hexagon:2611.01068v1

The Complexity of Choosing Local Bases for Zero-Error Quantum Readout

Viktor Diakin
Viramelis Laboratory, independent research

25 November 2026

Abstract

A graph compactly specifies a quantum memory. Its edges carry parities of vertex bits. Some parities are recorded as exact classical flags, while each remaining pair of edges selects one of four pure qubit states in two mutually unbiased bases. We ask whether one copy identifies the vertex word up to global inversion when each party fixes its measurement after learning the flags. Choosing suitable bases is Σ_2^{P}-complete. Checking specified bases is coNP-complete. Both results hold on simple bipartite graphs of maximum degree three whose data edges form a matching and whose flags form disjoint paths. Every instance in both reductions has a perfect adaptive decoder using a single additional broadcast bit to choose the data measurements. For unrestricted collective readout, the complementary problem is Conflict-Free Cut with paired edges, which gives a published source of its hardness. We also give a finite decision tree for adaptive local readout and report two exhaustive enumerations through six labelled vertices.

Keywords. Quantum state discrimination, polynomial hierarchy, conflict-free cuts, local measurements, zero-error decoding.

1 The input and the decision problems

A single early measurement can reveal the basis needed by an entire group of qubits. We study the complexity of deciding whether a network permits all data measurements to be chosen without such an early outcome. The source of computational difficulty is that an explicitly given graph can describe exponentially many messages, all of which must be decoded exactly.

Let $G = (V, E)$ be a finite undirected graph, $n = |V| \geq 2$. Edge occurrences are distinct. Intermediate descriptions may have parallel occurrences. Designated flag edges $F \subseteq E$ are singletons. The remaining edges are partitioned into p pairs (e_j, f_j). Pair membership is additional input and may connect edges without a common vertex. The input lists vertices, endpoints, flags and pairs explicitly. A message is an element of $X = \mathbb{F}_2^V/\langle \mathbf{1} \rangle$; fix a representative whose root bit is zero. Every message has positive prior probability. On an edge $e = vw$ write $d_e(x) = x_v + x_w$ in $\mathbb{F}_2$. A flag stores $|d_e(x)\rangle$. A data register for edge values (a, c) stores the rank-one density operator

$$R_{ac} = \tfrac{1}{2}\Big[I + \big((-1)^a X_{\mathrm{P}} + (-1)^c Z_{\mathrm{P}}\big)/\sqrt{2}\Big], \qquad B_s = \{R_{0,s}, R_{1,1+s}\}, \quad s \in \mathbb{F}_2. \tag{1}$$

Here $X_{\mathrm{P}}, Z_{\mathrm{P}}$ are Pauli matrices. Conditioned on a message, the records form a tensor product. This is the square geometry of a 2-to-1 qubit random access code [7, 8]. The objective is recovery of the whole message. The model assigns no dynamics or thermodynamic meaning to its labels.

Flags and access. Exact flag outcomes are public before the data measurements. Write $K = \{h \in X : d_e(h) = 0 \ \forall e \in F\}$. Each attainable transcript determines a known coset $x^0 + K$. A representative x^0 is obtained by propagating bits on a spanning forest of the flags. Conjugating qubit j by $Z_{\mathrm{P}}^{d_{e_j}(x^0)} X_{\mathrm{P}}^{d_{f_j}(x^0)}$ removes the offset in (1). Thus a decoder for K works on every transcript, and every exact decoder works on the zero transcript. In K labels are constant on each flag component, with $k = \dim K = c(V, F) - 1$. A central classical decoder may collect all outcomes. No pre-shared entanglement is available.

- $\mathsf{CHECK}(G, b)$ asks whether the specified product of binary measurements B_{b_j} succeeds on K with certainty. $\mathsf{FIX}(G)$ asks whether any local measurements, fixed simultaneously after the flags, succeed. Lemma 2 reduces this to $\exists b\, \mathsf{CHECK}(G, b)$.

- $\mathsf{LOCC}_{\mathrm{fin}}$ allows a finite number of rounds of local quantum operations and classical communication, one data qubit per party, and arbitrary local ancillas. $\mathsf{COL}(G)$ allows an arbitrary joint measurement.

- A completely predetermined apparatus fixes its measurements before receiving the flags. For a triangle with data (a, c) and flag $a+c$, the flag chooses B_{a+c} perfectly. An apparatus fixed before that announcement would have to separate opposite pairs in both bases on one qubit.

The restrictions concern dependence of settings on classical information. The work of inferring the final word from all recorded outcomes is unrestricted.

Theorem 1 (Complexity of fixed readout)**.** *Under deterministic polynomial-time many-one reductions,* FIX *is* Σ_2^{P}*-complete and* CHECK *is* coNP*-complete. The reductions produce simple bipartite physical graphs of maximum degree three. The flags are vertex-disjoint paths, allowing trivial paths, and the data edges form a matching. Every output graph, on both yes and no inputs, admits exact one-copy local readout with at most one additional broadcast bit after the flags and before all remaining data measurements.*

2 Which message pairs must a basis separate?

Put $\ell_j = d_{e_j}|_K$, $m_j = d_{f_j}|_K$, $t_j = \ell_j + m_j$, and $W(h) = \{j : \ell_j(h) = m_j(h) = 1\}$. The eigenstates in B_s are labelled by their first bit. Measuring a state of parity s in B_s reveals that bit. In the other basis both outcomes have probability $1/2$.

Lemma 2 (Pair separation)**.** COL *holds if and only if* $W(h) \neq \varnothing$ *for all nonzero* $h \in K$*. A fixed apparatus exists if and only if some* $b \in \mathbb{F}_2^p$ *satisfies*

$$\forall h \in K \setminus \{0\} \quad \forall u \in K \qquad \bigvee_{j \in W(h)} [\, b_j = t_j(u) \,]. \tag{2}$$

For given b, h*, failure at some* u *is equivalent to* $\mathbf{1} + b_{W(h)} \in \operatorname{im} T_{W(h)}$*, where* $T_W : u \mapsto (t_j(u))_{j \in W}$*.*

Proof. On site j, changing zero, one or both encoded bits gives squared overlap $1, 1/2$ or 0. Overlaps of products multiply. Exact discrimination of pure states requires mutual orthogonality: an effect identifying one state with probability one and another with probability zero has them in orthogonal eigenspaces. Orthogonal projectors completed on the complement attain the converse.

Two product outcome laws have disjoint support exactly when some factor separates them perfectly. For finite outcomes use the Cartesian product of supports. For general outcomes use positive densities relative to the sum of the local laws for the finite ensemble and take the product dominating measure. A qubit POVM separating an antipodal pair refines that

sharp basis: a positive effect with zero expectation on a ray vanishes on that ray. The same argument applies to its positive operator density almost everywhere. Such a measurement cannot separate an antipodal pair of the other basis. Replace each useful local instrument by its separating basis and choose either setting at the other sites. Shared randomness adds no power, since finitely many probability-one success events have intersection of full measure. For $u, u + h$ the separating sites are now precisely those specified in (2). Conversely each such site separates their classical outcomes. Failure at every indicated site gives the stated linear system, including the empty system. □

Failure of CHECK has a polynomial certificate (h, u). Verification tests the flags, a nonconstant difference and all p registers. For FIX guess b and universally quantify (h, u), evaluating the disjunction in polynomial time. These are the upper bounds in Theorem 1. When $K = 0$ every setting succeeds. When $K \neq 0$ and $p = 0$ every setting fails.

3 A circuit reduction with a readable control qubit

Start with $\exists x \in \mathbb{F}_2^m \, \forall y \in \mathbb{F}_2^r \; \phi(x, y)$ for an explicitly given Boolean circuit. This problem is Σ_2^{P}-complete [1, 2]: unroll the polynomial-time predicate in the class definition into a polynomial-size circuit with the input instance hard-wired. If $m = 0$, add one ignored existential bit. An empty universal block ranges over the singleton empty tuple. We allow constant and single-wire circuits.

3.1 Generating the sets that a colouring must split

Introduce $D = \{\star\} \cup \{v_i^0, v_i^1 : 1 \leq i \leq m\}$. Use the nonempty sets

$$\begin{aligned} P_i &= \{v_i^0, v_i^1\}, \\ H(x, y) &= \{\star, v_1^{x_1}, \ldots, v_m^{x_m}\} \qquad \text{when } \phi(x, y) = 0. \end{aligned} \tag{3}$$

A proper colouring uses both colours on every set. It gives the members of P_i different colours, so there is a unique x_i^* for which $v_i^{x_i^*}$ has the colour of $\star$. A violating y for x^* would generate a monochromatic $H(x^*, y)$. Conversely, if $\forall y \, \phi(x^*, y)$, give $\star$ and all $v_i^{x_i^*}$ colour zero and their partners colour one. Every generated $H(x, y)$ has a label of colour one since $x \neq x^*$. This establishes equivalence without listing the exponentially many possible sets.

3.2 Implementing the generator by pairs of edges

The circuit will act on the bits of a *message difference* h. Fix distinct logical vertices A, S, B; gauge $h_A = 0$. The intended case is $h_S = 1$, $h_B = 0$. Use the relation

$$R(i, j; k, l) \quad \Longleftrightarrow \quad (h_i = h_j) \;\vee\; (h_k = h_l).$$

Implement one relation by two distinct qubits with identical logical pairs $((i, j), (k, l))$. If R holds, neither belongs to $W(h)$. Otherwise both belong and share the same $t_j(u)$, so opposite settings separate this h for every u. Install a guard $R(z, A; z, S)$ for every logical vertex z except A, S, including B, inputs, unused inputs and auxiliary wires. For $h_S = 0$ the guards force the zero word. For $h_S = 1$ all guards hold.

A separate *data* qubit, the control, uses $((A, B), (A, B))$. Its zero-layer state is R_{dd}, $d = u_A + u_B$. Its canonical setting is 0. It handles $h_A \neq h_B$ for every u and stays out of $W(h)$ when $h_A = h_B$. Its value has not been announced among the flags.

For $h_S = 1$, A, S supply constants $0, 1$. For $z = \mathrm{mux}(w, a, b)$ with output a at selector one and b at zero, impose

$$R(w, A; z, a), \qquad R(w, S; z, b). \tag{4}$$

At $w = 1$ the first enforces $z = a$; at $w = 0$ the second enforces $z = b$. Topological evaluation gives the unique gate values for fixed inputs. Negation is $\mathrm{mux}(w, A, S)$ and conjunction is $\mathrm{mux}(w, a, A)$; translation takes linear size for a fixed bounded-fanin basis. Repeated operands, constant outputs and unused wires cause no exceptional assignment. A gate with identical branches can alias their common wire. If an edge is (z, z) its form is zero; any relation containing that equality is tautological and may be kept. The physical construction below assigns each occurrence its own two terminals.

In *pair mode* choose an index i with $m - 1$ selector bits: take the first one, defaulting to index m. Let q_i be the resulting one-hot indicator. Keeping a running prefix bit for the absence of earlier ones uses $O(m)$ multiplexers. For $m = 1$ set $q_1 = 1$. A mode bit v selects between pair mode (0) and violation mode (1). Impose $R(v, A; \phi, A)$, which says $v = 0$ or $\phi = 0$. The marked outputs are

$$H_\star = v, \quad H_{v_i^0} = \mathrm{mux}(v, \neg x_i, q_i), \quad H_{v_i^1} = \mathrm{mux}(v, x_i, q_i). \tag{5}$$

In pair mode the marked set is P_i. In violation mode it is $H(x, y)$ for a falsifying assignment. Every such set is attained by a consistent wire assignment. Outputs may share a wire; label occurrences are distinct.

For each $a \in D$ append one *label register* $((A, H_a), (B, H_a))$. When $h_A = h_B = 0$ it belongs to $W(h)$ exactly when $h_{H_a} = 1$. The required basis on any orientation word u is

$$t_a(u) = (u_A + u_{H_a}) + (u_B + u_{H_a}) = u_A + u_B. \tag{6}$$

The relations filter the differences h under consideration. They place no preparation constraint on u. The zero flag layer allows both values of $u_A + u_B$ for each fixed generated h.

3.3 The two implications, including arbitrary auxiliary settings

Lemma 3. *A proper colouring yields an accepted vector of settings. The label settings of every accepted vector form a proper colouring.*

Proof. Given a proper colouring, use it on label registers, B_0 on the control and B_0, B_1 on every doubled relation. For any $h \neq 0$ and any u, the control separates if $h_B = 1$. A violated guard, gate or mode constraint also separates by its doubled relation. In the remaining case $h_S = 1$ by nonzeroness and the guards. Evaluation marks one of the sets in (3). Its two colours include $u_A + u_B$, so a marked label separates. This exhausts (2).

Conversely, fix *arbitrary* auxiliary settings. For each P_i set $h_A = h_B = 0, h_S = 1$, choose pair mode, select i, take any x, y and evaluate the circuit. For each generated $H(x, y)$ use the same constants, violation mode, its generating x, y and any selectors. This gives a nonzero h with all relations true and the control absent from $W(h)$. Exactly the marked labels remain. If their common colour is c, choose $u_A = 0, u_B = 1 - c$, with all other logical u-bits zero. Then u and $u + h$ are legal, distinct and unseparated, whatever the auxiliary settings. Hence an accepted vector splits each set. Reading the labels of the star colour recovers x^* from the preceding colouring proof. □

This gives $\phi \mapsto G_\phi$ for FIX. For CHECK start from the coNP-complete circuit tautology problem, equivalently the complement of circuit satisfiability after output negation. Given $\psi(y)$ put $m = 1$ and $\phi(x, y) = \psi(y)$. Supply the canonical control and doubled settings and label colours $(b_\star, b_{v_1^0}, b_{v_1^1}) = (0, 0, 1)$. If ψ is a tautology only P_1 occurs and it is split. If $\psi(y_0) = 0$, the assignment $x = 0, y = y_0$ generates the colour-zero set $\{\star, v_1^0\}$. Choosing $u_A = 0, u_B = 1$ refutes the supplied vector. This is a many-one reduction to verification itself, with both implications independent of an oracle for synthesis.

4 A physical graph, an adaptive protocol and an example

Lemma 4 (Physical realization)**.** *A logical description with L vertices and p registers, allowing repeated occurrences and loops, has a simple bipartite realization of size $O(L+p)$ and maximum degree three. On the zero flag layer its message space, edge forms and setting certificates are exactly the logical ones.*

Proof. Create two new physical terminals for each of the $2p$ data-edge occurrences; join them by that data edge and assign their bipartition colours $0, 1$. Label the terminals by the respective logical endpoints. For each logical vertex z, order its q terminals. If $q = 0$, create one isolated representative. If $q = 1$, its only terminal suffices. Otherwise join consecutive terminals by flag paths of length two for equal colours and length three for opposite colours. Use fresh internal vertices, alternating colours along each path. All vertices on these paths represent z.

There are $4p$ terminals and at most $2(4p)$ internal vertices besides isolated representatives. The flag graph is a union of paths. Each terminal is incident to exactly one data edge and at most two flags; each internal vertex has degree two. Each edge has opposite colours at its ends. Data edges are a matching; every flag segment uses a fresh internal vertex, so no flag duplicates a data edge or another flag. For a logical loop the two terminal ends are distinct physical vertices on the same flag path. Multiple occurrences of the same logical edge have disjoint terminal pairs.

Choose the physical gauge at a representative of A. Copying each logical bit to all vertices representing it gives a linear bijection onto the zero flag layer: flag equalities force exactly this constancy and identify no distinct logical vertices. Each physical data difference is the pullback of the logical one, including zero for a loop. The bijection preserves $W(h)$ and $t_j(u)$ in both directions for every h, u. The other flag transcripts are locally equivalent by the Pauli correction in Section 1. Decoding all physical labels, including inserted vertices, follows from logical decoding and the known transcript. Construction and translation of a supplied b take polynomial time. Data edges in a common register need not share a physical endpoint. □

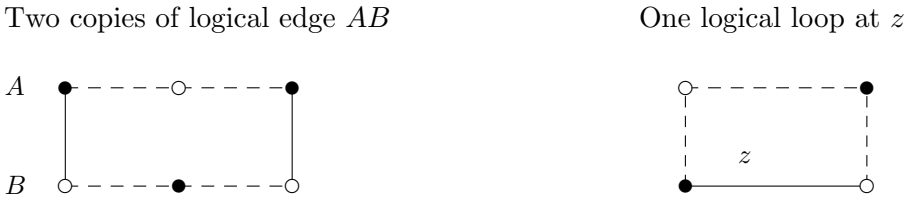


Figure 1: Two local examples of the physical realization. Solid segments are data edges and dashed segments are exact flags. Filled and open circles give the two bipartition colours. At left each horizontal path is one logical variable and the two vertical data edges form the control register. At right the flag path identifies all four vertices, so the solid data edge has zero difference on the zero layer.

What one broadcast bit buys. After the flag transcript and local corrections, measure the control in B_0. Broadcast its deterministic outcome $s = u_A + u_B$ once to all parties needing it. Measure every label register in B_s; use opposite bases on the copies of every relation. All remaining settings can be chosen simultaneously. Messages with different s already have disjoint control outcomes. For two distinct messages in one branch the difference has $h_A = h_B$. A violated relation separates via its two qubits. Otherwise there is a legal marked set; it contains at least one label and all its settings equal the required s. This works even if no proper colouring exists.

The unit counted is one bit on a public broadcast channel *in the measurement-selection phase following the flags*. The protocol also makes flags public and sends final local outcomes to the classical decoder. A point-to-point implementation may transmit a copy to several sites. Total communication and classical decoding cost are separate resources. Some yes instances need zero adaptive bits. The entire image of each reduction admits this protocol; deciding the promise itself is unnecessary. This completes Theorem 1.

Worked reduction. Let $\phi(x, y) = x \vee y = \mathrm{mux}(x, S, y)$, with one bit of each kind. The true witness is $x = 1$. There is one pair set, and only $x = y = 0$ generates a violation set. The two rows below are the edges of a path on three label vertices.

Generator event	$\star$	v_1^0	v_1^1	Required inequality
Pair mode		•	•	$b_0 \neq b_1$
Violation at $(0, 0)$	•	•		$b_\star \neq b_0$

Thus $(0, 1, 0)$ and $(1, 0, 1)$ are exactly the possible label colourings of a successful vector. For $(0, 0, 1)$ the violation row is monochromatic; choose h from that row and $u_B = 1, u_A = 0$ as in Lemma 3. All surviving label qubits are measured on the wrong axis. The archived compiler uses $q_1 = S$, the OR gate, the negation gate for x, and the two output selectors in (5). It realizes ten logical vertices, seventeen doubled relations, one control and three labels, giving 38 data qubits, 295 flag edges and 305 physical vertices. Testing all eight colourings against these physical edges gives the two above.

5 Collective readout and the prior cut problem

Corollary 5. COL *is* coNP*-complete. Hardness already holds without flags on simple connected graphs of maximum degree five. With flag paths allowed it also holds on simple bipartite physical graphs of maximum degree three with a matching of data edges.*

Proof. A nonzero word in K defines a nontrivial cut crossing no flag. It witnesses failure of COL exactly when the cut crosses at most one edge per pair. This is a polynomial certificate. Equivalently, some selection retaining the flags and one edge from each pair is disconnected: given the cut choose an edge outside it from every pair; conversely take a component of a disconnected selection. Nonorthogonal words remain nonorthogonal under any finite tensor power. When COL holds, two copies suffice with fixed settings B_0 on one and B_1 on the other, since one register separates each same-flag message pair in its own basis.

For hardness use the theorem of Rauch, Rautenbach and Souza [3]. On a connected simple host of maximum degree five, deciding whether some cutset is independent in a supplied 1-regular conflict graph on its edges is NP-complete. Use the very same host with no flags; the conflict edges partition its edge occurrences into pairs. If deleting an independent cutset disconnects the host, the edge boundary of a component is contained in that cutset and is independent. Conversely an independent nontrivial boundary itself disconnects the host. The complementary answers agree exactly. This is an identity reduction on flag-free inputs up to encoding the pairs. Lemma 4 transfers both answers to the flagged physical class. These COL reductions carry no promise of adaptive success. □

Direct reduction for the flagged formulation. For an independent check of the computational statement, here is a reduction directly from circuit unsatisfiability, without a conflict graph theorem. Use only logical constants A, S, the circuit inputs and one vertex per multiplexer output. Put *one* register, with edges $(i, j), (k, l)$, for each relation $R(i, j; k, l)$ in (4) and for each guard $R(z, A; z, S)$. Require the output wire o to be one by adding $R(o, S; o, S)$. In this construction a register is absent from $W(h)$ exactly when its relation is satisfied. A

satisfying circuit input extends by evaluation to a nonzero cut h with $h_A = 0, h_S = 1$ and $W(h) = \varnothing$. Conversely a nonzero witness with empty $W(h)$ obeys all guards, so $h_S = 1$; by the gate relations its input gives circuit output one. If $o = A$ the final relation excludes the sentinel-one case; if $o = S$ it is tautological. With no input wires the constant cases still have the two logical constant vertices, even if one needs an isolated representative in the lift. Thus the circuit is unsatisfiable exactly when the lifted graph belongs to COL. This proof does not alter the combinatorial priority above.

Contracting the flags gives the cut test on a logical multigraph, retaining edge occurrences, loops and isolated vertices. A disconnected graph with two or more components already has an empty witnessing boundary. A single logical vertex has no nontrivial cut. The flag-free hardness is a published combinatorial result. Kalyanasundaram and Kumar study related planar and bounded-degree cases [4, 5]; their November preprint gives a polynomial algorithm for matching conflicts at host maximum degree four. Flag contraction can increase degree. Indeed, on a flag-free subcubic graph on $n > 4$ vertices any transversal retains $p \le 3n/4 < n - 1$ edges, so COL always fails. It is essential to specify which graph satisfies the degree bound. A rainbow cut gives the same obstruction when each pair is a colour class. Asking for rainbow separation of every designated pair of vertices has different quantifiers [6].

6 Adaptive measurements, separability and limits

A separable POVM has effects that are sums of positive product operators. For a fixed finite set of guessed labels, $\overline{\mathsf{LOCC}_{\mathrm{fin}}}$ is the closure of finite-round LOCC measurements in the finite-dimensional space of effects. Positive priors make vanishing average error equivalent to vanishing error on every label. The inclusions $\mathsf{LOCC}_{\mathrm{fin}} \subseteq \overline{\mathsf{LOCC}_{\mathrm{fin}}} \subseteq \mathsf{SEP} \subseteq \mathsf{ALL}$ follow from refinement into product Kraus operators and closedness of the finite-dimensional separable cone [12, 13]. The finite decision tree below addresses the first class.

Theorem 6 (Two-basis decision tree)**.** *Given candidates $L \subseteq K$ and available sites J, a finite exact local decoder exists precisely under the following recursion. Accept for $|L| \le 1$. Otherwise choose $j \in J, s \in \mathbb{F}_2$ and require acceptance on* both *children*

$$L_o = \{u \in L : t_j(u) \neq s \text{ or } \ell_j(u) = o\}, \qquad J' = J \setminus \{j\}, \qquad o = 0, 1. \tag{7}$$

Reject if there is no successful choice. From $(K, \{1, \dots, p\})$ the depth is at most p, and the graph decision problem lies in PSPACE*.*

Proof. For sufficiency implement the chosen projective measurement; (7) is exactly its positive Born support. For necessity the residual words must be orthogonal if there is more than one candidate. Theorem 2 of Duan, Xin and Ying [9] gives a protocol measuring individual qubits projectively in an adaptive order for any LOCC-distinguishable orthogonal multiqubit ensemble. Its hypothesis applies to each finite residual list; coincident residual states already preclude identification of those two labels. On a product word all other qubits are unchanged by a local measurement. A projective basis different from both alphabet bases has both outcomes with positive probability for every candidate: a qubit ray has a unique orthogonal ray. Pick either successful continuation, omit that measurement and supply its known postmeasurement state locally if needed. In an alphabet basis the surviving alternatives are precisely (7); the measured factor is again a known local state. Induction gives the tree. For space store a depth-first branch of at most p triples (j, s, o), polynomial-size counters, and the current candidate while enumerating K to decide whether at most one survives. Recompute supports from the branch. The workspace is polynomial. □

Shifts on five vertices. Put flags on 03 and 14. Pair 01 with 12, 24 with 02, and 23 with 34. On the zero flag layer their values are $(a, b), (b, c), (c, a)$, $c = a + b$. The rotation $R_y(-\pi/4)$

takes rays $00, 11, 01, 10$ of (1) to $0, 1, +, -$ respectively. A further Z_{P} interchanges $+, -$ and preserves $0, 1$ as rays. Up to order the words become

$$|000\rangle, \qquad |1{+}{-}\rangle, \qquad |{-}1{+}\rangle, \qquad |{+}{-}1\rangle.$$

Each string denotes a tensor product. This is the three-qubit Shifts family [10]. Each party is the unique orthogonality site for one pair in the $0, 1$ basis and another in the $+, -$ basis. A first local effect E must have vanishing off-diagonal element in each of these bases to preserve orthogonality on each surviving pair. If it eliminates a word, positivity gives the same zero. Thus E is scalar. One may refine outcomes into individual Kraus operators: error contributions are nonnegative and the continuation can ignore the refined label. Each nonzero M obeys $M^\dagger M = cI$ and is proportional to a locally reversible isometry on the input support. A finite successful protocol has no first informative step. This realizes a known obstruction on the graph.

Here a stronger closure statement follows from a separate published criterion [14, Thm. 3]. On each party the unique-site pairs and their reversals supply the dyads $|0\rangle\langle 1|$, $|1\rangle\langle 0|$, $|+\rangle\langle -|$, $|-\rangle\langle +|$. They span the three-dimensional traceless subspace: the first two span the off-diagonal matrices and the sum of the last two is Z_{P}. Cohen's local-span hypothesis is satisfied at all three sites, so Shifts cannot be read perfectly in $\overline{\mathsf{LOCC}_{\mathrm{fin}}}$. Restricting an exact decoder, or an approximating sequence, to the four zero-flag messages would give a Shifts reader. Its fixed flag states can be prepared locally. The same argument applies to the prism with flags $03, 14, 25$ and pairs $(01, 45), (12, 35), (02, 34)$; its zero layer has the same three factors.

A separable measurement nevertheless reads all four Shifts words. Use the rank-one projectors of this complete product basis:

$$000, \quad 1{+}{-}, \quad {-}1{+}, \quad 1{-}{-}, \quad 10{+}, \quad 01{-}, \quad 001, \quad {+}1{+}. \tag{8}$$

Every two entries have an orthogonal local pair; eight orthonormal vectors exhaust dimension eight. Decode the first three outcomes to themselves and the last five to $|+-1\rangle$, orthogonal to the first three. Coarse-graining preserves separability. To cover every flag layer, tensor the locally corrected per-layer effects with their product flag projectors and collect outcomes by message. This is the product-completion construction of [11] in explicit coordinates. Thus finite-LOCC failure and collective success in the census do not imply SEP failure. The remaining members of that class have not been classified here for closure or SEP.

7 Context, priority and open questions

Stockmeyer's hierarchy and Hemaspaandra's alternation-bounded classifications [1, 2] explain the quantified source problem. They do not identify which quantified predicates a graph of edge differences can realize. Conflict-Free Cut [3] identifies exactly the collective obstruction on unflagged paired hosts. Its witness is a single cut. For fixed readout one must further exhibit a reference word whose bases simultaneously miss all witnessing registers, and for synthesis one setting vector must work across all reference words and cuts. The compiler above realizes this extra alternation despite the shared incidence structure.

One-way LOCC can be studied by graphs of orthogonality among explicitly indexed states [18]; in our input a vertex indexes a hidden bit and a cut indexes a difference between states. In discrimination with postmeasurement information [15, 16] the time at which classical information is revealed changes decoding ability. The distinction is operationally relevant to our flags and control. Croke and Barnett [17] study communication-dependent discrimination of product ensembles. Neither communication savings alone nor the use of incompatible bases constitutes the complexity result. Noncommutative graphs [19] concern zero-error channel communication, and many-copy restricted hypothesis testing [20] concerns a different asymptotic objective.

Our search included quantified Boolean constraint languages, implicit set splitting, conflict-free and rainbow cuts, robust connectivity, nonadaptive product-basis discrimination, postmeasurement information and graph methods for LOCC. The closest exact identification concerned COL. In the sources examined through 25 November 2026 we did not identify a theorem giving Σ_2^{P}-hardness of graph-specified fixed-basis synthesis with the one-broadcast promise. This is the scope of a literature comparison, not a claim to worldwide priority. The complexity of the finite adaptive decision problem beyond its space upper bound, and the complexity of FIX when the two edges of each register must share an endpoint, remain open here. The pure product and exact-flag assumptions matter: full-rank noise destroys the orthogonality needed for exact discrimination in a single flag layer.

A Structure and an explicit-list comparison

Write $T : u \mapsto (t_j(u))_j$. Settings differing by $T(v)$, $v \in K$, are equivalent by substituting $u + v$ in (2). Enumerating $2^{p-\operatorname{rank} T}$ representatives, then $2^k - 1$ differences and one linear system per difference yields a constructive exact algorithm. There is also an algorithm fixed-parameter tractable in k: after contraction at most $1 + \binom{k+1}{2}$ distinct edge forms remain. Registers of the same unordered pair type have the same required basis and the same set of separating differences. If the type occurs at least twice, fixing two copies in opposite bases covers everything that type can ever separate. Thus it suffices to enumerate settings on singleton types, at most $O((k+1)^4)$ of them. This gives $2^{O((k+1)^4)} \operatorname{poly}(|G|)$ time, including the search over differences. The polynomial algorithm for all compact instances would imply $\mathsf{P} = \Sigma_2^{\mathrm{P}}$.

In one structured class, registers with identical forms (q_i, q_i) can all be set to zero. Suppose every remaining register has the same nonzero parity t on K, with t spanned by the q_i. Let $L = K \cap \bigcap_i \ker q_i$. Every difference outside L has been covered. On L the two forms of each remaining register agree, while $t(u)$ takes both values on K. Condition (2) therefore requires each colour class of remaining forms $\ell_j|_L$ to span L^*. This is equivalent to two disjoint bases of the represented binary matroid: take a basis from each spanning class, or extend two disjoint bases by colouring leftover elements arbitrarily. Matroid union with rank computations over $\mathbb{F}_2$ solves this case in polynomial time [21]; at $L = 0$ empty bases suffice.

Explicit lists. Suppose instead that the input explicitly enumerates arbitrary words over the four local rays, without advance side information. Checking b is polynomial: for each pair of words form a clause over settings, using one literal per antipodal site with its required basis. Synthesis is NP-complete even for an orthogonal list. For a 3-CNF with nonempty clauses, no tautologies and distinct variables in each clause, create two words per clause. Give both the same binary clause-index tag on additional sites in B_0. For each variable occurring positively (negatively) in that clause place antipodal rays in B_1 (B_0); for all absent variables put the same ray on both words. Within each clause the words are orthogonal; different clauses have orthogonal tags. A satisfying assignment selects the variable bases and B_0 on tags, separating every pair. Conversely the identical tags cannot help on the two words of a clause, so any successful settings satisfy every clause. This is a polynomial construction and the explicit pair test proves membership. Providing the full list changes the input size and changes the classification.

B A direct finite-round proof for the four-ray alphabet

For completeness we give a proof of the normal form in Theorem 6 that does not invoke the projective-reduction theorem. The argument also explains why a local filter need not itself be projective. Fix a finite set of product words over the alphabet. Refine local instruments into single Kraus outcomes and include local ancillas in their output spaces. The probability of

a transcript on a word factors into expectations of the cumulative input effects $F_i = M_i^\dagger M_i$ at its sites. Refinement preserves exactness since errors sum nonnegatively. Ignore zero-trace effects. Normalizing each F_i by its trace gives a Bloch ball coordinate. At the next action by party i, this coordinate is the convex average of its children, all other coordinates staying fixed. Conversely any such decomposition of F_i is implementable there: if $F_i = \sum_r F_{ir}$, use $F_i^{-1/2} F_{ir} F_i^{-1/2}$ on its support, transporting by the isometry in the polar decomposition of M_i and completing arbitrarily on the unused output subspace. Scale factors affect branch probabilities but not which messages survive.

Let D_0 be the Bloch ball, with centre O, and let Q be the square formed by the four alphabet rays. A normalized local effect kills a pure input exactly when it is the antipodal vertex of that input. The set of terminal tuples that leave at most one possible message is a finite union of boxes (Cartesian products). To describe a box, choose a possible label to keep and for each other label choose a site that kills it. A free coordinate is D_0 and a prescribed coordinate is one of the four vertices; inconsistent choices give an empty box. Empty candidate sets may always be accepted.

Induct on allowed rounds. From child boxes used in a split on site i, form a parent box by intersecting their factors at every other site and taking the convex hull of their site-i factors. Every point in this box admits such a split, discarding zero weights, by the local implementation above. Conversely a successful split belongs to one such parent box. The factors that can arise are D_0, the empty set and convex hulls of subsets of the four vertices and O. This family is closed under the required operations: proper polygons lie in the square and their boundary segments lie on its sides or diagonals, whose only intersections are vertices and the centre. A hull involving D_0 is D_0. Thus each bounded-depth winning set is a finite union of compact boxes, and a witnessing construction may be retained for each. Even a general outcome space gives no further locations: at fixed other coordinates the allowed child points lie in a finite union of compact factors; their barycentre is a finite convex combination by Carathéodory's theorem. Null branches for the finite list may be discarded; a set of children killing every candidate is already terminal. Binary mixtures can implement any finite convex combination, though the following simplification needs only two outcomes.

Take a winning construction at a tuple of centres or vertices, as at the all-centre root. If the tuple already lies in some child box, follow that child's construction without measuring. Otherwise the changing coordinate cannot be a vertex, which is extreme in the ball. It is O, no participating child factor contains O, and their joint convex hull does. None is D_0; a subset of the four corners containing no opposite pair lies on one side of the square. The union must therefore contain opposite vertices. Measure their basis, moving the centre to each of those vertices with its corresponding child construction. Unchanged coordinates satisfy the required intersections. Induction on the height of the witnessing construction yields a successful tree. A coordinate fixed to a vertex never requires a further informative move. After at most p such moves the resulting physical protocol has precisely the form (7). This argument treats finite height; passage to a closure would require a separate result.

C Small graphs, exact counts and reproducibility

For each simple labelled host include all subgraphs, flags and unordered pair partitions of the remaining edges. Isolated vertices are allowed. Table 1 counts layouts, not isomorphism orbits. The first three classes have exact collective readout on one copy. The third class fails finite LOCC; SEP and closure are not tested by these counts.

An analytic check through four vertices. For $q = \binom{n}{2}$ possible edges, the number of layouts satisfies $N_0 = 1$, $N_1 = 2$, $N_q = 2N_{q-1} + (q-1)N_{q-2}$: a marked edge is absent, a flag or paired with another edge. Disconnected hosts fail COL. On a tree only the all-flag case

n	All layouts	FIX	$\mathsf{LOCC}_{\mathrm{fin}} \setminus \mathsf{FIX}$	$\mathsf{COL} \setminus \mathsf{LOCC}_{\mathrm{fin}}$	Fails COL
2	2	1	0	0	1
3	14	7	0	0	7
4	499	218	24	0	257
5	123 109	46 878	9 870	1 224	65 137
6	266 906 858	88 900 744	28 644 060	2 613 570	146 748 484

Table 1: Exact classifications by the least resource among the specified classes. The final column stays impossible at every finite copy number.

succeeds; selecting one edge from a pair loses connectivity. On a connected unicyclic host, the transversal condition permits at most one pair and requires its two edges to lie on the cycle. In the zero flag layer they have the same form, as the flag path components identify their endpoints across the two arcs. Thus a cycle of length c contributes $1 + \binom{c}{2}$ successful layouts including all flags. On two and three vertices the totals are one and seven. On four vertices there are 16 labelled trees, three four-cycles and twelve triangles with a pendant vertex, giving $16 + 3 \cdot 7 + 12 \cdot 4 = 85$ successful layouts.

In a diamond (K_4 minus an edge) every edge lies on a cycle. With zero or one pair this gives $1 + 10$ successes: if the flags alone do not connect, both paired edges cross the same bipartition since deleting either one leaves the host connected. With two pairs there is one flag. A flag joining the two degree-three vertices leaves forms x, x, y, y for independent x, y. Testing the two coordinate cuts forces the identical forms to be paired, giving one fixed success. A peripheral flag leaves forms $x, x, y, x + y$. The cut $h_x = 0, h_y = 1$ forces y and $x + y$ to be paired. The unique resulting layout reads x from (x, x) then y in B_x. For that cut only the second register separates, and it requires both settings as u_x varies. Hence there are four adaptive-only layouts per diamond, six labelled diamonds, and 16 total collective successes per diamond.

On K_4, zero or one pair gives $1 + 15$ fixed successes since removing two edges leaves a connected graph. Two pairs leave two flag edges and a one-dimensional K. At least three of the four remaining edges cross the unique nontrivial cut (three for adjacent flags and four for disjoint flags). Any pairing puts two crossing edges together; their forms agree on K, giving a separator in B_0. This accounts for $\binom{6}{2} \cdot 3 = 45$ further successes. Three pairs with no flags cannot cover all four singleton cuts, since a pair of distinct simple edges covers at most one. Altogether there are $85 + 6 \cdot 16 + 61 = 242$ collective successes on four vertices. Over 74 labelled hosts on two through four vertices the 515 layouts thus split into 226 fixed successes, 24 requiring adaptive one-copy readout and 265 failures at every finite copy number.

When the zero setting fails. Take flags $12, 34$ and registers $(01, 02), (03, 13), (04, 14)$. The zero-layer labels are $(0, a, a, c, c)$, with local pairs $(a, a), (c, a + c), (c, a + c)$. Changing a alone forces setting 0 on the first register. Changing only c requires at least one of the last two settings to equal each possible a. Conversely these conditions suffice for every difference. Exactly $001, 010$ work. This example is minimal in both vertices and edges for a simple host. To see the edge bound, first observe that for $p \leq 2$ every feasible layout admits $b = 0$. A zero row of T is optimally set to zero. One nonzero parity row, or two independent ones, can be shifted to zeros by $T(K)$. If the two rows are equal and nonzero, choose h with both parities one; $W(h)$ is empty. The preceding enumeration shows zero suffices for all feasible layouts with $n \leq 4$. At $n = 5$, a potential counterexample with at most seven edges must have $p = 3$, seven edges and one flag, since a spanning transversal needs $|E| - p \geq n - 1$. There are eight messages and eight binary outcome strings, so each message must have a singleton support. The zero word lies in B_0 at every site, forcing all three settings to zero. For $n \geq 6$, $p \geq 3$ already implies $|E| \geq 8$. These facts give both minima.

What was reproduced. The first full enumeration used C++ and orbits of flag graphs with cut and linear-support tests. The second used Python, orbits of the footprint and pairing of data edges, and positive Born supports of pairs of words. Orbit sizes restored every labelling. Both full output files agree in all four columns and in the number of fixed successes requiring a nonzero vector (630 at five vertices, 2 247 300 at six). They were written in the same AI-assisted project and constitute two algorithms, not two independent research groups.

For this revision we unpacked the archived code, rebuilt C++17 and recomputed its census through five vertices. We reran both Python implementations through four, the named examples, 88 solver queries on 24 circuit truth tables and 16 small discriminations with a separate numerical optimizer. Each integer output reproduced the archived result in its tested range. The largest reproduced primal–dual gap was 1.286×10^{-7} and all 16 intervals overlapped their archived counterparts within the numerical tolerance. The noisy tests are regression checks, outside the hypotheses of the exact complexity theorems. We compared the independently generated archived six-vertex outputs; the 266 906 858 layouts were not re-enumerated for this revision. New tests checked all eight label colourings of the worked circuit, five constant or degenerate fixed-reading circuits, the direct collective reduction on all 22 truth tables with zero through two inputs, and the product-basis identities in (8). The supplement includes the original archive unchanged, its SHA-256 digest, an English replay guide, the verification script and its machine-readable output. No quantum device was used.

Research responsibility and AI assistance

Viktor Diakin initiated the research, formulated the questions, defined the scope of the hardware-free investigation and coordinated the exchange and review of results. AI systems contributed substantially to the construction, proofs, computations, literature review and writing. The following attribution records specific contributions reported in the research log and by the project lead.

- **Astra (OpenAI, GPT-6.2)** developed the graph model, original criteria, complexity reduction, adaptive analysis, enumeration software and drafts. The final proof expansion, reference comparison, targeted replay and English preparation form part of this AI-assisted workflow.
- **Mira (OpenAI, GPT-6.2)** audited the earlier criteria, performed a separate small combinatorial enumeration, verified the five-vertex setting example and maintained the research record.
- **Gemini Neon High (Google)** supplied the five-vertex counterexample to zero settings, identified the fixed-flag subensemble issue for the prism, raised the complexity question and proposed checking collective-readout hardness. The latter has the published cut-theoretic antecedent credited in Corollary 5.
- **DeepSeek** (model version unrecorded) reviewed necessity and sufficiency of the earlier fixed-reading criterion, model assumptions and limitations, and the proposed validation controls. The lead also reports a review of the complexity-stage manuscript.
- **Qwen4-Max (Alibaba)** independently gave the small-layout classification and the 6×4 diamond count, and commented on purity, Shifts and references. The lead also reports a review of the complexity-stage manuscript.

The lead reports that the Gemini, DeepSeek and Qwen reviews of the preceding complexity-stage draft found no counterexample within its stated scope. These are model reviews of a preceding draft. They do not certify the final revision or constitute human peer review.

References

[1] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science* 3 (1976), 1–22. Publisher.

[2] E. Hemaspaandra. Dichotomy theorems for alternation-bounded quantified Boolean formulas. arXiv:cs/0406006.

[3] J. Rauch, D. Rautenbach and U. S. Souza. On conflict-free cuts: algorithms and complexity. *Information Processing Letters* 187 (2025), 106503. Publisher; arXiv:2311.01077.

[4] S. Kalyanasundaram and S. Kumar. Conflict-free cuts in planar and 3-degenerate graphs with 1-regular conflicts. Preprint, May 2026. arXiv:2605.12068.

[5] S. Kalyanasundaram and S. Kumar. Conflict-free cuts: exploring the boundaries of tractability. Preprint, November 2026. arXiv:2611.15846.

[6] Z. Huang and X. Li. Hardness results for rainbow disconnection of graphs. arXiv:1811.11939.

[7] A. Ambainis, A. Nayak, A. Ta-Shma and U. Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. arXiv:quant-ph/9804043.

[8] R. W. Spekkens, D. H. Buzacott, A. J. Keehn, B. Toner and G. J. Pryde. Preparation contextuality powers parity-oblivious multiplexing. *Physical Review Letters* 102 (2009), 010401. Publisher.

[9] R. Duan, Y. Xin and M. Ying. Locally indistinguishable subspaces spanned by three-qubit unextendible product bases. *Physical Review A* 81 (2010), 032329. Publisher; arXiv:0708.3559.

[10] C. H. Bennett, D. P. DiVincenzo, T. Mor, P. W. Shor, J. A. Smolin and B. M. Terhal. Unextendible product bases and bound entanglement. *Physical Review Letters* 82 (1999), 5385–5388. Publisher.

[11] D. P. DiVincenzo, T. Mor, P. W. Shor, J. A. Smolin and B. M. Terhal. Unextendible product bases, uncompletable product bases and bound entanglement. *Communications in Mathematical Physics* 238 (2003), 379–410. arXiv:quant-ph/9908070.

[12] R. Duan, Y. Feng, Y. Xin and M. Ying. Distinguishability of quantum states by separable operations. arXiv:0705.0795.

[13] E. Chitambar, D. Leung, L. Mančinska, M. Ozols and A. Winter. Everything you always wanted to know about LOCC (but were afraid to ask). arXiv:1210.4583.

[14] S. M. Cohen. Local approximation for perfect discrimination of quantum states. *Physical Review A* 107 (2023), 012401. Publisher; arXiv:2207.03573.

[15] M. A. Ballester, S. Wehner and A. Winter. State discrimination with post-measurement information. arXiv:quant-ph/0608014.

[16] C. Carmeli, T. Heinosaari and A. Toigo. State discrimination with postmeasurement information and incompatibility of quantum measurements. arXiv:1804.09693.

[17] S. Croke and S. M. Barnett. Difficulty of distinguishing product states locally. *Physical Review A* 95 (2017), 012337. Publisher.

[18] D. W. Kribs, C. Mintah, M. Nathanson and R. Pereira. Vector representations of graphs and distinguishing quantum product states with one-way LOCC. arXiv:1912.10483.

[19] R. Duan, S. Severini and A. Winter. Zero-error communication via quantum channels, noncommutative graphs, and a quantum Lovasz theta function. arXiv:1002.2514.

[20] H.-C. Cheng, A. Winter and N. Yu. Discrimination of quantum states under locality constraints in the many-copy setting. arXiv:2011.13063.

[21] J. Edmonds. Minimum partition of a matroid into independent subsets. *Journal of Research of the National Bureau of Standards B* 69B (1965), 67–72. NIST original.