When Matchgate Base Collapse Fails:
A Qutrit Trichotomy and Unbounded Exact Width

Chenghua Liu[1] and Boning Meng[2]

[1]Institute of Software, Chinese Academy of Sciences, Beijing, China
[2]University of Regensburg, Regensburg, Germany
liuch.russell@gmail.com, mengboning2013@gmail.com

Abstract

Holographic algorithms solve planar counting problems by encoding each value of a finite domain into several Boolean matchgate wires. Base collapse asks whether every exact representation can be compressed to a number of wires per edge bounded only by the domain size. Chen (STOC 2016) and Xia (STOC 2016) established broad positive collapse theorems under full-rank and related structural hypotheses. Together, these works left open whether a domain-only collapse bound survives when several locally deficient signatures must share one encoding. We resolve this open problem negatively, already on a three-state (qutrit) domain. For every $a \geq 4$, we construct a five-label integer-valued qutrit language of maximum arity a with an exact rational matchgate presentation, yet every exactly equivalent label-preserving presentation requires common width $\Omega(2^{a/2}/a)$, even if its domain, base, tensors, and complex weights may all change. The construction avoids the usual local degeneracies, so the obstruction is genuinely simultaneous. Algebraic geometry then drives an exhaustive trichotomy explaining the boundary of collapse: the gadget closure either separates into rays, is Gaussian-mobile and collapses to width at most three, or is confined to a rigid plane-plus-ray flag containing our unbounded family. Pure-spinor geometry forces compression in the mobile case, while dimension bounds for matchgate varieties and Zariski avoidance place exact integer data outside every low-width algebraic image. Thus one geometric framework explains both why collapse occurs and why it can fail without bound.

Acknowledgments

OpenAI Codex's principal substantive contribution to the mathematical development was to help identify, introduce, and formulate the algebraic-geometric viewpoint used in the proofs. In particular, it suggested treating the Zariski closures of low-width star-table images as $\mathbb{Q}$-defined algebraic varieties, using finite Zariski avoidance to select an integer obstruction table, and applying pure-spinor incidence and common-annihilator geometry to obtain low-rank matchgate covers. The authors checked, revised, and integrated these suggestions into the final arguments.

1 Introduction

Holographic algorithms turn finite-domain planar counting problems into Boolean networks of planar matchgates, where the global sum is evaluated by Pfaffians and weighted perfect matchings. Matchgates first arose in Valiant's study of classically simulable quantum circuits and later became the computational engine of holographic algorithms [21–23]. In broad Boolean counting-CSP settings they account precisely for the tractability created by planarity [10]. Understanding their representations is therefore part of understanding why this form of exact cancellation works at all.

The change of representation is mediated by a *base.* A width-t base assigns every original domain value a vector in a 2^t-dimensional Boolean space: operationally, each domain edge is replaced by t Boolean matchgate wires. Crucially, the same encoding must work at every occurrence of every label. Extra wires may therefore be mere coordinate redundancy, or they may store compatibility information that is visible only when many labelled tensors are represented simultaneously.

The central collapse question is whether every matchgate-realizable labelled language on a q-element domain admits a common presentation whose width is bounded by a function of q alone. We answer it negatively: already for $q = 3$, minimum exact common width is unbounded. This resolves the uniform-width aspect of the multi-signature problem left open after Cai and Fu's full-rank-generator theorem [5] and Chen's all-domain full-rank theorem [11]. Xia supplied further collapse certificates but explicitly isolated this gap: base directions unused by one signature may still be needed by another, and the attainable common collapse size was unknown [24, end of Section 5].

If a domain-only bound existed, arbitrarily large ambient matchgate spaces would add no exact expressive power at fixed domain size, and the search for holographic algorithms could be restricted to bounded-width bases. Our result shows instead that simultaneous compatibility can carry irreducibly large information. The question is therefore about the intrinsic size of a shared representation, not merely about shortening one displayed base matrix.

Matchgate signatures also have two complementary descriptions. They are boundary tensors of planar perfect-matching gadgets, and, after Boolean indices are read as occupation modes, they are Pfaffian or fermionic-Gaussian tensors, subject to the usual parity conventions [2, 18]. Base collapse can accordingly be viewed as a structured common-mode compression problem for an exact Gaussian tensor language. Ordinary linear algebra alone does not answer it: a qutrit fits inside two Boolean bits, but an entire labelled qutrit language need not become matchgate under one small common encoding.

1.1 The collapse question and our answer

We measure representation size by *minimum exact common width*: the fewest Boolean wires per domain edge among all matchgate presentations that preserve the finite label families, their arities and port orders, and the exact value of every planar labelled instance. A competing presentation may change the finite domain, the common base, every coordinate tensor, and all complex weights. It may not replace a label by a gadget or merely preserve polynomial-time computability. The precise notions appear in Section 2.3.

Informal Theorem (Unbounded exact width on a qutrit domain)**.** *For every $k \geq 2$ there is a five-label integer-valued qutrit language $\mathcal{F}_k \mid \{B_k\}$ whose unique right label B_k has arity $k + 2$. It has an exact rational matchgate presentation, but its minimum exact common width is*

$$\Omega\left(\frac{2^{k/2}}{k}\right).$$

The displayed presentation can simultaneously be chosen with a full-row-rank base, support-essential left and right families, rank exactly two at every port of every nonunary primitive, and a connected two-center gadget whose boundary tensor is nonfactorable across the two centers.

There are only five labels: the three unary qutrit pins, binary disequality on states $0, 1$, and one right label B_k of arity $k + 2$. Thus the obstruction is not a proliferation of unrelated primitives.

Moreover every left label is symmetric in its ports, so each singleton left subproblem is compatible with the type of symmetry used in Xia's singleton collapse [24, Theorem 5.3]. The large width is genuinely simultaneous: it is the requirement that the whole labelled language share one exact representation that becomes costly.

The result rules out any width bound depending only on the domain size in this unrestricted model. It is stronger than a lower bound against compressing the displayed base, because it survives arbitrary exact re-presentation. It does not assert computational hardness: the high-arity right label grows with k, and our equivalence notion is stricter than gadget or polynomial-time equivalence. The full statement and quantitative inequality are given in Theorem 3.1, and the domain-only consequence is Corollary 3.2.

1.2 Where can non-collapse occur?

The negative result leaves a structural question: why can width hide in a three-state language at all? Consider actual connected planar gadgets whose dangling edges are incident with left vertices; their internal vertices may use labels from either side. At an exposed port, the *mode support* is the span of all qutrit vectors obtainable there as the other ports vary. A port is deficient when this span misses at least one qutrit direction, so a nonzero support is either a one-dimensional *ray* or a two-dimensional *plane*. Bi-essentiality says that the mode supports of the primitive left labels, and separately those of the induced right labels, collectively see all three qutrit directions.

After the base is applied, matchgate parity singles out two endpoint rays P^+, P^- inside every realized plane P; these endpoints need not themselves occur as gadget supports. The exceptional *monomial flag* has all realized planes equal to one plane P, contains a transverse realized ray L, and confines every realized ray to P^+, P^-, or L.

Informal Theorem (Support-geometric trichotomy)**.** *For a full-row-rank, bi-essential, port-deficient qutrit presentation, exactly one of the following occurs.*

(I) Ray-separable. *No plane is realized, and every nonzero positive-arity all-left boundary tensor is a pure tensor, that is, it factors algebraically into unary vectors.*

(II) Gaussian-mobile. *A plane is realized, but the supports are not confined to one flag. The whole labelled language has an exact common presentation of width at most three; if two distinct planes occur, the width is at most two.*

(III) One-plane flag. *All realized planes equal P and all realized rays lie among P^+, P^-, L. Every nonzero positive-arity all-left boundary tensor is a Boolean matchgate core on P times pure-ray factors; a nullary boundary tensor is a scalar.*

The third branch contains a family whose minimum exact widths are unbounded.

This is the significance of the trichotomy. Once a plane-bearing support configuration escapes every monomial flag—through a second plane or additional transverse rays—Gaussian geometry produces one bounded-width encoding shared by every label. The flag still does not force large width for each individual language, and the ray-only branch is not yet a width classification. The branches are properties of a displayed presentation, whereas the collapse obtained in branch (II) is a new presentation of the same labelled language. Our unbounded family is constructed inside branch (III). There the flag rigidifies the all-left closure without controlling the residual right table. The label B_k can store a 2^k-entry residual table, which we choose outside every relevant low-width image. The precise statement is Theorem 3.3.

1.3 Proof architecture and technical ideas

The negative construction and the positive boundary use five main ideas.

1. *Composable-support geometry for base collapse.* Mode supports cannot increase when planar gadgets are composed. Primitive port deficiency therefore constrains an entire gadget closure by a small incidence geometry of realized rays and planes.

2. *Pure-spinor hulls as common-cover certificates.* The base sends realized rays to pure matchgate rays and realized planes to rank-two Gaussian row spaces. Two distinct planes, or two distinct pure rays outside one realized plane, have a low-codimension common annihilator. This produces a rank-four or rank-eight matchgate cover—one Gaussian row space containing the whole base row space—and hence collapse.

3. *All-arity flag-normal-form decoding.* In the one-plane branch, rank-one ports can be stripped off and the remaining plane-supported tensor can be decoded through one fixed pair of parity endpoints. This yields a Boolean matchgate core times pure-ray factors uniformly at every arity.

4. *Universal block-Pfaffian interpolation.* Any everywhere-nonzero Boolean table can be interpolated on two repetition codewords per block by a single exact matchgate. A multiplicative Möbius decomposition turns the table entries into independent Pfaffian controls and supplies the large-width presentation.

5. *A representation-independent star-table dimension obstruction.* The variety of s-wire matchgate signatures has dimension at most $1 + \binom{s}{2}$. Closed labelled stars extract the same 2^k-entry table from every exactly equivalent presentation, regardless of its new domain or base. Zariski avoidance then chooses an integer table outside every lower-width star-table image and gives the quantitative bound.

The fourth idea realizes arbitrary information; the fifth proves that this information cannot be hidden in fewer common modes. The first three explain why the construction must be placed in the fixed flag and why nearby support geometries collapse. The two proof lanes and their locations in the paper are summarized below.

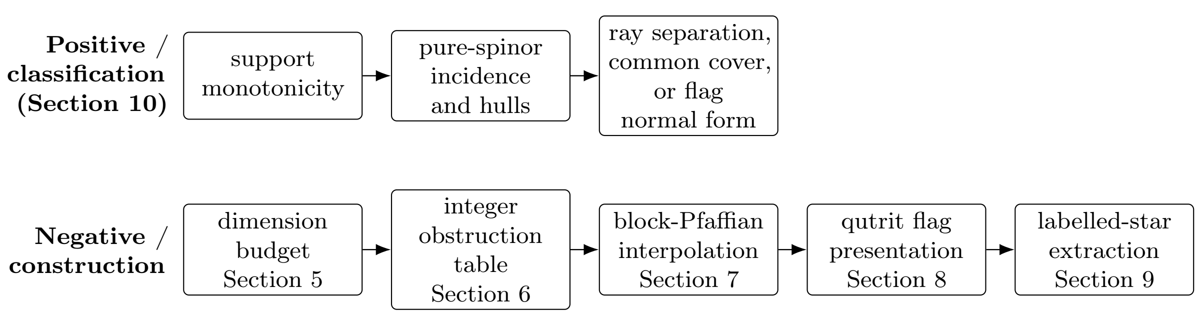

Figure 1: Conceptual roadmap. The upper lane classifies support geometry; the lower lane constructs information and proves that no exactly equivalent presentation can compress it below the stated width.

1.4 Relation to previous work

The collapse line begins with Cai and Lu's Boolean-domain results [7, 8]. Fu and Yang treated simultaneous realizability over rank-two bases [12]. Later work studied blockwise-symmetric signatures and holographic algorithms over general domains [13, 14].

The closest positive results delimit our regime. The broad formulation of Cai and Fu's domain-three collapse is used under their standing full-rank-generator hypothesis; in the degenerate alternative their conclusion is a direct computational simplification, rather than a domain-only bound on minimum exact common width [5]. Chen's all-domain theorem likewise covers the regime of a full-rank original signature [11]. Xia supplied several further sufficient certificates: a full-rank realizer, a bounded-rank common matchgate cover, or a singleton symmetric left family [24]. Our displayed presentations lie outside these regimes: their bases have full row rank and the two induced families are support-essential, but every port of every nonunary qutrit primitive has rank exactly two, and the simultaneous left family is not a singleton.

Xia also identified the precise compatibility gap that motivates our work. His singleton proof may discard a part of the base unused by one signature, but that part can remain necessary for another; after a two-signature example outside his hypotheses, he wrote that no characterization of matchgate covers was known and that the achievable collapse size was unknown [24, end of Section 5]. Under exact same-instance equivalence, Corollary 3.2 shows that the attainable width cannot be bounded as a function of the domain size alone. This answers the uniform-width aspect of Xia's question, not the characterization of matchgate covers or the optimal collapse size of each individual language.

Our structural result also identifies the boundary within this port-deficient qutrit regime. Support mobility yields an exact common presentation of width at most three—and at most two when two distinct planes occur—whereas the one-plane flag contains languages of unbounded minimum exact common width.

Nontrivial negative results for holographic algorithms did exist before this work. Cai and Choudhary used matchgate identities to rule out holographic templates for natural problem generalizations [3], and Fu, Yang, and Yin proved nonrealizability of higher-domain equality under holographic transformations [13]. Those results exclude a specified realization. Here every language is already matchgate-realizable; the question is how small any exact common realization can be. To our knowledge, Theorem 3.1 gives the first unbounded lower bound on minimum common matchgate-base width under exact labelled equivalence, even over a fixed three-element domain.

Algebraic geometry is likewise not new to matchgate theory. Matchgate identities define algebraic varieties, basis manifolds have been used for simultaneous realizability, and spinor/Pfaffian geometry has supplied a conceptual language for holographic algorithms [2,4,9,19]. Bradley also used dimension counts to prove lower bounds on the graph size needed to realize standard signatures [1]; that resource is the size of an individual matchgate, rather than the width of a common base. The new role here is to turn that geometry into a lower bound robust under arbitrary exact re-presentation: pure-spinor incidence controls the collapse side, while star-table varieties and Zariski avoidance control the non-collapse side. In this sense algebraic geometry is a central method of the paper, though not a first use of the subject in holographic algorithms.

1.5 Organization

Section 2 gives the minimum terminology needed to read the formal statements in Section 3. Section 4 then develops the proof machinery. Sections 5–9 build the obstruction and prove the lower bound; Section 10 proves the trichotomy, and Section 11 collects the open directions.

2 Statement-Level Preliminaries

This section isolates only the notions needed to parse the main results. The complete coordinate conventions, ordered planar operations, and proof-oriented definitions are developed in Section 4. A reader interested first in the statements may therefore proceed directly from this section to Section 3.

2.1 Signatures, planar instances, and local rank

Let D be a finite nonempty domain, take $\mathbb{N} = \{1, 2, \ldots\}$, and put $D_3 = \{0, 1, 2\}$. Write e_a for the coordinate vector of $a \in D$ and δ_a for the corresponding unary pin. A signature of arity a is a tensor $F : D^a \to \mathbb{C}$ whose ordered arguments are its *ports*. All families are finite and labelled: two labels remain distinct even if their coordinate tensors agree. A bipartite language $\mathcal{F} \mid \mathcal{G}$ places labels from $\mathcal{F}$ and $\mathcal{G}$ on the two sides of a planar embedded bipartite graph. Its value is the sum over all edge assignments of the product of the incident tensor entries. Exact labelled equivalence below preserves this value on the same embedded, ordered, labelled graph.

At port j, the *mode flattening* $\mathrm{Mat}_j(F)$ is the matrix with the jth coordinate as row index and all other coordinates as column index. Its image $S_j(F)$ is the *mode support*, and its rank is the *mode rank*. A port is *deficient* when this rank is smaller than $|D|$, and a signature is *non-full-rank* when every port is deficient. Thus a qutrit port is deficient exactly when its rank is at most two. A family is *support-essential* when the sum of all of its mode supports is the full domain space (three-dimensional for a qutrit); a bipartite language is *bi-essential* when this holds on both sides. A nonzero tensor is *completely decomposable*, or *pure*, when it is a tensor product of unary vectors. Its *Schmidt rank* across a partition of the ports is the rank of the corresponding two-block flattening. These notions and the left/right variance convention are formalized in Sections 4.1 and 4.2.

2.2 Common matchgate presentations

An exact s-wire matchgate signature is the boundary tensor of a finite weighted planar perfect-matching gadget with s ordered Boolean ports; the set of such tensors is denoted MG_s. A width-t base on D is a matrix

$$M \in \mathbb{C}^{|D| \times 2^t}.$$

It replaces every domain edge by one ordered block of t Boolean wires. For a left signature F of arity a, the Boolean lift is $FM^{\otimes a}$. A right domain signature is written $G = M^{\otimes b} H$, where H is its Boolean preimage and b is its block arity. A *valid common matchgate presentation* $(\mathcal{F}, \mathcal{H}, M)$ requires every $FM^{\otimes a}$ and every H to be an exact matchgate in the prescribed block order, with the same base M used for all labels. Its induced domain language is $\mathcal{F} \mid M\mathcal{H}$. The presentation is rational when all displayed coordinates and realizing edge weights may be chosen in $\mathbb{Q}$. Full definitions and the ordered contraction convention appear in Sections 4.3 and 4.6.

An r-input, t-output exact matchgate matrix is the matrix obtained by viewing r ordered Boolean ports of an exact matchgate as inputs and t as outputs, with the fixed planar output reversal. A *rank-2^r common matchgate cover* of M is such a matrix P satisfying

$$\mathrm{row}(M) \subseteq \mathrm{row}(P).$$

A common cover can be decoded to a width-r presentation. The matrix, cover, and decoding conventions are given in Section 4.4.

2.3 Exact labelled equivalence and minimum width

Definition 2.1 (Exact labelled equivalence)**.** Let $\mathcal{F} \mid \mathcal{G}$ be a finite-domain planar language over D, and let $\mathcal{P} \mid \mathcal{R}$ be one over an arbitrary finite domain E. They are *exactly equivalent* if there are arity-preserving bijections between the two left label sets and between the two right label sets such that every planar embedded $\mathcal{F} \mid \mathcal{G}$ instance has exactly the same complex value after replacing each label by its counterpart on the same embedded graph. Incidences, marked first ports, and linearized cyclic port orders are unchanged.

This is the planar, explicitly port-ordered version of Xia's *result equivalence*, which likewise uses label bijections and equality on the same instance graph [24, Section 2.3]. This notion does not permit replacing a label by a gadget, changing an arity, preserving values only up to a scalar, or merely retaining polynomial-time computability. It does permit an alternative presentation to change the finite domain, base, coordinate tensors, and exact complex weights. Thus it is a language-level notion, whereas ranks, supports, and flags associated with one displayed triple are presentation-relative.

Definition 2.2 (Minimum exact common width)**.** For a finite-domain language $\mathcal{F} \mid \mathcal{G}$ admitting at least one valid presentation, define $\mathrm{bw}(\mathcal{F} \mid \mathcal{G})$ to be the minimum integer $r \geq 0$ for which some valid common width-r presentation over some finite nonempty domain induces a labelled language exactly equivalent to $\mathcal{F} \mid \mathcal{G}$.

No rank condition is imposed on an alternative base or its domain signatures. At width zero the Boolean block alphabet is the singleton $\{0,1\}^0 = \{\epsilon\}$, so declared block arity still records the induced domain arity. The minimum ranges over all such alternative presentations, not merely submatrices or coordinate changes of a displayed base. A *uniform* width bound controls every language in a stated class by one function of the specified parameters; it is stronger than a pointwise assertion about each fixed finite language.

2.4 Realized supports and monomial flags

Fix a qutrit presentation with rank $M = 3$. A connected planar gadget is *all-left* when every dangling edge is incident with a left vertex. At a distinguished dangling port of a nonzero all-left boundary tensor, its mode support is a *realized support*. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be the sets of realized rays and planes, respectively. For every realized plane P, the transformed block support PM is a rank-two exact-matchgate row space. Its even- and odd-parity lines pull back to two distinguished rays of P, denoted P^+ and P^-. A pair (P, L), where $P \in \mathcal{S}_2$ and $L \in \mathcal{S}_1$ is transverse to P, is a *monomial flag* when

$$\mathcal{S}_2 = \{P\}, \qquad \mathcal{S}_1 \subseteq \{P^+, P^-, L\}.$$

In the flag branch, a *Boolean matchgate core on P times pure-ray factors* means that the plane-supported ports are encoded, through one fixed basis of the two parity endpoints, by a Boolean matchgate tensor, while every remaining port contributes a unary factor on one of the three allowed rays. The exact gadget, rerooting, and normal-form conventions are given in Section 4.7.

3 Main Results

3.1 Unbounded exact width

For $a, b \in D_3$, let E_{ab} denote the qutrit matrix unit, and put $X_{01} := E_{01} + E_{10}$; thus X_{01} is binary disequality on states $0, 1$ and vanishes whenever one input is 2.

Theorem 3.1 (Unbounded exact width with a flag presentation)**.** *For every integer $k \geq 2$, there is a finite integer-valued domain-three language*

$$\mathcal{F}_k = \{\delta_0, \delta_1, \delta_2, X_{01}\}, \qquad \mathcal{G}_k = \{B_k\}.$$

The right label B_k has arity $k+2$, which is therefore the maximum arity of the language. Put

$$\mu_k := \mathrm{bw}(\mathcal{F}_k \mid \mathcal{G}_k). \tag{1}$$

Then the following properties hold.

(i) *It has an exact rational matchgate presentation with a full-row-rank* 3*-by-*2^{2^k+3} *base, of width* 2^k+3.

(ii) *Every port of* X_{01} *and* B_k *has flattening rank exactly two, whereas the pins have rank one.*

(iii) *The language is support-essential on both sides. The single right tensor* B_k *has two link-port supports* $\operatorname{span}(e_0, e_1)$ *and* k *hard-port supports* $\operatorname{span}(e_0, e_2)$*—the link ports are its first two ports and the hard ports its remaining* k*—so it is support-essential by itself.*

(iv) *Two copies of* B_k *occur in a connected planar gadget whose boundary tensor has Schmidt rank exactly two across the hard ports of the first copy versus those of the second copy. Thus the construction is not a product of isolated stars.*

(v) *The minimum exact width satisfies*

$$2^k \leq 2\left(1+\binom{\mu_k}{2}\right)+1+\binom{k\mu_k}{2}. \tag{2}$$

Consequently,

$$\mu_k \geq \left\lceil \sqrt{\frac{2^k-3}{1+k^2/2}} \right\rceil = \Omega\left(\frac{2^{k/2}}{k}\right). \tag{3}$$

The presentation in part (i) *may be chosen to lie in alternative* (III) *of Theorem 3.3.*

The quantification in Theorem 3.1 is unrestricted: a competing presentation may change the finite domain, the common base, and every coordinate tensor, provided that the original labels, arities, ordered ports, and all planar instance values are preserved exactly. Thus the lower bound is not merely against compressions or invertible coordinate changes of the displayed base. The construction gives $\mu_k \leq 2^k+3$; no matching upper bound is claimed.

Corollary 3.2. *Under exact labelled equivalence, there is no function* $h\colon \mathbb{N} \to \mathbb{N}$ *such that, for every finite nonempty domain* D *and every valid presentation* $(\mathcal{F}, \mathcal{H}, M)$ *on* D *for which every signature in the induced domain families* $\mathcal{F} \cup M\mathcal{H}$ *is integer-valued, has positive arity, and is non-full-rank, one has* $\operatorname{bw}(\mathcal{F} \mid M\mathcal{H}) \leq h(|D|)$*. This remains false even when the displayed presentation is over* $\mathbb{Q}$*; equivalent presentations in the definition of* bw *may still use arbitrary exact complex weights.*

3.2 Support-geometric trichotomy

The lower bound is complemented by an exhaustive classification of the qutrit support geometries relevant to base collapse.

Theorem 3.3 (Support-geometry trichotomy and an unbounded flag family)**.** *Let* $(\mathcal{F}, \mathcal{H}, M)$ *be a valid common matchgate presentation on a qutrit domain with* $\operatorname{rank} M = 3$*, and let* $\mathcal{B} = M\mathcal{H}$ *be the induced right family. Assume that* $\mathcal{F} \mid \mathcal{B}$ *is bi-essential and every port flattening of every domain tensor in* $\mathcal{F} \cup \mathcal{B}$ *has rank at most two. Let* $\mathcal{S}_1, \mathcal{S}_2$ *be the supports realized at distinguished ports of actual connected all-left planar gadgets. Exactly one of the following alternatives holds.*

(I) Ray-separable closure: $\mathcal{S}_2 = \varnothing$. *Every nonzero positive-arity all-left boundary tensor is completely decomposable; a nullary tensor is a scalar.*

(II) Gaussian-mobile closure: $\mathcal{S}_2 \neq \varnothing$ *and no monomial flag exists. Then* M *has a common matchgate cover of rank at most eight, and yields an exactly equivalent common presentation of width at most three. If* $|\mathcal{S}_2| \geq 2$*, the cover has rank four and the width is at most two.*

(III) One-plane flag closure: *a monomial flag* (P, L) *exists. Every nonzero positive-arity all-left boundary tensor has the Boolean-core-times-rays form of Section 2.4, using one fixed endpoint-basis isomorphism* $\iota : \mathbb{C}^2 \to P$ *for the entire closure; a nullary boundary tensor is a scalar.*

The class of languages admitting an alternative (III) *presentation has no uniform bound on minimum exact common width. The rank, essentiality, and connected two-center properties in Theorem 3.1 may hold simultaneously.*

Remark 3.4 (Structural trichotomy versus width characterization). The alternatives are presentation-relative because P^+ and P^- are defined through the displayed base. The bounded conclusions in alternative (II) are nevertheless language-level: they construct a smaller exact presentation of the same labels. "No uniform bound" quantifies over the family $(\mathcal{F}_k \mid \{B_k\})$; every individual finite language has finite width. A flag is necessary for our constructed family to escape Gaussian collapse but is not sufficient for large width, since small-width flag presentations also exist. Alternative (I) also has no width classification. A pointwise characterization requires a language invariant that captures the surviving right-side data independently of any chosen flag presentation.

Theorem 3.1 is proved in Sections 5–10, Corollary 3.2 in Section 9, and Theorem 3.3 in Section 10.

4 Technical Preliminaries

4.1 Sets, fields, tensors, and ranks

We write $\mathbb{C}$, $\mathbb{Q}$, and $\mathbb{Z}$ for the fields of complex and rational numbers and the ring of integers, respectively. Unless another field is named, all vector spaces, tensor products, spans, and matrix ranks are over $\mathbb{C}$. We use

$$\mathbb{N} = \{1, 2, \ldots\}, \qquad \mathbb{Z}_{\geq 0} = \{0, 1, 2, \ldots\}, \qquad \mathbb{Z}_{>0} = \mathbb{N}.$$

For a positive integer a, write $[a] = \{1, \ldots, a\}$, and put $[0] = \varnothing$. The Boolean domain is $\{0, 1\}$, and a *qutrit* means a three-element domain, normally identified with $D_3 = \{0, 1, 2\}$. For a subfield $K \subseteq \mathbb{C}$, write $K^\times = K \setminus \{0\}$ and let $\overline{\mathbb{Q}}$ be the field of complex algebraic numbers.

For $x \in \{0, 1\}^s$, define

$$\operatorname{supp}(x) := \{i \in [s] : x_i = 1\}, \qquad |x| := |\operatorname{supp}(x)|,$$

and call $|x|$ the *Hamming weight*. The symbol $\oplus$ denotes bitwise exclusive-or, $\operatorname{dist}(x, y) := |x \oplus y|$ is Hamming distance, and $\bar{x} := x \oplus 1^s$ is the bitwise complement. For $T \subseteq [s]$, $\mathbf{1}_T \in \{0, 1\}^s$ is its indicator word. The words 0^s and 1^s are the all-zero and all-one repetition words; for $s = 0$ both equal the empty word ϵ. If I is an arbitrary finite index set, $\{0, 1\}^I$ means the set of functions $I \to \{0, 1\}$, and $\sigma(i)$ denotes the bit at i. For a proposition $\mathcal{A}$, the indicator $\mathbf{1}[\mathcal{A}]$ is 1 when $\mathcal{A}$ is true and 0 otherwise. For a coordinate array A, we also write $\operatorname{supp}(A) := \{\omega : A(\omega) \neq 0\}$ for its coordinate support; the argument makes this distinct from the support of a bit word.

If D is a finite set, a *signature of arity* a over D is a function $F \colon D^a \to \mathbb{C}$, and we write $\operatorname{arity}(F) = a$. A nullary signature is a complex scalar. A signature is *rational* or *integer-valued* when all of its coordinates lie in $\mathbb{Q}$ or $\mathbb{Z}$, respectively. All signature families below are finite labelled families: labels remain distinct even if two labels carry the same coordinate tensor. A *primitive signature* is a tensor attached to a single labelled vertex, as opposed to a tensor obtained by contracting a multi-vertex gadget; its argument positions are its *ports*.

Let $V_D = \mathbb{C}^D$ with its standard basis. Write e_a for the coordinate vector indexed by $a \in D$ and e_a^* for its dual covector. More generally, e_x denotes the coordinate vector indexed by a word x, so $e_x(y)$ is 1 for $y = x$ and 0 otherwise. The tensor product $u_1 \otimes \cdots \otimes u_a$ is the multilinear array whose coordinates are the products of the corresponding coordinates; $M^{\otimes a}$ is the a-fold Kronecker power applying M independently to the a ordered ports. The zeroth tensor power is the scalar identity on $\mathbb{C}$, and an empty tensor product equals 1. The symbols span, im, $\oplus$, and $\bigoplus$ denote linear span, image, direct sum, and an iterated direct sum. We write A^{T} for ordinary transpose, with no complex conjugation, and $\operatorname{row}(A)$, $\operatorname{col}(A)$, and $\operatorname{rank} A$ for the row space, column space, and matrix rank.

For tensor-variance bookkeeping, we regard a signature placed on the left side of a bipartite instance as an element of $(V_D^*)^{\otimes a}$ and one placed on the right as an element of $V_D^{\otimes a}$. Both have the same coordinate-array notation, and all pairings in this paper are multilinear over $\mathbb{C}$, never Hermitian. A *ray*, *line*, and *plane* mean respectively a one-dimensional subspace (the first two words are synonymous here) and a two-dimensional subspace. If $K \leq U$ are finite-dimensional, $\operatorname{codim}_U K := \dim U - \dim K$.

For $a \geq 1$ and $j \in [a]$, the *one-versus-rest flattening*, also called the *mode-j flattening*, of F is the matrix

$$\operatorname{Mat}_j(F) \in \mathbb{C}^{|D| \times |D|^{a-1}}, \qquad \operatorname{Mat}_j(F)_{u,\mathbf{v}} = F(v_1, \ldots, v_{j-1}, u, v_j, \ldots, v_{a-1}). \tag{4}$$

Its image

$$S_j(F) := \operatorname{im} \operatorname{Mat}_j(F) \tag{5}$$

is the *mode support* or *port support*, and $\operatorname{rank} \operatorname{Mat}_j(F) = \dim S_j(F)$ is the corresponding *mode rank*. With the variance convention above, a left mode support is viewed in V_D^* and a right mode support in V_D. A port is *deficient* when its mode rank is less than $|D|$. When $|D| = n$, we call F *full rank* if $\operatorname{rank} \operatorname{Mat}_j(F) = n$ for some j, and *non-full-rank* if every port is deficient. This is Chen's domain-signature notion of full rank [11]; it is unrelated to the matrix rank of a holographic base. A presentation is *port-deficient* when every port of every primitive signature in both induced domain families is deficient.

A *minimal mode decomposition* at port j is an expression

$$F = \sum_{a=1}^{d} u_a \otimes F_a, \qquad d = \operatorname{rank} \operatorname{Mat}_j(F),$$

after placing port j first in the notation, with both (u_a) and (F_a) linearly independent. It is precisely a rank factorization of the mode-j flattening.

A family of left signatures is *support-essential* when the sum of all of its mode supports is V_D^*; the analogous condition for a right family uses V_D. A bipartite problem, or a presentation of it, is *bi-essential* when both induced domain families are support-essential. Zero tensors contribute the zero subspace to these sums.

A nonzero positive-arity tensor is *completely decomposable*, or a *pure tensor*, if it is an elementary tensor $u_1 \otimes \cdots \otimes u_a$; equivalently, every mode rank is one. For a bipartition $A \sqcup B = [a]$, the *Schmidt rank across $A \mid B$* is the matrix rank obtained by flattening the A coordinates against the B coordinates. Thus Schmidt rank one is exactly factorization as one tensor on A times one tensor on B. An *outer product* is such a rank-one bipartite tensor. A *slice* of a tensor is the lower-arity tensor obtained by fixing a specified subset of its coordinates.

Finally, $u_k = \Omega(v_k)$ means that there are constants $c > 0$ and k_0 such that $|u_k| \geq c|v_k|$ for every $k \geq k_0$. For $q = p/r \in \mathbb{Q}$ in lowest terms, with $r > 0$, define its *height* by

$$H(q) := \max\{|p|, r\}.$$

The height of a finite rational table is the maximum height of its entries. If H is this table height, then $1 + \log_2 H$ controls the maximum entry bit length up to constant factors. We call an indexed family of tables *entrywise closed-form* when one uniform formula gives each entry directly from its index, rather than selecting the table by existence or enumeration. A table is *everywhere nonzero* or *nowhere zero* when none of its entries is zero.

4.2 Planar Holant networks and gadgets

Let $\mathcal{F}$ and $\mathcal{G}$ be finite labelled signature families over the same domain D. We call the ordered pair $\mathcal{F} \mid \mathcal{G}$ a *bipartite signature language* or *planar Holant problem*. A planar bipartite $\mathcal{F} \mid \mathcal{G}$ instance consists of a planar embedded bipartite graph (U, V, E), a label from $\mathcal{F}$ on every vertex of U, and a label from $\mathcal{G}$ on every vertex of V. At each vertex of positive degree we specify a first incident edge and then follow the cyclic order induced by the embedding; a nullary vertex has the unique empty port order. Thus every vertex has a linear port order compatible with its planar cyclic order. A vertex label must have arity equal to the degree of that vertex, and its arguments follow this ordered list of ports. Here a planar embedding is a fixed crossing-free embedding; it determines an outer face and a cyclic order at every vertex. We write $E(w)$ for the resulting linearly ordered list of edges incident with w. The value of the instance I is

$$\mathrm{Holant}_I(\mathcal{F} \mid \mathcal{G}) = \sum_{\sigma\colon E\to D} \prod_{u\in U} F_u(\sigma|_{E(u)}) \prod_{v\in V} G_v(\sigma|_{E(v)}). \tag{6}$$

Here F_u and G_v denote the coordinate tensors assigned by the labels, and $\sigma|_{E(w)}$ is the tuple of incident edge values read in the specified linear port order. Formula (6) is the full tensor contraction of the network.

A *planar gadget* is the same kind of finite tensor network embedded in a closed disc, with some uncontracted half-edges ending on the boundary of the disc; these are its *dangling* or *boundary edges*. Their cyclic order on the boundary, together with a marked first edge, is the specified linear boundary order. Contracting all internal edges produces the *boundary tensor*, whose ports occur in that order. If e is a boundary edge, $j(e)$ denotes its position in this order. The vertex tensors are the gadget's primitive signatures. A gadget is *connected* when its underlying graph after deleting the dangling half-edges is connected. It is *all-left* when every dangling edge is incident with a left vertex. The *all-left gadget closure* is the collection of boundary tensors of actual connected all-left gadgets; it is not their linear span. A *star gadget* has one central vertex; each central port is either contracted with a unary leaf or left as a dangling edge. It is *closed* when it has no dangling edges, equivalently, when its boundary arity is zero and every central port is contracted with a unary leaf. A *pin* is a unary coordinate signature $\delta_a = e_a^*$ on the left, or e_a on the right; *pinning a port* means contracting it with that unary tensor. Ordered planar composition means gluing compatible consecutive boundary ports without changing their inherited cyclic orders. The Boolean *disequality* signature is $\mathrm{NEQ}(x, y) = 1$ for $x \neq y$ and 0 otherwise.

4.3 Planar matchgates, Pfaffians, and ordered operations

A *perfect matching* of a finite graph is a set of edges incident with every vertex exactly once. For an edge-weighted graph $\mathcal{N}$, let $w(e)$ denote the weight of edge e and put

$$\mathrm{PerfMatch}(\mathcal{N}) := \sum_{P \text{ a perfect matching of } \mathcal{N}} \prod_{e\in P} w(e), \tag{7}$$

with value 0 if no perfect matching exists and value 1 on the empty graph. A *planar matchgate* $\mathcal{M}$ is a finite edge-weighted planar graph with *external vertices* $v_1, \dots, v_s$ occurring in that cyclic order on its outer face, with v_1 marked as first; all other vertices are *internal.* Under the deletion convention, its s-bit signature $G\colon \{0,1\}^s \to \mathbb{C}$ is

$$G(x) = \mathrm{PerfMatch}(\mathcal{M} - \{v_i : x_i = 1\})\,, \tag{8}$$

where deletion removes the indicated external vertices and all incident edges. We fix this convention throughout. The other common convention is obtained by globally complementing every index bit. Let MG_s denote the set of *exact* s-bit planar-matchgate signatures in the stated external order: "exact" means equality with the finite weighted graph's coordinates, not approximation or a limiting construction. We set $\mathrm{MG}_0 = \mathbb{C}$.

For $s \geq 0$, let

$$\rho_s(x_1, \ldots, x_s) = (x_s, \ldots, x_1), \qquad \rho_0(\epsilon) = \epsilon.$$

Lemma 4.1 (Reflection reversal)**.** *If $G \in \mathrm{MG}_s$, then the reversed signature*

$$G^{\mathrm{rev}}(x) := G(\rho_s(x))$$

also belongs to MG_s*.*

Proof. For $s = 0$ there is nothing to prove. Reflect an exact planar realization of G in the plane and mark the reflected copy of the old last external vertex as the new first external vertex. The reflected external order is the reverse of the old order, while the weighted graph and every perfect-matching weight are unchanged. Its deletion signature in the new order is therefore $G \circ \rho_s$. □

The *matchgate identities*, also called the quadratic Grassmann–Plücker relations in these coordinates, are the following equations. For $\alpha, \beta \in \{0,1\}^s$, list $\mathrm{supp}(\alpha \oplus \beta) = \{p_1 < \cdots < p_\ell\}$; then

$$\sum_{j=1}^{\ell} (-1)^j G(\alpha \oplus \mathbf{1}_{\{p_j\}}) G(\beta \oplus \mathbf{1}_{\{p_j\}}) = 0. \tag{9}$$

They are necessary and sufficient for exact planar-matchgate realizability and imply parity [4, 6].

A matrix A is *skew-symmetric* when $A^{\mathsf{T}} = -A$. If $S = \{i_1 < \cdots < i_{2m}\}$ has even size, its principal submatrix is $A[S]$ and its *Pfaffian* is

$$\mathrm{Pf}(A[S]) := \sum_{\pi} \mathrm{sgn}(\pi) \prod_{q=1}^{m} A_{u_q, v_q}. \tag{10}$$

Here π ranges over partitions of S into pairs. For each π, write each pair as $\{u_q, v_q\}$ with $u_q < v_q$, and index the pairs so that $u_1 < \cdots < u_m$. The increasing order on S is the displayed sequence $(i_1, \ldots, i_{2m})$. Let σ_π be the unique permutation of $[2m]$ satisfying

$$(u_1, v_1, \ldots, u_m, v_m) = (i_{\sigma_\pi(1)}, i_{\sigma_\pi(2)}, \ldots, i_{\sigma_\pi(2m)}).$$

Then

$$\mathrm{sgn}(\pi) := \mathrm{sgn}(\sigma_\pi) = (-1)^{N_\pi},$$

where N_π is the number of pairs of positions $a < b$ for which the ath entry of $(u_1, v_1, \ldots, u_m, v_m)$ is larger than its bth entry. For example, if $S = \{1, 2, 3, 4\}$ and $\pi = \{\{1, 3\}, \{2, 4\}\}$, then the associated sequence $(1, 3, 2, 4)$ has one inversion, so $\mathrm{sgn}(\pi) = -1$. Each product is a *perfect-matching monomial.* We use $\mathrm{Pf}(A[\varnothing]) = 1$.

For skew-symmetric matrices A and B of even orders, their block direct sum in concatenated coordinate order satisfies

$$\mathrm{Pf}(A \oplus B) = \mathrm{Pf}(A)\,\mathrm{Pf}(B).$$

More generally, if P_σ is the permutation matrix of a coordinate permutation σ, then, for every even-order skew-symmetric matrix C,

$$\mathrm{Pf}(P_\sigma C P_\sigma^{\mathsf{T}}) = \det(P_\sigma)\,\mathrm{Pf}(C) = \mathrm{sgn}(\sigma)\,\mathrm{Pf}(C).$$

Swapping consecutive coordinate chunks of sizes p and q consists of pq transpositions and therefore contributes the sign $(-1)^{pq}$. Consequently, permuting even-sized chunks introduces no sign.

We use the following two standard consequences of the identities.

1. Every nonzero $G \in \mathrm{MG}_s$ has one parity: all nonzero coordinates $G(x)$ have the same parity of $|x|$.
2. For a skew-symmetric matrix $A \in \mathbb{C}^{s\times s}$, define its *principal-Pfaffian signature*

$$\Gamma_A(x) = \begin{cases} \mathrm{Pf}(A[\mathrm{supp}(x)]), & |x| \text{ is even,} \\ 0, & |x| \text{ is odd,} \end{cases} \qquad \mathrm{Pf}(A[\varnothing]) = 1, \tag{11}$$

   Then $\Gamma_A \in \mathrm{MG}_s$. By sufficiency of the matchgate identities, Γ_A has an exact planar realization in the stated external order; A is a Pfaffian parameter matrix and need not itself be planar. To reconcile conventions explicitly, if $\widetilde{\Gamma}_A(x) := \mathrm{Pf}(A[[s] \setminus \mathrm{supp}(x)])$, then $\Gamma_A(x) = \widetilde{\Gamma}_A(\bar{x})$. The support- and deletion-Pfaffian conventions therefore differ by the one global bit complement already fixed above. The matchgate identities are invariant under this common XOR shift of all coordinate indices, so Γ_A satisfies them in the same coordinate order and the sufficiency theorem realizes it exactly in the deletion convention.

The all-zero signature belongs to MG_s: one may adjoin an isolated internal vertex, which destroys every perfect matching. Every standard basis vector e_x belongs to MG_s as a tensor product of unary pins. If $G \in \mathrm{MG}_s$ and $\lambda \in \mathbb{C}$, then $\lambda G \in \mathrm{MG}_s$; take the disjoint union with an internal edge of weight λ. More generally, exact matchgates are closed under disjoint union/tensor product, pinning an external bit, and ordered planar contraction: the last operation glues compatible consecutive external vertices, sums over their common Boolean values, and retains the inherited order on the unglued vertices. These are the closure operations invoked later.

Lemma 4.2 (Field-preserving matchgate realization)**.** *Let $K \subseteq \mathbb{C}$ be a subfield, and fix an order of s coordinates. If $G \in K^{\{0,1\}^s}$ satisfies the matchgate identities in that order, then G is exactly the deletion signature of a planar matchgate whose edge weights lie in K and whose external vertices occur in the prescribed order.*

Proof. For a nonzero signature this is the field-preserving sufficiency construction for the matchgate identities of Cai and Gorenstein [6, Section 5 and Theorem 5.1]. Their normalization and restoration steps use only field operations, and their planar crossover gadget—the fixed matchgate implementing a signed exchange of two ordered wires—uses weights in $\{+1, -1\}$. Their construction places the labels $1 < \cdots < s$ clockwise; naming those labels by the fixed coordinate order gives the prescribed external order here. The zero signature is realized over K by including an isolated internal vertex. Nullary signatures are scalars and are realized by a single internal edge of the given weight (with the zero case handled as above). $\square$

4.4 Matchgate matrices, covers, and fermionic geometry

Let $G \in \mathrm{MG}_{r+t}$ have its first r ports designated as inputs and its last t as outputs. Its *r-input, t-output exact matchgate matrix* is the 2^r-by-2^t matrix P defined by

$$P(x,y) := G(x, \rho_t(y)), \qquad x \in \{0,1\}^r, \quad y \in \{0,1\}^t, \tag{12}$$

where ρ_t is the fixed planar reversal defined in Lemma 4.1. With this convention, gluing an ordered output block to an ordered input block is ordinary matrix multiplication. A square such matrix is an *input* or *output matchgate transformation* according to the side on which it is composed; it is *invertible* in the ordinary linear-algebraic sense. The inverse of an invertible exact matchgate transformation is again an exact matchgate transformation, by the ordered canonical-form theorem [4, 24]. Here an *ordered canonical form* means the normalization of an exact matchgate matrix, by invertible input and output matchgate transformations, to disjoint active input–output wires with all inactive modes pinned; every later use states the particular consequence needed. A Boolean *wire* or *mode* is one binary tensor index; an ordered consecutive group of wires is a *block*.

The matchgate matrix rank theorem says that if P is a nonzero a-input, b-output exact matchgate matrix, then

$$\operatorname{rank} P = 2^d$$

for a unique integer d with $0 \le d \le \min\{a,b\}$. In the ordered canonical form, d is the number of active input–output wires [24, Theorem 4.1]. Conversely, the canonical matchgate with d active wires and all remaining modes pinned shows that, for fixed a, b, every d in this range occurs. Thus, when the input and output arities are unrestricted, d may be any nonnegative integer. For $0 \le r \le t$, the *standard inclusion* $J_{r,t}$ is the exact matchgate matrix that copies the r input bits to r active output wires and pins the remaining $t-r$ inactive output wires to zero. In the fixed matrix convention,

$$J_{r,t}(x,y) = \mathbf{1}[y = (x, 0^{t-r})], \qquad J_{r,t} J_{r,t}^{\mathsf{T}} = I_{2^r}.$$

In its planar realization, the first $t-r$ physical output ports are zero-pinned and r nested unit edges join input i to the physical output port carrying logical bit y_i. Thus the physical output order is $y_t, \ldots, y_1$, exactly as required by (12). A fixed-mode *flip* exchanges the two Boolean values on that mode and is itself an invertible matchgate transformation: replace the corresponding unit equality edge in the nested identity realization by a two-edge unit path with one internal vertex. Its two-terminal deletion signature is 1 exactly on 01 and 10, so its ordinary matrix is the bit flip on that mode.

Following Xia, if $M \in \mathbb{C}^{n \times 2^t}$ has rank 2^r, a *full-rank matchgate realizer* of M is a matrix $A \in \mathbb{C}^{2^s \times n}$, for some $s \ge 0$, such that AM is an s-input, t-output exact matchgate matrix of rank 2^r [24, Definition 5.1]. More generally, a *rank-2^r common matchgate cover* of a width-t base M is an r-input, t-output exact matchgate matrix P of rank 2^r satisfying

$$\operatorname{row}(M) \subseteq \operatorname{row}(P). \tag{13}$$

Its *cover width* is r. It is called common because the same row space covers the single base shared by every label. A *Gaussian cover* is the same object; "Gaussian" here is the standard fermionic synonym for matchgate/Pfaffian structure, not a probabilistic term. A *rank-2^r Gaussian hull matrix* of subspaces $R_1, \ldots, R_m$ is such an exact matchgate matrix whose row space contains $R_1 + \cdots + R_m$; its row space is the corresponding *Gaussian hull*. When that row space contains $\operatorname{row}(M)$, the matrix is a common cover of M.

For parity bookkeeping, put

$$\mathcal{E}_t := \operatorname{span}\{e_x : |x| \text{ is even}\}, \qquad \mathcal{O}_t := \operatorname{span}\{e_x : |x| \text{ is odd}\}.$$

A subspace is *parity-invariant* when it is the direct sum of its intersections with $\mathcal{E}_t$ and $\mathcal{O}_t$. If R is the row space of a rank-two exact matchgate matrix, the canonical form gives

$$R = R^+ \oplus R^-, \qquad R^+ := R \cap \mathcal{E}_t, \quad R^- := R \cap \mathcal{O}_t, \qquad \dim R^+ = \dim R^- = 1.$$

A nonzero vector is *parity-homogeneous* when it lies entirely in $\mathcal{E}_t$ or entirely in $\mathcal{O}_t$. The two lines R^+, R^- are its *parity endpoints.* When a domain plane P maps to $R = PM$, their pullbacks in P are denoted P^+, P^-.

We now fix the pure-spinor terminology used in the support argument. Let $f_1, \dots, f_t$ be the standard basis of $W = \mathbb{C}^t$, and identify $\mathbb{C}^{\{0,1\}^t}$ with the exterior algebra $\Lambda^\bullet W$ by

$$e_x \longleftrightarrow f_{i_1} \wedge \cdots \wedge f_{i_m} \quad \text{when } \operatorname{supp}(x) = \{i_1 < \cdots < i_m\}.$$

Thus $e_\varnothing = e_{0^t}$ is the *vacuum*, a vector supported on words of weight one is a *one-particle vector*, and $\wedge$ is the alternating exterior product. The space $W \oplus W^*$ carries the symmetric bilinear form

$$\langle w + \phi, w' + \phi' \rangle := \phi(w') + \phi'(w)$$

and acts on $\Lambda^\bullet W$ by the Clifford action

$$(w + \phi) \cdot u := w \wedge u + \iota_\phi u,$$

where ι_ϕ is contraction by ϕ. For nonzero u, define its *annihilator*

$$\operatorname{Ann}(u) := \{z \in W \oplus W^* : z \cdot u = 0\}. \tag{14}$$

A subspace $K \leq W \oplus W^*$ is *isotropic* when the displayed form vanishes on $K \times K$; it is maximal isotropic when it has the maximum possible dimension t. A nonzero u is a *pure spinor* when $\operatorname{Ann}(u)$ is maximal isotropic. In the coordinates above, satisfying the matchgate identities in the fixed ordered basis is equivalent to being a pure spinor. Together with Lemma 4.2, a nonzero vector is therefore an exact matchgate vector if and only if it is a pure spinor [4, 6, 16]. For a vector space U, its projective space $\mathbb{P}(U)$ is the set of one-dimensional subspaces; $[u]$ denotes the point $\mathbb{C}u$. We call $\mathbb{C}u$ a *pure matchgate ray* and define $\operatorname{Ann}(\mathbb{C}u) := \operatorname{Ann}(u)$, which is independent of the chosen nonzero generator. The projective set of pure matchgate rays of one fixed parity is the *spinor variety*; a *pure pencil* is a projective line contained in that variety. A *common annihilator* of several pure vectors or rays is the intersection of their annihilators. The contractions $\iota_{f_i^*}$ are the *standard annihilation operators*; their common kernel on $\Lambda^\bullet W$ is the intersection of their ordinary linear kernels.

A linear *isometry* of $W \oplus W^*$ is an invertible linear map that preserves the displayed symmetric bilinear form. The *Witt extension theorem* says that an isometry between two isotropic subspaces extends to an isometry of the whole nondegenerate space. The group of all such isometries is denoted $\mathrm{O}(W \oplus W^*)$, and its determinant-one subgroup is $\mathrm{SO}(W \oplus W^*)$. The *Pin/Spin (Clifford) lift* is the induced invertible action on $\Lambda^\bullet W$ intertwining the Clifford action; a determinant-reversing isometry uses the Pin action, or equivalently one fixed-mode flip in addition to a Spin action. Thus, if $g \in \mathrm{O}(W \oplus W^*)$ and U_g is a lift, then, up to the immaterial nonzero scalar in the choice of lift,

$$U_g\big(z \cdot u\big) = (gz) \cdot U_g(u). \tag{15}$$

The next lemma supplies the ordered-planar realization statement needed later; it is not being folded into the terminology.

Lemma 4.3 (Ordered transpose and Clifford-lift realization)**.** *The following statements use the matrix convention* (12).

(a) *If P is an r-input, t-output exact matchgate matrix, then P^{T} is a t-input, r-output exact matchgate matrix.*

(b) *For $g \in \mathrm{O}(W \oplus W^*)$, the ordinary coefficient matrix of any Pin lift U_g in the ordered exterior-monomial basis is an invertible t-input, t-output exact matchgate matrix. The same is true of U_g^{-1}.*

Proof. For (a), let G realize P, so $P(x, y) = G(x, \rho_t(y))$. If $r + t = 0$, transposition changes nothing and the claim is immediate. Otherwise reflect the embedded graph and mark the reflected copy of the old last external port as the new first port. In the resulting prescribed order its signature is

$$H(y, z) := G\big(\rho_r(z), \rho_t(y)\big).$$

Reflection preserves planarity, weights, and deletion-signature values. Moreover,

$$H\big(y, \rho_r(x)\big) = G\big(x, \rho_t(y)\big) = P(x, y) = P^{\mathsf{T}}(y, x),$$

so the reflected matchgate realizes exactly P^{T}. Notice that this argument does not assert that a bare unsigned wire reversal is a matchgate transformation.

For (b), define the coefficient convention by

$$U_g(e_y) = \sum_{x \in \{0,1\}^t} (U_g)_{x,y} e_x.$$

Take ordered copies $W_{\mathrm{in}}, W_{\mathrm{out}}$ of W and, in the ordered graded-tensor identification $\Lambda(W_{\mathrm{in}} \oplus W_{\mathrm{out}}) \cong \Lambda W_{\mathrm{in}} \mathbin{\widehat{\otimes}} \Lambda W_{\mathrm{out}}$, put

$$\Omega_t := \sum_{y \in \{0,1\}^t} e_y^{\mathrm{in}} \wedge e_{\rho_t(y)}^{\mathrm{out}}.$$

Here $\widehat{\otimes}$ denotes the graded tensor product, with all input modes ordered before all physical output modes. This is the signature of the nested-unit-edge realization of $J_{t,t} = I_{2^t}$ described above, and hence is a pure spinor. First assume $g \in \mathrm{SO}(W \oplus W^*)$ and take its even Spin lift. In the ordered factorization above its action on doubled spinors is literally $U_g \widehat{\otimes} I$: input creation and contraction act on the first factor, while contraction of an input covector with the second factor vanishes. Since the lift is even, it intertwines the doubled Clifford action with $g \oplus I$. Therefore

$$\mathrm{Ann}\big((U_g \widehat{\otimes} I)\Omega_t\big) = (g \oplus I)\, \mathrm{Ann}(\Omega_t),$$

which is again maximal isotropic. Consequently

$$\Omega_g := (U_g \widehat{\otimes} I)\Omega_t$$

is pure and therefore satisfies the matchgate identities. Its ordered coordinates obey

$$\Omega_g\big(x, \rho_t(y)\big) = (U_g)_{x,y},$$

so Lemma 4.2 realizes the ordinary matrix of U_g as an exact planar matchgate in precisely the prescribed input/output order.

If $\det g = -1$, fix a standard determinant-reversing mode reflection s with $s^2 = I$ whose lift is the exact fixed-mode flip recorded above, and put $h = sg \in \mathrm{SO}(W \oplus W^*)$. Up to a nonzero scalar, a Pin lift of $g = sh$ is the ordered composition of that fixed-mode flip with the Spin lift of h. Exact matchgate matrices are closed under this ordinary planar composition and under nonzero rescaling. This proves the assertion for every Pin lift. Applying the same argument to g^{-1} proves it for U_g^{-1}. $\square$

An exact-matchgate row space is also called a *Gaussian row space*, and an invertible matchgate transformation a *Gaussian transformation*. In the Grassmann description, symbols such as $\xi_1, \dots, \xi_s, \eta_1, \dots, \eta_t$ are anticommuting input and output variables. A *pivot chart* chooses a nonzero matrix coordinate, uses fixed input/output flips to place it at the vacuum coordinate, and divides by that coordinate; *pivot-normalized* means that this pivot is 1. In such a chart a matchgate matrix has a skew quadratic parameter matrix in the standard Pfaffian/Grassmann chart [4, 24]

$$\mathcal{A} = \begin{pmatrix} A_{\mathrm{in}} & B \\ -B^{\mathsf{T}} & D_{\mathrm{out}} \end{pmatrix}$$

and coefficient-generating polynomial

$$\exp\left(\frac{1}{2}\xi^{\mathsf{T}} A_{\mathrm{in}} \xi + \xi^{\mathsf{T}} B \eta + \frac{1}{2}\eta^{\mathsf{T}} D_{\mathrm{out}} \eta\right).$$

The variables anticommute, so the exponential is a finite exterior-algebra polynomial; the coefficient of each input–output monomial is the corresponding matrix entry. If the pivot output row is the vacuum, all its two-output coordinates vanish and hence $D_{\mathrm{out}} = 0$. Multiplication on the input Grassmann algebra by the inverse input-only exponential removes A_{in} by ordinary invertible row operations. More explicitly, for $q(\xi) = \frac{1}{2}\xi^{\mathsf{T}} A_{\mathrm{in}} \xi$, the map

$$L_{-q} : p(\xi) \longmapsto \exp(-q(\xi)) \wedge p(\xi)$$

is an invertible linear map with inverse L_q. Applied coefficientwise in η, its ordinary coefficient matrix left-multiplies the matchgate matrix and hence leaves its row space unchanged. No exact-matchgate claim about L_{-q} is needed: since the exterior elements involved are even,

$$\exp(-q) \wedge \exp(q + \xi^{\mathsf{T}} B \eta) = \exp(\xi^{\mathsf{T}} B \eta),$$

and the matrix on the right is independently exact by the same principal-Pfaffian chart. Its coefficient matrix, up to the fixed signs dictated by (12), is the *exterior-compound matrix*

$$\bigoplus_{q=0}^{\min(s,t)} \wedge^q B.$$

If $d = \operatorname{rank} B$, then

$$\operatorname{rank}\left(\bigoplus_{q=0}^{\min(s,t)} \wedge^q B\right) = \sum_{q=0}^{d} \binom{d}{q} = 2^d. \tag{16}$$

The nontrivial Clifford-lift realization has been isolated in Lemma 4.3; the remaining geometric inputs are stated with citations where they enter the proofs.

4.5 Algebraic geometry and transcendence conventions

For a field $K \subseteq \mathbb{C}$, affine N-space $\mathbb{A}_K^N$ has coordinate ring $K[Y_1, \dots, Y_N]$, the polynomial ring of its coordinate functions. An *ideal* is an additive subgroup of a ring closed under multiplication by arbitrary ring elements. If $I \subseteq K[Y_1, \dots, Y_N]$ is an ideal, write

$$V_{\mathbb{C}}(I) := \{y \in \mathbb{C}^N : p(y) = 0 \text{ for every } p \in I\}.$$

A set X cut out in $\mathbb{A}_K^N$ by an ideal $I \subseteq K[Y_1, \ldots, Y_N]$ is called a K-defined *affine algebraic set*; its complex points are $X(\mathbb{C}) := V_{\mathbb{C}}(I)$. A subset of $\mathbb{C}^N$ is *K-defined Zariski closed* if it equals $X(\mathbb{C})$ for such an X. The *Zariski closure* $\overline{S}^{\mathrm{Zar}}$ is the smallest Zariski-closed set containing S; S is *Zariski dense* when that closure is the whole ambient space, and a closed subset is *proper* when it is not the whole ambient space. A closed set is *irreducible* when it is not a union of two proper closed subsets, and an *irreducible component* is a maximal irreducible closed subset. We use *affine algebraic set* and *affine variety* interchangeably and do not require irreducibility.

The *dimension* of an affine algebraic set is its Krull dimension: the largest length of a strict chain of irreducible closed subsets, equivalently the maximum dimension of its irreducible components; in particular $\dim \mathbb{A}_K^N = N$. A *principal open set* has the form $D(f) = \{x : f(x) \neq 0\}$, and $\mathbb{G}_m := \mathbb{A}^1 \setminus \{0\}$. A *regular map* or *morphism* has coordinate functions that are regular: locally they are quotients of polynomials with nonvanishing denominators. On a closed affine algebraic set, polynomial coordinate functions give the special case used below; on $\mathbb{G}_m$, Laurent polynomials—finite sums of integer powers of the coordinate—are also regular. We use the standard facts that

$$\dim \overline{\phi(X)}^{\mathrm{Zar}} \leq \dim X$$

for a morphism ϕ, that product dimensions add, and that extending the ground field from $\mathbb{Q}$ to $\mathbb{C}$ does not change dimension.

A polynomial is *homogeneous of degree d* when every one of its monomials has total degree d; a homogeneous degree-two polynomial is a *quadric*. A projective algebraic set in $\mathbb{P}(U)$ is the common zero set of homogeneous polynomials, and its *homogeneous vanishing ideal* consists of all homogeneous polynomials vanishing on it. Saying that this ideal is *generated by quadrics* means that every polynomial in it is a polynomial combination of degree-two generators. Restricting a quadric to a *projective line* $\mathbb{P}(L)$, where L is a two-dimensional vector subspace, gives a homogeneous degree-two polynomial in two line coordinates; three distinct zeros force that restriction to vanish identically.

If $X = V_{\mathbb{C}}(I) \subseteq \mathbb{C}^m$ and a polynomial map $\phi : X \to \mathbb{C}^N$ has target coordinates $Y_1, \ldots, Y_N$, its *graph ideal* is

$$J := I + \langle Y_j - \phi_j(\mathbf{U}) : j \in [N] \rangle \subseteq K[\mathbf{U}, \mathbf{Y}],$$

and $J \cap K[\mathbf{Y}]$ is the corresponding *elimination ideal*. The elimination theorem identifies the Zariski closure of the image with $V_{\mathbb{C}}(J \cap K[\mathbf{Y}])$. A *Gröbner basis* is a generating set whose leading monomials generate the leading-monomial ideal for a fixed *term order*, meaning a multiplicative well-order of monomials; the leading monomial is the largest monomial in that order. An elimination-order basis computes this intersection. We use only the elementary fact that a Gröbner basis over $\mathbb{Q}$ remains one after extending coefficients to $\mathbb{C}$.

For $S \subseteq \mathbb{C}$, $\mathbb{Q}(S)$ denotes the smallest subfield of $\mathbb{C}$ containing $\mathbb{Q} \cup S$. Elements $a_1, \ldots, a_m$ are *algebraically independent over* $\mathbb{Q}$ if no nonzero polynomial in $\mathbb{Q}[X_1, \ldots, X_m]$ vanishes at them. The *transcendence degree* $\operatorname{trdeg}_{\mathbb{Q}} K$ is the maximum cardinality of an algebraically independent subset of K; the *coordinate field* of a tensor is the field generated by all of its coordinates. The *compositum* of subfields is the smallest field containing all of them, and transcendence degree is subadditive under finite composita. By *finite Zariski avoidance* we mean choosing a point outside a finite union of proper Zariski-closed sets.

For the number-theoretic argument, the *square-class group* $\mathbb{Q}^\times/(\mathbb{Q}^\times)^2$ identifies nonzero rationals whose quotient is a rational square. A *multiquadratic field* is a field $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_m})$ obtained by adjoining finitely many square roots. These terms occur only in the proof of the explicit transcendental table.

4.6 Bases and valid common presentations

For an integer $t \geq 0$, a width-t base from Boolean blocks to a finite domain D is a matrix

$$M \in \mathbb{C}^{|D| \times 2^t}. \tag{17}$$

Rows are indexed by D and columns by t-bit strings. This is Xia's orientation [24]; Chen uses the transpose orientation [11].

A row of the form e_b^T encodes its domain state by the Boolean *block codeword* $b \in \{0,1\}^t$. For several blocks we abbreviate

$$e_{x_1,\dots,x_m} := e_{x_1} \otimes \cdots \otimes e_{x_m}.$$

If $G = M^{\otimes b} H$, then H is a *Boolean preimage* of the domain signature G under M. A *base collapse* is the replacement of a valid width-t presentation by a valid presentation of the same problem at smaller width. A *holographic algorithm* uses such a presentation to replace the domain tensor network by a Boolean matchgate network and hence by weighted perfect matching; the present paper studies exact representability and does not infer an algorithmic running time from width alone.

We view a left domain signature $F\colon D^a \to \mathbb{C}$ as a row tensor and a right Boolean signature $H\colon (\{0,1\}^t)^b \to \mathbb{C}$ as a column tensor. We call b the *block arity* of H and write $\operatorname{barity}(H) = b$; its ordinary Boolean arity is bt. Their blockwise transforms are

$$\widehat{F} = FM^{\otimes a}, \qquad \widehat{F}(x_1,\dots,x_a) = \sum_{\mathbf{i}\in D^a} F(\mathbf{i}) \prod_{j=1}^{a} M_{i_j,x_j}, \tag{18}$$

$$G = M^{\otimes b} H, \qquad G(\mathbf{i}) = \sum_{x_1,\dots,x_b\in\{0,1\}^t} H(x_1,\dots,x_b) \prod_{j=1}^{b} M_{i_j,x_j}. \tag{19}$$

Every x_j is a consecutive t-bit block, and the displayed block order is the external linear order compatible with the planar cyclic order. This is *block-major order*: list every wire of block 1, then every wire of block 2, and so on. The traditional names for the transformed left signatures $FM^{\otimes a}$ and the Boolean right preimages H are *generators* and *recognizers*, respectively. *Simultaneous realizability* means that the same base makes all of these labelled tensors exact matchgates. A *full-rank generator* is a generator whose matrix obtained by choosing one Boolean block as output and all remaining blocks as inputs has rank equal to $\operatorname{rank} M = 2^r$; writing that matrix as AM, its left factor A is then a full-rank matchgate realizer in the sense of Section 4.4. This phrase is used only in the historical discussion. A *singleton symmetric left family* contains one signature invariant under every permutation of its ports; a *blockwise-symmetric* Boolean signature is invariant under permutations of its equal-sized Boolean blocks. These historical terms occur only in the literature discussion. Equations (18)–(19) also fix the tensor variance used throughout the paper.

For a Boolean column tensor Q on k consecutive r-bit blocks and Boolean row tensors $g_1,\dots,g_k$ on one r-bit block, define the multilinear, non-Hermitian contraction

$$\operatorname{Contr}(Q; g_1,\dots,g_k) := \sum_{x_1,\dots,x_k\in\{0,1\}^r} Q(x_1,\dots,x_k) \prod_{j=1}^{k} g_j(x_j). \tag{20}$$

No complex conjugation occurs. For $r \geq 1$, the contraction is represented by the closed block-star tensor network in Figure 2. Each arm contains r distinct ordered Boolean wires, whose values form $x_j = (x_{j,1},\dots,x_{j,r}) \in \{0,1\}^r$. These entries are separate summation indices; in particular, there is no copying constraint $x_{j,1} = \cdots = x_{j,r}$. The schematic is drawn for $k \geq 3$; when $k = 1$ or 2, only the corresponding distinct arms remain.

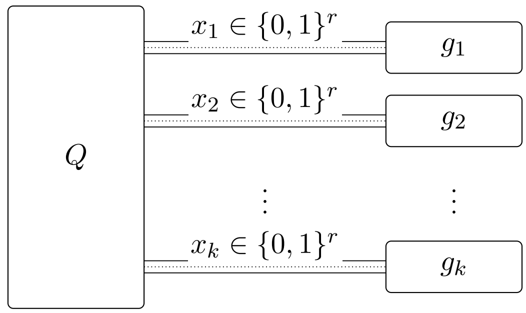

Figure 2: The closed block-star tensor network representing $\mathrm{Contr}(Q; g_1, \dots, g_k)$. The central tensor Q has k consecutive r-port blocks, and the jth arm connects g_j to block j. The within-arm order is fixed in the paragraph below. Each three-stroke bundle schematically represents the r distinct wires of an arm and is not a literal wire count; intermediate arms are suppressed.

For $r = 0$, every block is the unique empty word, no wires are present, and the contraction reduces to $Q \prod_{j=1}^{k} g_j$.

We use the identity block-order convention in (20): the jth block and its internal coordinate order in Q contract with g_j in that same order. There is one planar-order point to check. Noncrossing physical gluing reads the ports of the leaf in the reverse order from the central block. Realize on the jth arm the reflected exact matchgate $g_j^{\mathrm{rev}} = g_j \circ \rho_r$, which is exact by Lemma 4.1. If the central coordinate is x_j, the leaf is physically read as $\rho_r(x_j)$ and contributes $g_j^{\mathrm{rev}}(\rho_r(x_j)) = g_j(x_j)$. Thus the algebraic same-order contraction in (20) is an exact planar matchgate contraction whenever Q and the g_j are exact matchgates.

In Xia's output-port convention, the same scalar is written, after the coordinate reindexing $Q^\rho(x_1, \dots, x_k) := Q(\rho_r(x_1), \dots, \rho_r(x_k))$, as

$$\sum_{x_1,\dots,x_k \in \{0,1\}^r} Q^\rho(x_1, \dots, x_k) \prod_{j=1}^{k} g_j(\rho_r(x_j)).$$

Here Q^ρ is only a fixed renaming of coordinates for comparison of conventions; we do not assert that an arbitrary blockwise reversal of Q is itself a matchgate operation. The renaming uses exactly the same coordinate sets and hence does not alter any coordinate-field or transcendence-degree bound. We suppress these fixed permutations below.

We also use partial contraction. For $J \subseteq [k]$, define the tensor on the uncontracted blocks by

$$\mathrm{Contr}_J\big(Q; (g_j)_{j\in J}\big)\big((x_\ell)_{\ell\notin J}\big) := \sum_{(x_j)_{j\in J}} Q(x_1, \dots, x_k) \prod_{j\in J} g_j(x_j). \tag{21}$$

In an expression such as $\mathrm{Contr}(Q; g_1, g_2, -, \dots, -)$, each dash marks an uncontracted block and the surviving blocks retain their original order. The reflected-leaf argument above and ordered planar matchgate closure from Section 4.3 show that such a partial contraction is again an exact matchgate whenever all participating tensors are exact matchgates.

Definition 4.4 (Valid holographic presentation)**.** Let $\mathcal{F}$ be a finite labelled family of domain-D row signatures, let $\mathcal{H}$ be a finite labelled family of Boolean column signatures with specified block arities, and let M be a width-t base. The triple $(\mathcal{F}, \mathcal{H}, M)$ is a *valid common planar holographic presentation* if

$$FM^{\otimes\,\mathrm{arity}(F)} \in \mathrm{MG}_{t\,\mathrm{arity}(F)} \quad \text{for every } F \in \mathcal{F}, \qquad H \in \mathrm{MG}_{t\,\mathrm{barity}(H)} \quad \text{for every } H \in \mathcal{H},$$

in the prescribed blockwise cyclic orders. Its induced domain problem is

$$\mathcal{F} \mid M\mathcal{H}, \qquad M\mathcal{H} := \{M^{\otimes \operatorname{barity}(H)} H : H \in \mathcal{H}\},$$

with labels inherited from $\mathcal{H}$. The same base M is used for every label on both sides.

The presentation is *over a subfield* $K \subseteq \mathbb{C}$ if every coordinate of M, every domain tensor, and every Boolean preimage lies in K, and the required exact planar matchgates can be realized with edge weights in K. We use *valid presentation* and *valid common matchgate presentation* as abbreviations for the object in Definition 4.4. A statement made *labelwise* means separately for each original label, without merging or replacing the label family.

Associativity of tensor contraction gives the holographic identity

$$\operatorname{Holant}_I(\mathcal{F} \mid M\mathcal{H}) = \operatorname{Holant}_I(\mathcal{F}M \mid \mathcal{H}) \tag{22}$$

for every corresponding planar instance, where $\mathcal{F}M$ denotes the family of transforms in (18). The cyclic order is unchanged.

4.7 Realized supports, Gaussian mobility, and flags

Fix a qutrit presentation $(\mathcal{F}, \mathcal{H}, M)$ with $V = \mathbb{C}^3$ and a width-t base M of row rank three. Consider the all-left gadget closure of its induced problem $\mathcal{F} \mid M\mathcal{H}$. For a connected all-left gadget Γ with nonzero boundary tensor T_Γ, and a distinguished boundary edge e, the subspace

$$S_e(\Gamma) := S_{j(e)}(T_\Gamma) \le V^*$$

is a *realized support*. Thus "realized" always means that the support occurs at an actual port of an actual connected planar gadget, not merely in a linear combination of tensors. Define

$$\mathcal{S}_d := \{S_e(\Gamma) : \dim S_e(\Gamma) = d\}, \qquad d \in \{1, 2\}. \tag{23}$$

The *support geometry* of the closure is the pair $(\mathcal{S}_1, \mathcal{S}_2)$ together with the containment relations among its realized rays and planes and, for each realized plane, the two parity endpoints induced by M. *Cyclic rerooting* at a distinguished t-wire Boolean block means rotating the boundary order until that block comes first and designating it as the *input* block; the wires inside the block retain their inherited order. To check the coordinates, let s wires remain, write the physical block flattening of the rerooted matchgate tensor G as

$$F_G(z, w) := G(z, w), \qquad z \in \{0,1\}^t, \quad w \in \{0,1\}^s,$$

and let

$$A_G(z, y) := G\big(z, \rho_s(y)\big)$$

be its t-input, s-output exact matchgate matrix. Thus $A_G = F_G R_s$ for the column-permutation matrix R_s induced by ρ_s, and

$$\operatorname{col}(F_G) = \operatorname{col}(A_G) = \operatorname{row}(A_G^\mathsf{T}). \tag{24}$$

We call $\operatorname{col}(F_G)$ the *block support* at the distinguished block; it is the support of the whole block, rather than of one Boolean wire. By Lemma 4.3, A_G^T is exact. Hence the reversal acts only on the complementary column index and never on the distinguished block. Consequently a realized line maps literally to a pure matchgate ray, and a realized plane P maps literally to the rank-two Gaussian row space PM, with pullback parity endpoints P^+, P^- as defined in Section 4.4.

Definition 4.5 (Monomial flag and connection-primality)**.** Suppose $P \in \mathcal{S}_2$ and $L \in \mathcal{S}_1$ with $L \not\subset P$. The pair (P, L) is a *monomial flag* if

$$\mathcal{S}_2 = \{P\}, \qquad \mathcal{S}_1 \subseteq \{P^+, P^-, L\}.$$

The all-left gadget closure is *connection-prime* if its realized supports span V^*, $\mathcal{S}_2$ is nonempty, and it admits no monomial flag.

An *adapted ray basis* for a flag (P, L) is an ordered basis of V^* obtained by choosing nonzero generators of P^+, P^-, and L. Changing those generators only rescales the three coordinates.

The names used in the trichotomy are now literal. A *ray-only* or *ray-separable* closure has $\mathcal{S}_2 = \varnothing$. A *Gaussian-mobile* closure has a realized plane but no monomial flag; the proofs show that its moving planes or rays generate a constant-rank Gaussian cover. A *one-plane flag closure* admits a monomial flag. A *Boolean matchgate core on P times pure-ray factors* is a tensor obtained by encoding an m-bit signature $h \in \mathrm{MG}_m$ in an ordered endpoint basis (p_0, p_1) with $p_0 \in P^+$ and $p_1 \in P^-$ at the plane-supported ports, and tensoring nonzero generators of P^+, P^-, or L at the remaining ports. The exact coordinate formula is proved in Lemma 10.9. *Rank-one stripping* is the elementary step of factoring off such a ray at a line-supported port. For a fixed adapted basis, the *residual right tables* are the labelled coordinate arrays obtained from induced right tensors after contracting selected ports with fixed pure-ray factors and expressing all remaining plane-supported ports in the dual endpoint coordinates. Their port order and labels are retained. This phrase names the right-side data left uncontrolled by the all-left flag normal form; it does not introduce a new gadget operation or a new notion of equivalence.

5 Matchgate Parameter Bounds and Star-Table Varieties

The technical framework of Sections 4.3 and 4.5 turns the width question into a parameter-counting problem. We begin the negative lane of Figure 1 by asking how much independent information a width-r labelled matchgate star can encode. We first bound the parameter dimension available to a single matchgate and then propagate that bound through sampled block contractions. The resulting $\mathbb{Q}$-defined star-table closures are the low-width algebraic loci that Section 6 will avoid.

For $s \geq 0$, set

$$\delta(s) := 1 + \binom{s}{2}.$$

Lemma 5.1 (Coordinate transcendence bound)**.** *For every integer $s \geq 0$ and every $G \in \mathrm{MG}_s$,*

$$\operatorname{trdeg}_{\mathbb{Q}} \mathbb{Q}(G(x) : x \in \{0,1\}^s) \leq \delta(s). \tag{25}$$

Proof. For $s = 0$, the signature consists of one scalar, so the claim is immediate. The zero signature is also immediate. Suppose therefore that $s \geq 1$ and $G \neq 0$. Choose a pivot $X \in \{0,1\}^s$ with $c := G(X) \neq 0$.

Matchgate parity implies that $G(Y) = 0$ whenever $\mathrm{dist}(X, Y)$ is odd. The standard pivoted Pfaffian chart for the matchgate identities [6] gives fixed signs $\eta_{ij} \in \{+1, -1\}$ such that

$$a_{ij} = \eta_{ij} \frac{G(X \oplus \mathbf{1}_{\{i,j\}})}{c} \qquad (1 \leq i < j \leq s) \tag{26}$$

are the upper-triangular entries of a skew-symmetric matrix A_X. For every even subset $T \subseteq [s]$, there is another fixed sign $\eta_T \in \{+1, -1\}$ for which

$$G(X \oplus \mathbf{1}_T) = \eta_T c \, \mathrm{Pf}(A_X[T]). \tag{27}$$

Here $\mathbf{1}_T$ is the indicator bit string of T, and $\oplus$ is bitwise exclusive-or. The signs depend only on the pivot X, the indicated index set, and the ordered-coordinate convention, not on the coordinate values of G.

Equations (26) and (27) show that every coordinate of G lies in

$$\mathbb{Q}\big(c, \{G(X \oplus \mathbf{1}_{\{i,j\}}) : 1 \leq i < j \leq s\}\big).$$

This field has at most $1 + \binom{s}{2}$ generators over $\mathbb{Q}$, which proves (25). The argument uses a pivot of the particular nonzero signature and treats both parity components; the zero signature was handled separately. $\square$

The next proposition packages the information bound in the form needed for labelled closed-star instances. We call L the *sampler alphabet*, the family $(g_\ell)_{\ell \in L}$ the *samplers*, and T the corresponding *sampled star table*. When $L = \{0, 1\}$, (g_0, g_1) is the *sampler pair*, and the displayed contraction is a *star-table representation*.

Proposition 5.2 (Sampled star-table bound)**.** *Let L be a finite set of size $q \geq 1$, and let $k \geq 1$ and $r \geq 0$. Suppose*

$$g_\ell \in \mathrm{MG}_r \quad (\ell \in L), \qquad Q \in \mathrm{MG}_{kr}\,.$$

Define a table $T\colon L^k \to \mathbb{C}$ by

$$T(\ell_1, \ldots, \ell_k) = \mathrm{Contr}(Q; g_{\ell_1}, \ldots, g_{\ell_k}), \tag{28}$$

where the jth r-bit tensor factor is contracted with the jth consecutive block of Q. Then

$$\mathrm{trdeg}_{\mathbb{Q}}\, \mathbb{Q}(T(\boldsymbol{\ell}) : \boldsymbol{\ell} \in L^k) \leq q\delta(r) + \delta(kr). \tag{29}$$

Proof. Let K be the field generated over $\mathbb{Q}$ by all coordinates of $\{g_\ell : \ell \in L\}$ and Q. Every coordinate in (28) is a polynomial with integer coefficients in these generators, and hence

$$\mathbb{Q}(T(\boldsymbol{\ell}) : \boldsymbol{\ell} \in L^k) \subseteq K.$$

Transcendence degree is subadditive under field composita. Applying Lemma 5.1 to the q unary-block signatures and to Q gives

$$\mathrm{trdeg}_{\mathbb{Q}}\, K \leq \sum_{i=1}^{q} \delta(r) + \delta(kr) = q\delta(r) + \delta(kr),$$

as required. $\square$

We now promote the same parameter count to an algebraic-geometric statement. This is the step that permits an integer obstruction even though integer coordinates have transcendence degree zero.

For $s \geq 1$ and $\nu \in \{0, 1\}$, let $\mathfrak{M}_s^\nu \subseteq \mathbb{A}_{\mathbb{Q}}^{2^s}$ be the affine algebraic set obtained by imposing the parity-ν coordinate equations $Y_x = 0$ whenever $|x| \not\equiv \nu \pmod 2$, together with the matchgate identities in the fixed external order. All these equations have integer coefficients. Set

$$\mathfrak{M}_s := \mathfrak{M}_s^0 \cup \mathfrak{M}_s^1.$$

This is a $\mathbb{Q}$-defined closed algebraic set; its defining ideal may be taken to be the intersection of the two parity-component ideals. Necessity and sufficiency of the matchgate identities, together with parity, imply

$$\mathfrak{M}_s(\mathbb{C}) = \mathrm{MG}_s, \tag{30}$$

where the zero tensor belongs to both parity components [6]. Separately, set $\mathfrak{M}_0 := \mathbb{A}_{\mathbb{Q}}^1$, so that $\mathfrak{M}_0(\mathbb{C}) = \mathrm{MG}_0 = \mathbb{C}$.

Lemma 5.3 (Dimension of the matchgate variety)**.** *For every $s \geq 0$,*

$$\dim \mathfrak{M}_s \leq \delta(s) = 1 + \binom{s}{2}. \tag{31}$$

Proof. The assertion is immediate for $s = 0$. For $s \geq 1$, dimension is unchanged under the field extension $\mathbb{Q} \subseteq \mathbb{C}$, so we work after base change to $\mathbb{C}$.

Fix one parity component and a bit string X of that parity. Let U_X be the principal open subset on which the coordinate indexed by X is nonzero. For fixed pivot-dependent signs $\eta_{X,T} \in \{+1, -1\}$, define a regular map

$$\psi_X : \mathbb{G}_m \times \mathbb{A}^{\binom{s}{2}} \longrightarrow \mathbb{A}^{2^s}$$

by

$$\psi_X(c, A)(X \oplus \mathbf{1}_T) = \begin{cases} \eta_{X,T}\, c\, \mathrm{Pf}(A[T]), & |T| \text{ even}, \\ 0, & |T| \text{ odd}, \end{cases} \qquad \eta_{X,\varnothing} = 1, \tag{32}$$

where A ranges over the skew-symmetric $s \times s$ matrices. This map is defined over $\mathbb{Q}$. The pivoted Pfaffian chart (26)–(27) shows that every point of U_X lies in the image of ψ_X. Hence

$$\dim U_X \leq \dim \left(\mathbb{G}_m \times \mathbb{A}^{\binom{s}{2}}\right) = 1 + \binom{s}{2}.$$

The finitely many U_X cover all nonzero points of the chosen parity component; adjoining the zero point does not increase dimension. Taking the maximum over the two parity components proves the claim. □

Fix $k \geq 1$ and $r \geq 0$. On

$$\mathfrak{X}_{k,r} := \mathfrak{M}_r \times \mathfrak{M}_r \times \mathfrak{M}_{kr}, \tag{33}$$

define the $\mathbb{Q}$-morphism

$$\Phi_{k,r} : \mathfrak{X}_{k,r} \longrightarrow \mathbb{A}^{2^k}_{\mathbb{Q}}.$$

Index the target coordinates by bit strings $\boldsymbol{z} = (z_1, \ldots, z_k) \in \{0,1\}^k$ and set

$$[\Phi_{k,r}(g_0, g_1, Q)](\boldsymbol{z}) := \mathrm{Contr}(Q; g_{z_1}, \ldots, g_{z_k}) \qquad (\boldsymbol{z} \in \{0,1\}^k). \tag{34}$$

Every output coordinate is a polynomial with integer coefficients in the source coordinates. This formula also covers $r = 0$: the three matchgate signatures are then scalars and every block has the unique empty word. Identifying $\mathbb{A}^{2^k}_{\mathbb{Q}}(\mathbb{C})$ with the table space $\mathbb{C}^{\{0,1\}^k}$, let

$$\mathcal{S}_{k,r} := \Phi_{k,r}(\mathfrak{X}_{k,r}(\mathbb{C})) \subseteq \mathbb{C}^{\{0,1\}^k}, \qquad \mathfrak{Z}_{k,r} := \overline{\mathcal{S}_{k,r}}^{\mathrm{Zar}} \subseteq \mathbb{C}^{\{0,1\}^k}. \tag{35}$$

Proposition 5.4 (Algebraic star-table bound)**.** *The variety $\mathfrak{Z}_{k,r}$ is defined over $\mathbb{Q}$ and satisfies*

$$\dim \mathfrak{Z}_{k,r} \leq D_{k,r} := 2\delta(r) + \delta(kr). \tag{36}$$

In particular, if $D_{k,r} < 2^k$, then $\mathfrak{Z}_{k,r}$ is a proper $\mathbb{Q}$-defined Zariski-closed subset of $\mathbb{C}^{\{0,1\}^k}$.

Proof. By Lemma 5.3,

$$\dim \mathfrak{X}_{k,r} \leq 2\delta(r) + \delta(kr).$$

The dimension of the Zariski closure of the image of a morphism is at most the dimension of its source, proving (36).

It remains to verify that $\mathfrak{Z}_{k,r}$ is defined over $\mathbb{Q}$. Let $\mathbf{U}$ denote all source-coordinate variables and let $\mathbf{Y} = (Y_{\boldsymbol{z}})_{\boldsymbol{z}\in\{0,1\}^k}$ be the target coordinates. Choose a $\mathbb{Q}$-coefficient ideal

$$I_{k,r} \subseteq \mathbb{Q}[\mathbf{U}]$$

whose complex zero set is $\mathfrak{X}_{k,r}(\mathbb{C})$. In $\mathbb{Q}[\mathbf{U}, \mathbf{Y}]$, form the graph ideal

$$J_{k,r} := I_{k,r} + \left\langle Y_{\boldsymbol{z}} - [\Phi_{k,r}(\mathbf{U})](\boldsymbol{z}) : \boldsymbol{z} \in \{0,1\}^k \right\rangle$$

and its elimination ideal

$$E_{k,r} := J_{k,r} \cap \mathbb{Q}[\mathbf{Y}].$$

After extending scalars from $\mathbb{Q}$ to $\mathbb{C}$, the elimination theorem gives

$$\mathfrak{Z}_{k,r} = V_{\mathbb{C}}(E_{k,r}).$$

Equivalently, an elimination-order Gröbner basis computed over $\mathbb{Q}$ remains a Gröbner basis after extension to $\mathbb{C}$. Thus $\mathfrak{Z}_{k,r}$ is defined over $\mathbb{Q}$. If $D_{k,r} < 2^k$, its dimension is smaller than that of the ambient affine space, so it is proper. $\square$

6 Two Obstruction Tables

The parameter budget of Section 5 reduces the next step to an avoidance problem: find a 2^k-entry table outside every relevant low-width star-table closure. Using the tensor and algebraic-geometric conventions of Sections 4.1 and 4.5, we obtain two nowhere-zero witnesses whose one-versus-rest flattenings all have rank two: a closed-form transcendental table and a positive-integer table selected by finite Zariski avoidance. The integer witness drives the rational qutrit construction, while the transcendental witness gives an explicit companion. The next section addresses the complementary question of realizing these hard tables exactly at a larger finite width.

6.1 A closed-form transcendental table

We next construct a table with the maximum possible coordinate transcendence degree 2^k. The only transcendence-theoretic input is the following standard consequence of the Lindemann–Weierstrass theorem.

Lemma 6.1. *Let $p_1, \ldots, p_N$ be distinct primes. Then*

$$\exp(\sqrt{p_1}), \ldots, \exp(\sqrt{p_N})$$

are algebraically independent over $\mathbb{Q}$.

Proof. The prime square classes are independent in $\mathbb{Q}^\times/(\mathbb{Q}^\times)^2$: unique factorization shows that a product $\prod_i p_i^{\epsilon_i}$ with $\epsilon_i \in \{0,1\}$ is a square only when every $\epsilon_i = 0$. Hence the multiquadratic field $\mathbb{Q}(\sqrt{p_1}, \ldots, \sqrt{p_N})$ has the automorphisms that independently change the sign of any one square root; an automorphism here is a bijection preserving addition, multiplication, and every rational number. Applying such an automorphism to a putative rational linear relation and subtracting shows that $\sqrt{p_1}, \ldots, \sqrt{p_N}$ are linearly independent over $\mathbb{Q}$.

Suppose that a nonzero polynomial $P \in \mathbb{Q}[X_1, \ldots, X_N]$ vanished at the displayed exponentials. Expanding its monomials would give

$$\sum_{\mathbf{m}\in\mathcal{S}} c_{\mathbf{m}} \exp\left(\sum_{i=1}^{N} m_i\sqrt{p_i}\right) = 0, \tag{37}$$

for a finite nonempty $\mathcal{S} \subseteq \mathbb{Z}_{\geq 0}^N$ and nonzero rational coefficients $c_{\mathbf{m}}$. Linear independence of the square roots makes the algebraic exponents in (37) pairwise distinct, including the zero exponent if P has a constant monomial. The Lindemann–Weierstrass theorem says that exponentials of distinct algebraic numbers are linearly independent over $\overline{\mathbb{Q}}$ [20]. This contradicts (37). $\square$

Fix $k \geq 2$, and let $p_1, \ldots, p_{2^k}$ be the first 2^k primes. For $z = (z_1, \ldots, z_k) \in \{0,1\}^k$, define its one-based lexicographic index by

$$\operatorname{ind}_k(z) := 1 + \sum_{j=1}^{k} z_j 2^{k-j} \in [2^k], \tag{38}$$

and set

$$f_k^{\mathrm{tr}}(z) := \exp\Big(\sqrt{p_{\operatorname{ind}_k(z)}}\Big). \tag{39}$$

Every entry is positive and nonzero. By Lemma 6.1,

$$\operatorname{trdeg}_{\mathbb{Q}} \mathbb{Q}(f_k^{\mathrm{tr}}(z) : z \in \{0,1\}^k) = 2^k. \tag{40}$$

Proposition 6.2 (Explicit transcendental obstruction)**.** *Every one-versus-rest flattening of f_k^{tr} has rank two. Moreover, if $r \geq 0$, $g_0, g_1 \in \mathrm{MG}_r$, and $Q \in \mathrm{MG}_{kr}$ satisfy*

$$f_k^{\mathrm{tr}}(z) = \operatorname{Contr}(Q; g_{z_1}, \ldots, g_{z_k}) \qquad (z \in \{0,1\}^k),$$

then

$$2^k \leq 2\delta(r) + \delta(kr). \tag{41}$$

Proof. For each port, choose two distinct assignments to the other $k-1$ coordinates. The corresponding 2-by-2 determinant is the difference of two distinct degree-two monomials in four different entries of f_k^{tr}. It is nonzero because those entries are algebraically independent by (40). Thus every mode flattening has rank two.

Apply Proposition 5.2 with the sampler alphabet $L = \{0,1\}$, then use (40). $\square$

6.2 A positive-integer table

We use the following elementary density fact.

Lemma 6.3 (Density of the positive integer grid)**.** *For every $N \geq 1$, the set $\mathbb{Z}_{>0}^N$ is Zariski dense in $\mathbb{C}^N$.*

Proof. We induct on N. A univariate polynomial vanishing at every positive integer is zero. For $N > 1$, write a polynomial vanishing on $\mathbb{Z}_{>0}^N$ as a polynomial in its last variable. After fixing any positive integer values of the first $N-1$ variables, the resulting univariate polynomial has infinitely many roots, so every coefficient vanishes on $\mathbb{Z}_{>0}^{N-1}$. The induction hypothesis makes all coefficient polynomials zero. $\square$

Proposition 6.4 (Positive-integer obstruction)**.** *For every $k \geq 2$, there exists a table*

$$f_k^{\mathrm{int}} \colon \{0,1\}^k \longrightarrow \mathbb{Z}_{>0} \tag{42}$$

such that, for every $r \geq 0$, $g_0, g_1 \in \mathrm{MG}_r$, and $Q \in \mathrm{MG}_{kr}$, the identities

$$f_k^{\mathrm{int}}(z) = \operatorname{Contr}(Q; g_{z_1}, \ldots, g_{z_k}) \qquad (z \in \{0,1\}^k) \tag{43}$$

imply

$$2^k \leq 2\delta(r) + \delta(kr). \tag{44}$$

The conclusion allows the representing matchgates to have arbitrary complex weights. In addition, every one-versus-rest flattening of f_k^{int} has rank two.

Proof. Let

$$R_k := \{r \in \mathbb{Z}_{\geq 0} : D_{k,r} < 2^k\}.$$

This set is finite because $D_{k,r} = 2\delta(r) + \delta(kr)$ tends to infinity with r. Moreover, $D_{k,0} = 3 < 2^k$ for $k \geq 2$, so $0 \in R_k$. For each $r \in R_k$, Proposition 5.4 makes $\mathfrak{Z}_{k,r}$ a proper $\mathbb{Q}$-defined closed subset of $\mathbb{C}^{\{0,1\}^k}$. Its defining ideal therefore contains a nonzero polynomial $P_{k,r} \in \mathbb{Q}[(Y_{\boldsymbol{z}})_{\boldsymbol{z}\in\{0,1\}^k}]$. The finite product

$$P_k := \prod_{r \in R_k} P_{k,r}$$

is nonzero. For every port j, choose two distinct assignments $u_j, v_j \in \{0,1\}^{k-1}$ and let Δ_j be the corresponding symbolic 2-by-2 flattening minor. Each Δ_j is a nonzero polynomial, so

$$P_k \prod_{j=1}^{k} \Delta_j$$

is nonzero. By Lemma 6.3, choose $a = (a_{\boldsymbol{z}})_{\boldsymbol{z}\in\{0,1\}^k} \in \mathbb{Z}_{>0}^{2^k}$ on which this product does not vanish, and set $f_k^{\mathrm{int}}(\boldsymbol{z}) = a_{\boldsymbol{z}}$. The nonzero minor $\Delta_j(a)$ gives rank at least two in mode j; the two-row shape gives rank at most two. Thus every mode rank is exactly two.

Suppose (43) holds at width r. Then f_k^{int} lies in the exact image $\mathcal{S}_{k,r}$ and hence in $\mathfrak{Z}_{k,r}$. Because $\mathcal{S}_{k,r}$ was defined using all complex points of the matchgate parameter space, this implication places no algebraicity or field-of-definition restriction on the representing matchgates. If $D_{k,r} < 2^k$, then $r \in R_k$, which would force $P_{k,r}(f_k^{\mathrm{int}}) = 0$, contrary to the choice of a. Therefore $D_{k,r} \geq 2^k$, proving (44). □

Remark 6.5. The integer table is existential but effective in principle: for each of the finitely many relevant widths one may compute elimination equations, then enumerate positive integer tables until their product does not vanish. The proof supplies no useful running-time, height, or bit-length bound. By contrast, f_k^{tr} has a simple closed formula but uses transcendental constants. Keeping both versions separates these two tradeoffs cleanly.

7 Pfaffian Interpolation on Two Block Codewords

The tables of Section 6 force every matchgate-star representation to obey the width inequality derived in Section 5, but a counterexample also requires an exact presentation at some finite width. Using the Pfaffian and ordered-block conventions of Sections 4.3 and 4.6, we solve this realization problem in a field-preserving form: every nowhere-zero Boolean table can be interpolated by one exact matchgate on the two repetition codewords in each block. This construction supplies the Boolean preimage that will be packaged into a qutrit presentation in the next section.

Theorem 7.1 (Block-Pfaffian interpolation)**.** *Let $K \subseteq \mathbb{C}$ be a subfield, let $k \geq 1$, let $f\colon \{0,1\}^k \to K^\times$, and put $t = 2^k$. There is an exact matchgate signature $H_f \in \mathrm{MG}_{kt}$ with coordinates in K, on k consecutive t-bit blocks, such that, for*

$$c_0 := 0^t, \qquad c_1 := 1^t,$$

one has

$$H_f(c_{z_1},\dots,c_{z_k}) = f(z_1,\dots,z_k) \qquad \textit{for every } z\in\{0,1\}^k. \tag{45}$$

More precisely, the construction produces a named skew-symmetric matrix $A_f \in K^{kt\times kt}$ whose principal-Pfaffian signature is normalized by

$$\Gamma_{A_f}(c_{z_1},\dots,c_{z_k}) = \frac{f(z)}{f(0^k)}, \qquad \Gamma_{A_f}(c_0,\dots,c_0) = 1, \tag{46}$$

and $H_f = f(0^k)\Gamma_{A_f}$. The blocks and all wires within them occur in the specified external linear order. Moreover, H_f has an exact planar matchgate realization whose edge weights all lie in K.

The construction has one independent Pfaffian component for each nonempty $S \subseteq [k]$. Its principal Pfaffian will equal a multiplier y_S when every block in S is selected and will equal 1 otherwise. Disjointness then multiplies these contributions and implements a multiplicative Möbius expansion of f.

Proof. We first decompose the table multiplicatively. For each nonempty subset $S \subseteq [k]$, define $y_S \in K^\times$ recursively under inclusion by

$$y_S := \frac{f(\mathbf{1}_S)/f(0^k)}{\displaystyle\prod_{\varnothing\neq U\subsetneq S} y_U}, \tag{47}$$

where $\mathbf{1}_S \in \{0,1\}^k$ is the indicator of S. The denominator is nonzero. The recursion determines the multipliers in increasing order of $|S|$. Induction on $|T|$ gives

$$\prod_{\varnothing\neq S\subseteq T} y_S = \frac{f(\mathbf{1}_T)}{f(0^k)} \qquad (T\subseteq[k]), \tag{48}$$

with the empty product interpreted as 1.

We now define a skew-symmetric matrix whose principal Pfaffians implement the left side of (48). In logical block $i \in [k]$, create an adjacent ordered pair of modes

$$(a_{S,i}, b_{S,i})$$

for every subset $S \subseteq [k]$ containing i. A fixed block contains $2\cdot 2^{k-1} = 2^k = t$ modes. Order the blocks by increasing i; within each block, order the pairs lexicographically by S, with $a_{S,i}$ immediately before $b_{S,i}$; explicitly, subsets are compared by their indicator words $\mathbf{1}_S$ in lexicographic order.

Let A_f be the skew-symmetric matrix indexed by these kt ordered modes. All its entries lie in K. All unspecified upper-triangular entries are zero, and different subsets S have no entries between their modes. For a singleton $S = \{i\}$, set

$$A_f[a_{S,i}, b_{S,i}] = y_S. \tag{49}$$

For $S = \{i_1 < \cdots < i_s\}$ with $s \ge 2$, set the baseline entries

$$A_f[a_{S,i_q}, b_{S,i_q}] = 1 \qquad (1\le q\le s), \tag{50}$$

the consecutive alternate entries

$$A_f[b_{S,i_q}, a_{S,i_{q+1}}] = 1 \qquad (1\le q< s), \tag{51}$$

and the closing entry

$$A_f[a_{S,i_1}, b_{S,i_s}] = y_S - 1. \tag{52}$$

The entries in the reverse direction are fixed by skew-symmetry.

Fix one subset S and suppose a set $T \subseteq [k]$ of logical blocks is selected in full. The selected modes in the S-component are precisely the pairs indexed by $S \cap T$. If $|S| = 1$, then (49) immediately gives contribution y_S when $S \subseteq T$ and contribution 1 otherwise. Assume henceforth that $|S| \geq 2$. The baseline and alternate positions in (50)–(52) form one formal alternating cycle on the $2|S|$ modes. The Pfaffian expansion may be viewed on this cycle; if $y_S - 1 = 0$, the alternate matching term is simply zero. If $S \nsubseteq T$, deleting the unselected two-vertex pairs breaks this cycle into disjoint even paths. Each such path has its baseline edges as its unique nonzero perfect-matching monomial, so the component Pfaffian is 1; the same holds when no pair is selected because $\mathrm{Pf}(A_f[\varnothing]) = 1$.

If $S \subseteq T$, all modes of the component remain. The cycle has exactly two formal perfect-matching terms, one possibly zero. In the local order

$$a_{S,i_1}, b_{S,i_1}, \ldots, a_{S,i_s}, b_{S,i_s},$$

the baseline matching consists of adjacent pairs and has Pfaffian sign $+1$. The alternate matching consists of the pair in positions $(1, 2s)$ and the adjacent pairs $(2,3), (4,5), \ldots, (2s-2, 2s-1)$, so it is noncrossing—no two paired intervals interleave in the displayed linear order—and also has sign $+1$. Its weight is $y_S - 1$. Thus the full S-component contributes $1 + (y_S - 1) = y_S$. This reasoning also allows $y_S = 1$, in which case the second term is zero; a zero-weight edge may be omitted from the eventual planar matchgate.

Let Γ_{A_f} be the principal-Pfaffian signature from (11). Its coordinates are Pfaffian polynomials with integer coefficients in the entries of A_f, so they lie in K. The principal-Pfaffian relations imply the matchgate identities; by Lemma 4.2, Γ_{A_f} has an exact planar realization over K in the stated order. For a logical assignment $z \in \{0,1\}^k$, let $T = \{i : z_i = 1\}$ and evaluate Γ_{A_f} on the block word $(c_{z_1}, \ldots, c_{z_k})$. Regrouping the selected modes by S permutes only adjacent two-coordinate chunks $(a_{S,i}, b_{S,i})$. Every transposition of two chunks has sign $(-1)^{2\cdot 2} = +1$, and regrouping preserves increasing i-order within a fixed S. Pfaffian factorization over the disjoint components therefore introduces no global sign. The component analysis and (48) give

$$\Gamma_{A_f}(c_{z_1}, \ldots, c_{z_k}) = \prod_{\varnothing \neq S \subseteq T} y_S = \frac{f(z)}{f(0^k)}. \tag{53}$$

Finally, multiply the whole matchgate signature by the nonzero scalar $f(0^k) \in K$, for example by taking a disjoint internal edge of that weight. The resulting exact matchgate signature H_f is realized over K and satisfies (45). The matrix order was block-major throughout, so this construction also proves the asserted external linear order. □

Remark 7.2. The construction specifies the target signature exactly through (11), (47), (49)–(52), and the final scalar multiplication by $f(0^k)$. The field-preserving matchgate-identity sufficiency theorem supplies an exact planar realization over the same coefficient field. It is applied to the Pfaffian tensor Γ_{A_f}; the parameter matrix A_f itself is not asserted to be planar. No claim of polynomial-size encoding or polynomial-time construction is intended: the theorem concerns exact representability.

8 The Controlled Domain-Three Presentation

The interpolation theorem of Section 7 realizes the positive-integer obstruction from Section 6 on Boolean blocks, but it does not yet produce a qutrit language with the structural properties required by Theorem 3.1. Using the tensor-transform, exact-matchgate, and common-base conventions of Sections 4.1, 4.3 and 4.6, we add marker and control modes and a full-row-rank qutrit base, while retaining the hard table as an exact restriction of the resulting right tensor. We then verify the local ranks, two-sided support-essentiality, and connected two-center coupling. This completes the construction side; the next section asks whether an arbitrary equivalent presentation can compress it.

Let $f = f_k^{\mathrm{int}}$ be the positive-integer obstruction table from Proposition 6.4, put

$$t_0 := 2^k, \qquad T_k := t_0 + 3,$$

and use the normalized skew matrix A_f named in Theorem 7.1; normalized means precisely (46). Put $c_0 := 0^{t_0}$ and $c_1 := 1^{t_0}$. Thus, for $z \in \{0,1\}^k$, $\Gamma_{A_f}(c_{z_1}, \dots, c_{z_k}) = f(z)/f(0^k)$.

One qutrit block consists of T_k Boolean wires in the order

$$(\text{marker};\ t_0 \text{ selectors};\ \text{two controls}).$$

Define the three block codewords

$$b_0 = (0, 0^{t_0}, 00), \qquad b_1 = (1, 0^{t_0}, 00), \qquad b_2 = (0, 1^{t_0}, 11). \tag{54}$$

Since t_0 is even, b_0, b_2 have even parity and b_1 has odd parity. With e_x denoting a Boolean coordinate vector, set

$$M_k = \begin{pmatrix} e_{b_0}^{\mathsf{T}} \\ e_{b_1}^{\mathsf{T}} \\ e_{b_2}^{\mathsf{T}} \end{pmatrix} \in \mathbb{Q}^{3 \times 2^{T_k}}. \tag{55}$$

The rows are distinct coordinate matchgates, so M_k has row rank three.

The left family contains the pins $\delta_0, \delta_1, \delta_2$. For $a, b \in D_3$, let the qutrit matrix unit E_{ab} be defined by $E_{ab}(u, v) = 1$ when $(u, v) = (a, b)$ and 0 otherwise, and define the binary signature

$$X_{01} := E_{01} + E_{10}.$$

Their transforms are exact matchgates: the pins become $e_{b_a}^{\mathsf{T}}$, and

$$X_{01} M_k^{\otimes 2} = e_{b_0, b_1}^{\mathsf{T}} + e_{b_1, b_0}^{\mathsf{T}}, \tag{56}$$

which is the odd two-bit disequality matchgate on the two marker modes, tensored with fixed zero pins.

It remains to construct the right Boolean preimage. It has $k+2$ consecutive blocks: two *link blocks*, followed by k *hard blocks*. Fix one designated hard block $j \in [k]$, and set

$$\alpha = 1, \qquad \beta = 2, \qquad w(0) = \alpha, \quad w(2) = \beta.$$

Construct a skew matrix $\widehat{A}_f$ on these $(k+2)T_k$ modes as follows.

1. On the selector modes of the hard blocks, use A_f.
2. Give zero rows and columns to every hard marker and to every nonmarker mode of the two link blocks.

3. On the control pair of every nondesignated hard block, put one isolated 2-by-2 skew block of weight 1: the entry from the first control to the second is 1, its reverse is -1, and all other entries incident with that pair are zero.

4. Let p, q be the link markers and a, b the controls of the designated hard block. In the external order $p < q < a < b$, put

$$\widehat{A}_f[a,b] = 1, \quad \widehat{A}_f[p,q] = \alpha, \quad \widehat{A}_f[p,a] = 1, \quad \widehat{A}_f[q,b] = \alpha - \beta, \tag{57}$$

with every other entry involving these four modes equal to zero.

All other entries between the listed components are zero. Let $\widehat{H}_k$ be $f(0^k)$ times the principal-Pfaffian signature of $\widehat{A}_f$, and define

$$\mathcal{F}_k := \{\delta_0, \delta_1, \delta_2, X_{01}\}, \qquad \mathcal{H}_k := \{\widehat{H}_k\}, \qquad B_k := M_k^{\otimes(k+2)} \widehat{H}_k. \tag{58}$$

Proposition 8.1 (Exact restriction and rational presentation)**.** *The triple $(\mathcal{F}_k, \mathcal{H}_k, M_k)$ is a valid width-T_k exact matchgate presentation over $\mathbb{Q}$, with a full-row-rank base. Its right tensor is integer-valued. The only nonzero domain coordinates of B_k have link states in $\{0,1\}$ and hard states in $\{0,2\}$. For $x \in \{0,1\}^k$, writing $2x = (2x_1, \ldots, 2x_k)$, one has*

$$\begin{aligned} &B_k(0,0,2x) = f(x), && B_k(1,1,2x) = w(2x_j) f(x), \\ &B_k(0,1,2x) = 0, && B_k(1,0,2x) = 0. \end{aligned} \tag{59}$$

Proof. The left transforms were verified in (56). The scaled principal-Pfaffian signature $\widehat{H}_k$ has rational coordinates and, by Lemma 4.2, an exact planar realization with rational weights in the prescribed order. Hence the presentation is valid over $\mathbb{Q}$.

A principal-Pfaffian coordinate indexed by a codeword selects exactly the matrix rows and columns at its 1-bits. Hence a link state 2 selects isolated link selectors and controls, while a hard state 1 selects an isolated hard marker; either choice gives Pfaffian zero. The four-mode control block satisfies

$$\mathrm{Pf}(\widehat{A}_f[\varnothing]) = 1, \quad \mathrm{Pf}(\widehat{A}_f[\{a,b\}]) = 1, \quad \mathrm{Pf}(\widehat{A}_f[\{p,q\}]) = \alpha, \quad \mathrm{Pf}(\widehat{A}_f[\{p,q,a,b\}]) = \alpha - (\alpha - \beta) = \beta.$$

Selecting exactly one link marker has odd parity and gives zero. The hard selector blocks contribute $f(x)/f(0^k)$ by (53); the global scale restores $f(x)$. In each surviving case, all added selected pieces have even size, so regrouping introduces no Pfaffian shuffle sign. This proves (59) and the coordinate-support restriction. Since f and w are positive integers, B_k is integer-valued. $\square$

Proposition 8.2 (Deficiency and essentiality)**.** *Every port of X_{01} and B_k has flattening rank exactly two, while the pins have rank one. The language is support-essential on both sides, and B_k alone is support-essential on the right. Its link-port support is* $\mathrm{span}(e_0, e_1)$ *and its hard-port support is* $\mathrm{span}(e_0, e_2)$.

Proof. The assertion for X_{01} is immediate. At either link port of B_k, the state-0 and state-1 rows have disjoint nonzero supports, while the state-2 row vanishes. Thus the link support is exactly $\mathrm{span}(e_0, e_1)$.

At every hard port, the state-1 row vanishes. At a nondesignated port, the chosen nonzero 2-by-2 minor of the corresponding flattening $\mathrm{Mat}_j(f)$ makes the state-0 and state-2 rows independent. At the designated port, proportionality of those rows by a scalar $\lambda \in \mathbb{C}$ would, after fixing the other hard coordinates to some $x' \in \{0,1\}^{k-1}$, give the following. Let $x^{(a)} \in \{0,1\}^k$ be the extension of x'

having value a at coordinate j. The link slice $(0,0)$ gives $f(x^{(1)}) = \lambda f(x^{(0)})$, while the slice $(1,1)$ gives $\beta f(x^{(1)}) = \lambda\alpha f(x^{(0)})$. Since f is nowhere zero, this would force $\alpha = \beta$, a contradiction. Hence every hard support is exactly $\mathrm{span}(e_0, e_2)$.

These two planes span the right qutrit space. On the left, the three pins span the dual space. This proves support-essentiality and every rank claim. □

For $z \in \{0,2\}^k$, define the relabelled table

$$\widetilde{f}(z) := f(z/2), \qquad z/2 := (z_1/2, \ldots, z_k/2) \in \{0,1\}^k. \tag{60}$$

Proposition 8.3 (A genuinely coupled two-center gadget). *Write the two link ports of each copy, in their prescribed local cyclic order, as ℓ_1, ℓ_2, followed by the hard ports $h_1, \ldots, h_k$. Join ℓ_2^L to ℓ_1^R through one copy of X_{01}, and join ℓ_1^L to ℓ_2^R through another copy. Draw the first length-two path above the two centres and the second below them. Thus the link* names *are cross-paired, but the two drawn paths do not cross. At each centre the interior corner of the resulting four-cycle is the empty cyclic sector from ℓ_1 to ℓ_2; the complementary sector contains $h_1, \ldots, h_k$ and lies in the outer face. Both copies therefore retain their prescribed local port orders. Marking h_1^L first, the gadget boundary order can be taken to be*

$$h_1^L, \ldots, h_k^L, h_1^R, \ldots, h_k^R.$$

On hard assignments $z, y \in \{0,2\}^k$, its boundary tensor is

$$K_f(z,y) = \widetilde{f}(z)\widetilde{f}(y)\big(w(z_j) + w(y_j)\big), \tag{61}$$

Its Schmidt rank across $z \mid y$ is exactly two.

Proof. Let a and b be the common link states at the left and right centres, respectively. Cross-pairing the link names does not permute the arguments at either centre. Each B_k forces its two link states to agree, and the symmetry of X_{01} means that both copies impose the same condition $a \neq b$. Exactly the two pairs $(a,b) = (0,1)$ and $(1,0)$ survive, giving (61). The full boundary matrix is the sum of two outer products

$$[\widetilde{f}(z)w(z_j)]_z[\widetilde{f}(y)]_y^{\mathsf{T}} + [\widetilde{f}(z)]_z[\widetilde{f}(y)w(y_j)]_y^{\mathsf{T}},$$

and every entry involving hard state 1 vanishes. Its rank is therefore at most two. Multiplication on the left and right by the invertible diagonal matrices with entries $\widetilde{f}(z)^{-1}$ and $\widetilde{f}(y)^{-1}$ preserves rank. After these scalings and restricting to the two values of z_j, y_j, the remaining matrix is

$$\begin{pmatrix} 2\alpha & \alpha+\beta \\ \alpha+\beta & 2\beta \end{pmatrix},$$

whose determinant is $-(\alpha-\beta)^2 \neq 0$. The Schmidt rank is exactly two, so the gadget cannot factor as $U(z)V(y)$ into one tensor for each centre. □

For comparison, retain the closed-form table as a companion construction. Let

$$K_k^{\mathrm{tr}} := \mathbb{Q}\big(f_k^{\mathrm{tr}}(z) : z \in \{0,1\}^k\big). \tag{62}$$

Repeat the controlled construction (54)–(58) with f_k^{tr} in place of f_k^{int}, using the same base M_k, left family $\mathcal{F}_k$, and control weights $\alpha = 1, \beta = 2$. Denote the resulting Boolean preimage and right tensor by $\widehat{H}_k^{\mathrm{tr}}$ and B_k^{tr}.

Proposition 8.4 (Closed-form companion presentation)**.** *The triple* $(\mathcal{F}_k, \{\widehat{H}_k^{\mathrm{tr}}\}, M_k)$ *is a valid common width-*T_k *presentation over* K_k^{tr}*. The coordinate restriction* (59)*, every rank and essentiality conclusion of Proposition 8.2, and the two-center conclusion of Proposition 8.3 all hold with* $f_k^{\mathrm{tr}}, B_k^{\mathrm{tr}}$ *in place of* f, B_k*.*

Proof. Apply Theorem 7.1 over K_k^{tr} and repeat the proofs of Propositions 8.1 to 8.3. Those arguments use only nonvanishing and selected nonzero 2-by-2 minors of the one-versus-rest flattenings of the table, together with $\alpha \neq \beta$. The explicit table has those properties by Proposition 6.2. $\square$

Recall the definition of μ_k from (1). The presentation in Proposition 8.1 shows

$$\mu_k \leq T_k = 2^k + 3. \tag{63}$$

In particular, μ_k is a well-defined finite integer.

For the companion problem, put

$$\mu_k^{\mathrm{tr}} := \mathrm{bw}(\mathcal{F}_k \mid \{B_k^{\mathrm{tr}}\}). \tag{64}$$

By Proposition 8.4, $\mu_k^{\mathrm{tr}} \leq 2^k + 3$.

9 The Lower Bound against Arbitrary Equivalent Presentations

The preceding section supplies a valid width-(2^k+3) presentation, and therefore only an upper bound on its minimum exact common width. A competing exactly equivalent presentation may change its finite domain, common base, coordinate tensors, and complex weights. Using the star-contraction conventions of Section 4.6 and the equivalence notion of Section 2.3, we exploit the preservation of closed labelled planar instances: a fixed family of labelled stars recovers the same obstruction table from every competing presentation. The obstruction of Section 6 then yields the quantitative lower bound and Corollary 3.2.

Fix any finite domain E and any problem $\mathcal{P} \mid \mathcal{R}$ over E that is exactly equivalent to $\mathcal{F}_k \mid \{B_k\}$ in the sense of Definition 2.1. Suppose it has a valid common width-r presentation $(\mathcal{P}, \mathcal{Q}, N)$, so its induced right family is $\mathcal{R} = N\mathcal{Q}$ labelwise, with base

$$N \in \mathbb{C}^{|E| \times 2^r}. \tag{65}$$

Let $P_i \in \mathcal{P}$ be the unary left label corresponding to δ_i. Let $B_k' \in \mathcal{R}$ be the arity-$(k+2)$ right label corresponding to B_k. By validity of the new presentation, there is a Boolean block signature $Q_B \in \mathrm{MG}_{(k+2)r}$ satisfying

$$N^{\otimes(k+2)} Q_B = B_k', \tag{66}$$

and the transformed unary signatures

$$g_i := P_i N \in \mathrm{MG}_r \qquad (i \in D_3). \tag{67}$$

Neither N nor these domain signatures are assumed to have any particular rank. Contract the first two consecutive blocks of Q_B with two copies of g_0. Exact matchgates are closed under ordered planar contraction, so the result is a signature

$$Q_0 := \mathrm{Contr}(Q_B; g_0, g_0, -, \ldots, -) \in \mathrm{MG}_{kr}\,. \tag{68}$$

Here each dash denotes an uncontracted block, so Q_0 retains the final k blocks in their original order.

For each $x = (x_1, \dots, x_k) \in \{0,1\}^k$, consider the closed planar star whose centre is labelled B_k, whose two link leaves are labelled δ_0, and whose jth hard leaf is labelled δ_{2x_j}. Its original value is $f_k^{\mathrm{int}}(x)$ by (59). Exact labelled equivalence gives the same value after replacing these labels by $B'_k, P_0, P_{2x_1}, \dots, P_{2x_k}$. Equations (66)–(68) and associativity in the fixed block order give

$$\begin{aligned} f_k^{\mathrm{int}}(x) &= \left(P_0 \otimes P_0 \otimes \bigotimes_{j=1}^{k} P_{2x_j} \right) B'_k \\ &= \mathrm{Contr}(Q_B; g_0, g_0, g_{2x_1}, \dots, g_{2x_k}) \\ &= \mathrm{Contr}(Q_0; g_{2x_1}, \dots, g_{2x_k}). \end{aligned} \tag{69}$$

Set $\widetilde{g}_0 := g_0$ and $\widetilde{g}_1 := g_2$. Then (69) is exactly a star-table representation with sampler pair $(\widetilde{g}_0, \widetilde{g}_1)$. Applying Proposition 6.4 therefore gives

$$\begin{aligned} 2^k &\le 2\delta(r) + \delta(kr) \\ &= 2\left(1 + \binom{r}{2}\right) + 1 + \binom{kr}{2} \\ &= 3 + r(r-1) + \frac{kr(kr-1)}{2}. \end{aligned} \tag{70}$$

This derivation uses no coordinate of the arbitrary domain E, no inverse or rank hypothesis on N, and no relationship between N and the displayed base M_k. It therefore applies to every presentation allowed in Definition 2.2, including $r = 0$.

Finally, for every integer r occurring in the minimization set of Definition 2.2, (70) gives the first inequality below; the second inequality holds for every integer $r \ge 0$:

$$2^k - 3 \le r(r-1) + \frac{kr(kr-1)}{2} \le r^2 \left(1 + \frac{k^2}{2}\right). \tag{71}$$

Every such width r must consequently satisfy

$$r \ge \left\lceil \sqrt{\frac{2^k - 3}{1 + k^2/2}} \right\rceil. \tag{72}$$

Taking the minimum over all equivalent presentations proves (2) and (3). Together with (63) and Propositions 8.1 to 8.3, this proves parts (i)–(v) of Theorem 3.1. The final assertion that the presentation in part (i) lies in alternative (III) is established in the proof of Theorem 3.3 in Section 10.

Corollary 9.1 (Explicit transcendental companion)**.** *For the closed-form problem $\mathcal{F}_k \mid \{B_k^{\mathrm{tr}}\}$ of Proposition 8.4, every valid common width-r presentation whose induced problem is exactly equivalent to it, even with arbitrary complex weights, satisfies*

$$2^k \le 2\delta(r) + \delta(kr). \tag{73}$$

Consequently,

$$\mu_k^{\mathrm{tr}} \ge \left\lceil \sqrt{\frac{2^k - 3}{1 + k^2/2}} \right\rceil = \Omega\left(\frac{2^{k/2}}{k}\right). \tag{74}$$

Thus the same unbounded-width result also has a simple closed-form witness, at the cost of transcendental signature values.

Proof. Fix a candidate presentation $(\mathcal{P}, \mathcal{Q}, N)$ from the corollary. Let P_i be the unary labels corresponding to δ_i, let Q_B be the Boolean preimage of the right label corresponding to B_k^{tr}, and put

$$g_i := P_i N \in \mathrm{MG}_r, \qquad Q_0 := \mathrm{Contr}(Q_B; g_0, g_0, -, \ldots, -) \in \mathrm{MG}_{kr} .$$

The same closed labelled stars used in (69), now with B_k^{tr}, give

$$f_k^{\mathrm{tr}}(x) = \mathrm{Contr}(Q_0; g_{2x_1}, \ldots, g_{2x_k}) \qquad (x \in \{0,1\}^k).$$

With $\widetilde{g}_0 := g_0$ and $\widetilde{g}_1 := g_2$, this is a star-table representation with sampler pair $(\widetilde{g}_0, \widetilde{g}_1)$. The explicit obstruction in Proposition 6.2 proves (73). Substituting the definition of δ and repeating (71) proves the displayed lower bound after taking the minimum over all equivalent presentations. □

Proof of Corollary 3.2. If such a function h existed, applying it to the presentations from Proposition 8.1 would give $\mu_k \le h(3)$ for every k. This contradicts $\mu_k \to \infty$. □

10 The Positive Boundary: Support Geometry and Collapse

Section 9 completes the quantitative negative argument by showing that the constructed qutrit language cannot be compressed to uniformly bounded width under arbitrary exact re-presentation. We now turn to the positive/classification lane of Figure 1: under the qutrit, full-row-rank, bi-essential, and port-deficient hypotheses, which realized support geometries force a constant-rank common cover, and what structural form remains when they do not? The proof proceeds in the three stages indicated in the roadmap. Its first stage is support monotonicity, which propagates primitive port deficiency through the all-left gadget closure and carries the resulting supports faithfully through the full-row-rank base.

We use the tensor, gadget, ordered-matchgate, and support vocabulary formalized in Sections 4.1, 4.2, 4.4 and 4.7. Thus $V = \mathbb{C}^3$, M is a width-t qutrit base of row rank three, and $\mathcal{S}_1, \mathcal{S}_2$ are the realized supports of (23). These are supports of actual connected all-left gadgets, not of their linear span. Cyclic input rerooting and the fixed reversal on the complementary output wires are those of Section 4.7, so ordinary matrix multiplication below is ordered planar composition and no internal wire permutation is implicit. The parity endpoints P^+, P^-, monomial flags, and connection-primality are those of Section 4.4 and Definition 4.5.

The first observation records why primitive deficiency persists at a dangling boundary edge.

Lemma 10.1 (Boundary-support monotonicity)**.** *Let Γ be a planar gadget, and let a dangling edge e be incident with port i of a primitive left tensor F. If T_Γ is the boundary tensor of Γ, then*

$$S_e(\Gamma) = S_{j(e)}(T_\Gamma) \subseteq S_i(F).$$

Proof. Choose a minimal mode decomposition

$$F = \sum_{a=1}^{d} u_a \otimes F_a, \qquad S_i(F) = \mathrm{span}\{u_1, \ldots, u_d\}.$$

Contracting the rest of the gadget changes only the tensors multiplying the u_a. Thus

$$T_\Gamma = \sum_{a=1}^{d} u_a \otimes R_a$$

for suitable residual boundary tensors R_a, proving the inclusion. □

Using the block-support convention fixed in Section 4.7, for a tensor G on q consecutive t-bit blocks let $S_i^{\mathrm{blk}}(G) \le \mathbb{C}^{\{0,1\}^t}$ be the image of the flattening that treats the whole ith block as its row index and all other blocks as its column index. It differs from the support of one individual Boolean mode.

Lemma 10.2 (Transformed block supports)**.** *Suppose $q \ge 1$, $T \in (V^*)^{\otimes q}$ is nonzero, and M has row rank three. If T has mode support $U = S_i(T) \le V^*$ at port i, then*

$$S_i^{\mathrm{blk}}(TM^{\otimes q}) = UM := \{uM : u \in U\}.$$

In particular, the two supports have the same dimension.

Proof. Write a minimal mode decomposition

$$T = \sum_{a=1}^{d} u_a \otimes T_a, \qquad U = \operatorname{span}\{u_1, \dots, u_d\},$$

with the T_a linearly independent. Full row rank makes both $u \mapsto uM$ and the corresponding tensor-power map injective. Hence

$$TM^{\otimes q} = \sum_{a=1}^{d} (u_a M) \otimes (T_a M^{\otimes (q-1)})$$

is again a minimal mode decomposition, with block support UM. □

The preceding support lemmas reduce the gadget closure to an incidence problem among realized rays and planes, while full row rank carries these supports without loss of dimension into exact matchgate block row spaces. We next use pure-spinor geometry to ask when such Gaussian pieces must lie inside one low-rank matchgate row space. The following hull, decoder, and common-annihilator lemmas turn the relevant incidence relations into rank-four or rank-eight common-cover certificates.

Lemma 10.3 (Rank-two Gaussian hull)**.** *Let $R_1, R_2 \le \mathbb{C}^{\{0,1\}^t}$ be row spaces of exact matchgate matrices. Suppose*

$$\dim R_1 = \dim R_2 = 2, \qquad R_1 \ne R_2, \qquad R_1 \cap R_2 \ne 0.$$

Then there is a two-input exact matchgate matrix P of rank four such that

$$R_1 + R_2 \subseteq \operatorname{row}(P).$$

Proof. Let parity refer to the t ordered output wires. Xia's rank-two canonical form has one nonzero row-space dimension in each output parity, so

$$R_i = R_i^+ \oplus R_i^-, \qquad \dim R_i^+ = \dim R_i^- = 1.$$

The intersection is parity invariant. Since two distinct planes with nonzero intersection meet in a line, choose a nonzero parity-homogeneous $u \in R_1 \cap R_2$. In either original matchgate matrix an actual pinned input row spans that parity component, so u is, up to a nonzero scalar, an exact matchgate vector.

Apply an invertible output matchgate transformation C sending u to the vacuum $e_\varnothing$. This is the vector case of Xia's canonical form [24, Theorem 4.1]. Put $R_i' = R_i C$ and choose a matchgate matrix A_i with s_i input wires and row space R_i'. After input flips and a global rescaling, an actual input row is the pivot row $e_\varnothing$. In the pivot-Pfaffian chart of Section 4.4, write its s_i input and t output Grassmann variables as ξ and η. Let $A_{\mathrm{in},i}$, B_i, and D_i denote, respectively, the input–input, input–output, and output–output blocks of the skew quadratic parameter matrix. Up to the pivot scalar, the matchgate matrix has generating function

$$\exp\left(\frac{1}{2}\xi^\mathsf{T} A_{\mathrm{in},i}\xi + \xi^\mathsf{T} B_i\eta + \frac{1}{2}\eta^\mathsf{T} D_i\eta\right).$$

The pivot row is exactly the output vacuum, so its two-output coordinates (the coefficients of $\eta_p\eta_q$) vanish and $D_i = 0$. Put

$$q_i(\xi) := \frac{1}{2}\xi^\mathsf{T} A_{\mathrm{in},i}\xi.$$

Exterior multiplication by $\exp(-q_i)$ is an invertible linear operator on the input coefficient space, with inverse multiplication by $\exp(q_i)$. Applied coefficientwise in η, its ordinary coefficient matrix performs invertible row operations and therefore preserves the row space. Since q_i and $\xi^\mathsf{T} B_i\eta$ are even exterior elements,

$$\exp(-q_i) \wedge \exp\bigl(q_i + \xi^\mathsf{T} B_i\eta\bigr) = \exp(\xi^\mathsf{T} B_i\eta).$$

Thus no unproved boundary transformation is being used: the resulting coefficient matrix is independently an exact matchgate matrix by the principal-Pfaffian chart and has the same row space as A_i. It is the exterior-compound matrix of B_i, whose rank is $2^{\operatorname{rank} B_i}$. Since $\operatorname{rank} A_i = 2$, the cross matrix has rank one, and therefore

$$R_i' = \operatorname{span}\{e_\varnothing, a_i\}$$

for a nonzero one-particle vector a_i. The two a_i are independent because $R_1' \neq R_2'$.

Use a_1, a_2 as the rows of a 2-by-t cross matrix B_0. The coefficient matrix P_0 of $\exp(\xi^\mathsf{T} B_0\eta)$ is an exact two-input, t-output matchgate matrix by the Grassmann chart, and its exterior-compound row space is

$$\operatorname{row}(P_0) = \operatorname{span}\{e_\varnothing, a_1, a_2, a_1 \wedge a_2\}$$

and rank four. It contains $R_1' + R_2'$. Composing the ordered outputs with C^{-1} gives the exact matchgate matrix $P = P_0C^{-1}$ and

$$\operatorname{row}(P) \supseteq (R_1' + R_2')C^{-1} = R_1 + R_2.$$

□

We next make explicit the exact-labelled interface between a Gaussian row space and a smaller common base.

Lemma 10.4 (Full-row-rank matchgate decoder)**.** *If P is an r-input, t-output exact matchgate matrix of rank 2^r, then there is a t-input, r-output exact matchgate matrix D such that*

$$PD = I_{2^r}.$$

Proof. Xia's ordered canonical form [24, Theorem 4.1] gives invertible exact input and output matchgate transformations A, C such that

$$APC = J_{r,t} =: J,$$

where J is the standard inclusion that copies the r input modes to r active output modes and fixes the remaining $t - r$ output modes. In the ordered matrix convention fixed above, $JJ^\mathsf{T} = I_{2^r}$, and the standard projection J^{T} is exact by Lemma 4.3. Concretely it is the reflected nested-edge realization: its r output ports inherit precisely the required reversed physical order, the r active wires are paired by the reflected unit edges, and the other $t-r$ input wires carry the reflected zero-pins. Therefore

$$D = CJ^{\mathsf{T}}A$$

is an exact ordered matchgate composition and

$$PD = A^{-1}JJ^{\mathsf{T}}A = I_{2^r}.$$

□

Lemma 10.5 (Exact-labelled cover compression)**.** *Let M be a domain base with t Boolean wires per block. Suppose P is an exact matchgate matrix of rank 2^r, with r input wires and the same t ordered output wires, such that*

$$\operatorname{row}(M) \subseteq \operatorname{row}(P).$$

Then every valid presentation using M has a presentation of the same domain tensors, with the same labels, arities, and port orders, using a width-r base.

Proof. Write $M = NP$, with N having 2^r columns, and let D be the decoder in Lemma 10.4. For a right Boolean preimage H, put $b := \operatorname{barity}(H)$ and $H' = P^{\otimes b}H$. Matchgate closure under ordered planar composition gives $H' \in \mathrm{MG}_{rb}$, and associativity gives, for the same right label,

$$M^{\otimes b}H = N^{\otimes b}H'.$$

For a left label F, put $a := \operatorname{arity}(F)$. Its old lift is

$$FM^{\otimes a} = (FN^{\otimes a})P^{\otimes a}.$$

Attach D to every consecutive output block. Exact matchgate closure and $PD = I_{2^r}$ give

$$(FM^{\otimes a})D^{\otimes a} = FN^{\otimes a},$$

so the new left lift is an exact matchgate. Thus N is the required width-r base. This decoder form of Xia's cover-compression argument [24, Theorem 5.2] is performed separately for every label. No label-dependent scalar, arity change, or port permutation occurs. □

Lemma 10.6 (Isotropic-kernel cover)**.** *Let $K \le W \oplus W^*$ be an isotropic subspace of dimension $t-r$, where $0 \le r \le t$. Then there is an r-input, t-output exact matchgate matrix P_K of rank 2^r such that*

$$\operatorname{row}(P_K) = \{u \in \Lambda^{\bullet}W : z \cdot u = 0 \textit{ for every } z \in K\}.$$

Proof. Write

$$\widehat{L}_K := \{u \in \Lambda^{\bullet}W : z \cdot u = 0 \text{ for every } z \in K\}.$$

By Witt extension, choose $g \in \mathrm{O}(W \oplus W^*)$ carrying K to

$$K_0 := \operatorname{span}\{f^*_{r+1}, \dots, f^*_t\} \le W^*.$$

Let U_g be a Pin lift, acting on column spinors. The intertwining identity (15) shows that

$$U_g\widehat{L}_K = \widehat{L}_{K_0}.$$

Put $C := U_g^{\mathsf{T}}$. Both C and $C^{-1} = (U_g^{-1})^{\mathsf{T}}$ are exact square matchgate transformations by Lemma 4.3. If u is written as a coefficient column, then

$$u^{\mathsf{T}} C = (U_g u)^{\mathsf{T}},$$

so, in row coordinates, $\widehat{L}_K C = \widehat{L}_{K_0}$.

The Clifford actions of K_0 are the standard annihilation operators, and their joint kernel is exactly $\Lambda^\bullet \operatorname{span}(f_1, \ldots, f_r)$, of dimension 2^r. In the fixed ordered coordinates this is $\operatorname{row}(J_{r,t})$. Therefore

$$P_K := J_{r,t} C^{-1}$$

is an exact r-input, t-output matchgate matrix of rank 2^r, and in fact

$$\operatorname{row}(P_K) = \widehat{L}_{K_0} C^{-1} = \widehat{L}_K.$$

This also realizes directly the common-annihilator dimension formula used here; compare [15, Theorem 1(1)]. □

With support monotonicity and the Gaussian hull and compression machinery in place, we return to the realized support geometry and perform the final separation in Figure 1. A ray-only closure forces complete factorization; sufficient motion among planes or pure rays produces a rank-four or rank-eight common cover; and the remaining plane-bearing configuration is a one-plane monomial flag. In the flag case incidence no longer forces collapse, so we instead derive a uniform all-arity Boolean-core-times-rays normal form.

Theorem 10.7 (Connection-prime collapse)**.** *Let $(\mathcal{F}, \mathcal{H}, M)$ be a valid qutrit presentation, let $\mathcal{B} = M\mathcal{H}$ be its induced right family, and consider the all-left gadget closure of $\mathcal{F} \mid \mathcal{B}$. Assume* $\operatorname{rank} M = 3$*, $\mathcal{B}$ is support-essential and every port of every domain tensor in $\mathcal{F} \cup \mathcal{B}$ is deficient. (The displayed rank hypothesis is redundant under right support-essentiality, but makes the domain of Definition 4.5 explicit.) If this closure is connection-prime, then M has a common matchgate cover of rank at most eight. Hence the presentation collapses exactly and labelwise to width at most three. If two distinct planes occur in $\mathcal{S}_2$, the cover has rank four and the width is at most two.*

Proof. The stated rank condition also follows independently from right essentiality, since all induced right mode supports lie in $\operatorname{col}(M)$ and together span $\mathbb{C}^3$. Now Lemma 10.2 identifies every realized domain support U with the matchgate block support UM. If distinct $P, Q \in \mathcal{S}_2$ occur, then PM and QM are distinct rank-two matchgate row spaces in the three-dimensional space $\operatorname{row}(M)$. They meet in a line. The Gaussian-hull lemma puts their sum, hence $\operatorname{row}(M)$, in a rank-four matchgate row space. This is a common cover, and Lemma 10.5 gives width at most two.

Suppose instead that $\mathcal{S}_2 = \{P\}$. Connection-primality supplies a realized line $L \not\subset P$, whose image LM is a pure matchgate ray. Since there is no monomial flag, there is another realized line L' different from P^+, P^- and L; its image is likewise a pure matchgate ray. A realized line inside P can only be P^+ or P^-, because every exact matchgate vector has one parity. Thus L' lies outside P.

If $t \leq 3$, the standard matrix $J_{t,t} = I_{2^t}$ has row space $\mathbb{C}^{\{0,1\}^t}$ and rank $2^t \leq 8$. It is a common cover, so Lemma 10.5 proves the conclusion. Assume henceforth that $t \geq 4$.

The space

$$\operatorname{row}(M) = PM + LM$$

is parity invariant. Let $\sigma \in \{+, -\}$ be the sign for which $P^\sigma M$ has the same parity as LM, and write $-\sigma$ for the opposite sign. Choose nonzero generators

$$u \in P^\sigma M, \qquad v \in LM, \qquad w \in P^{-\sigma} M.$$

The first parity component is the plane $\mathrm{span}(u, v)$, while the other is the line $\mathbb{C}w$. The pure ray $L'M$ must have the first parity, since the opposite component contains only $P^{-\sigma}M \subset PM$. The choice of L' therefore implies

$$L'M = \mathbb{C}(au + bv), \qquad a, b \in \mathbb{C}^{\times}.$$

Thus the projective line through the three distinct pure points $[u], [v], [au + bv]$ contains three pure points. The homogeneous vanishing ideal of the spinor variety is generated by quadrics [16, Section 2.4]. Every restricted defining quadric has three zeros on this projective line and therefore vanishes identically. The whole line is a pure pencil, so the pure-pencil classification of lines on the spinor variety [17, Example 4.12] supplies a fixed isotropic $(t - 2)$-space K_{uv} contained in the annihilator of every point of the line. The two maximal isotropics $\mathrm{Ann}(u)$ and $\mathrm{Ann}(v)$ are distinct and belong to the same spinor component. In one component the codimension of their intersection in either maximal isotropic is even [16, Section 2.2]. Containment of K_{uv} makes that codimension at most two, while distinctness excludes zero; it is therefore exactly two. Hence

$$\dim(\mathrm{Ann}(u) \cap \mathrm{Ann}(v)) = t - 2.$$

We justify the second annihilator dimension through an explicit Clifford isometry. Put $R := PM$. By Witt extension, choose $g \in \mathrm{O}(W \oplus W^*)$ with $g\,\mathrm{Ann}(u) = W^*$, and let U_g be a Pin lift. In row coordinates set $C := U_g^{\mathsf{T}}$. The intertwining identity (15) shows that uC is a nonzero multiple of the vacuum $e_\varnothing$, and Lemma 4.3 shows that C and C^{-1} are exact matchgate transformations. Hence RC is a rank-two exact matchgate row space containing the vacuum. The pivot-chart argument in the proof of Lemma 10.3 applies verbatim and gives

$$RC = \mathrm{span}\{e_\varnothing, \alpha\}$$

for a nonzero one-particle vector $\alpha \in W$. A Pin lift either preserves parity globally or swaps it globally. Since u and w span the two opposite-parity lines of R, the vector wC spans $\mathbb{C}\alpha$. Now $\mathrm{Ann}(e_\varnothing) = W^*$, and

$$\mathrm{Ann}(e_\varnothing) \cap \mathrm{Ann}(\alpha) = \{\phi \in W^* : \phi(\alpha) = 0\},$$

a hyperplane of dimension $t - 1$. The same intertwining identity identifies this intersection with $g(\mathrm{Ann}(u) \cap \mathrm{Ann}(w))$. Since g is an isometry, it follows that

$$\dim(\mathrm{Ann}(u) \cap \mathrm{Ann}(w)) = t - 1.$$

Both intersections lie in the t-dimensional maximal isotropic space $\mathrm{Ann}(u)$. Their intersection has dimension at least $t - 3$ and annihilates every vector in $\mathrm{row}(M)$. Choose a $(t - 3)$-dimensional subspace

$$K \leq \mathrm{Ann}(u) \cap \mathrm{Ann}(v) \cap \mathrm{Ann}(w).$$

It is isotropic because $K \leq \mathrm{Ann}(u)$. Every vector in $\mathrm{row}(M)$ is annihilated by K, so Lemma 10.6 with $r = 3$ gives a rank-eight exact matchgate matrix whose row space contains $\mathrm{row}(M)$. This is a common cover, and Lemma 10.5 gives width at most three [24, Theorem 5.2]. $\square$

The exceptional case has an all-arity normal form. We first decode one rank-two block.

Lemma 10.8 (Rank-two block decoding). *Let $R \leq \mathbb{C}^{\{0,1\}^t}$ be the row space of a rank-two exact matchgate matrix, and write*

$$R = R^+ \oplus R^-, \qquad \dim R^+ = \dim R^- = 1,$$

where R^+ is even and R^- is odd. There are generators $g_0 \in R^+$, $g_1 \in R^-$ and one invertible output matchgate transformation C, depending only on R, such that

$$g_0 C = e_{0^t}^{\mathsf{T}}, \qquad g_1 C = e_{10^{t-1}}^{\mathsf{T}}.$$

Let $m \geq 0$. If an exact mt-wire matchgate tensor G belongs to $R^{\otimes m}$ and

$$G = \sum_{x \in \{0,1\}^m} h(x) g_{x_1} \otimes \cdots \otimes g_{x_m},$$

then h is an exact m-wire Boolean matchgate signature.

Proof. Xia's rank-two canonical form [24, Theorem 4.1], in the terminology of Section 4.4, gives an invertible output matchgate transformation that sends R to the row space of $J_{1,t}$, after fixed output flips if needed. Thus it sends R to

$$\operatorname{span}\{e_{0^t}^{\mathsf{T}}, e_{10^{t-1}}^{\mathsf{T}}\}.$$

The two parity lines map to the displayed coordinate lines. If they are interchanged, compose with a flip of the active output mode. Rescale the generators to obtain the stated equalities.

If $m = 0$, the assertion is the convention $\mathrm{MG}_0 = \mathbb{C}$. Assume $m \geq 1$.

Apply the same transformation to every consecutive block of G. Matchgate signatures are closed under this blockwise composition, and the result is

$$\sum_{x \in \{0,1\}^m} h(x) e_{x_1 0^{t-1}}^{\mathsf{T}} \otimes \cdots \otimes e_{x_m 0^{t-1}}^{\mathsf{T}}.$$

Pin the $t-1$ inactive wires of every block to zero, keeping the inherited cyclic order on the active wires. The remaining signature is exactly h, so matchgate closure under pinning proves $h \in \mathrm{MG}_m$. $\square$

Lemma 10.9 (Rank-one stripping and flag normal form)**.** *Assume $q \geq 1$ and M has row rank three, and let $T \in (V^*)^{\otimes q}$ be nonzero with $TM^{\otimes q} \in \mathrm{MG}_{qt}$.*

If every mode support of T is a line, then T is completely decomposable. More generally, suppose a plane $P \leq V^*$ has the property that PM is the row space of a rank-two exact matchgate matrix, let P^+, P^- be its pullback parity endpoints, and suppose a line $L \not\subset P$ satisfies

$$S_i(T) \in \{P, P^+, P^-, L\} \qquad \text{for every port } i.$$

Choose the generators g_0, g_1 supplied by Lemma 10.8 for $R = PM$, and choose $p_0 \in P^+$ and $p_1 \in P^-$ so that $p_a M = g_a$ for $a \in \{0,1\}$. Put $\iota(e_a) = p_a$ for $a \in \{0,1\}$. If

$$I = \{i \in [q] : S_i(T) = P\},$$

then there are nonzero q_j spanning $S_j(T)$ for $j \notin I$ and an exact Boolean matchgate $h_T \in \mathrm{MG}_{|I|}$ such that, in the original port order,

$$T = \sum_{\sigma \in \{0,1\}^I} h_T(\sigma) \bigotimes_{j=1}^{q} v_j(\sigma), \qquad v_j(\sigma) = \begin{cases} p_{\sigma(j)}, & j \in I, \\ q_j, & j \notin I. \end{cases} \tag{75}$$

Here h_T uses the cyclic order inherited by the ports in I. If $I = \varnothing$, it is a nullary scalar. Necessarily $|I| = 0$ or $|I| \geq 2$; in the latter case every mode of h_T has rank two. Each $q_j M$, as well as $p_0 M, p_1 M$, is an exact pure matchgate vector.

Proof. If $S_i(T)$ is a line, minimality of the mode support gives $T = q_i \otimes T'$ after placing port i first only in the notation. Tensoring by a nonzero vector does not change any other mode support, so the operation may be repeated at every rank-one port. If every port has rank one, induction on q gives complete decomposability.

For any stripped factor q_j, pin all other Boolean wires of the nonzero lift $TM^{\otimes q}$ to a joint coordinate on which the residual factor is nonzero. The remaining block is a nonzero scalar multiple of q_jM, so q_jM is an exact matchgate vector. The vectors p_0M, p_1M are the parity endpoints of a rank-two matchgate row space and are exact matchgate vectors by canonical form.

In the second situation, stripping the ports outside I leaves a core $T_I \in P^{\otimes |I|}$. For each $j \notin I$, choose a bit string y_j with $(q_jM)(y_j) \neq 0$. Pinning the corresponding Boolean block of $TM^{\otimes q}$ to y_j leaves a nonzero scalar multiple of $T_IM^{\otimes |I|}$. Thus the core lift is an exact matchgate. Expand T_I uniquely in the basis (p_0, p_1) of P; its coefficient tensor is h_T. Applying Lemma 10.8 with $R = PM$ proves $h_T \in \mathrm{MG}_{|I|}$. Finally, injectivity of $p \mapsto pM$ shows that a rank-one mode of h_T would make the corresponding mode support of T a line rather than P. Hence every retained mode has rank two. □

We now prove the trichotomy stated in Section 3.

Proof of Theorem 3.3. The rank assumption is in fact also forced by right essentiality, because every induced right mode support lies in $\mathrm{col}(M)$. By Lemma 10.1, every nonzero boundary support in the all-left closure has dimension one or two. Since the primitive left family occurs among the one-vertex gadgets, left essentiality gives

$$\sum_{S \in \mathcal{S}_1 \cup \mathcal{S}_2} S = V^*. \tag{76}$$

If $\mathcal{S}_2 = \varnothing$, every port of every nonzero boundary tensor has rank one, so Lemma 10.9 gives alternative (I). Suppose $\mathcal{S}_2 \neq \varnothing$. If no monomial flag exists, the closure is connection-prime by (76), and Theorem 10.7 proves alternative (II). Otherwise a monomial flag (P, L) exists and gives the support inclusions in alternative (III). Every positive-arity boundary tensor has an exact lift by planar contraction of the transformed primitive matchgates, so Lemma 10.9 applies with the same plane, rays, and encoding; a nullary boundary tensor is already a scalar. The alternatives are mutually exclusive and exhaustive according to whether $\mathcal{S}_2$ is empty and, when nonempty, whether a monomial flag exists.

For the final assertion, take the displayed presentation $(\mathcal{F}_k, \mathcal{H}_k, M_k)$ from Proposition 8.1. Every exposed edge of an all-left gadget is incident with either a unary pin, whose support is one of $\mathbb{C}e_0^*, \mathbb{C}e_1^*, \mathbb{C}e_2^*$, or a port of X_{01}, whose support is

$$P = \mathrm{span}(e_0^*, e_1^*).$$

Boundary-support monotonicity confines the resulting support to one of those lines or to P. A realized line inside P maps under M_k to a pure matchgate ray inside the parity-split row space PM_k and hence is one of its two parity endpoints $\mathbb{C}e_0^*, \mathbb{C}e_1^*$. Since X_{01} realizes P and δ_2 realizes $L = \mathbb{C}e_2^*$, the displayed presentation lies in alternative (III). Part (v) of Theorem 3.1 makes μ_k tend to infinity, and Propositions 8.2 and 8.3 supply the remaining operational properties. This also proves the final flag-placement assertion of Theorem 3.1. □

The same geometry has a useful binary-context consequence.

Corollary 10.10 (Binary-context alternative)**.** *Let $(\mathcal{F}, \mathcal{H}, M)$ be a valid qutrit presentation, let $\mathcal{B} = M\mathcal{H}$ be support-essential, assume* $\operatorname{rank} M = 3$*, and consider the all-left gadget closure of*

$\mathcal{F} \mid \mathcal{B}$. Suppose one realized rank-two plane P and one realized line $L \not\subset P$ span V^*. As recorded in Section 4.7, LM is then a pure matchgate ray. Choose nonzero generators $p_+ \in P^+$, $p_- \in P^-$, and $\ell \in L$; the ordered triple (p_+, p_-, ℓ) is an adapted ray basis. A binary context *means the boundary tensor of a connected all-left gadget with two ordered boundary ports. Then at least one of the following holds:*

(a) *M has a rank-four common matchgate cover;*

(b) *M has a rank-eight common matchgate cover;*

(c) *every binary context is partial monomial in the adapted ray basis, meaning that its* 3*-by-*3 *domain matrix has at most one nonzero entry in every row and column.*

Consequently, one non-partial-monomial binary context forces width at most three.

Proof. Let A be the domain matrix of a binary context, expressed in the chosen adapted basis. The stated full row rank also follows from support-essentiality of $\mathcal{B}$, because every induced right mode support is contained in $\operatorname{col}(M)$. Its Boolean lift is an exact matchgate matrix, and this full row rank preserves the matrix rank of A. Since $\operatorname{rank} A \leq 3$ while a nonzero exact matchgate matrix has power-of-two rank, the only possibilities are ranks zero, one, and two. Let $\sigma \in \{+, -\}$ be the sign for which $P^\sigma M$ has the same parity as LM.

If $\operatorname{rank} A = 1$, write $A = cxy^\mathsf{T}$ with $c \in \mathbb{C}^\times$ and nonzero coordinate columns $x, y \in \mathbb{C}^3$. Full row rank of M makes both xM and yM nonzero. Choose $\beta \in \{0,1\}^t$ with $(yM)(\beta) \neq 0$ and pin the second Boolean block of the exact lift to the coordinate β. The remaining block is the nonzero scalar multiple $c(yM)(\beta)\, xM$, and hence $\mathbb{C}(xM)$ is a pure matchgate ray. Choosing α with $(xM)(\alpha) \neq 0$ and pinning the first block proves the same for $\mathbb{C}(yM)$. Thus the local rays $\mathbb{C}x$ and $\mathbb{C}y$ map to pure matchgate rays. Call a ray *noncoordinate* when it is not one of $\mathbb{C}p_+, \mathbb{C}p_-, \mathbb{C}\ell$. If either local ray is noncoordinate, parity places it as a third pure point on the same-parity line through P^σ and L. The pure-pencil and common-annihilator argument in Theorem 10.7 gives a rank-eight cover. Otherwise A is a scalar multiple of a matrix unit—a matrix with one nonzero coordinate—and is partial monomial.

If $\operatorname{rank} A = 2$, its row and column mode supports $S_1(A), S_2(A)$ are planes that map to rank-two matchgate row spaces. If either plane differs from P, Lemma 10.3 gives a rank-four cover. Otherwise both supports equal P. Expanding the lift in the opposite-parity endpoint rays P^+, P^-, global matchgate parity allows only the two diagonal entries or only the two antidiagonal entries. Hence A is partial monomial. Rank zero is trivial. $\square$

11 Discussion and Future Directions

The main results separate two sources of representation size. Local domain rank and port rank control the geometry visible on one edge, but they do not control how much labelled information must be coordinated across many ports. Support mobility eliminates that information by producing a small common Gaussian cover. A fixed plane-plus-ray flag leaves residual right-side data on which the minimum exact width can grow. The natural next step is to make this distinction intrinsic to the language rather than to one displayed presentation.

11.1 From a support trichotomy to a language-level width theory

The central open problem is a pointwise characterization of minimum exact width. The residual right tables defined from a flag presentation are useful coordinates, but are not themselves invariants:

another presentation may change the flag, the domain, or the table. One seeks a presentation-independent invariant that yields a bounded-width realization when small and quantitative lower bounds, valid for every exact presentation, when large. The labelled star-table obstruction is one such lower-bound certificate for our family, but not yet a general characterization.

Two boundary cases sharpen the question. Alternative (I) of Theorem 3.3 forces every nonzero positive-arity all-left boundary tensor to factor, yet the minimum exact width of languages in this branch is unclassified. Is there a constant C such that every language admitting an alternative (I) qutrit presentation has minimum exact width at most C, or is there such a family whose minimum widths tend to infinity because of right-side data? At the other boundary, Corollary 10.10 shows that one non-partial-monomial binary context forces width at most three. Is there a function $c(a, m)$ such that, for qutrit presentations with primitive arity at most a and at most m labels, failure of ray separability or flag confinement is always witnessed by a planar context with at most $c(a, m)$ vertices and boundary ports? For rational or explicitly represented algebraic presentations, can the remaining universal confinement conditions be decided from the primitive signatures?

11.2 Bounded arity and the source of unbounded information

Our family uses a fixed number of labels and stores its growing table in the increasing arity of B_k. Does bounded arity alone force a bound $h(q, a)$ on the minimum exact width of every exactly matchgate-realizable finite labelled language on a q-element domain whose labels have arity at most a? If not, does a uniform bound $h(q, a, m)$ hold once the number of labels is also fixed? A counterexample to this latter statement would require a storage mechanism different from the growing-arity table used here.

For rational or integer languages, an effective version is equally important. With domain size q, arity bound a, label bound m, and coefficient bit length B, can one give a useful quantitative or computable bound $h(q, a, m, B)$? Merely knowing that a finite maximum exists after all four parameters are fixed would not explain the representation theory; the goal is a structural and effective bound.

11.3 Robustness under broader equivalence notions

Exact labelled equivalence preserves the value of every ordered labelled planar instance under the labelwise correspondence. A first relaxation would allow each label to be replaced by a uniformly bounded exact planar gadget while retaining a common base. One can then ask whether the star-table lower bound survives, perhaps after changing its sampling contexts. More general gadget interreducibility may alter arities and duplicate labels, so a width measure must account for the size and uniformity of the substitutions as well as the number of base wires.

Polynomial-time interreducibility is a still broader question and requires a specified uniform input model for growing languages and coefficients. The present theorem makes no claim under either bounded-gadget equivalence or polynomial-time interreducibility: it is a representation lower bound, not a complexity lower bound. Establishing which part of the obstruction is invariant under bounded gadgets would be the most direct next test.

11.4 Explicit and quantitative obstructions

The main obstruction uses a positive-integer table chosen by Zariski avoidance, but the proof supplies neither a simple formula nor a useful height bound. The transcendental companion has a closed formula and the same width lower bound, at the cost of leaving the algebraic coefficient setting. A natural target is a uniformly generated positive-integer or rational family with per-entry bit length polynomial in k and construction time polynomial in the total output length.

There is also a large quantitative gap. The current construction gives

$$\Omega\left(\frac{2^{k/2}}{k}\right) \le \mu_k \le 2^k + 3.$$

It is unknown whether the dimension obstruction can be strengthened, whether the block-Pfaffian realization can be compressed, or whether the true width lies at an intermediate scale. Resolving this gap would distinguish the intrinsic representational cost of the table from the overhead of the present universal interpolation.

11.5 Beyond qutrits and intrinsic compression

The qutrit case is special because every deficient nonzero mode support is a ray or a plane. Larger domains admit intermediate-dimensional supports and richer incidence patterns among Gaussian row spaces. A higher-domain theory would likely organize these supports through Grassmannian or spinorial incidence data and determine which kinds of mobility force a domain-controlled common cover. For each fixed domain size, does this incidence geometry admit a finite stratification analogous to Theorem 3.3? If such stratifications exist, do their types eventually stabilize as the domain grows, or do genuinely new exceptional strata continue to appear?

From the free-fermionic viewpoint, the same question asks when a common exact Gaussian tensor representation admits simultaneous mode compression. This is an algebraic reformulation, not a claim about physical qubit compression, but it may suggest invariants that are invisible in the original base coordinates.

Thus the trichotomy identifies support mobility as the local mechanism that forces compression, while the construction shows that a fixed flag can preserve labelled information of arbitrarily large exact cost. The remaining challenge is to recognize that information intrinsically—from the language rather than from a chosen presentation.

AI Disclosure: We used OpenAI Codex to assist with developing and checking mathematical arguments, reviewing the literature, revising the exposition and organization, and preparing the LaTeX manuscript. The tool materially affected Sections 1–11. More details on the affected proof architecture and technical components can be found in Section 1.3. The authors verified the correctness and originality of all content, including references.

References

[1] William F. Bradley. A recursive definition of the holographic standard signature. arXiv preprint arXiv:0911.2907 [cs.CC], 2009. https://arxiv.org/abs/0911.2907.

[2] Sergey Bravyi. Contraction of matchgate tensor networks on non-planar graphs. In Kazem Mahdavi and Deborah Koslover, editors, *Advances in Quantum Computation*, volume 482 of *Contemporary Mathematics*, pages 179–211. American Mathematical Society, 2009. https://doi.org/10.1090/conm/482/09419. Author version: https://arxiv.org/abs/0801.2989.

[3] Jin-Yi Cai and Vinay Choudhary. Some results on matchgates and holographic algorithms. *International Journal of Software and Informatics*, 1(1):3–36, 2007. Preliminary version in ICALP 2006: https://doi.org/10.1007/11786986_61.

[4] Jin-Yi Cai, Vinay Choudhary, and Pinyan Lu. On the theory of matchgate computations. *Theory of Computing Systems*, 45(1):108–132, 2009. https://doi.org/10.1007/s00224-007-9092-8.

[5] Jin-Yi Cai and Zhiguo Fu. A collapse theorem for holographic algorithms with matchgates on domain size at most 4. *Information and Computation*, 239:149–169, 2014. https://doi.org/10.1016/j.ic.2014.10.002. Author version: https://arxiv.org/abs/1305.1409.

[6] Jin-Yi Cai and Aaron Gorenstein. Matchgates revisited. *Theory of Computing*, 10(7):167–197, 2014. https://doi.org/10.4086/toc.2014.v010a007.

[7] Jin-Yi Cai and Pinyan Lu. Basis collapse in holographic algorithms. *Computational Complexity*, 17(2):254–281, 2008. https://doi.org/10.1007/s00037-008-0249-x.

[8] Jin-Yi Cai and Pinyan Lu. Holographic algorithms: The power of dimensionality resolved. *Theoretical Computer Science*, 410(18):1618–1628, 2009. https://doi.org/10.1016/j.tcs.2008.12.047.

[9] Jin-Yi Cai and Pinyan Lu. Holographic algorithms: From art to science. *Journal of Computer and System Sciences*, 77(1):41–61, 2011. https://doi.org/10.1016/j.jcss.2010.06.005.

[10] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms with matchgates capture precisely tractable planar #CSP. *SIAM Journal on Computing*, 46(3):853–889, 2017. https://doi.org/10.1137/16M1073984. Author version: https://arxiv.org/abs/1008.0683.

[11] Sitan Chen. Basis collapse for holographic algorithms over all domain sizes. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pages 776–789. ACM, 2016. https://doi.org/10.1145/2897518.2897546.

[12] Zhiguo Fu and Fengqin Yang. Holographic algorithms on bases of rank 2. *Information Processing Letters*, 114(11):585–590, 2014. https://doi.org/10.1016/j.ipl.2014.06.006.

[13] Zhiguo Fu, Fengqin Yang, and Minghao Yin. On blockwise symmetric matchgate signatures and higher domain #CSP. *Information and Computation*, 264:1–11, 2019. https://doi.org/10.1016/j.ic.2018.09.012.

[14] Zhiguo Fu and Jin-Yi Cai. Holographic algorithms on domains of general size. *Theory of Computing Systems*, 67(3):417–436, 2023. https://doi.org/10.1007/s00224-022-10088-7.

[15] Carlos Batista. Pure subspaces, generalizing the concept of pure spinors. *Journal of Geometry and Physics*, 81:117–127, 2014. https://arxiv.org/abs/1310.0104.

[16] Laurent Manivel. On spinor varieties and their secants. *SIGMA*, 5:078, 2009. https://arxiv.org/abs/0904.0565.

[17] Joseph M. Landsberg and Laurent Manivel. On the projective geometry of rational homogeneous varieties. *Commentarii Mathematici Helvetici*, 78:65–100, 2003. https://doi.org/10.1007/s000140300003.

[18] Richard Jozsa and Akimasa Miyake. Matchgates and classical simulation of quantum circuits. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 464(2100):3089–3106, 2008. https://doi.org/10.1098/rspa.2008.0189. Author version: https://arxiv.org/abs/0804.4050.

[19] Joseph M. Landsberg, Jason Morton, and Serguei Norine. Holographic algorithms without matchgates. *Linear Algebra and its Applications*, 438(2):782–795, 2013. https://doi.org/10.1016/j.laa.2012.01.010. Author version: https://arxiv.org/abs/0904.0471.

[20] Sever Angel Popescu. A simple and self-contained proof for the Lindemann–Weierstrass theorem. In Jordi Guàrdia, Nicuşor Minculete, Diana Savin, Montserrat Vela, and Abdelkader Zekhnini, editors, *New Frontiers in Number Theory and Applications*, Trends in Mathematics, pages 349–366. Birkhäuser, Cham, 2024. https://doi.org/10.1007/978-3-031-51959-8_16.

[21] Leslie G. Valiant. Quantum circuits that can be simulated classically in polynomial time. *SIAM Journal on Computing*, 31(4):1229–1254, 2002. https://doi.org/10.1137/S0097539700377025.

[22] Leslie G. Valiant. Accidental algorithms. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, pages 509–517. IEEE Computer Society, 2006. https://doi.org/10.1109/FOCS.2006.7.

[23] Leslie G. Valiant. Holographic algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008. https://doi.org/10.1137/070682575.

[24] Mingji Xia. Base collapse of holographic algorithms. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pages 790–799. ACM, 2016. https://doi.org/10.1145/2897518.2897560. Author version: https://arxiv.org/abs/1511.01230.