerate finite queries and decide their threshold comparisons. Over algebraic gate data, parameterize a decoder by its Choi matrix, positive semidefiniteness and trace constraints. Quantify over input purifications of bounded dimension. Squared fidelity is real semialgebraic: introduce positive semidefinite variables for the square roots in F, with polynomial equations $X^2 = A$, and take the squared trace of the outer square root. Separate real and imaginary parts. Quantifier elimination over real closed fields, Basu–Pollack–Roy, Theorem 2.77 [32], decides the finite sentence, including threshold equality: all coefficients are real algebraic and the constraints are polynomial equalities or inequalities, exactly its hypotheses. Strict failure at one query semidecides the complement of $f_* \geq 1 - \varepsilon$, proving membership in Π_1^0. On this template the simpler rational formula (19) also supplies that enumeration directly.

At zero tolerance (19) says all finite words have $p(v) = 0$. With nonnegative transitions this is exactly absence of a positive-probability path from the start to an accepting state. Finite graph reachability decides it. Independently the template lies inside Theorem 6's class. Both arguments establish exact decidability. □

The PFA gap technique is standard, with earlier strict-cutpoint and robust reductions by Nasu–Honda, Paz, and Condon–Lipton [21, 22, 23]. Blondel–Tsitsiklis [54] also transfer PFA emptiness to rational matrix-product boundedness. The application here appears new through the exact optimal-fidelity identity (19): it is a fixed-dimensional recovery-fidelity realization of a known probabilistic-automaton gap. Douglas–Mghirbi [61] study the separate question of the exact spectrum of rational PFA values. Theorem 15 proves a uniform approximation obstruction; it does not classify individual values on this fixed-dimensional recovery template.

For measure-once automata with algebraic *unitary* letters and a final projection, strict-cutpoint emptiness is decidable by Blondel et al., Theorem 3.2 [27], and effective algebraic-group methods are given by Derksen et al., Theorem 1 [28]. Non-strict finite-word tests have a different status. The addressed port matrices in Section 6 are contractions, so that unitary hypothesis fails. Section 7 makes the input-control distinction precise: an inert passive excitation has an algebraic value, whereas fresh inputs steering dissipative memory along one ray give a robust undecidable gap without requiring broadcast to the tested footprints. Whether dissipation is necessary for these codes remains open.

6 Single-excitation propagation and the exact root threshold

We now construct a homogeneous rational depth-one tree code whose fidelity infimum is uniformly computable but whose comparison at one exact threshold is Π_1^0-complete. The continuing amplitude at each spatial port is transformed by one matrix. Internal alternatives reaching that port add as amplitudes. The cell nevertheless forwards a spent quantum probe and its old forwarding wire. In particular, its no-arrival probe state is two-dimensional and

records the logical sector in nonorthogonal polarizations. These features are part of the construction, not side information omitted from the decoder experiment.

6.1 Rational scattering and a polarized no-arrival state

Use the binary 18-state rational PFAs of Import 1 in Section 5. Transpose their matrices into column convention and denote them T_0, T_1. Let π be the rational initial distribution, e the accepting coordinate and $\lambda = 1/18$. For $w = i_1 \cdots i_k$ set

$$\pi_w = T_{i_k} \cdots T_{i_1} \pi, \qquad p(w) = e^T \pi_w. \tag{21}$$

Thus π_w is nonnegative with coordinate sum one. Its entries will now be amplitudes, not sampled transition probabilities. In a 21-dimensional active space X, define rational data

$$\begin{gathered} y = (\pi, 6\lambda/5, 2\lambda/5), \quad N = y^T y, \quad s = \frac{2}{1+N}, \quad t = \frac{1-N}{1+N}, \\ \xi = (sy, t), \qquad D_i = \operatorname{diag}(T_i, 1, 1, 1). \end{gathered} \tag{22}$$

The two extra coordinates of y are markers, and the last coordinate of ξ is neutral. Since $s^2 N + t^2 = 1$, ξ is a rational unit vector. Fix $\alpha = 1/16$, $\beta = 1/2$, and $h_0 = 3/5$.

Lemma 16 (A full-domain single-excitation cell)**.** *There is an effective homogeneous rational depth-one isometry on*

$$D = a \otimes A \otimes E \otimes J, \qquad (\dim a, \dim A, \dim E, \dim J) = (7, 22, 44, 85), \qquad \dim D = 575960, \tag{23}$$

with a fixed basis root cap and a fresh bulk input $B = \mathbb{C}^7$, having the following specified operations.

1. *$A = \mathbb{C}|\mathrm{vac}\rangle \oplus X$ scatters its occupied component into continuation ports zero and one with amplitude matrices $\alpha D_0, \alpha D_1$, a terminal port four with amplitude matrix βI, and a residual port three with a rational matrix M of at most 84 rows. Every other scattering output in a given occupied component is in its fixed vacuum.*
2. *E consists of 22 real rotation planes, complexified as Hilbert spaces: a two-dimensional no-arrival plane and 21 active planes. Initially only direction zero in each plane is populated. For logical sectors P_+, P_-, of ranks three and four, the cell applies I or U to the existing E, respectively. The block of U on a plane is*

$$R(c, d) = \begin{pmatrix} c & -d \\ d & c \end{pmatrix}. \tag{24}$$

   The accepting plane uses $(c, d) = (7/25, 24/25)$, the two marker planes use $(4/5, 3/5)$, and all other planes (the 17 non-accepting state planes, the neutral plane and the no-arrival plane) use $(3/5, 4/5)$.
3. *The whole rotated incoming E is embedded into child zero's J. The incoming J is forwarded intact to child one's J. The scatterer's residual occupies child three's J, and its terminal occupies direction zero in the active planes of child four's E. All unused factors are fixed. After the controlled rotation, the same G of Lemma 2 encodes incoming a and fresh B into the five data factors.*

The cell is isometric for arbitrary joint inputs on all the declared factors, including incompatible controls and external references.

Proof. Each column of T_i has Euclidean norm at most one, so $\|T_i\|^2 \le \|T_i\|_F^2 \le 18$, and $\|D_i\|^2 \le 18$. Hence

$$Q_0 = I - \beta^2 I - \alpha^2(D_0^T D_0 + D_1^T D_1) \succeq (3/4 - 36/256)I \succ 0. \tag{25}$$

Rational LDL^T elimination gives $Q_0 = L_0 \Lambda L_0^T$ with positive rational diagonal pivots. For a pivot a/b, Lagrange's four-square theorem, Hardy–Wright, Theorem 369 [31], gives four integers whose squares sum to ab. Dividing by b gives four rational numbers with square sum a/b. Multiply these by the corresponding row of L_0^T and stack the rows. The resulting rational matrix M has at most $4 \cdot 21 = 84$ rows and $M^T M = Q_0$. Enumeration of the four-square witnesses terminates.

For $x \in X$, send x to the sum of its four occupied spatial components with matrices $\alpha D_0, \alpha D_1, \beta I, M$. The ranges are mutually orthogonal by spatial occupation; their squared norms sum to $\|x\|^2$ by (25). Send the incoming vacuum to the joint vacuum, orthogonal to every occupied output. This proves an isometry on the whole A input, including superpositions with vacuum. The residual uses the 84 directions orthogonal to the padding vector of J, and the terminal uses the 21 direction-zero active vectors in E. On vacuum, those destinations are their padding and no-arrival vectors.

All three rotation pairs in the statement have $c^2 + d^2 = 1$. Controlled rotation on $B \otimes E$ is therefore unitary. Its output E has an isometric embedding into child zero's 85-dimensional J; the old 85-dimensional J goes to a separate child-one factor. None of these slots is used by the scatterer. Compose these disjoint embeddings, the scatterer, and G after the controlled rotation. Every component is isometric on its full input space, proving the claim without a restriction to reachable states. A consumed probe never reenters A or E at a later refinement.

Initialize old data and J to their fixed padding, E to $|\mathrm{vac}, 0\rangle$, and A to the occupied vector ξ. To express this as one fixed basis cap, choose a rational orthogonal matrix taking a designated occupied basis vector to ξ. The Householder reflection in their difference does so unless they coincide, in which case use the identity. Identify physical input active coordinates by this matrix and continuing output coordinates by its inverse; apply the same identification in every cell. Terminal and residual embeddings use the corresponding input-coordinate change. All entries remain rational and all dimensions unchanged. Number the designated occupied basis vector as the root's zero vector. This gives the fixed basis-cap convention in the theorem. □

The no-arrival plane must be explicit. Its two logical settings give

$$|\mathrm{vac}, 0\rangle, \qquad U|\mathrm{vac}, 0\rangle = \frac{3}{5}|\mathrm{vac}, 0\rangle + \frac{4}{5}|\mathrm{vac}, 1\rangle. \tag{26}$$

Both vectors lie in the same no-arrival plane, with overlap 3/5. This fixed nonorthogonal quantum record of the logical sector is forwarded on J. More generally $(U + U^\dagger)/2 \succeq (7/25)I_E$, so for any carrier state, even with a purification, the two probe outputs have overlap $\mathrm{Tr}(\rho_E U)$ with real part at least 7/25. No perfectly distinguishable setting record is created. This statement does not assert absence of a quantum record. Likewise, forwarding the old J preserves an intact quantum wire, rather than copying a classical transition label.

6.2 What the receiver can recover

Lemma 17 (Unconditional carrier state and optimal witnesses)**.** *For the initialized cell of Lemma 16, a nonvacuum terminal can reach only a vertex $v = w4$, with $w \in \{0,1\}^*$, where its direction-zero marginal is*

$$\rho_E = (1 - \|x_w\|^2)|\mathrm{vac}, 0\rangle\langle \mathrm{vac}, 0| + |x_w\rangle\langle x_w|, \quad x_w = \beta\alpha^{|w|}(s\pi_w, 6s\lambda/5, 2s\lambda/5, t). \tag{27}$$

At all other vertices the incoming E is $|\mathrm{vac}, 0\rangle$. Set $z_v = \mathrm{Tr}(\rho_E U)$ and $p_v = (1 + z_v)/2$. For every query at every cutoff, $f(V_n; v, R_n(v, j)) \geq p_v$. At birth, child-zero loss attains this bound at $v = w4$ and at the root. In particular, the root value is $f_0 = 4/5$, and

$$z_{w4} = \frac{3}{5} + \frac{8}{25}s^2\beta^2\alpha^{2|w|}(\lambda^2 - p(w)^2), \qquad z_v = \frac{3}{5} \quad (v \neq w4). \tag{28}$$

Here f_0 denotes the root query's value, not a cutoff numbered zero.

Proof. Conditional on one occupied continuation, its amplitude acquires exactly αD_i and every other scattering output is vacuum. Induction along a binary address gives the vector in (27) when it exits through port four. The alternative has the excitation in orthogonal other spatial modes, so tracing those modes leaves the stated vacuum/live mixture. No excitation can return from J. All other incoming terminal factors are fixed no-arrival vectors. These are unconditional marginals: there is no division by an arrival probability.

Only direction zero in each plane is populated before its pulse, so z_v is real. Relative to the neutral cosine $3/5$, acceptance has deviation $-8/25$ and each marker has deviation $1/5$. The sum of squared marker amplitudes before the common scale is $((6/5)^2 + (2/5)^2)\lambda^2 = (8/5)\lambda^2$. The acceptance squared amplitude is $p(w)^2$. Therefore

$$z_{w4} = \frac{3}{5} + s^2\beta^2\alpha^{2|w|}\left[-\frac{8}{25}p(w)^2 + \frac{1}{5}\frac{8}{5}\lambda^2\right],$$

which proves (28). It also shows $z_v > 0$; alternatively the positive Hermitian part of U gives $z_v \geq 7/25$.

For a lower decoder at a later cutoff invert the four full surviving subtree isometries on their images and discard later bulk factors. Such inverses extend to channels off their images, and inversion preserves correlations between subtrees. Tracing a full lost subtree removes its isometry and traces its input. We may consequently work with the four birth shares. Invert G on those data shares and discard its old logical a and all control wires.

The terminal E before the pulse has joint marginal $\rho_E \otimes \sigma_{BL}$, even with entangled spectators. To verify this, commute every earlier data operation past all later continuing scatterers: it acts only on data, bulk inputs and an already detached terminal wire. Initially the passive network is independent of all bulk/reference inputs. Once the earlier detached wires, earlier data and spectator bulk inputs are traced, trace preservation on these systems leaves the same EBL marginal as omitting their operations. This remains true when the target was entangled with those spectators. The inverse of G returns B with its pulse correlations even when the old a and E are correlated; discarding a leaves this same marginal.

Writing $Z = P_+ - P_-$, tracing the probe preserves the two diagonal sector blocks and multiplies off-diagonal blocks by the real z_v. Thus the decoder implements

$$p_v\sigma_{BL} + (1 - p_v)(Z \otimes I_L)\sigma_{BL}(Z \otimes I_L). \tag{29}$$

On a purification, its ideal pure overlap is at least p_v. Discarding the extra purification cannot decrease fidelity. This proves the lower bound simultaneously for all losses, cutoffs and input states.

For the upper bound at the birth of $v = w4$, fix spectators in pure product states and entangle the target with a two-level reference on one basis vector from each sector. The incoming old data a and terminal E are independent in this experiment. Indeed every strict ancestor on the addressed path is the root or has a binary address and has fixed terminal vacuum, even when its continuing A is occupied. Its data pulse therefore acts independently of that continuing A. The old J is routed and never read by data or continuing scatterers. Tracing off-path outputs and inducting down the binary path proves the asserted independence. At $v = w4$ the incoming A is vacuum and incoming J is padding, because its final digit is four. On losing child zero the entire spent E is inaccessible, and every surviving control is a fixed ancillary state. This accounts for all available quantum side information.

The retained data arise by sending the dephased target and an independent old datum through G and losing one share. Any decoder therefore induces a channel $\mathcal{A}$ on the dephased target. Denote the two relative-sign Bell projectors on the chosen sector pair by Φ_+, Φ_-. Put

$$q_\pm = \mathrm{Tr}[\Phi_+(\mathcal{A} \otimes \mathrm{id})(\Phi_\pm)].$$

These are nonnegative and $q_+ + q_- \le 1$: their input sum is $|00\rangle\langle 00| + |11\rangle\langle 11|$ on the chosen pair, and its overlap after the channel is half the sum of two diagonal output probabilities, each at most one. Since $p_v \ge 1/2$, the recovered Bell overlap is at most $p_v q_+ + (1-p_v)q_- \le p_v$. This matches the lower bound. At the root the incoming E is vacuum, and the other scattered outputs, though mutually entangled, are jointly independent of the root data and its pulse. Surviving controls add only a fixed ancillary state. The same upper bound therefore holds at the root, where $p_v = (1 + 3/5)/2 = 4/5$. □

Theorem 18 (Single-excitation exact root threshold)**.** *On the explicit single-excitation depth-one tree-code template of Lemma 16, with one amplitude matrix per scattering port, rational entries, fixed basis cap and radial dimension* 575960*, deciding whether the optimal all-scale recovery fidelity equals its root value* $4/5$ *is* Π^0_1*-complete. Equivalently,* $f_* \ge 4/5$ *is complete on this template. Nevertheless* f_* *is uniformly computable from the finite instance. For every fixed rational* $c \ne 4/5$*, deciding* $f_* \ge c$ *is decidable, uniformly when* c *is also an input promised different from* $4/5$*. There is no positive-gap hardness conclusion for this template.*

Proof. Lemma 17's lower bound applies to every query, and its birth queries attain each bound needed for the infimum. Consequently, with $C = (4/25)s^2\beta^2 > 0$,

$$f_* = \min\left\{\frac{4}{5},\ \inf_{w\in\{0,1\}^*}\left[\frac{4}{5} + C\alpha^{2|w|}(\lambda^2 - p(w)^2)\right]\right\}. \tag{30}$$

This formula concerns the actual optimized worst-case subsystem fidelity. Its terminal-state calculation is unconditional and its upper witnesses bound every decoder.

All $p(w)$ are nonnegative and $\lambda > 0$. Thus the term at w is below $4/5$ exactly when $p(w) > \lambda$; its prefactor is strictly positive at every finite word. The root query ensures $f_* \le 4/5$. Nonemptiness for Import 1 therefore reduces to strict failure of $f_* = 4/5$. To specify the many-one degree, Neary's Lemma 9 and Theorem 11 [37] effectively encode halting

by a binary tag system and a five-pair PCP instance. Use the corrected terminal marker in Rote, Section 2.8, footnote 3, followed by Theorem 1(c), Section 3.2 [25]. The mandatory initial domino is incorporated into the rational initial distribution. These finite rational binary 18-state automata meet (21)'s hypotheses, so the code is an effective reduction from nonhalting to equality at the root. Conversely, enumerate binary words and their rational $p(w)$ to enumerate a strict failure. This gives Π_1^0 membership and completeness. The finite algebraic-query argument in Theorem 15 gives the same upper bound without using the closed formula.

For computability, let U_m be the minimum of $4/5$ and the finitely many bracketed values in (30) for $|w| < m$. Since $p(w), \lambda \in [0, 1]$, each omitted value differs from $4/5$ by at most $C\alpha^{2m}$. Hence

$$U_m - C\alpha^{2m} \leq f_* \leq U_m. \tag{31}$$

The constants are known rationals and $\alpha = 1/16$. Given any positive rational precision, increase m until $C\alpha^{2m}$ is smaller; the finite minimum is exactly computable. This is a uniform terminating value algorithm, including instances with $f_* = 4/5$.

If $c > 4/5$, the root query decides the threshold negatively. If $c < 4/5$, choose m with $C\alpha^{2m} < 4/5 - c$. Every word of length at least m then has its value strictly above c. The comparison is equivalent to $U_m \geq c$, a finite rational calculation including equality. The same procedure also decides equality to or strict comparison with a rational $c \neq 4/5$: the tail is strictly above c, so a value at or below it is attained among finitely many words. The undecidable case is therefore the exact root threshold, despite uniform value computability. No uniform positive deficit is established by a witness; its known amplitude factor decays geometrically. □

6.3 The role of the no-arrival plane and the known mechanism

The two-dimensional no-arrival plane in (26) is substantive. In the analogous template with a one-state no-arrival space, a unitary probe acts there by a scalar phase. The controlled scalar phase on B can be undone on the recovered target, so the root value is one. Equality of the all-scale value to that root value is then exact recovery, decidable by Theorem 6. A one-state vacuum therefore cannot support the asserted $4/5$ root-threshold theorem. The construction above rotates and forwards its no-arrival polarization at every site and retains old J inputs intact. It has one continuing amplitude matrix per port but is not described as devoid of quantum records.

The mechanism of representing stochastic languages by linear amplitudes and embedding letter contractions into quantum evolution is established. Turakainen [34, 56] relates generalized automata to stochastic languages. Yakaryilmaz–Say, CSR 2009 [35], characterizes the unbounded-error languages of one-way Kondacs–Watrous automata as exactly the stochastic languages (Theorem 3 of its preprint version); the underlying quantum-automaton model is due to Kondacs–Watrous [36]. These are precedents for the mechanism, not imports asserting the recovery theorem. Here the rational completion (25), the no-arrival polarization, the all-spectator lower decoder and the Bell upper witness prove the required code statement. Its application to this recovery experiment appears new.

There is no conflict with the strict-cutpoint decision for measure-once unitary letters [27,

Theorem 3.2]. The addressed matrices αD_i are lossy contractions with internal mixing. An isometric dilation of the full scatterer does not make a single occupied-port submatrix unitary. The attenuation remains in the unconditioned fidelity as $\alpha^{2|w|}$, and its effective tail is exactly what makes (30) computable. Section 5 gives a robust gap and an infimum that is not uniformly approximable within 12/49. Section 7 places these facts in the input-control classification: inert passive single excitations have algebraic values, while worst-case fresh inputs steering dissipative memory along one ray yield a robust undecidable gap. Broadcast to the tested footprints is unnecessary; necessity of dissipation remains open.

7 Input control and recovery values

The relevant distinction is whether fresh bulk inputs can steer the continuing excitation. We prove an algebraic value theorem for a single excitation whose word is determined by its spatial address, and an undecidable separated comparison for a dissipative memory programmed by fresh inputs along one ray. The latter needs no unbounded distribution of a fixed source channel to disjoint recovery footprints. The passive arrival-tail bound below follows from a Gram identity, rather than from no-cloning.

7.1 A passive excitation and its finite recovery channel

Let X be a finite-dimensional occupied space, $A = \mathbb{C}|\emptyset\rangle \oplus X$, and let E contain a unit no-arrival vector e_0 orthogonal to an active terminal space. All matrices and cap vectors have rational real and imaginary parts. On the occupied space specify continuation matrices A_0, A_1, A_2, a residual matrix M, and a terminal matrix T, with

$$\sum_{i=0}^{2} A_i^\dagger A_i + M^\dagger M + T^\dagger T = I_X. \tag{32}$$

The splitter sends an occupied vector into orthogonal modes: $A_i u$ in child i's live register, Mu in child three's archive, and Tu in child four's active E register. It sends vacuum to joint vacuum. Missing modes are zero. At a site apply a specified rational unitary Q to fresh $B = \mathbb{C}^7$ and incoming E, forward the entire spent E into child zero's archive and incoming archive into child one's, using disjoint slots, and encode old datum a and postinteraction B with G of Lemma 2. Archive dimensions accommodate these basis embeddings and the residual mode; archives never control later operations. The cap has occupied unit vector ξ and incoming $E = e_0$. The Gram identity, orthogonal vacuum and disjoint forwarding make the cell isometric on its entire input space. In particular state-independent routing is not assumed.

For $w = i_1 \cdots i_k \in \{0,1,2\}^*$ set $A_w = A_{i_k} \cdots A_{i_1}$, $y_w = A_w \xi$, $p_w = \|T y_w\|^2$, and

$$\rho_w = (1 - p_w)|e_0\rangle\langle e_0| + |Ty_w\rangle\langle Ty_w|, \qquad N_\rho(b) = \mathrm{Tr}_E[Q(b \otimes \rho)Q^\dagger]. \tag{33}$$

Let $g(\rho)$ be the optimal worst-case entanglement fidelity of N_ρ after a channel $B \to B$, with a seven-dimensional reference, and put $b_0 = g(|e_0\rangle\langle e_0|)$.

Lemma 19 (Finite reduction for a passive excitation)**.** *For the cells just specified, with the recovery experiment* (17)*,*

$$f_* = \inf\Big(\{b_0\} \cup \{g(\rho_w) : w \in \{0,1,2\}^*\}\Big), \qquad |g(\rho_w) - b_0| \le p_w. \tag{34}$$

Each displayed finite value is an effectively specified real algebraic number.

Proof. Tracing every other excitation mode removes its cross terms with a chosen terminal, so the unconditional incoming state at $w4$ is exactly (33). All other incoming probes are vacuum; a spent probe cannot feed back. The excitation dynamics do not depend on bulk inputs. Invert each surviving child's full descendant isometry on its range, extending the inverse off the range by a fixed replacement state. Four recovered data shares return the postinteraction B, including its correlations with arbitrary spectators, by Lemma 2. Discarding controls and applying an optimizer for g therefore gives the corresponding lower bound at every cutoff and every loss. Pure-state tests with a seven-dimensional reference suffice: purify, Schmidt-compress, and discard the added purification using fidelity monotonicity.

For the converse use the birth cutoff at $w4$, erase child zero, and fix spectators to pure product states. Strict prefixes of w contain only continuation digits, so all their incoming probes are fixed vacuum and their data operations are independent of the splitter. At $w4$ the incoming live register is vacuum; its entire spent probe leaves through the erased child. The incoming archive is forwarded to child one, but is independent of the terminal probe and fresh target under this product-spectator witness: along continuation-only prefixes it can contain only earlier vacuum probes and the fixed cap padding. Every other surviving control is likewise target-independent. Conditional on the reduced probe, the retained data are obtained by a channel from $N_{\rho_w}(B)$ and an independent old datum. Hence any boundary decoder induces a channel $B \to B$, and cannot exceed $g(\rho_w)$. The same birth witness at the root gives b_0. This proves the infimum identity without discarding any side information in the upper bound.

For fixed decoder and pure test state the score lies in $[0,1]$ and is affine in ρ. Writing $\rho_w = (1-p_w)\rho_0 + p_w\sigma_w$ bounds its change from the vacuum score by p_w, uniformly in both optimizing variables. Taking their extrema preserves that bound. Finally use the decoder's unnormalized Choi matrix, its positive semidefinite and trace constraints, and unit-sphere coordinates for the pure input. The overlap with the pure ideal is polynomial in these real and imaginary coordinates and in the algebraic probe entries. The compact sentence $\exists J_{\mathcal{D}}\ \forall\psi:\ s(J_{\mathcal{D}},\psi,\rho) \ge t$ defines exactly $t \le g(\rho)$. Tarski quantifier elimination [64], in the effective form of Basu–Pollack–Roy, Theorems 2.77 and 14.16 [32], applies: there are finitely many variables, polynomial constraints and effective real-algebraic coefficients. Quantifier elimination and real-root isolation compute the unique endpoint as algebraic data. □

Theorem 20 (Algebraic value with inert no-arrival state)**.** *For every rational passive single-excitation cell above satisfying*

$$Q(b \otimes e_0) = (Vb) \otimes e_0 \quad \textit{for all } b \in B \tag{35}$$

for a fixed unitary V, one algorithm outputs the exact real-algebraic number f_* and decides every comparison $f_* \ge c$, $c \in \mathbb{Q}$. The hypotheses are finite matrix tests. No fresh bulk input steers the excitation, and continuation matrices may depend on its state.

Proof. Put $D_T = T^\dagger T$, $\Phi(Z) = \sum_i A_i^\dagger Z A_i$ and $S_n = \sum_{k=0}^{n-1} \Phi^k(D_T)$. Equation (32) gives $0 \le D_T \le I - \Phi(I)$ and hence

$$0 \le S_n \le S_{n+1} \le I, \qquad S_n \le I - \Phi^n(I).$$

The increasing traces converge. For $m \ge n$, $\|S_m - S_n\| \le \mathrm{Tr}(S_m - S_n)$, so S_n converges in norm to H and $\Phi^n(D_T) \to 0$. Compute the cyclic real space $K = \mathrm{span}\{D_T, \Phi D_T, \dots\}$ by exact linear algebra, stopping at its earliest dependency; its dimension is at most $(\dim X)^2$. This space is Φ-invariant. Every basis vector tends to zero under iteration, so $I - \Phi$ is injective on K and is invertible there. Solve $(I - \Phi)H = D_T$ on K (use $H = 0$ if $K = 0$). This computes H exactly and avoids inversion on persistent dark subspaces in the full matrix space.

Expanding the word products gives the computable tail

$$t_n = \sum_{|w| \ge n} p_w = \xi^\dagger (H - S_n)\xi = \xi^\dagger \Phi^n(H)\xi \downarrow 0. \tag{36}$$

This summability follows from the positive Gram budget. It holds even when some $\|A_i\| = 1$. For $m_n = \min(b_0, \{g(\rho_w) : |w| < n\})$, Lemma 19 gives

$$\min(m_n, b_0 - t_n) \le f_* \le m_n. \tag{37}$$

The width is at most t_n, so there is a uniform approximation algorithm. This is the standard finite-dimensional reachability idea; compare Ying–Feng–Yu–Ying, Theorem 8 [57]. The derivation here supplies the required terminal effect directly.

It remains to decide equality at $b_0 = 1$, which follows from (35). A channel $N : B \to B$ has corrected worst-case fidelity one exactly when it is unitary conjugation. Indeed a score one on a maximally entangled vector implies $\mathcal{R}N = \mathrm{id}$ by Choi injectivity. For Kraus operators L_j, K_i of $\mathcal{R}, N$, each $L_j K_i$ is a scalar multiple of I, since the identity Choi support has rank one. Some scalar is nonzero. That L_j is invertible, so all K_i are proportional to its inverse; trace preservation makes the resulting single normalized Kraus matrix unitary. The converse uses its inverse. Equivalently the Choi matrix of N has rank one.

Realify $y = A_w \xi$. The entries of ρ_w are polynomials of degree at most two in y, so every real or imaginary part of a two-by-two minor of the Choi matrix of N_{ρ_w} has degree at most four. Vanishing of all these minors means rank at most one, and trace preservation excludes zero. Form the finite vector $h(y)$ of real monomials of degrees zero through four. Substitution by each realified A_i is an explicit linear map on h. Close the span of $h(\xi)$ under these maps by adjoining independent images until its dimension stops growing. Each minor vanishes at every word exactly when its linear functional vanishes on this finite basis. Include the empty word. This decides whether every $g(\rho_w) = 1$, hence whether $f_* = 1$.

If the test succeeds output one. Otherwise some finite value is below one; enumerate n until $1 - t_n > m_n$. The search terminates by the strict deficit and (36), and (37) then proves $f_* = m_n$. Output its algebraic representation and compare it exactly with any rational threshold. Zero ports, zero arrivals and dark subspaces are covered by the same algorithm. Polynomial-orbit identities and matrix-semigroup closures are standard tools; see Hrushovski–Ouaknine–Pouly–Worrell [58]. Their general closure algorithm is not needed for this bounded-degree test. The application to this recovery experiment appears new. □

7.2 A polarized readout with unitary continuation

Theorem 21 (Unitary continuation and two sectors)**.** *In the passive class suppose each nonzero continuation matrix is $A_i = \sqrt{q_i}U_i$, $q_i > 0$, with unitary U_i. Let $B = B_+ \oplus B_-$ have two nonzero sectors, $Q = P_+ \otimes I + P_- \otimes R$, and $h = \langle e_0, Re_0\rangle \in [0,1]$ be real. Assume the compression of the unitary R to $\operatorname{im} T$ is Hermitian positive semidefinite. For rational input data, f_* again has a uniformly computable exact algebraic description, and every rational threshold is decidable.*

Proof. Put $D = T^\dagger T$, $C = T^\dagger R T \succeq 0$. The opposite-sector coherence is multiplied by

$$z_w = \operatorname{Tr}(\rho_w R) = h + y_w^\dagger (C - hD) y_w \in [0,1].$$

Thus $N_{\rho_w} = p\,\mathrm{id} + (1-p)Z(\cdot)Z$, where $Z = P_+ - P_-$ and $p = (1+z_w)/2$. The identity decoder guarantees fidelity at least p by purification. For the upper bound choose one vector in each sector and their equal Bell superposition with a two-dimensional reference. For any decoder with Kraus L_j, write $a_j = \langle +|L_j|+\rangle$, $b_j = \langle -|L_j|-\rangle$. Its Bell overlap is

$$\frac{1}{4}\sum_j (|a_j|^2 + |b_j|^2 + 2z_w \Re(a_j \bar{b}_j)) \le \frac{1+z_w}{2},$$

since trace preservation bounds each of $\sum |a_j|^2, \sum |b_j|^2$ by one and Cauchy–Schwarz bounds the cross sum. Hence $g(\rho_w) = (1+z_w)/2$ and $b_0 = (1+h)/2$.

For words without a zero arm write $A_w = a_w U_w$, $a_w > 0$. Equality $f_* = b_0$ is equivalent to nonnegativity of $(U_w\xi)^\dagger (C - hD)(U_w\xi)$ for every such word; zero arms give b_0. Realify the finitely many algebraic unitary U_i into orthogonal matrices. Derksen–Jeandel–Koiran, Lemma 2, Algorithm 1 and Theorem 9 [28], compute the group Zariski closure: the generators are invertible square matrices over an effective number field. Their Section 2.2 and Blondel et al., Theorem 3.1 and the algebraic extension following Theorem 3.2 [27], identify its real orthogonal locus with the compact Euclidean monoid closure. Real points, rather than the complex closure alone, are the optimization domain. Continuity and density turn the sign test into nonnegativity of a quadratic form on that computed compact real algebraic set. Quantifier elimination decides it. Related effective closure bounds are given by Nosan et al. [59].

If the sign test succeeds output b_0; otherwise use the tail and finite-minimum stopping test in (37) to output the attained algebraic value below b_0. This proves the claimed uniform algorithm. It does not replace lossy addressed contractions by unitary matrices: normalization is used only in a sign test, where its strictly positive scalar cancels. □

7.3 Fresh inputs steering a dissipative memory on one ray

The next construction keeps full memory weight on the zero-child ray. Dissipation refers to its reduced memory dynamics after detached emissions are traced; the entire cell is isometric. Its instruction word is chosen by fresh bulk inputs in the worst-case quantifier of (17). It uses a known finite-state memory mechanism [26]; the realization as optimal subsystem recovery requires the following explicit construction. Classical-state quantum automata provide related finite-control models [60].

Lemma 22 (Lazy instructions and a rational memory dilation)**.** *From the rational 39-state, four-letter PFA of Import 2 with value v, one can effectively construct five rational stochastic instructions on 78 states with the same supremum v, and rational isometries $V_i : \mathbb{C}^{78} \to \mathbb{C}^{78} \otimes \mathbb{C}^{391}$ whose reduced memory populations follow those instructions for every input matrix. The emitted register has a common identity component of amplitude* $6/7$*.*

Proof. Use banks zero and one of the original states, accepting the original accepting state in either bank. For $i = 1, \ldots, 4$, let E_i transfer bank zero to bank one applying the original stochastic matrix P_i, leaving bank one fixed. Let E_5 reset bank one to bank zero with unchanged label, leaving bank zero fixed. Each E_i is stochastic and idempotent. Set $S_i = (7/8)I + (1/8)E_i$. Expanding any instruction product as a convex combination of I, E_i products gives acceptance at most v: its bank trajectory is independent of the internal state, and its effective original letters form a finite PFA word. Conversely

$$S_i^r = (7/8)^r I + (1 - (7/8)^r)E_i.$$

Alternate each desired original letter with a reset and repeat each instruction r times. For any fixed finite word the resulting product tends to its exact bank-transfer product as $r \to \infty$. The empty word is included. Thus the suprema agree; no witness length is assumed computable.

Put $c = 6/7$. Use one environment direction $|*\rangle$ and 78 orthogonal five-dimensional source blocks. For each letter choose rational vectors

$$e_{st} = c\,\mathbf{1}_{s=t}|*\rangle + \eta_{st}, \qquad \eta_{st} \in \text{source block } s, \qquad \|\eta_{st}\|^2 = S_i(s,t) - c^2\mathbf{1}_{s=t}.$$

Every positive remainder is strictly positive, since $S_i(s,s) \geq 7/8 > 36/49$. Choose a common positive rational coordinate smaller in square than all positive remainders, and complete each remaining norm by four nonnegative rational squares, using Hardy–Wright, Theorem 369 [31], on ab for a rational $a/b \geq 0$. Zero edges use zero vectors. Finite enumeration supplies all amplitudes. Set $V_i|s\rangle = \sum_t |t\rangle e_{st}$.

Each column has norm one. Distinct source columns are orthogonal: their source blocks are disjoint, and a common destination cannot have a star term from both. This proves full-domain isometry. The same calculation shows that the diagonal output blocks on any input, including operator-valued spectator entries, are $\sum_s S_i(s,t)X_{ss}$. Off-diagonal input units $|s\rangle\langle u|$, $s \neq u$, reduce to $c^2|s\rangle\langle u|$. Within one source block distinct feasible successors have positive overlap through the chosen common coordinate; the dilation need not leave perfectly distinguishable successor records. The star component is exactly cI on memory, and $1 + 78 \cdot 5 = 391$. □

Theorem 23 (One-ray dissipative-memory gap)**.** *There is an effective family of homogeneous rational depth-one fivefold tree isometries, with bulk dimension seven and fixed radial dimension*

$$d = 7 \cdot 158 \cdot 7^{80} \quad (\textit{about } 5 \times 10^{70}), \tag{38}$$

for which distinguishing $f_* \geq 1/6$ from $f_* = 1/49$ is undecidable. Every instance satisfies $1/49 \leq f_* \leq 9/49$, so the high side lies in $[1/6, 9/49]$. The quantifier over worst-case fresh bulk inputs supplies the instruction word. This theorem asserts no hardness reduction with those inputs frozen.

Proof. Apply Imports 1 and 2 of Section 5 with old threshold $1/18$, $d = 18$ and new threshold $1/2$. The output has rational transitions, 39 states, four letters and a deterministic start; its value is at most $1/2$ or is one, undecidably. These are precisely Lemma 22's hypotheses. Use its two-bank memory and five lazy instructions.

Terminal noise. For $b \in \mathbb{Z}_7$, $z \in \mathbb{F}_2^3$, let U_{bz} multiply basis label $j \in \{0, \dots, 6\}$ by $(-1)^{z \cdot \mathrm{bin}(j)}$ and then shift it by b modulo seven. Averaging these 56 signed permutations kills off-diagonal units by character orthogonality and sends diagonal units to $I/7$. Their uniform mixture is $\mathcal{T}(X) = \mathrm{Tr}(X)I/7$.

A key has 168 labels (b, z, m), $m = 1, 2, 3$. For an accepting memory label prepare χ_1 with amplitudes $(6, 4, 2)/56$ for each (b, z); for other labels prepare $\chi_0 = |0, 0, 1\rangle$. The norm is one since $56(36 + 16 + 4)/56^2 = 1$. Memory-controlled preparation is isometric. Apply the key-controlled U_{bz} to the fresh B. Losing the key implements $\mathcal{T}$ in the accepting case and identity in the rejecting case. This preparation preserves memory populations, also as joint diagonal blocks with untouched future bulk systems.

Cell and staggered archives. The control has phases P, T, F, O, with 78 states in each of P, T and one state in each of F, O. Only child zero receives the continuing phase; every nonzero child has tag O. The cap starts at the deterministic bank-zero label in P. Divide B into six fixed nonempty basis sectors of ranks $(2, 1, 1, 1, 1, 1)$, with five instruction labels and stop. At P, an instruction applies its V_i, preserving the whole sector of B, and sends its memory in P to child zero. Stop sends the unchanged memory in T. A fresh raw symbol stores the instruction and emitted register, or a separate stop marker. At T perform the terminal pulse, store both memory label and spent key in raw, and send F to child zero. At F, O preserve B while writing its sector label to raw and continue the respective tag. Use disjoint raw subspaces for all these phase choices. At most one terminal occurs on any history.

One raw symbol has dimension $r = 7^5 = 16807$: five blocks of 391 dimensions, a stop marker, $78 \cdot 168$ terminal labels, and six markers for each of F, O require $15072 \le r$. Encode old datum a and the post-pulse B with G. All new raw is initially private to child zero.

At each input keep a previous raw symbol and four waiting lists of lengths $4, 3, 2, 1$, all in r-dimensional symbols. Encode the previous raw and fixed basis padding with $G^{\otimes 5}$ on their five base-seven digits. Emit its initial share and the next share from each waiting list. These five symbols form Z of dimension r^5. The remaining lists, of lengths $4, 3, 2, 1$, together with the current raw, are the next eleven delay/raw symbols in child zero; other children have fixed delay padding. Finally encode the entire incoming spatial archive A and Z by $G^{\otimes 25}$, digitwise, into the five child archives, each of dimension r^5. Spatial archives never control the live memory or delays. Initialize all data, delay and archive factors to fixed basis vectors.

Four shares of G recover both inputs by Lemma 2. A single share is a replacement channel, including with any reference: tracing four evaluations makes bra and ket polynomials identical, killing distinct logical matrix units, while summing the constant coefficient makes each remaining evaluation uniform. The identity holds on both encoded inputs jointly. Thus it holds for any number of parallel copies. Every delay operation is a fresh isometric encoding, a reordering, or an encoding of the entire remaining input. Distinct input phases have orthogonal outputs, including T, F, which share the outgoing tag but have disjoint raw labels. Within P, sector preservation makes the controlled updates isometric on superpositions. The data and spatial encoders act on disjoint slots. This proves isometry on the full input

space, including arbitrary waiting-list contents and entangled inputs. Entries and embeddings are rational. The radial factors are a of dimension seven, control of dimension 158, eleven r-dimensional delay/raw symbols, and archive of dimension r^5. Their product is (38). The same cell is repeated everywhere, and its depth per refinement is one.

Lower decoder with arbitrary spectators. At any query invert each surviving whole subtree on its range, using the trace-preserving extension described in Proposition 14. Recover the post-pulse B from four data shares and discard old data and all controls. Six-sector pinching $\mathcal{D}_6$ satisfies $\mathcal{D}_6(\sigma) \succeq \sigma/6$, also with references: write a test vector as $x = \sum_j x_j$ in the sectors and use $\| \sum_j \sqrt{\sigma} x_j \|^2 \le 6 \sum_j \| \sqrt{\sigma} x_j \|^2$. At an ordinary P, F, O step, tracing its instruction/phase-labelled raw gives exactly this pinching on its ideal subnormalized target block. Each within-sector update is an isometry on memory and preserves the target within that sector.

For clarity the population induction at a terminal uses this trace order. Discard old data and spatial archives; commute the trace of every detached raw/delay history past later operations, since none controls memory. At each earlier program step trace its controlling bulk after resolving its orthogonal instruction sectors. The lemma updates the memory's diagonal blocks by S_i with arbitrary operator-valued entries on the untouched future target and reference. Earlier terminal keys and their targets, when both traced, preserve the other marginal after the controlled preparation; no off-diagonal entry can contribute to a subsequent population. Resolve the stop and phase histories by their orthogonal raw labels. Equivalently project all earlier fresh program bulks into their sectors in the original joint input, then trace those bulks and all unqueried systems. This gives positive subnormalized target/reference operators σ_w for prefixes w ending in stop. Their accepting and rejecting blocks are exactly $p(w)\sigma_w$ and $(1 - p(w))\sigma_w$. Other phase histories have their own ideal blocks. All ideal blocks sum to the desired target/reference marginal by trace preservation. This establishes the assertion for coherent stop choices and arbitrarily entangled bulk inputs, without choosing a decoder after learning the word.

A terminal block is therefore recovered as

$$(1 - p(w))\sigma_w + p(w)(\mathcal{T} \otimes \mathrm{id})(\sigma_w).$$

The map $\mathcal{T} - \mathrm{id}\,/49$ is completely positive: its unnormalized Choi matrix is $I \otimes I/7 - |\Omega\rangle\langle\Omega|/49 \succeq 0$, where $\|\Omega\|^2 = 7$. If all $p(w) \le \lambda$, this block dominates $(1 - 48\lambda/49)\sigma_w$. If the PFA value is at most $1/2$, terminal blocks dominate $25\sigma_w/49$, and all other blocks dominate $\sigma_w/6$. The unconditional output dominates the desired state divided by six. The implication $\omega \succeq t\sigma \Rightarrow F(\sigma, \omega) \ge t$ follows by comparing the eigenvalues of $\sqrt{\sigma}\,\omega\sqrt{\sigma} \succeq t\sigma^2$ and summing their square roots. Hence $f_* \ge 1/6$. With $p \le 1$ the same argument gives $f_* \ge 1/49$ for every instance.

Upper witnesses. Choose the fresh program bulks as basis-sector instructions for any lazy word, then stop, and place an independent Bell target/reference pair at the following T. Erase child zero at its birth. The current raw, including memory and spent key, is lost. Tracing its memory label removes cross-label terms. Conditional on acceptance, the pre-pulse key is the pure χ_1, independent of the old archives, old data and Bell pair. Tracing it depolarizes the target, leaving the reference independent of everything retained, even if old archives reveal the memory label. Any decoder has Bell overlap $1/49$ in that branch and at most one otherwise.

The witness gives $f_* \le 1 - 48p(w)/49$. Lemma 22 makes $p(w)$ approach one in the value-one case, hence $f_* = 1/49$. The two promises are separated by $1/6 - 1/49 = 43/294 > 0$, proving undecidability.

For the universal upper bound choose an O site, an independent Bell target and loss of child zero at birth. Its old controls and archives are independent of the fresh target; its sector label is lost, so its effective target channel is six-sector pinching. For sector projectors P_j of ranks r_j and decoder Kraus matrices L_ℓ, the Bell overlap is

$$\frac{1}{49}\sum_{j,\ell} |\operatorname{Tr}(L_\ell P_j)|^2 \le \frac{1}{49}\sum_j r_j \sum_\ell \|P_j L_\ell P_j\|_{\mathrm{HS}}^2 \le \frac{1}{49}\sum_j r_j^2 = \frac{9}{49}.$$

The initial inequality is Hilbert–Schmidt Cauchy–Schwarz, and the next uses $\sum_\ell L_\ell^\dagger L_\ell = I$ and that compression cannot increase the indicated norm. Independent side information does not improve this bound, since composing with its fixed preparation still gives a channel on B. Thus $f_* \le 9/49$ always. The $25/49$ above is a terminal lower bound; it is the high-side constant only for the recorded template of Section 5, not for this one-ray theorem. All witnesses use ordinary fresh inputs, without heralding or postselection. □

7.4 A probability budget for disjoint recovery footprints

Fix one finite-depth source system, its encoding into a source site's radial registers, and a preparation of all other inputs independent of the source. Those other inputs may be mutually entangled. Two source states are compared in this same experiment. A receiver realizes a fixed channel Λ when a deterministic decoder on its entire region induces Λ, also with references. The prescribed channel may be noisy; its two selected outputs must have trace distance $\Delta > 0$, where $d_{\mathrm{tr}}(\rho,\sigma) = \|\rho - \sigma\|_1/2$. This is the usual marginal-channel notion of compatibility [62]; it is distinct from broadcasting arbitrary noncommuting states [63].

Proposition 24 (At most $2/\Delta$ nonexceptional footprints)**.** *For the one-ray cells of Theorem 23, at a common cutoff and for a fixed pair of source inputs, at most $2/\Delta$ pairwise disjoint nonexceptional recovery footprints can have output trace distance at least Δ. This also bounds disjoint footprints realizing a fixed channel whose outputs on that pair have contrast Δ, and subsets of disjoint bounding footprints. Nonexceptional excludes at most six early side slices after the source preparation, at most one footprint containing the final spine leaf, and at most one footprint containing all children of an off-spine source site. Thus the total bound is $2/\Delta + O(1)$, uniformly in cutoff and the common future preparation. It is a bound per input pair, not a bound obtained by varying the source pair from receiver to receiver.*

Proof. A whole child departing from the source's zero ray has fixed tag and delays and one data share and one archive share. The tensor product of the single-share replacement channels proved above resets their joint inputs, including with all untouched future systems and a reference. Consequently every whole descendant interval strictly below such a departure has a constant source marginal after any refinement. A footprint containing all children of an off-spine source is explicitly excluded; disjointness permits at most one such footprint. For the other footprints follow that source's zero-child continuation, where the same replacement identity applies after every departure. A footprint anchored on the active ray and losing a

nonzero child contains the final spine leaf, so at most one such footprint is disjoint from the others. The remaining potentially sensitive footprints are side slices $R_n(0^k, 0)$ relative to the source ray.

Trace all siblings from the preceding spine step, which lie outside a given side slice. The entering data a and spatial archive A are then independent maximally mixed factors, jointly independent of the machine, delays and future spectators: their two encoders induce a tensor product of replacement channels. Neither factor is read before its next encoding. After flushing the delay line, its five emitted symbols Z are also fixed maximally mixed factors independent of the current target, phase and future systems. Indeed they are one share each from five different raw encodings. Each encoding acts on a detached raw input without feedback, so commute it past all later operations on memory and future bulks when computing this slice. Earlier emitted shares could return only through old A, which was just replaced. The remaining delayed shares enter child zero and are traced. Thus, even for entangled raw inputs, the effective map is a tensor product of five single-share replacement channels. Six steps suffice to flush arbitrary injected delay contents; the initial slices are precisely the finite exceptions above.

Let $p_k^{(x)}$ be the unconditional probability of phase T at step k for source input $x \in \{0, 1\}$. When child zero's current phase/raw is traced, distinct phase histories decohere. On every nonterminal block the current data is merely pinched into six sectors. Replace the terminal block by that same pinching to define a baseline. Each positive terminal block has trace $p_k^{(x)}$, so its two possible trace-preserving outputs differ in trace distance by at most $p_k^{(x)}$. The baseline depends only on the unconditional marginal of the current fresh bulk and untouched future systems. This marginal is source-independent: all prior operations are trace preserving on its complement. The old-data and archive factors have already been replaced. Subsequent encoding, refinement within the side subtrees and any receiver processing cannot increase trace distance. Therefore the entire nonexceptional footprint obeys

$$d_{\mathrm{tr}}(\omega_k^{(0)}, \omega_k^{(1)}) \le p_k^{(0)} + p_k^{(1)}. \tag{39}$$

For each input, $\sum_k p_k^{(x)} \le 1$. The probability remaining in P decreases by the stop flow into T, and the next step sends all T weight permanently into F. Telescoping gives the inequality, including possible initial T weight. Orthogonal phase/raw labels justify this probability balance with coherent or entangled fresh controls. Sum (39) over a disjoint family. Its total contrast is at most two, so at most $2/\Delta$ of its nonexceptional slices can each have contrast at least Δ. A decoder realizing Λ requires at least the contrast of its outputs by trace-distance contraction. The same argument covers a subset of a bounding footprint. Add the explicitly excluded finite exceptions to obtain the total bound. No no-cloning or general no-broadcasting theorem is used in this probability-budget proof. □

Pooled disconnected receivers remain open. These codes do serve receivers that pool different side slices. Park the memory in either of two bank-one basis states and repeat one ordinary instruction, which leaves that bank fixed. Each step emits two distinct nonorthogonal raw states: their source-block components differ and their overlap is $c^2 = 36/49$. This realizes a fixed nonconstant noisy channel on the classical source label. Five consecutive later side slices supply all five temporal shares of one such raw symbol. Each slice's four archive shares

recover Z and hence the specified temporal share; pooling those slices inverts the raw encoding. Choose emission times at least six steps apart to obtain arbitrarily many disjoint such pooled receivers at a common late cutoff, using subtree inverses as needed. Proposition 24 concerns each individual tested footprint. A necessity statement for unrestricted disconnected receivers remains open.

Input control and dissipation. Theorem 20 gives an exactly computable value when the single excitation is unaffected by fresh inputs and the no-arrival state is inert. Theorem 23 gives a robust undecidable gap when worst-case fresh inputs steer a dissipative memory that keeps full weight on one ray. Broadcast to the tested footprints is not needed. Is dissipation necessary for this gap? For algebraic unitary letters a measure-once automaton value is computable by Derksen–Jeandel–Koiran [28, Theorem 1 and Section 2.2], as used in Theorem 21. A necessity theorem for these recovery codes has not been proved.

8 Triangular disk isometries with matching rules

8.1 Whole-leg geometry and local identities

Use the vertex convention $\{p, z\}$: p edges per face and z edges at a vertex. We reserve χ for Hilbert dimension. In this section $(p, z) = (3, 32)$. A physical site is one complete dangling planar leg of dimension χ, irrespective of any tensor factors used to describe it. Each vertex has an independent bulk input of dimension χ. A vertex-centered cutoff $D_n(o)$ is the graph-distance ball of radius n; $n = 0$ is the isolated seed. All internal planar edges are contracted. A valid disk map has Gram $c_n I$ with $c_n > 0$, and is normalized by this one scalar. No cap or postselection is applied to the physical boundary.

Here is the fan geometry. The seed has 32 private neighbors. A nonseed row vertex of type a has one parent, two tangent neighbors and 29 outward spokes; type b has two consecutive parents, two tangents and 28 outward spokes. Across an old rim edge the new triangle has one shared b child. Between its two shared extreme children a parent has respectively 27 or 26 private a children. Every row is a simple cycle and shared children are isolated. The boundary construction glues the missing regular triangles in these fans. It gives nested topological disks; all completed interior stars have valence 32. Developing the resulting complete simply connected regular-polygon surface into the hyperbolic plane gives the regular tiling. Completeness follows, for example, because the star of a face contains a uniform metric neighborhood of that face, so a finite-length path crosses only finitely many successive star shells. No further vertex identifications occur. Every new vertex has a parent and edges skip no row, proving that these are graph-distance balls. These fan counts agree with the triangular inflation in Basteiro et al., Eq. (11) [20].

At birth, order a vertex's l outward legs cyclically after its inward fan and divide them at $\lfloor jl/5 \rfloor$, $j = 0, \ldots, 5$. A query deletes one group and retains the others. At the seed use its 32 legs and any rotation of the ordered grouping. At later cutoffs each old spoke owns a contiguous descendant block: a private child's entire outward fan belongs to its parent spoke; at a shared child split the 28 outgoing legs equally between the clockwise-earlier and later parents. These are ownership labels for the observation, not extra contractions. The

footprint of a vertex is the union of its birth spokes' descendant blocks. Recovery is (3) with all other bulk inputs arbitrary, quantified over every disk, every included vertex and every allowed group loss.

We use the following *qudit extension* of Steinberg et al., Definition 2 and Fig. 5 [17]. The vertex tensor A is cyclic in its 32 planar slots, and B on an edge is symmetric unitary. The seed encodes one χ-level input into 32 χ-level planar systems, with subspace-code distance exactly two. The one-vertex flattening has the logical and one planar leg as inputs. The two-vertex flattening joins neighboring tensors by B and takes, from each vertex, its logical leg and a planar leg immediately adjacent to the joined bond, on the same side of that bond, as inputs. Both flattening Grams must be positive scalar identities; each diagram may be normalized to unit Gram. This is a matrix-identity convention, with no stabilizer or qubit hypothesis. It does not require the entire annulus to be isometric on its full inner rim. Appendix A proves these identities for the construction, including distance exactly two.

Cheng's angular-uniformity criteria [49] provide a geometric comparison for the distinction between local isometric identities and full hyperinvariance. Evenbly's state networks [18], the ququart and state-dependent recovery of Steinberg–Feld–Jahn [19], and the qubit-code identities of Steinberg et al. [17] have separate hypotheses; the qudit extension above declares exactly which identities are used.

8.2 An effective history tensor

Hyperbolic domino undecidability is established by Margenstern [50], and Carette–Moutot [51] formulate tensor-valued quantum Wang tilings. Together with the matching-rule precedents [13, 14, 15, 16], these supply standard simulation and tensor techniques. The contribution here is seeded disk-isometric recovery hardness using standard matching rules, with independent bulk inputs and a witness on erased whole legs. Grogan's explicit seeded half-plane construction from Glazer's machine [53] is an independence precedent in a different geometry and predicate; no independence claim for recovery is inferred. The following construction gives the raw finite matching table before its physical whole-leg realization. Describing finite port labels and a coefficient for every admissible tuple specifies a tensor effectively.

Lemma 25 (Rooted and root-free histories)**.** *From a blank-tape machine one can compute a finite cyclic matching table on $\{3, 32\}$, with nonnegative raw amplitudes and a logical involution Z, having these properties.*

1. *A configuration has at most one internal root. It may instead be root-free or empty. At a root x it chooses a bit t and one permanent arm, and acts on its logical input by Z^t; at all other vertices its logical action is identity.*
2. *Whole boundary records determine every applied logical phase, for any machine.*
3. *If the machine never halts, each retained region determines the phase at its queried vertex using two retained private rays.*
4. *If the machine halts, for the concentric target o and every sufficiently late cutoff there are two rooted histories, with opposite bits and the same permanent arm, whose boundary records differ only in one dangling payload inside an allowed erased sector, at least two ports from that sector's ends.*

Proof. **Frames and their uniqueness.** A root frame has every spoke outward and requires each neighbor's unique parent to point back. An ordinary frame has one or two consecutive inward ports, tangents immediately outside them, and the remaining outward fan. A private, non-extreme outward spoke requires an a child with its parent back along that edge. An extreme spoke requires a b child with its other parent on the appropriate side in the common triangle. Tangent orientation is fixed by handedness. Port roles and expected remote frames are finite labels; matching compares complementary roles and their data. With a true root, the actual distance frames satisfy these rules.

For uniqueness, retract the outer row of $D_{r+1}(o)$ to $D_r(o)$, taking a private child to its parent and a shared child to its clockwise-earlier parent, fixing the old disk. Adjacent new vertices have equal or adjacent selected parents, and a new-to-old edge maps to a point or an edge. The retraction is 1-Lipschitz on graph distances. Retracting any shortest path from a larger ball proves that each disk is an isometric subgraph. From an internal root x choose an in-disk shortest path to any vertex. Induction along outward steps forces exactly the true distance frame about x at that vertex. A second root is impossible. Root-free partial frames remain allowed; none is presumed to extend beyond the disk.

Straight arms and a finite computation table. Mark all 32 straight rays from a root and distinguish one permanently. At an a carrying an incoming mark, continue it through the opposite outgoing port; forbid marks at b vertices. The opposite port is strictly inside the private fan. A chain of opposite equal hyperbolic edges is a geodesic and therefore a graph shortest path. Induction gives the root's rays. In a root-free configuration a mark backtraces along the same straight line to the boundary. Two noncollinear such lines whose backtraces meet at an internal vertex v identify a root at v, because an ordinary vertex carries at most one marked line. The same test excludes any different root.

Compile the machine into a one-sided tape machine that, on simulated halt, returns to its permanent left-end cell and enters a stationary ACCEPT state there. A two-sided tape can be interleaved on two tracks of a one-sided tape with finite control. The new machine accepts exactly on the original halting inputs. In each root sector place an independent copy of this tape, whose home cell is on its clockwise-initial marked ray. Distance one initializes the headed blank with no update. At each later row an active marker sends its updated candidate to its straight private child. A home also sends a new blank marker to the private port immediately preceding that continuation. Other outward requests are inactive; b vertices require both parent requests inactive. In cyclic sector order this preserves the old tape and appends one blank before the next home's marker. Thus the tapes are arbitrarily long and disjoint.

On tangent gaps carry right-going and left-going candidate tape letters and home bits. Fillers match each component through. An active marker fixes its departing messages to its own candidate and reads the arriving neighbor candidates. At home ignore the left neighbor; just before home use an unheaded blank to the right, excluding moves across sector boundaries. The radius-one head rule then computes the next tape state. Inconsistent triples can be disallowed or completed by a fixed rule. On a complete rooted shell each gap's messages are fixed by its two endpoint markers, including consecutive markers. Hence at radius $r \geq 1$ each home has executed $r - 1$ transitions. These are simultaneous local constraints on neighboring labels, not an angular isometric sweep.

Payload and visible phases. At the root choose $t \in \{0,1\}$ and apply a fixed diagonal involution Z with both signs on the fresh input. Carry t only along marked arms, never on tangents or unmarked requests. At a nonpermanent home, blank its outgoing payload on ACCEPT or if it was already blank; otherwise continue it. The permanent arm always carries t. Computation and marks do not depend on t. Add an all-empty sector, with no phase, whose labels cannot match live ones. The local table has coefficient one for each canonically prescribed tuple; duplicate hidden descriptions can be grouped with their positive multiplicity.

The whole boundary sees all straight exits, so two noncollinear lines identify a root and its permanent arm gives its bit. Root-free histories have no logical phase. This proves item 2, including any coincidences of full boundary tuples.

For a vertex v at depth h in a disk of radius n, a lost group has at most six outward spokes (seven at the seed). At a nonseed discard that group, both extremes, and two ports from each end of the at most two retained runs. At least $l - 6 - 2 - 8 \geq 12$ eligible spokes remain. Choose two nonopposite private directions. Their straight continuations remain private and outward relative to o, since at the next a the opposite port is about 14 slots from the extremes. At their terminal vertices the two neighboring ports on either side also belong to the retained footprint. The full ball of radius $n - h$ about v lies in $D_n(o)$ by the triangle inequality, so every complete computation shell needed to determine these outgoing payloads is present. In a nonhalting history rooted at v both payloads are t. Their backtraced marks identify v; other roots and root-free histories cannot imitate them. Thus they determine the target phase, with default zero otherwise. This includes the birth cutoff with no propagation steps, proving item 3.

If halt occurs, choose the physical center as root and a permanent arm in the interior of one of the lost seed groups. Beyond return-home acceptance all other payloads are blank. Histories with $t = 0, 1$ then agree on all other raw legs. The arm's terminal direction and its two neighbors on either side remain in the erased descendant sector. The deterministic forward construction provides both histories at every such cutoff, proving item 4. □

Theorem 26 (Exact triangular disk recovery)**.** *Exact all-cutoff sector recovery is* Π_1^0*-complete on an explicit effective template of uniform cyclic qudit tensor codes on* $\{3, 32\}$ *satisfying the one- and two-vertex identities stated above. Every complete disk is a nonzero scalar isometry on all bulk inputs. The common bulk/whole-leg dimension* χ *is finite and cutoff-independent within an instance, and unbounded across instances. The tensor entries are algebraic. This conclusion does not assume full-rim isometric annuli.*

Proof. Apply Lemma 25 and the physical construction in Lemmas 28–30 of Appendix A. The latter uses five routed polynomial shares for each raw label, an independent cyclic base encoder G_n, and ancillary straight Bell wires. Every internal edge has at least three complete routes, so the physical disk is exactly the raw contraction followed by independent boundary encodings and fixed Bell factors, up to a nonzero scalar. In particular its physical erasures are the prescribed whole-leg erasures. All entries are effective algebraic numbers. The appendix verifies the local identities and the seed-only sectors giving distance two.

For $n \geq 1$, undo the boundary coding in thought and suppress the fixed Bell factors. Let x range over all exposed raw tuples, and let $a_x \geq 0$ count histories with that tuple. By

Lemma 25, applied phases $t_v(x)$ are common to every such history. At most one is nonzero. Thus, up to a common scalar,

$$V_n|\psi\rangle = \sum_x a_x|x\rangle \otimes G_n \left(\prod_v Z_v^{t_v(x)}\right)|\psi\rangle. \tag{40}$$

The raw basis is orthogonal, $G_n^\dagger G_n$ is scalar identity, and the empty assignment exists. The disk Gram is therefore positive scalar identity. At $n = 0$ the additional auxiliary sectors have distinct raw tuples and base isometric images with normalized logical-dependent Bell factors; their individual Grams and their sum are scalar identity. Normalization does not select or reject a computation.

For nonhalting, use the independent base decoder of Lemma 29 on the retained base factors. It recovers the target coherently with arbitrary spectator inputs. The two retained ray records in Lemma 25 are fully accessible through their five-share codes by the margin property in Lemma 28. Read them coherently and undo $Z_v^{t_v}$ on the recovered target. On every term of (40) its target phase has now vanished; phases on spectators do not depend on the target value. The decoder dilation sends $|i, h\rangle$ to $|i\rangle|\eta_h\rangle$, independent of i; isometry makes different η_h orthogonal. At the isolated seed the retained raw records also distinguish the auxiliary sectors. A lost group has length at most seven and contains no antipodal pair. After base recovery, undo each logical-dependent straight-pair phase on a retained half, using $(M \otimes I)|\Omega\rangle = (I \otimes M^T)|\Omega\rangle$. This is a coherent controlled unitary. The junk is then independent of the target in every sector. Discarding it gives exact recovery, including the seed cutoff.

For halting use the two concentric histories of Lemma 25, beyond acceptance. All five shares of their differing permanent raw label lie in the erased physical region. After the known local decoding and basis change, let O interchange these two raw basis labels on their span and be zero elsewhere. This Hermitian operator has norm one and is an observable on E alone. Fix spectators in $+1$ eigenstates and compare opposite Z eigenstates of the target. On tracing other labels, each common completion y contributes, up to the same positive normalization,

$$2a_{x_0,y}a_{x_1,y}\left[1 - (-1)^{t_v(x_0,y)+t_v(x_1,y)}\right] \tag{41}$$

to the difference of expectations. Every term is nonnegative, and the exhibited pair makes one positive. The erased states differ. Lemma 1's necessity excludes exact recovery from R. This is an inaccessible-region witness in the actual normalized code, without a boundary condition or conditioning on an event.

The map from machines to the finite tensor table is computable, with nonhalting exactly the passing case. Finite disks can be generated and contracted effectively; their Gram equations and each recovery query are decidable over algebraic numbers by Lemma 1. Enumerating failures gives membership in Π_1^0 on the template. If invalid disks are to count as failures in a larger language, enumerate a non-scalar or zero disk Gram too. This proves completeness with the stated scope. □

Proposition 27 (An annular kernel)**.** *The triangular tensors in Theorem 26 obey the specified local one- and two-vertex identities but their innermost isolated annulus is not a scalar isometry on the full inner-rim space. The same is true after fixing every logical input to the same basis vector.*

Proof. In Appendix A's auxiliary sectors, J_+ matches only J_- across a raw bond, and neither matches live or empty labels. Choose the nonzero rim vector obtained from the isolated seed's J_+ component for a fixed logical basis input. The cable contraction identity is linear, so feeding this vector to the annulus recovers precisely raw matching with the center fixed to J_+. A nonzero term would require alternate signs around a triangle through the center. This is impossible. The annulus kills that nonzero vector for any new bulk inputs. The complete ordinary disk remains nonzero by its empty sector. Therefore the annulus has a kernel and cannot be scalar isometric on its full domain.

Freezing all logical legs to the same unit basis vector restricts the local scalar-identity Grams to the remaining planar inputs, preserving those identities, cyclicity and symmetric unitary edges. The same seed vector is still nonzero and still killed. This also gives a state-network counterexample to the full-rim interpretation of the sentence following Steinberg et al., Definition 1 [17], asserting an isometric map along any radial direction, for $p = 3$ in the qudit setting. It does not concern the separate bounded gate decompositions on $\{5, 4\}$. □

The annular kernel prevents full-input radial isometry, so these codes are not hyperinvariant in Evenbly's 2017 sense [18].

9 Interpretation and remaining boundaries

For the exact codes and geometries specified here, the tested observable is access to a local bulk quantum subsystem from a prescribed boundary subregion. Such access permits boundary representatives of its observables by pulling them back through the decoder. The geometric intervals are fixed by the experiment; no metric is inferred from an entropy calculation. This use of quantum error correction is motivated by bulk reconstruction in HaPPY and related codes [2, 1, 19], but our theorems do not assume their perfect-tensor, entropy or continuum hypotheses.

The exact decidability proof has three separate requirements: full-input isometry for each local circuit gate, a contracting bound for the endpoint neighborhoods, and finite-state control of their locations. The seam violates the bounded communication radius. The growing-depth schedule violates its uniform bound. The growing-dimensional template violates the fixed finite patch space. The triangular matching construction preserves initialized disk information while violating full-rim annular isometry. The approximate recorded template preserves all circuit hypotheses and tests an all-word inequality in place of zeroness. These statements concern the specified experiments and finite descriptions.

The approximation results turn on input control. A single excitation whose addressed matrices are independent of fresh inputs has a summable terminal-arrival budget. With an inert no-arrival state its value is exactly algebraic and every rational threshold is decidable (Theorem 20); the polarized example of Theorem 18 retains only an exact-root obstruction. Worst-case fresh inputs can instead steer a dissipative memory that keeps full weight on one ray, producing the robust gap of Theorem 23. Broadcast to the tested footprints is not needed: Proposition 24 gives at most $2/\Delta$ nonexceptional receivers per source input pair, with its stated finite exceptions.

Open questions. Is dissipation necessary for the one-ray gap? A measure-once value with algebraic unitary letters is computable by Derksen–Jeandel–Koiran [28]; a necessity theorem for these codes is not proved. The case of unrestricted pooled disconnected receivers is also open. These examples serve such receivers by distributing temporally separated shares of a spent raw symbol, so the individual-footprint bound does not apply to their unions.

For the physical valences of the pentagon HaPPY graph $\{4, 5\}$ and the four-planar-leg Evenbly graph $\{5, 4\}$, the literal five-sector birth experiment is degenerate: some birth footprints have only two or one spokes. Proposition 35 proves the obstruction, and Corollary 34 proves the buffered alternative. Under that alternative, the bounded annular circuits treated in Appendix C have decidable exact recovery. This does not classify arbitrary tensor presentations on those graphs.

Several boundaries remain open within the questions posed here. Full-rim isometric annuli on the symmetry-fixed regular tilings, without a bounded circuit decomposition, have no classification in this paper. Neither do fixed-dimensional stabilizer restrictions of the triangular hardness template, generic coherent approximate templates, or stability of the exact hard instances under perturbation. Approximate hardness for the recorded tree gives a uniform decision obstruction at each fixed tolerance; it is not a noise threshold or a dynamical statement. No result here concerns quantum gravity itself.

A Physical realization of the triangular matching table

We give the tensor factors, slot ordering, edge matching, and local identities used in Theorem 26. All dimensions in this appendix are finite functions of the input transition table. Their largeness is allowed across instances, never across cutoffs of one instance.

Lemma 28 (Five routes and whole-leg decoding)**.** *Each raw half-edge label of the history table has a fixed physical realization, cyclic at vertices and with a symmetric unitary edge matrix, such that every disk contraction equals the raw contraction followed by independent boundary isometries and fixed Bell factors, up to a nonzero scalar. A raw dangling spoke at least two ports inside a retained or lost run has all of its boundary shares on that side of the physical partition.*

Proof. For an edge uv, use five simple routes from u to v: the edge, the two length-two routes through its triangular tips, and the complementary arc in the 32-cycle of neighbors of each tip, omitting uv. Their initial ports at either end are $e, e-1, e+1, e-2, e+2$ in cyclic order. The sign reverses on edge reversal. Each physical edge carries a distinct tensor factor for each occurrence of a route through it. Intermediate vertices connect that route's factors by a two-leg equality tensor. A route has length at most 31; the finite number of route occurrences meeting a port can therefore be computed in a radius-33 flag neighborhood. Order these slots by route type, endpoint flags and relative local path words, using the lesser description for a reversed occurrence. Use coordinates relative to the actual port, so a cyclic vertex rotation preserves transmitted basis values. Midpoint rotation interchanges edge endpoints and induces an involutive slot permutation. Its matrix is symmetric unitary.

Pad the finite raw alphabet to $\mathbb{F}_s$, a power of 97, reserving distinct live, empty and two auxiliary symbols. Unused labels have zero coefficient. Raw compatibility is an involution P on labels: it exchanges complementary input/output or tangent roles, their expected frames and message data. Redundant compatibility checks can be made local by including expected remote fields in a port label. Extend P on unused labels by identity. Set

$$L = (I + P)/2 + i(I - P)/2.$$

Since $P = P^T = P^{-1}$, L is unitary and $L^T L = P$. Applying L at both ends turns ordinary bilinear contraction into raw P matching. On an exposed index it is a known local basis change.

Encode a raw label j by

$$C|j\rangle = \frac{1}{s} \sum_{a,b \in \mathbb{F}_s} |f(\infty), f(1), f(-1), f(2), f(-2)\rangle, \qquad f(t) = j + at + bt^2, \quad f(\infty) = b. \tag{42}$$

Assign ∞ to the direct route, ± 1 to the short routes and ± 2 to the long routes. Reversal exchanges the signs and leaves every codeword invariant by $a \mapsto -a$. Three extended evaluations determine f. For each fixed j, any two selected evaluations are a nonsingular linear function of (a, b): at two distinct nonzero finite points its determinant is $t_1 t_2 (t_2 - t_1)$, and for ∞, t it is $\pm t$. Thus for a set S of at most two shares,

$$\mathrm{Tr}_{S^c}(C|j\rangle\langle k|C^\dagger) = \delta_{jk}\, I_S / s^{|S|}. \tag{43}$$

Off-diagonal terms vanish by interpolation on the complementary three coordinates. The code is real, so the same equation describes bilinear contraction of two encoders on at least three equal coordinates, leaving identity wires on the missing coordinates and exact raw matching. This is an algebraic tensor identity, not a decoding assumption.

Every internal disk edge has three complete routes. If it is not on the outer cycle both incident triangles are included, so the direct route and both short routes are complete. If it is on the outer cycle, its inward tip lies on the preceding row and all of that tip's neighbors lie in the disk; the direct, inward short and inward complementary-link routes are complete. These cover all internal edges. Equation (43) therefore contracts them to raw P matching times independent Bell factors and nonzero constants. A raw edge with exactly one endpoint inside emits its full code, with shares at their respective initial exits. Other route pieces have two exposed ends and are independent Bell wires. This proves the claimed whole boundary factorization. The endpoint ports of a raw dangling spoke are within two slots of it, so two-port margins ensure all five exits lie in the same retained or lost run. There is no boundary measurement or removal of part of a physical site. □

Lemma 29 (Independent cyclic base code)**.** *There is an effective cyclic base tensor with a χ-dimensional bulk leg and 32 Q-dimensional planar factors, with $\chi = Q\chi_{\mathrm{aux}}$, such that each complete base disk is a positive scalar isometry and each prescribed retained footprint recovers its target subsystem at every cutoff. Here χ_{aux} is the dimension of all nonbase factors on one physical leg.*

Proof. All auxiliary factors can be chosen to have dimensions powers of 97, including one extra straight-wire factor of dimension 97. Choose a power Q of 97 with $Q^3/32 \geq Q\chi_{\mathrm{aux}}$. Since $32 \mid 96$, $\mathbb{F}_Q$ contains a primitive 32nd root ω. Encode three field coefficients $u = (u_0, u_1, u_2)$ by equal-amplitude evaluation at $1, \omega, \ldots, \omega^{31}$ of

$$f(t) = \sum_{i=0}^{13} r_i t^i + u_0 t^{14} + u_1 t^{15} + u_2 t^{16}, \qquad \text{amplitude } Q^{-7} \text{ for each } r \in \mathbb{F}_Q^{14}. \tag{44}$$

For a subset S of at most 14 evaluations the complement has at least 18, hence determines all 17 coefficients. Off-diagonal logical matrix units therefore vanish on tracing the complement. For a fixed u, evaluation on S is a surjective linear map of r by the Vandermonde rank, with constant-size fibers. Thus its partial trace is $\delta_{uv} I_S / Q^{|S|}$. In particular the tensor, with logical input and any at most 14 planar inputs, has scalar-identity Gram.

Cyclic rotation of evaluations multiplies each coefficient by its fixed power of ω and permutes the r's. It permutes logical basis triples in orbits of size at most 32. Their normalized orbit sums give at least $Q^3/32$ orthonormal pointwise invariant inputs. Restrict to the span of the earliest $\chi = Q\chi_{\mathrm{aux}}$ orbit sums in a computable order. The matrix-unit identities survive an arbitrary orthonormal restriction, and the resulting tensor is literally cyclic. The complete planar dimension is $Q\chi_{\mathrm{aux}} = \chi$, matching the bulk dimension. Coefficients of orbit sums and code normalizations are algebraic.

For disk isometry prepare a normalized Bell pair on every tangent bond of a new row. Each vertex can then be treated in parallel as a scalar isometry from its bulk, its one or two inward planar factors and both tangent factors: at most four planar inputs. Composition gives scalar isometry of the disk. This reasoning concerns the independent base factor, without imposing the raw matching constraints.

At birth an ordinary queried vertex has at least $28 - 6 - 2 = 20$ retained private spokes after the lost group and extreme shared spokes are omitted. Of its 32 ports at most 12 are unavailable, within the allowance 14. At the seed the bound is stronger. View each unavailable outgoing leg and a corresponding input argument as an unnormalized Bell pair. The scalar isometry from bulk plus unavailable ports to retained ports has an inverse on its image, which recovers the target while preserving arbitrary surrounding contractions. Its normalization is independent of its input. This proves the birth decoder with arbitrary spectators.

At later cutoffs follow only private child trees of the required spokes. A private child's entire outward fan belongs to its parent spoke. Inverting its tensor on its private outward spokes recovers its entering factor: only its inward port, two tangents and two extremes are unavailable, at most five. Repeat on disjoint private descendant trees. This recovers all required birth factors inside the prescribed footprint; unrelated inputs and tangent Bell factors can be discarded. Inverses extend to channels by projection off their images. The same tensor identity with an arbitrary surrounding state proves validity with entangled spectators and references, not merely for product inputs. □

Lemma 30 (Cyclicity, local identities and distance two)**.** *Tensor the raw history table, its routed realization and the base code with the following straight-wire factors and auxiliary sectors. The resulting tensor is cyclic, has a symmetric unitary edge, obeys the qudit one- and two-vertex identities of Section 8, and its isolated seed has distance exactly two. The auxiliary sectors vanish from every disk with $n \geq 1$.*

Proof. Pair antipodal planar ports locally by 16 normalized 97-level Bell vectors Ω. Internal straight factors contract by identity. In live and empty sectors these wires are undressed. Add raw sectors J_+, J_-, each with its own symbol on all 32 ports. They match only opposite signs across P and never live or empty symbols. Let $D = \operatorname{diag}(1, \zeta, \ldots, \zeta^{96})$, $\zeta^{97} = 1$ primitive, so $D^T = D$ and $\operatorname{Tr} D = 0$. Partition the logical basis into two nonempty parts and write its indicator $e(i) \in \{0, 1\}$. In either auxiliary sector put

$$G|i\rangle \otimes \bigotimes_{\text{antipodal pairs}} (I \otimes D^{e(i)})|\Omega\rangle.$$

Moving an operator between Bell halves transposes it; $D^T = D$ therefore makes these decorations cyclic even when a rotation exchanges halves. Base cyclicity and port-relative route ordering give literal cyclicity of the complete vertex. The edge is the symmetric unitary slot permutation of Lemma 28, identity on base and straight factors. All internal normalization factors are finite algebraic scalars.

To compute a one-vertex Gram, retain its logical reference and one planar port in its Choi tensor. Temporarily allow two adjacent planar ports. Each raw encoder contributes at most two of its five endpoint shares, since their initial ports are distinct. Equation (43) kills its off-diagonal raw matrix units and leaves identity on these shares, also after L. Distinct raw tuples therefore decohere. If several internal descriptions give the same tuple, their logical action agrees because a root bit is recorded on its permanent spoke; group their positive multiplicities. Per tuple the base marginal on the reference and selected ports is scalar identity by Lemma 29. Its input off-diagonals vanish, so the auxiliary basis-controlled Bell

phase cannot restore a logical off-diagonal. Adjacent selected ports contain no antipodal pair, hence the straight factors give independent identities for every logical value. Through-route factors may connect the two selected ports by a fixed Bell vector; for a single selected port all are identities. The one-port marginal is thus scalar identity on logical reference and port, exactly the one-vertex identity.

For the two-vertex identity at neighbors v, w, retain at v its bond to w and the adjacent input port towards the common triangle tip u; do the same at w on that side. Use the two-adjacent-port marginals just calculated and contract vw. A surviving through-wire connection between the two external ports would need three consecutive edges u, v, w, u of one route. Each of the five routes is simple, so this cannot occur. No route turns back at a port. Remaining contractions give identities on both external planar inputs and both logical references, times positive constants, including traces of identities or Bell pairs whose mate is unobserved. The local unitary edge does not change that Gram. This proves the two-vertex identity on either side. Normalize each flattening by its positive scalar; the disk tensor itself stays fixed.

Single-port error detection follows from the one-vertex identity, so distance is at least two. Two antipodal ports contain their straight pair. In live or empty tuples it is Ω for all inputs. Auxiliary tuples have positive seed weight independent of i and contribute $(I \otimes D^{e(i)})\Omega$. After tracing all unused shares the raw sectors decohere by (43). For $e(i) = 0$ the pair is always Ω; for $e(i) = 1$ it has positive weight orthogonal to Ω, since $\operatorname{Tr} D = 0$. Thus two ports detect a logical difference, proving distance exactly two in the subspace-code sense. Every larger disk contains a triangle, which cannot alternate J signs. Since J/live junctions are forbidden on the connected raw graph and Lemma 28 makes matching exact, no auxiliary assignment survives on those disks. Rescale the vertex once so that its bulk-to-planar seed map is an encoding isometry. All asserted conditions follow. □

B A normalized gap for triangular disks

An exact environmental distinction need not survive normalization with a fixed strength. In the unweighted construction there are $2z$ centered histories, $z = 32$, with distinct records. A genuine complete history rooted at each internal vertex can be restricted to the physical disk. Its root, bit and permanent arm remain identifiable on the full boundary. Thus the total squared history norm is at least $2z|D_n(o)|$, while the centered sector has mass at most $1/|D_n(o)|$. Removing its logical phase changes the normalized map in operator norm by at most twice the square root of that mass. The resulting map recovers the center from the independent base code, since other sectors act only on spectators. Accordingly the exact witness in Theorem 26 alone supplies no uniform fidelity deficit. We now change the positive local history weights and prove a gap for the changed entries.

B.1 Weights and boundary defects

Use canonical raw labels in which all outgoing requests are determined once the local frame, candidates and two tangent messages are known; there is no extra guess about a missing neighbor's outgoing data. Give each admissible tuple unit multiplicity. Let $A \geq 2$ bound the

number of raw half-edge labels, with any padding included, and put

$$K_0 = 100z^2A, \qquad r = A/K_0 = 1/(100z^2). \tag{45}$$

Multiply every live tuple by K_0 per outward port, K_0^{-1} per inward port and one per tangent. Use the all-outward rule at a root; keep empty and auxiliary tuples at weight one. This occurs before the basis change and five-share coding. On a matched internal edge the factors cancel. If a disk has B dangling legs, of which T are tangent and I inward in a live frame, its weight is

$$K_0^{B-T-2I}, \qquad \text{relative to } K_0^B \text{ it is at most } K_0^{-m}, \quad m = T + I. \tag{46}$$

Weights are positive, independent of bit and logical input. The physical factorization, phase determination and the local identity calculation in Appendix A are unchanged: the latter is diagonal in raw tuples and each scalar identity is merely multiplied by its positive squared weight. The auxiliary distance witness remains at the seed, and no auxiliary assignment survives for $n \geq 1$.

Lemma 31 (Rooted and root-free defect counts)**.** *For $D_n(o)$ with $n \geq 1$, a rooted frame with root x at distance h from o has $m \geq 4h$. For fixed x, m, at most $2zA^m$ compatible histories occur. For $x = o$ there are exactly $2z$ histories, with $m = 0$, distinct records and relative weight one. Root-free live histories have $m \geq 4n$, and their number with a given m is at most*

$$2z^{2n}(3zA)^m. \tag{47}$$

Proof. **An internal root.** On the physical rim S_n, distance to x has minimum $n - h$ and maximum $n+h$. Triangle inequalities give the bounds. Extend a shortest o-to-x path outward relative to o to reach the lower value. For the upper value, choose at o a private outward step in x's frame and continue straight through private children for n equal edges. This is outward relative to x and a hyperbolic geodesic, so its graph length cannot be shortened. For $h = 0$ both values are n. The cyclic height variation is at least twice the difference of maximum and minimum, hence at least $4h$.

Each changing rim edge has height difference one. At its higher endpoint, the edge to the lower neighbor is an inward port. The neighboring exterior port, towards the exterior shared triangle tip, is inward or tangent in this frame. The two incident physical rim edges at a higher endpoint give distinct such exterior ports. This injects changing edges into the m non-outward dangling ports, proving $m \geq 4h$.

For a fixed frame, arm and bit, specify the m non-outward boundary labels. Increasing-height propagation determines incoming candidates, then each local output. Cut tangent arcs are determined by their endpoint labels and internal markers. A complete tangent cycle bounds a whole sphere about the root; the disk fills its interior and contains its anchored home markers. Its messages are therefore determined too. This also handles incomplete outer shells and isolated cut-arc vertices. At most A^m choices remain, times z arms and two bits. At the center every complete shell is determined by the initial computation; all boundary ports point outward. This gives exactly $2z$ distinct records.

Root-free frames without an exterior extension. Assign an increment $+1$ along each outward edge, -1 inward and zero tangent. Each triangle has exactly one tangent edge. Two

tangent directions cannot be consecutive in a 32-port frame. An all-radial triangle would have a corner with adjacent inward and outward directions, forbidden because a tangent separates these fans. At the third vertex opposite a tangent the radial signs agree. Thus increments sum to zero around every triangle. Simple connectivity gives an integer height, up to translation.

Orient tangent edges by handedness and radial edges uphill. This directed graph is acyclic. A cycle would have constant height and contain a simple tangent cycle. Its inward side is consistent. If that side is inside, a triangle tip initiates a strictly descending path trapped in the finite interior, which cannot stop without a root; if the outward side is inside, an ascending path gives the same contradiction because every ordinary vertex has an outward neighbor. The bounded interior of a cycle in the disk has no exterior vertices, so touching the physical rim changes nothing.

Predecessors of a vertex consist of its one or two consecutive parents and the preceding tangent, all on at most three consecutive ports. At a rim vertex its actual disk neighbors lie in a contiguous fan of at most four ports. Its intersection with the predecessor arc is connected when nonempty, since the arcs together have length less than 32. Consecutive predecessors are joined by a triangle. In a topological ordering, all predecessors lie in one earlier connected component. Adding a vertex can therefore never merge two earlier components. The final graph is connected, so exactly one vertex has no predecessor.

A site without any in-disk radial parent lies on S_n. If its preceding tangent neighbor is inside, that neighbor is a rim neighbor at the extreme of the physical fan. The virtual inward triangle tip is the exterior shared child, and any second virtual parent lies still farther outside. Thus that preceding tangent neighbor is also parentless. Tracing such predecessors back reaches the unique predecessor-free vertex. Each site has at most one tangent successor, so the parentless sites form one directed consecutive rim segment.

Specify the start, length and direction of that segment (at most $2|S_n|^2 \le 2z^{2n}$ choices). Each of its sites has at most $2z$ choices of frame and consumes at least one non-outward dangling port. This frame data forces the remaining frames by radial propagation, since following parents backwards reaches the segment. All tangent components are cut, by the absence of cycles; their endpoint messages and the m non-outward labels determine histories as above. For fixed segment data the combined safe bound $(3zA)^m$ absorbs frame choices, label choices and the at most m segment length possibilities already bounded. This proves (47).

Finally follow strictly descending spokes from o to the initial exit, and separately strictly ascending spokes to an exit. Their terminal internal sites are on S_n, reached after at least n steps. The rim height range is at least $2n$. The same changing-edge injection gives $m \ge 4n$. This uses only the finite disk, not a hypothetical extension of its frame. □

Theorem 32 (Reweighted triangular late-cutoff gap)**.** *There is an effective reweighting of Theorem 26's template, preserving its qudit identities and all normalized disk isometries, such that nonhalting gives $f_n = 1$ at every cutoff and halting gives*

$$f_n \le 1 - \frac{1}{16384} = \frac{16383}{16384} \quad \textit{at every sufficiently late cutoff.}$$

The local dimension is unbounded across instances and fixed within each one. The same concentric target and one fixed sector loss witness all late failures of a halting instance.

Proof. Remove the common live factor K_0^B and the common base, route and straight-wire scalars. For $x \neq o$ there are at most z^h possible roots at distance h. With Lemma 31 and (46), the sum of absolute assignment amplitudes for all noncentral roots is bounded by

$$2z \sum_{h\geq 1} z^h \sum_{m\geq 4h} r^m = \frac{2z}{1-r} \sum_{h\geq 1} (zr^4)^h < 1. \tag{48}$$

The geometric tail factor $(1-r)^{-1}$ is included explicitly; dropping it is not an upper bound. For root-free histories, $z^{2n} \leq z^{m/2}$ for $m \geq 4n$, so their sum of absolute amplitudes is at most

$$2 \sum_{m\geq 4n} (3z^{3/2}r)^m < 1. \tag{49}$$

Both strict bounds follow directly from $z = 32$, $r = 1/(100z^2)$; for example $3z^{3/2}r = 3/(100\sqrt{z}) < 1/100$. Sums of absolute amplitudes bound Euclidean norms even if assignments have the same boundary tuple. The centered norm is $\sqrt{2z} = 8$. The separate empty tuple has relative amplitude at most one. Auxiliary sectors are absent for $n \geq 1$. The triangle inequality bounds the full relative norm by $8 + 1 + 1 + 1 = 11$, and its squared normalization by 121, independently of logical input. Whole records still determine all phases, so the disk Gram remains scalar identity.

Nonhalting decoders in Theorem 26 depend on supports and phases, which reweighting has not changed. They therefore remain exact at every cutoff, including the separately corrected auxiliary seed sectors.

After halting fix the center, an erased seed sector and a permanent arm in its interior. At every cutoff beyond return-home acceptance the two centered bit histories have relative amplitudes one and differ only in their erased terminal payload. Choose the erased observable O from (41), with norm one. For opposite target eigenstates and fixed $+1$ spectators, every common-completion contribution is nonnegative, and the displayed centered pair contributes at least $4/121$ after normalization. Thus their erased output states ρ_0, ρ_1 obey

$$D_E = \tfrac{1}{2}\|\rho_0 - \rho_1\|_1 \geq 2/121 > 1/64. \tag{50}$$

All five shares defining O are physically lost. No rare-outcome conditioning is used.

To convert this to a recovery bound, entangle the target's two opposite-eigenvalue basis states with a two-level reference and let $|\psi_i\rangle_{RE}$ be their pure encoded states. For any decoder the squared Bell overlap is at most

$$\begin{aligned} \frac{1 + \|\mathrm{Tr}_E |\psi_0\rangle\langle\psi_1|\|_1}{2} &= \frac{1 + \|\sqrt{\rho_0}\sqrt{\rho_1}\|_1}{2} \\ &\leq \frac{1 + \sqrt{1 - D_E^2}}{2} \leq 1 - D_E^2/4 < 1 - 1/16384. \end{aligned} \tag{51}$$

For completeness, the two diagonal contributions to Bell overlap are at most one each before division by four. The adjoint decoder is unital completely positive, so it contracts operator norm; pairing its image of a matrix unit with the cross matrix bounds each off-diagonal contribution by that matrix's trace norm. Polar decomposition, or maximizing overlap between purifications, gives the equality in (51). Choose purifications maximizing that overlap. Their

pure-state trace distance is the square root of one minus their squared overlap; partial trace contracts trace distance, giving $D_E \le \sqrt{1 - \|\sqrt{\rho_0}\sqrt{\rho_1}\|_1^2}$. Finally $\sqrt{1-x} \le 1 - x/2$. This proves all inequalities without an approximate-complementarity import. Since (17) minimizes over inputs, this one Bell input bounds f_n as stated at every late cutoff. □

C Buffered protocols and annular circuits

C.1 Finite endpoint control for unequal refinement lengths

The argument of Theorem 6 is not limited to five equal blocks. We record the extension needed to distinguish the triangular disk construction from common low-valence code circuits.

Proposition 33 (Expanding substitution controller)**.** *Suppose the actual physical open-leg types form a finite alphabet with a fixed finite cyclic seed and substitution σ, each parent emitting a consecutive block of between $m \ge 2$ and B children. Suppose each annulus has an effective bounded-depth, bounded-range local scalar-isometric circuit, with type-dependent local tables and effective coefficient field as in Theorem 6. Let a query's birth interval and loss endpoints lie within a fixed number C of fine cells of one parent, and let later endpoints follow exact descendant block starts. With any specified finite early-cutoff conventions, exact all-query recovery is decidable.*

Proof. Normalize a scalar gate $U^\dagger U = cI$ in the doubled pullback by division by c^2, which remains in the field and implements (10) without taking square roots. Use the same fresh-reference contractions as before. A fixed number of interval endpoints is assumed in the birth table, so a fixed number of joint collars suffices. Increase C to include $B, 1$ and the whole finite birth neighborhood. Let D_0, A_0, L be the round, anchor and lookup bounds in (8), including any fixed type-table context. Define

$$w = 10(L + C + B + 10), \qquad u = w + 2L + C + B + 10,$$
$$K = 10(B+1)(u + C + B + 10), \qquad H_1 = 20K + 10C + 100.$$

On the lifted line let $f(i)$ be the initial child of parent i. Then $f(i+1) - f(i) = |\sigma(a_i)|$. Fine offsets of size at most r about a center $f(i) + s$, $0 \le s \le C$, meet parent cells within $\lceil (r + C + B)/m \rceil + 1$. To cross h complete parent blocks requires at least mh fine sites, and the extra unit allows both partially crossed blocks. With two cut cushions the backward exception radius is at most

$$\left\lceil \frac{w + C + B}{m} \right\rceil + 3 + D_0 \le (w + C + B)/2 + 4 + D_0 \le w.$$

The chosen u covers all necessary anchors and tags. The largest fine offset needed is at most $B(u+1) + C + 1 < K$. The inverse window bound $\lceil (K + C + B)/m \rceil + 1 \le K$ also holds. These are the unequal-length counterparts of (9).

Store type windows of radius K about every labeled center, gaps capped at H_1, and the *whole type word* of each gap shorter than H_1. For a short circumference store its entire cyclic word. Windows alone would not determine the expansion of a short intervening gap, so the

gap words are necessary. An extension expands each short gap word to its morphic image, capping when its length reaches H_1. Long gaps remain long because $m \geq 2$. New windows are obtained by expanding stored parent windows and restricting to radius K around the selected child. The inverse bound proves that no unknown parent letters are needed. At a mark the only possibly long gap is outside the birth neighborhood; from a saturated circumference its length is at least $2H_1 - 2C \geq H_1$. Otherwise the exact old ring is stored. Resolve cyclic order, coincident centers and tags by actual indices on those small rings. A finite decoration whose updates use bounded type context can be included in the windows. Thus signatures and updates are finite and effective.

Generate an address by selecting a legal child from each parent block, then a birth mark and a loss from the finite table, then any number of extensions. Shared origins are resolved by the actual ownership rule. Every vertex's designated output port and its birth bulk input lie within the bounded mark neighborhood. Full swaps outside the collars pass through each gate; spectators outside give a common positive scalar. The target lies inside the mark patch. The same functional construction and marked subtraction as (16) therefore give a finite linear representation of the purity difference. The span algorithm in Theorem 6 terminates and decides every query, including the finitely contracted early exceptions. □

Corollary 34 (The two low-valence annular presentations)**.** *Exact all-cutoff recovery for any fixed finite buffered-interval protocol is decidable for the pentagon perfect-tensor code on its vertex graph* $\{4, 5\}$ *and for cyclic tensors on* $\{5, 4\}$ *satisfying the stated one- and two-vertex bulk identities, with algebraic entries or entries in the effective field of Theorem 6. Finite symmetry decorations are included when their flag tables and identities are effectively specified.*

Proof. For $p \geq 4$, ring vertices have type a with zero inward edges or b with one, two tangent edges, and $l = z - 2 - k$ outward spokes, $k = 0, 1$. An old fan emits the vertex word

$$(ba^{p-3})^{l-1}ba^{p-4}. \tag{52}$$

Between successive spokes of the same vertex a face has $p-3$ intervening new a vertices; across an old tangent edge it has $p - 4$. This proves the rule by gluing the missing regular polygons in the outward fans, with the same complete-development argument as in Section 8. The initial ring is $(ba^{p-3})^z$. An old spoke owns the new b it meets and the following intervening a's, with all their outward legs. It therefore emits a fixed finite word of actual open-leg types. For $\{4, 5\}$ a b already has two outgoing legs, giving minimum length two. For $\{5, 4\}$ the one-legged b after a spoke is followed by at least one two-legged a, again giving minimum length at least two. Maximum lengths are bounded by the finite fan table. A finite symmetry decoration updates by the fixed bounded flag walk across a fan, so it is local finite-state data.

It remains to exhibit full-input annular circuits. For $\{4, 5\}$ the perfect rank-six tensor is scalar isometric with any at most three inputs. Runs of b have length at most two by (52). Pick an a sink in each bounded a run and orient the tangent segments from a b source towards adjacent sinks. A b has at most one inward tangent in addition to its radial input, and an a sink has two tangent inputs and no radial one. Including the bulk gives at most three inputs per tensor. Directed paths have bounded length. Evaluating in that order, assigning unitary edge factors downstream, is a bounded local scalar-isometric circuit agreeing with the contraction on every input. The finite seed is handled directly.

For $\{5,4\}$ gaps between b's have one or two a's. In a two-a gap combine those a's into a two-vertex sink, with the exterior input legs adjacent to their common edge on the inner side. Make the b just before such a gap a one-vertex source using its radial input. In a one-a gap pair its preceding b with that a: the inputs are the radial spoke at b and the tangent on the right at a, which are adjacent to the common bond on its inner side. This pair sends information to the left. The source sends right to the long-gap sink and left through the consecutive short-gap pairs to the preceding sink. Every vertex lies in exactly one block. Old b's are isolated, old a's have two spokes, and at most two short gaps can be consecutive by (52); the initial ring has only long gaps. Thus the path depth and range are bounded. Every singleton or pair includes its independent bulk inputs in precisely the stated identity. Unit edge factors can be moved between complete input/output legs without changing the Gram. This is an all-input circuit, not an argument on only the disk image.

A fixed buffer, a fixed finite number of refinements before choosing a birth interval, or a bounded local offset in its endpoints gives exactly the birth table of Proposition 33. It now decides all queries in either presentation. The diagrams and fan notation agree with Steinberg et al., Figs. 4–5 and Eqs. (8)–(13) [17]; the circuit decompositions are proved above. □

Proposition 35 (The literal small-footprint obstruction)**.** *If a nontrivial logical subsystem has a birth footprint with at most two nonempty loss sectors and each may be erased, exact recovery from all their complements is impossible. Hence the literal five-sector experiment fails on* $\{4,5\}$ *and* $\{5,4\}$ *for every code with independent bulk dimension greater than one satisfying the disk promise.*

Proof. With one nonempty sector one allowed loss leaves no physical output depending on the target. With two, the complementary retained outputs for their separate losses are disjoint. Fix spectators and apply both decoders to these disjoint systems. They would copy every target pure state into two factors. If $0 < |\langle u, v\rangle| < 1$, an isometric dilation of such copying would give

$$|\langle u, v\rangle| = |\langle u, v\rangle|^2 |\langle e_u, e_v\rangle| \le |\langle u, v\rangle|^2,$$

a contradiction. In the $p \ge 4$ fan above a b has $z-3$ outward spokes. Such vertices occur in the initial ring; for $\{4,5\}$ this is two and for $\{5,4\}$ one. Dividing them into five groups yields at most that many nonempty sectors. This birth-cutoff obstruction proves the assertion. It makes no assertion about a widened buffered footprint. □

C.2 Fixed dimensions and numerical variations

The fixed-architecture principle below has direct methodological overlap with Shiraishi, Theorem 1, Corollary 2 and Section III [52], which use Hilbert's basis theorem for finite-size frustration-freeness equations and a dimension-dependent, uncomputable cutoff. Here the equations impose subsystem recovery on the disk-isometry promise. The Noetherian principle is standard; its application to these recovery equations appears new.

Proposition 36 (A nonuniform finite test at fixed architecture)**.** *Fix the tiling, discrete port dimensions, symmetry decoration and computable interval protocol, and vary only finitely many numerical tensor entries. On the promise of nonzero scalar-isometric disks, some finite*

set of exact query equations suffices for all-cutoff recovery for this architecture. This is an existence result without a uniform procedure for extracting its cutoff across architectures.

Proof. Use the unnormalized Gram equations (4) at every finite query. In the real and imaginary tensor entries they are polynomials with rational coefficients. Let I be the ideal generated by all these equations in the polynomial ring on the finitely many entry variables. The Hilbert basis theorem, Atiyah–Macdonald, Theorem 7.5 [33], makes this ideal finitely generated: the coefficient field $\mathbb{Q}$ is Noetherian and there are finitely many variables. Each finite generator is itself a finite polynomial combination of the original query equations. Taking their finite union gives a finite subset of original queries generating I. Its maximum cutoff tests all query equations simultaneously, on the disk-isometry promise. This does not say that a finite initial stretch of unchanged ideals certifies that no later generator will appear. Theorem 26 varies the discrete local dimension; a uniform extraction of these finite cutoffs for its templates would decide nonhalting, contradicting its reduction. □

A prescribed common computable level schedule is a statement about input presentation. If it carries a halting enumeration to local matching rules, that is external information encoded in the schedule. It is not an additional completeness theorem for a single autonomous repeated bounded-depth circuit. Likewise, the local identities alone do not classify full-rim isometric annuli without a bounded circuit description; Proposition 27 shows why disk normalization cannot supply that missing hypothesis.

References

[1] A. Almheiri, X. Dong and D. Harlow, Bulk locality and quantum error correction in AdS/CFT, JHEP **04** (2015), 163. arXiv:1411.7041.

[2] F. Pastawski, B. Yoshida, D. Harlow and J. Preskill, Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence, JHEP **06** (2015), 149. arXiv:1503.06237.

[3] D. Kretschmann, D. W. Kribs and R. W. Spekkens, Complementarity of private and correctable subsystems in quantum cryptography and error correction, Phys. Rev. A **78** (2008), 032330. arXiv:0711.3438.

[4] G. Vidal, Entanglement renormalization, Phys. Rev. Lett. **99** (2007), 220405, arXiv:cond-mat/0512165; Class of quantum many-body states that can be efficiently simulated, Phys. Rev. Lett. **101** (2008), 110501, arXiv:quant-ph/0610099.

[5] G. Evenbly and G. Vidal, Algorithms for entanglement renormalization, Phys. Rev. B **79** (2009), 144108. arXiv:0707.1454.

[6] M.-P. Schützenberger, On the definition of a family of automata, Information and Control **4** (1961), 245–270. doi:10.1016/S0019-9958(61)80020-X.

[7] W.-G. Tzeng, A polynomial-time algorithm for the equivalence of probabilistic automata, SIAM J. Comput. **21** (1992), 216–227. doi:10.1137/0221017.

[8] S. Kiefer, A. S. Murawski, J. Ouaknine, B. Wachter and J. Worrell, On the complexity of equivalence and minimisation for $\mathbb{Q}$-weighted automata, Logical Methods in Computer Science **9**(1) (2013), 8. arXiv:1302.2818.

[9] N. J. Ross and S. Wesley, Cutoff theorems for the equivalence of parameterized quantum circuits, MFCS 2025, LIPIcs **345** (2025), 85:1–85:19. doi:10.4230/LIPIcs.MFCS.2025.85.

[10] P. Chuklanov, Parameterized equivalence checking of quantum circuit generators via weighted automata, preprint (2026). arXiv:2611.12049v1.

[11] M. Kliesch, D. Gross and J. Eisert, Matrix product operators and states: NP-hardness and undecidability, Phys. Rev. Lett. **113** (2014), 160503. arXiv:1404.4466.

[12] R. Cleve, D. Gottesman and H.-K. Lo, How to share a quantum secret, Phys. Rev. Lett. **83** (1999), 648–651. arXiv:quant-ph/9901025.

[13] A. S. Kahr, E. F. Moore and H. Wang, Entscheidungsproblem reduced to the $\forall\exists\forall$ case, Proc. Natl. Acad. Sci. USA **48** (1962), 365–377. doi:10.1073/pnas.48.3.365.

[14] J. Kari, The tiling problem revisited (extended abstract), MCU 2007, LNCS **4664** (2007), 72–79. doi:10.1007/978-3-540-74593-8_6.

[15] J. Morton and J. Biamonte, Undecidability in tensor network states, Phys. Rev. A **86** (2012), 030301(R). arXiv:1205.3315.

[16] G. Scarpa, A. Molnár, Y. Ge, J. J. García-Ripoll, N. Schuch, D. Pérez-García and S. Iblisdir, Projected entangled pair states: Fundamental analytical and numerical limitations, Phys. Rev. Lett. **125** (2020), 210504. arXiv:1802.08214.

[17] M. Steinberg, J. Fan, R. J. Harris, D. Elkouss, S. Feld and A. Jahn, Far from perfect: Quantum error correction with (hyperinvariant) Evenbly codes, Quantum **9** (2025), 1826. arXiv:2407.11926v2.

[18] G. Evenbly, Hyperinvariant tensor networks and holography, Phys. Rev. Lett. **119** (2017), 141602. arXiv:1704.04229.

[19] M. Steinberg, S. Feld and A. Jahn, Holographic codes from hyperinvariant tensor networks, Nature Communications **14** (2023), 7314. arXiv:2304.02732.

[20] P. Basteiro, G. Di Giulio, J. Erdmenger, J. Karl, R. Meyer and Z.-Y. Xian, Towards explicit discrete holography: Aperiodic spin chains from hyperbolic tilings, SciPost Physics **13** (2022), 103. arXiv:2205.05693.

[21] M. Nasu and N. Honda, Mappings induced by PGSM-mappings and some recursively unsolvable problems of finite probabilistic automata, Information and Control **15** (1969), 250–273. doi:10.1016/S0019-9958(69)90449-5.

[22] A. Paz, *Introduction to Probabilistic Automata*, Academic Press, 1971, especially Theorem 6.17 and the bibliographical notes crediting Nasu and Honda.

[23] A. Condon and R. J. Lipton, On the complexity of space bounded interactive proofs, FOCS 1989, 462–467. doi:10.1109/SFCS.1989.63519.

[24] H. Gimbert and Y. Oualhadj, Probabilistic automata on finite words: Decidable and undecidable problems, ICALP 2010, LNCS **6199** (2010), 527–538. doi:10.1007/978-3-642-14162-1_44.

[25] G. Rote, Probabilistic finite automaton emptiness is undecidable for a fixed automaton, MFCS 2025, LIPIcs **345** (2025), 86:1–86:18. doi:10.4230/LIPIcs.MFCS.2025.86. Full proof of Theorem 7 in Appendix A of arXiv:2412.05198.

[26] D. Elkouss and D. Pérez-García, Memory effects can make the transmission capability of a communication channel uncomputable, Nature Communications **9** (2018), 1149. doi:10.1038/s41467-018-03428-0.

[27] V. D. Blondel, E. Jeandel, P. Koiran and N. Portier, Decidable and undecidable problems about quantum automata, SIAM J. Comput. **34** (2005), 1464–1473. doi:10.1137/S0097539703425861.

[28] H. Derksen, E. Jeandel and P. Koiran, Quantum automata and algebraic groups, J. Symbolic Computation **39** (2005), 357–371. doi:10.1016/j.jsc.2004.11.008.

[29] J. Eisert, M. P. Müller and C. Gogolin, Quantum measurement occurrence is undecidable, Phys. Rev. Lett. **108** (2012), 260501. arXiv:1111.3965.

[30] M. M. Wolf, T. S. Cubitt and D. Pérez-García, Are problems in quantum information theory (un)decidable?, preprint (2011). arXiv:1111.5425.

[31] G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, 5th ed., Oxford University Press, 1979, Theorem 369.

[32] S. Basu, R. Pollack and M.-F. Roy, *Algorithms in Real Algebraic Geometry*, 2nd ed., Springer, 2006, Theorem 2.77 (quantifier elimination over real closed fields).

[33] M. F. Atiyah and I. G. Macdonald, *Introduction to Commutative Algebra*, Addison-Wesley, 1969, Theorem 7.5.

[34] P. Turakainen, Generalized automata and stochastic languages, Proc. Amer. Math. Soc. **21**(2) (1969), 303–309. doi:10.1090/S0002-9939-1969-0242596-1.

[35] A. Yakaryilmaz and A. C. Cem Say, Languages recognized with unbounded error by quantum finite automata, CSR 2009, LNCS **5675** (2009), 356–367. doi:10.1007/978-3-642-03351-3_33. Preprint: arXiv:0809.0073v2.

[36] A. Kondacs and J. Watrous, On the power of quantum finite state automata, FOCS 1997, 66–75. doi:10.1109/SFCS.1997.646094.

[37] T. Neary, Undecidability in binary tag systems and the Post correspondence problem for five pairs of words, STACS 2015, LIPIcs **30** (2015), 649–661. doi:10.4230/LIPIcs.STACS.2015.649.

[38] P. A. Abdulla, Y.-F. Chen, M. Hečko, L. Holík, O. Lengál, J.-A. Lin and R. S. Thinniyam, Parameterized verification of quantum circuits, Proc. ACM Program. Lang. **10**, POPL (2026), article 70. doi:10.1145/3776712.

[39] H. Miller–Bakewell, Finite verification of infinite families of diagram equations, EPTCS **318** (2020), 27–52. arXiv:1904.00706.

[40] Q. Huang, L. Zhou, W. Fang, M. Zhao and M. Ying, Efficient formal verification of quantum error correcting programs, Proc. ACM Program. Lang. **9**, PLDI (2025), 1068–1093. doi:10.1145/3729293.

[41] A. J. Ferris and D. Poulin, Tensor networks and quantum error correction, Phys. Rev. Lett. **113** (2014), 030501. arXiv:1312.4578.

[42] T. Farrelly, R. J. Harris, N. A. McMahon and T. M. Stace, Tensor-network codes, Phys. Rev. Lett. **127** (2021), 040507. arXiv:2009.10329.

[43] C. H. Bennett, Logical reversibility of computation, IBM J. Res. Dev. **17** (1973), 525–532. doi:10.1147/rd.176.0525.

[44] N. Shiraishi and K. Matsumoto, Undecidability in quantum thermalization, Nature Communications **12** (2021), 5084. arXiv:2012.13889.

[45] T. Cubitt, D. Elkouss, W. Matthews, M. Ozols, D. Pérez-García and S. Strelchuk, Unbounded number of channel uses are required to see quantum capacity, Nature Communications **6** (2015), 7739. arXiv:1408.5115.

[46] C. Schön, E. Solano, F. Verstraete, J. I. Cirac and M. M. Wolf, Sequential generation of entangled multiqubit states, Phys. Rev. Lett. **95** (2005), 110503. arXiv:quant-ph/0501096.

[47] D. Kretschmann and R. F. Werner, Quantum channels with memory, Phys. Rev. A **72** (2005), 062323. arXiv:quant-ph/0502106.

[48] J. Kari, V. Salo and T. Worsch, Sequentializing cellular automata, Natural Computing **19** (2020), 759–772. doi:10.1007/s11047-019-09745-7.

[49] W. Cheng, Angular k-uniformity and the hyperinvariance of holographic codes, preprint (2025). arXiv:2506.06577.

[50] M. Margenstern, The domino problem of the hyperbolic plane is undecidable, Bull. EATCS **93** (2007), 220–237. arXiv:0706.4161.

[51] T. Carette and E. Moutot, Aperiodicity in quantum Wang tilings, Theoretical Computer Science (2026); author preprint (2024). Author version.

[52] N. Shiraishi, Frustration-freeness is decidable but unprovable, preprint, 2026. arXiv:2611.07337v1.

[53] C. Grogan, 975 Wang tiles whose seeded half-plane tiling is equivalent to Con(ZFC), public construction (2026), compiled from E. Glazer's K230. Construction and documentation.

[54] V. D. Blondel and J. N. Tsitsiklis, The boundedness of all products of a pair of matrices is undecidable, Systems & Control Letters **41** (2000), 135–140. Author version.

[55] C. Bény and O. Oreshkov, General conditions for approximate quantum error correction and near-optimal recovery channels, Phys. Rev. Lett. **104** (2010), 120501. arXiv:0907.5391.

[56] P. Turakainen, On languages representable in rational probabilistic automata, Ann. Acad. Sci. Fenn., Ser. A I **439** (1969), 1–10. Publisher version.

[57] S. Ying, Y. Feng, N. Yu and M. Ying, Reachability probabilities of quantum Markov chains, CONCUR 2013. arXiv:1304.0060.

[58] E. Hrushovski, J. Ouaknine, A. Pouly and J. Worrell, Polynomial invariants for affine programs, LICS 2018. arXiv:1802.01810.

[59] K. Nosan, A. Pouly, S. Schmitz, M. Shirmohammadi and J. Worrell, On the computation of the Zariski closure of finitely generated groups of matrices, ISSAC 2022, 129–138. arXiv:2106.01853.

[60] S. Zheng, D. Qiu, L. Li and J. Gruska, One-way finite automata with quantum and classical states, preprint (2011). arXiv:1112.2022.

[61] S. Douglas and N. Mghirbi, The values of rational probabilistic finite automata are the left-c.e. reals, Zenodo (2026). doi:10.5281/zenodo.24142961.

[62] T. Heinosaari and T. Miyadera, Incompatibility of quantum channels, J. Phys. A: Math. Theor. **50** (2017), 135302. arXiv:1608.01794.

[63] H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa and B. Schumacher, Noncommuting mixed states cannot be broadcast, Phys. Rev. Lett. **76** (1996), 2818–2821. arXiv:quant-ph/9511010.

[64] A. Tarski, *A Decision Method for Elementary Algebra and Geometry*, 2nd ed., University of California Press, 1951.