A Full Complexity Dichotomy for Complex-Valued Boolean Holant Problems

Chenghua Liu[1,*], Boning Meng[2,†], and Juqiu Wang[1,‡]

[1]Institute of Software, Chinese Academy of Sciences, Beijing, China
[2]University of Regensburg, Regensburg, Germany
[*]liuch.russell@gmail.com [†]mengboning2013@gmail.com [‡]wangjuqiu104@gmail.com

Abstract

We prove a complexity dichotomy for Boolean Holant problems defined by arbitrary finite sets of algebraic complex-valued signatures. The tractable cases are characterized by an explicit, decidable criterion.

Acknowledgments

Chenghua Liu is sincerely grateful to Zhengfeng Ji and Mingsheng Ying for their continued encouragement and generous support. Boning Meng and Juqiu Wang warmly thank Mingji Xia for his invaluable guidance and dedicated mentorship. The authors used OpenAI's ChatGPT in preparing this manuscript, including exploratory proof development, exact symbolic and finite-case checks, language editing, and LaTeX preparation. All companion verification code was written by OpenAI Codex and subsequently reviewed by the authors. All mathematical claims, computer-assisted certificates, and the final presentation were reviewed and finalized by the authors, who take full responsibility for the content.

Code Availability

The verification code for the finite calculations in Appendix A is available in the `certificates/` directory of the *GitHub repository: Full-Holant-Dichotomy*. The complete suite is run from the repository root with `python3 certificates/run_all_exact.py`. The `README.md` in that directory gives requirements and reproduction commands; `theorem-to-script-map.md` identifies the statements checked by each program and distinguishes supplementary checks from calculations used in the proof.

Contents

1 Introduction **3**
1.1 From counting CSP to the full Holant dichotomy 3
1.2 Our result: the full dichotomy 5
1.3 Proof organization 5
1.4 Related work 6

2 Preliminaries and the Tractable Boundary **7**
2.1 Algebraic notation . . . 7
2.2 Complexity and exact-input conventions . . . 7
2.3 Tensor and group conventions . . . 8
2.4 Graphs, matchings, and Walsh coordinates . . . 8
2.5 Gadgets and allowed operations . . . 9
2.6 Signatures, tensors, and Holant instances . . . 10
2.7 Holographic transformations and the K-basis . . . 11
2.8 Gadgets, interpolation, and factor reductions . . . 13
2.9 Zero and nullary preprocessing . . . 15
2.10 Charge and Hamming layers . . . 15
2.11 Bell signatures and the Reed–Muller core . . . 16
2.12 Exact finite verification . . . 17
2.13 Tractable signature sets and effective recognition . . . 17
2.14 Problem equivalence and normalization . . . 21
2.15 Binary-equality and quaternary dichotomies . . . 22

3 Auxiliary Reductions **23**
3.1 Binary Interpolation . . . 23
3.2 Single-occurrence Reductions and Nonzero Unaries . . . 23
3.3 Hardness from an Eulerian-orientation Signature . . . 25
3.4 Equalities and Generalized Equalities . . . 30

4 Obtaining a Mountain **31**

5 Mountain Height One **33**
5.1 Odd cyclic and dihedral transfer groups . . . 37
5.2 Contractions for the odd cyclic and dihedral groups . . . 40
5.3 The odd-dihedral case . . . 42
5.4 The odd-cyclic case . . . 48

6 Mountain Height Two **56**

7 Mountain Height Three **57**
7.1 Matching Products and Two-Copy Constructions . . . 57
7.2 Six-Port Loop and Two-Copy Contractions . . . 58

8 Mountain Height at Least Four **60**
8.1 Common Normalization and Phase Recovery . . . 61
8.2 Height at Least Five . . . 64
8.3 Height Four . . . 64

9 Synthesis and Proof of the Dichotomy **69**
9.1 Combining the Height Cases . . . 69
9.2 Proof of the Main Theorem . . . 69

10 Conclusion **70**

A Exact Finite Certificates **70**
A.1 Odd-dihedral and high-charge certificates 71

1 Introduction

A Holant problem asks for the value of a graph tensor-network partition function: each vertex is assigned a local function, called a signature, of its incident Boolean edge variables. One sums the product of these functions over all edge assignments. The framework includes weighted graph homomorphisms, counting CSPs, Eulerian orientations, and vertex models. A *holographic transformation* is an invertible simultaneous change of basis on the local tensors and wires [Val08, CLX09].

We prove a dichotomy for finite algebraic complex-valued Boolean Holant problems on general graphs. Our proof treats unrestricted all-even signature sets, in which every positive-arity signature has even arity and no auxiliary binary signature is prescribed, and combines this analysis with known dichotomies. We first review the graph-homomorphism, counting-CSP, and Holant classifications that precede this result.

1.1 From counting CSP to the full Holant dichotomy

Weighted graph homomorphisms (#GH) and counting constraint satisfaction problems (#CSP) provide earlier examples of counting dichotomies on general graphs. Their classifications developed in three directions: from Boolean to arbitrary finite domains, from symmetric to arbitrary local functions, and from Boolean-valued through nonnegative, signed-real, and complex weights. Cai and Chen's general-domain complex classification covers both problems [CC17]. Table 1 records these classifications. A dash indicates that no separate reference is listed for that case; the case is covered by a more general entry in the table.

Table 1: Development of general-graph dichotomies for weighted graph homomorphisms and counting CSP.

Weight range	Symmetric #GH	#GH	Boolean-domain #CSP	General-domain #CSP
$\{0, 1\}$	[DG00]	[DGP07]	[CH96]	[Bul13, DR13]
$\mathbb{Q}_{\geq 0}$	–	–	–	[BDG+12]
$\mathbb{R}_{\geq 0}$	[BG05]	[CC19]	[DGJ09]	[CCL16]
$\mathbb{R}$	[GGJT10]	–	[BDG+09]	–
$\mathbb{C}$	[CCL13]	[CC17]	[CLX14]	classification [CC17]; decidability [LM26b]

The last cell distinguishes classification from uniform recognition. Cai and Chen's 2017 theorem classified every fixed algebraic-complex language, but its tractability side was expressed by three conditions over an unbounded generated family of partial sums. In 2026, Liu and Meng gave a total exact algorithm deciding those conditions from the finite input language, proved that Block Orthogonality alone implies the other two, and also decided the corresponding degree-multiple counting-CSP criterion [LM26b]. This makes the tractability criterion decidable from a finite, exactly encoded algebraic complex-valued language on any finite domain.

For Holant on larger domains, classifications are available under further restrictions on the signatures. A domain-three theorem treated one symmetric ternary signature with all unary signatures available [CLX13]; its journal version appeared in 2026 [CLX26]. Subsequent classifications covered special edge-colouring signatures and restricted domain-three and domain-four settings [CGW16b, LFC25]. In 2025, Cai and Ihm proved a dichotomy for real-valued symmetric signatures on domain three with all unary signatures available, allowing an arbitrary finite signature set [CI25].

On the Boolean domain, the Holant framework was introduced in 2009 [CLX09]. Classifications first treated symmetric signatures or assumed access to specified unary signatures. The latter assumptions were reduced from all unary signatures (Holant*), to four distinguished unaries (Holant$^+$), to the two unary functions fixing a bit to zero or one (Holant^c). Later work treated signature sets containing any nonzero odd-arity signature (Holant$^{\text{odd}}$). The symmetric and arbitrary-signature classifications are separated in Table 2; the dash convention is the same as in Table 1.

Table 2: Development of Boolean-domain Holant dichotomies on general graphs.

(a) Symmetric signature sets

Weight range	Holant*	Holant^c	Unrestricted Holant
$\mathbb{R}_{\geq 0}$	–	–	–
$\mathbb{R}$	–	[CLX09]	[HL16]
$\mathbb{C}$	[CLX09]	[CHL12]	[CGW16a]

(b) Arbitrary signature sets

Weight range	Holant*	Holant$^+$	Holant^c	Holant$^{\text{odd}}$	Unrestricted Holant
$\mathbb{R}_{\geq 0}$	–	–	–	–	[LW18]
$\mathbb{R}$	–	–	[CLX18]	[CFS20b]	[SC20]
$\mathbb{C}$	[CLX11]	[Bac17]	[Bac21, Bac25]	[MWXZ25, GSS26]	This paper; with binary disequality: [LM26a]

Classifications without auxiliary unary signatures were obtained for symmetric complex-valued signatures [CGW16a], arbitrary nonnegative signatures [LW18], and arbitrary real-valued signatures [SC20]. These results address restrictions on symmetry or weights that are absent from the present theorem. Other dichotomies treat a single ternary signature on 3-regular graphs, the six-vertex and eight-vertex models, and weighted Eulerian orientations with arrow-reversal symmetry [YHF24, CFX18, CF23, CFS20a].

In 2025, Meng, Wang, and Xia proved an $\mathsf{FP}^{\mathsf{NP}}$ versus #P-hard dichotomy for weighted Eulerian orientations [MWX25]. Meng, Wang, Xia, and Zheng combined this result with tensor decomposition to obtain the same classification for complex Boolean Holant with a nonzero odd-arity signature [MWXZ25]. In 2026, Guan, Shao, and Shi gave a polynomial-time algorithm based on linear programming for the former $\mathsf{FP}^{\mathsf{NP}}$ cases, yielding FP versus #P-hard dichotomies for both problems [GSS26].

In 2026, Liu and Meng classified complex Boolean Holant with binary disequality $X(a, b) = \mathbf{1}[a \neq b]$ available. Their theorem allows an arbitrary finite algebraic complex-valued signature set and gives an explicit decidable tractability criterion [LM26a]. Their theorem has an equivalent formulation for K-Holant with binary equality $I(a, b) = \mathbf{1}[a = b]$ available, obtained by replacing equality and disequality wires [LM26a, Theorem 1.2]. We use this formulation below; K-Holant is defined in Section 2.7.

Xia also proposed a framework that uses finite groups associated with binary signatures to organize Boolean tensor-network classifications. It analyzes several branches under the assumption that all realizable nonzero quaternary signatures decompose [Xia26].

Our proof handles the unrestricted all-even case by combining Liu and Meng's binary-equality dichotomy and their dichotomy for tensor-prime even-parity quaternary signatures with both endpoint values nonzero [LM26a] with a height decomposition based on Hamming imbalance.

Together with the odd-arity dichotomy, this gives the full classification in the unrestricted complex cell of Table 2.

Independent work. We have been informed of at least the following related work, which was carried out independently of the present paper and brought to our attention before its public release. Yuan Huang informed us of work on complex-valued Holant problems on 4-regular graphs, for which a manuscript is being prepared for posting on arXiv. Jiayi Zheng informed us of unpublished work on the complex-valued Holant dichotomy in the case where the realizable binary signatures form an odd dihedral group. Mingji Xia informed us of work concerning the twisted Klein-group case and the C_1 case, under the condition that all quaternary signatures are decomposable. The results and proofs in the present paper were developed independently of these works.

1.2 Our result: the full dichotomy

The predicate $\mathsf{Tract}(\mathcal{F})$, defined in Section 2.13, is an equivalent formulation of the seven conditions in the odd-arity dichotomy of Meng, Wang, Xia, and Zheng [MWXZ25, Theorem 30]. Our result applies this criterion without an odd-arity assumption. Each condition concerns the signature set as a whole. In each holographic alternative, one basis must work simultaneously for the edge and every member of $\mathcal{F}$.

Theorem 1.1 (Full dichotomy theorem)**.** *Let $\mathcal{F}$ be a finite set of algebraic complex-valued Boolean signatures, given in the exact algebraic representation fixed in Section 2.2. The predicate $\mathsf{Tract}(\mathcal{F})$ is decidable, and the following complexity classification holds:*

$$\begin{aligned} \mathsf{Tract}(\mathcal{F}) \quad &\Longrightarrow \quad \mathrm{Holant}(\mathcal{F}) \in \mathsf{FP}, \\ \neg\mathsf{Tract}(\mathcal{F}) \quad &\Longrightarrow \quad \mathrm{Holant}(\mathcal{F}) \textit{ is } \#\mathsf{P}\textit{-hard} \end{aligned}$$

under polynomial-time Turing reductions.

The theorem allows arbitrary mixed arities, zero and tensor-decomposable signatures, and singular binary factors. Recognition is effective for the fixed exact algebraic encoding, while evaluation and all oracle reductions are polynomial in the input-instance size.

1.3 Proof organization

In the standard K-basis, ordinary equality edges become disequality X-edges. Each such edge contributes one zero and one one, so the *half-charge*—half the difference between the numbers of ones and zeros—is conserved. Cases supported only on one side of the central layer, or on a single layer, reduce to the known weighted Eulerian-orientation classification. Before constructing the height parameter, Section 3 proves the binary, unary, EO, and equality reductions used in the later cases. In the remaining two-sided, multilayer case, Section 4 shows that tensor-factor extraction and the charge argument either reach an already classified branch or produce a *mountain*. If d is the gcd of the nonzero half-charges supported by the current signature set, this is an arity-$2d$ tensor with nonzero all-zero and all-one coefficients and all remaining support on the central weight-d layer. This height d indexes the rest of the proof.

The first two heights are treated separately. At height one the mountain is binary. The singular and infinite-order subcases fall into classifications already available. In the remaining finite-order subcase, the symmetry split either exposes a nonsingular weighted equality, which can be normalized to I before applying the binary-equality dichotomy [LM26a, Theorem 1.2], stated here as Theorem 2.23, or reduces to an odd cyclic or odd dihedral form. Those two forms are resolved within the same height-one analysis, giving Theorem 5.1.

At height two the mountain is an eight-vertex signature whose two endpoint values are nonzero. If it does not factor nontrivially, Theorem 2.24 applies; if it does factor, one obtains a nonsingular weighted equality and returns to the binary-equality case. Together the two branches give Corollary 2.25.

Height three is settled by the loop and two-copy contractions in Theorem 7.1. For all heights at least four, Theorem 8.1 uses one relative port normalization and a global phase recovery. At heights at least five, the resulting matching-product structure gives a hard two-copy output directly. Height four is the exceptional eight-port case: a classification of the possible matchings leads to the Reed–Muller core and a Turing reduction adjoining binary equality. A final normalization argument transports each common tractable representation or hardness reduction back to the original coordinates.

The appendix records the exact finite certificates and their direct-use consequences for the odd-dihedral and high-charge arguments.

1.4 Related work

Beyond dichotomies for exact counting on general graphs, related work studies restricted graph classes, stronger lower bounds under #ETH, and algebraic criteria for equality of Holant values.

Restricting the input graph to be planar can make additional counting problems tractable. For Boolean counting CSP, Cai, Lu, and Xia classified real-valued symmetric constraints [CLX17]. Guo and Williams extended the classification to symmetric constraints with complex weights [GW20], and Cai and Fu removed the symmetry assumption [CF22]. In this planar model, the incidence graph is embedded in the plane and the argument order follows the cyclic order around each constraint vertex. Every case that is hard on general graphs but tractable in this planar setting admits a holographic reduction to matchgates. These reductions evaluate the partition function through weighted perfect-matching sums computed by the FKT algorithm.

Graph homomorphism problems are a subclass of counting CSP, so these Boolean-domain results also classify planar graph homomorphisms with complex symmetric weight matrices. For larger domains, Cai and Maran proved a dichotomy for real symmetric 3×3 matrices in 2023 [CM23]. Their 2024 preprint classified nonnegative real symmetric 4×4 matrices of full rank [CM24]. In a 2026 preprint, Cai, Maran, and Young obtained a dichotomy for nonnegative real symmetric matrices of arbitrary finite size whose diagonal entries are pairwise distinct [CMY26].

For planar Holant with symmetric complex-valued Boolean signatures, Cai, Fu, Guo, and Williams gave a classification that includes tractable cases beyond holographic reductions to the FKT algorithm [CFGW22]. In 2026, Cai, Fan, Shao, and Tang classified the eight-vertex model into three cases: tractable on general graphs, hard on general graphs but tractable on planar graphs, and hard even on planar graphs [CFST26].

Beyond planar graphs, classifications are also known for other graph restrictions. Thilikos and Wiederrecht characterized the minor-closed graph classes on which perfect matchings can be counted in polynomial time [TW24]. A class is minor-closed if it is closed under vertex deletion, edge deletion, and edge contraction. Meng and Pan's ESA 2026 paper gives dichotomies for complex-weighted Boolean counting CSP when the incidence graph excludes a fixed clique as a minor, including each fixed variable-occurrence bound $D \geq 3$ [MP26]. The incidence graph records variables and constraints as its two types of vertices. Their classifications show how the tractable cases depend on the excluded clique and the occurrence bound.

Another direction strengthens the running-time lower bounds within the hard cases. Under the counting exponential-time hypothesis (#ETH), which rules out counting satisfying assignments of 3-CNF formulas in time subexponential in the number of variables, Brand, Dell, and Roth proved dichotomies for the Tutte polynomial and unweighted Boolean counting CSP [BDR19]. Liu extended the Boolean counting-CSP classification to algebraic complex weights, including bounded-occurrence versions [Liu24].

For bipartite Holant, Cai and Young proved that two pairs of finite signature sets give equal values on all corresponding labelled graphs if and only if their orbit closures under holographic transformations intersect [CY26].

2 Preliminaries and the Tractable Boundary

2.1 Algebraic notation

We use standard algebraic notation. All rings are commutative with identity. The fields $\mathbb{Q}, \mathbb{R}, \mathbb{C}$ denote the rationals, reals, and complex numbers, and $F^\times = F \setminus \{0\}$ for a field F. We also use $\mathbb{F}_3 = \mathbb{Z}/3\mathbb{Z}$. For a finite set V, F^V denotes the functions $V \to F$. We write $[m] = \{1, \ldots, m\}$, with $[0] = \varnothing$, and $\mathrm{Mat}_{m\times n}(F)$ for matrices over F, abbreviated to $\mathrm{Mat}_n(F)$ when $m = n$. Transposes are denoted by M^{T}; they do not include complex conjugation. The group of invertible matrices is $\mathrm{GL}_n(F)$, and I_n is its identity.

2.2 Complexity and exact-input conventions

A *counting problem* is a function from finite encoded instances to exact numbers. The class FP consists of functions computable exactly by a deterministic algorithm in time polynomial in the input length. The class $\#\mathsf{P}$ consists of the functions that count accepting computation paths of a nondeterministic polynomial-time Turing machine. A (possibly complex-valued) counting problem is $\#\mathsf{P}$*-hard* if every function in $\#\mathsf{P}$ has a polynomial-time Turing reduction to it. We use *tractable* to mean “in FP” and *hard* to mean “$\#\mathsf{P}$-hard” under this convention.

An *oracle algorithm* for a problem B may request exact values of B on instances that it constructs. For counting problems A, B, $A \leq_{\mathrm{T}} B$ means that such an algorithm computes A, makes only polynomially many calls to B, constructs instances of polynomial size, and otherwise uses polynomial time. We write $A \equiv_{\mathrm{T}} B$ when both directions hold. Known nonzero algebraic rescalings and exact algebraic postprocessing are allowed. The notation $\mathsf{FP}^{\mathsf{NP}}$ permits an NP decision oracle. Recognition of a finite exact signature list is required to halt with the correct answer; it is not asserted to run in polynomial time in that list. In every tractable alternative, any shared basis, sign or class must work simultaneously for the entire signature set. Throughout,

$$\overline{\mathbb{Q}} := \{\alpha \in \mathbb{C} : p(\alpha) = 0 \text{ for some } 0 \neq p \in \mathbb{Q}[t]\} \qquad \text{and} \qquad \mathbb{F}_2 := \mathbb{Z}/2\mathbb{Z}$$

are respectively the algebraic complex numbers and the two-element field. We fix an *exact algebraic encoding* of $\alpha \in \overline{\mathbb{Q}}$ by its monic irreducible polynomial over $\mathbb{Q}$ and a rectangle in $\mathbb{C}$ whose vertices have rational real and imaginary parts, containing α and no other complex root of that polynomial; the rectangle is called isolating. A *number field* is a finite extension $L/\mathbb{Q}$, and $[L : \mathbb{Q}]$ is its dimension as a $\mathbb{Q}$-vector space. Equality, zero tests, conjugation, and field arithmetic are exact. The signature set is fixed when a Holant problem is evaluated; running time is measured in the encoded instance size.

For later comparison, a weighted Boolean *counting constraint satisfaction problem* (counting CSP) has Boolean variables, local constraint functions, and partition function equal to the sum, over all variable assignments, of the product of all constraint values. The notation $\#\mathrm{CSP}^q$ used below imposes the additional degree-multiple condition specified in Section 2.6.

2.3 Tensor and group conventions

For a nonzero vector, tensor or matrix v over F, write $[v] = F^\times v$ for its projective class. In particular, $\mathrm{PGL}_n(F) = \mathrm{GL}_n(F)/(F^\times I_n)$. A projective line consists of the classes in a two-dimensional linear span. We use the standard terms *Segre variety* for the projective classes of nonzero simple tensors, *ruling* for a linear family with all but one tensor factor fixed, and *secant* for a line through two distinct points. A pencil is a family $q_s = A - sB$; its parameter range is specified at each use.

Tensor products use the convention $V^{\otimes 0} = F$, with empty product 1. A nonzero tensor has rank one when it is a product of one vector on each of its specified factors. For two factors this is matrix rank one. We say a divides f, with cofactor b, when $f = a \otimes b$ up to port permutation and nonzero scalar redistribution. Every tensor is taken to divide the zero tensor. A proper tensor factor has positive arity strictly smaller than its parent.

The notation $\operatorname{span} S$ denotes linear span. An affine subspace is a translate $a + L$ of a linear subspace; its direction is L, and an affine plane has dimension two. For nonempty S, $\operatorname{aff}(S) = a + \operatorname{span}\{x - a : x \in S\}$, where $a \in S$. For $x, y \in \mathbb{F}_2^V$, the dot product $x \cdot y = \sum_{v \in V} x_v y_v$ is taken in $\mathbb{F}_2$. The notation $x_{\widehat{pq}}$ deletes coordinates p, q and keeps the inherited order.

We use $H \leq G$ for a subgroup, $H \triangleleft G$ for a normal subgroup, and $[G : H]$ for its index. An involution has order two. Our finite-group notation is

$$C_m, \qquad D_{2m} = \langle r, s \mid r^m = s^2 = 1,\ srs = r^{-1} \rangle, \qquad V_4 = C_2 \times C_2,$$

where C_m is cyclic of order m and D_{2m} has order $2m$. The symmetric and alternating groups are S_m and A_m. For $S \subseteq G$, the centralizer is $C_G(S) = \{g \in G : gs = sg \text{ for every } s \in S\}$, and the center is $Z(G) = C_G(G)$. If ζ is a primitive m-th root of unity, then $[\mathbb{Q}(\zeta) : \mathbb{Q}] = \varphi(m)$, where φ is Euler's totient.

Angle brackets denote generated subgroup in a group context, linear span in a vector-space context, and the tensor closure $\langle \mathscr{C} \rangle$ for a signature class, as defined in Section 2.6. The Frobenius pairing is defined in Section 2.8.

2.4 Graphs, matchings, and Walsh coordinates

Graphs may have parallel edges and loops; each loop contributes two to $\deg_G(v)$. A perfect matching of a finite set V is a partition into pairs. The mate of a vertex is the other endpoint of its matching edge. An incidence graph joins two families according to the stated incidence relation, and an alternating path alternates two specified edge types. We write K_n for the complete graph and $\binom{V}{k}$ for the k-element subsets of V.

A binary linear code is a subspace of $\mathbb{F}_2^V$, and an affine code is a translate of one. An information set consists of coordinates on which restriction is a bijection onto the corresponding binary coordinate space. For a represented binary matroid, independence means linear independence of the representing vectors; its basis-exchange graph joins bases differing by one exchanged element.

For $x \in \mathbb{F}_2^V$, put $\mathrm{wt}(x) = |\{v : x_v = 1\}|$. A Hamming layer fixes $\mathrm{wt}(x)$; parity means $\mathrm{wt}(x) \bmod 2$. The operations $\oplus$ and $\bigoplus$ denote coordinatewise addition in $\mathbb{F}_2$, while $\bar{x}$ is the bitwise complement. A portwise bit flip by $b \in \mathbb{F}_2^V$ sends f to $f^{\oplus b}(x) = f(x \oplus b)$. Bit-flip/port-permutation orbits use these flips and permutations of argument positions. Binary words are ordered lexicographically with $0 < 1$.

For an integer $r \geq 0$ and $f : \mathbb{F}_2^r \to \overline{\mathbb{Q}}$, viewed as the lexicographically ordered column array $(f(x))_{x \in \mathbb{F}_2^r}$, its support is $\operatorname{supp}(f) = \{x : f(x) \neq 0\}$. We fix the unnormalized Walsh matrix and the normalized Walsh coordinates

$$P = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad \widetilde{f} = (P^{-1})^{\otimes r} f. \tag{1}$$

Thus inversion is $f = P^{\otimes r} \widetilde{f}$. In coordinates,

$$f(x) = \sum_{y \in \mathbb{F}_2^r} (-1)^{x \cdot y} \widetilde{f}(y).$$

Consequently, if $\widetilde{f}$ has even-parity support then $f(\bar{x}) = f(x)$, and if it has odd-parity support then $f(\bar{x}) = -f(x)$. An overline on a complex scalar means complex conjugation. At ambient arity r, we write $\mathbf{1}_{\text{even}}$ and $\mathbf{1}_{\text{odd}}$ for the indicator tensors of even and odd parity; the arity is inferred from context. References below to *Fourier inversion* mean precisely this inverse Walsh formula.

Given a finite set V, a perfect matching M of V, and $\epsilon \in \mathbb{F}_2^M$, the corresponding *matching code* is

$$C(M, \epsilon) = \{x \in \mathbb{F}_2^V : x_u \oplus x_v = \epsilon_{\{u,v\}} \text{ for every } \{u, v\} \in M\}.$$

After choosing an orientation of every edge of M, write $\vec{M}$ for the resulting set of ordered pairs. A nonzero *matching product* on V is a function $f : \mathbb{F}_2^V \to \overline{\mathbb{Q}}$ of the form

$$f(x) = \lambda \prod_{(u,v) \in \vec{M}} b_{uv}(x_u, x_v), \qquad \lambda \in \overline{\mathbb{Q}}^{\times}, \quad b_{uv} : \mathbb{F}_2^2 \to \overline{\mathbb{Q}}.$$

For the all-one choice of ϵ, the set $C(M, \epsilon)$ is

$$\{x \in \mathbb{F}_2^V : x_u + x_v = 1 \text{ for every } \{u, v\} \in M\},$$

where this sum is over the integers. It has one independent binary choice per matching edge. Any coefficient function on these choices is only a coordinate representation; support of this form does not by itself imply a matching-product factorization.

2.5 Gadgets and allowed operations

A gadget is a finite tensor network over the allowed signatures, with labelled external ports. Its output is the tensor obtained by summing over all internal indices. The native edge is equality in ordinary coordinates and disequality in K-coordinates. A *direct* construction uses one fixed finite gadget. Connecting a binary signature at a port, joining two copies, and contracting two ports are specified by their wiring or matrix formulas when used.

A slice fixes some arguments of a tensor. Such a restriction, a linear combination or a Walsh-coordinate calculation does not alone give a gadget. Pinning requires an available point-mass unary. Interpolation instead gives an oracle reduction by solving a linear system from exact evaluations. The convention for adjoining signatures under these reductions is stated in Section 2.8; it does not identify general Turing reductions with direct gadgets.

A signature set is *preprocessed* after removing zero and nullary signatures as in Section 2.9. Tensor-factor extraction and global changes of basis are used only under the hypotheses of the results stated below.

2.6 Signatures, tensors, and Holant instances

Recall that $\overline{\mathbb{Q}}$ is the field of algebraic complex numbers and $\mathbb{F}_2$ is the binary field. For an integer $r \geq 0$, an arity-r Boolean signature is a function $f : \{0,1\}^r \to \overline{\mathbb{Q}}$, identified with its lexicographically ordered column tensor in $(\overline{\mathbb{Q}}^2)^{\otimes r}$. Its *ports* are its labelled argument positions, and their order is part of the data. We write $\operatorname{arity}(f) = r$ and $\operatorname{supp}(f) = \{x : f(x) \neq 0\}$, and use $\mathbf{1}[\mathsf{P}]$ for the indicator of a proposition P. A signature set is *all-even* when each positive-arity member has even arity. The terms nullary, unary, binary, ternary, quaternary, six-ary, and eight-ary mean arity $0, 1, 2, 3, 4, 6, 8$, respectively; an r-port tensor has arity r. For an ordered binary signature B, we identify B with the matrix $(B(a,b))_{a,b \in \mathbb{F}_2}$, with port 1 indexing rows and port 2 indexing columns. All binary-vector arithmetic is over $\mathbb{F}_2$, with $\oplus$ used when bitwise addition should be explicit. The symbols $\mathbf{0}^r$ and $\mathbf{1}^r$ denote the all-zero and all-one words, respectively; at $r = 0$, both denote the unique empty word. More generally, $\mathbf{0}^V, \mathbf{1}^V \in \mathbb{F}_2^V$ are the constant words on a finite index set V. The values $f(\mathbf{0}^r)$ and $f(\mathbf{1}^r)$ are the two *endpoint coefficients*. For $x \in \mathbb{F}_2^V$ and $A \subseteq V$, $x_{\widehat{A}}$ denotes the restriction of x to $V \setminus A$, in the inherited order when V is ordered. For scalars a, b, we use

$$\operatorname{diag}(a,b) = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}.$$

For an integer $r \geq 0$, let S_r denote the symmetric group on $[r]$. For an arity-r signature f and $\tau \in S_r$, $(\tau f)(x) = f(x_{\tau^{-1}(1)}, \ldots, x_{\tau^{-1}(r)})$. For signatures $f : \mathbb{F}_2^r \to \overline{\mathbb{Q}}$ and $g : \mathbb{F}_2^s \to \overline{\mathbb{Q}}$ on disjoint ordered ports, and $x \in \mathbb{F}_2^r, y \in \mathbb{F}_2^s$, define $(f \otimes g)(x,y) = f(x)g(y)$. The signature f is *symmetric* when $\tau f = f$ for every $\tau \in S_r$, equivalently when $f(x)$ depends only on $\operatorname{wt}(x)$. For $T \subseteq [r]$, $\mathbf{1}_T \in \{0,1\}^r$ denotes the incidence vector of T; in particular $\mathbf{1}_\varnothing = \mathbf{0}^r$. We write $e_j = \mathbf{1}_{\{j\}}$ for the j-th standard basis vector whenever the ambient arity is clear. A nonzero positive-arity signature is *tensor-prime* if, up to port permutation and nonzero scalar, it cannot be written as $a \otimes b$ with both a and b of positive arity. A nonzero positive-arity signature that is not tensor-prime is *tensor-decomposable*. Prime factorization is unique up to these operations and factor order [CFS20a, full version, Lemma 2.13]. For any class $\mathscr{C}$ of signatures, its tensor closure $\langle \mathscr{C} \rangle$ consists of tensor products of members of $\mathscr{C}$, up to port permutation and nonzero scalar. By convention, it also contains zero and the arity-zero empty product.

For an integer $r \geq 0$ and $a \in \{0,1\}^r$, define the *point-mass signature* $\delta_a : \{0,1\}^r \to \overline{\mathbb{Q}}$ at a by

$$\delta_a(x) = \mathbf{1}[x = a].$$

For an integer $r \geq 1$ and $a \in \{0,1\}^r$, a *weighted generalized equality* is a signature

$$c_0 \delta_a + c_1 \delta_{\bar{a}}, \qquad c_0, c_1 \in \overline{\mathbb{Q}}^{\times},$$

where $\bar{a}$ is the bitwise complement. It is *pure* when $a = \mathbf{0}^r$. Portwise bit flips make it pure, but are used only when actually available or inside a stated global transform. We abbreviate generalized equality by GE. For $p, q \in \overline{\mathbb{Q}}$, in arity two, $p\delta_{00} + q\delta_{11} = \operatorname{diag}(p,q)$ is a *weighted binary equality*, and $p\delta_{01} + q\delta_{10}$ is a *weighted binary disequality*; either is nonsingular when $pq \neq 0$.

For a signature set $\mathcal{F}$, a signature grid over $\mathcal{F}$ is a finite multigraph $G = (V, E)$, with loops allowed, together with a signature $f_v \in \mathcal{F}$ satisfying $\operatorname{arity}(f_v) = \deg_G(v)$ and an ordering, of that length, of the incident edge-ends at every vertex. An edge assignment $\sigma : E \to \{0,1\}$ therefore induces an ordered word x_v^σ at v. For the resulting grid Ω,

$$\operatorname{Holant}_\Omega = \sum_{\sigma : E \to \{0,1\}} \prod_{v \in V} f_v(x_v^\sigma).$$

The problem Holant($\mathcal{F}$) asks for this value on an input grid over $\mathcal{F}$. Throughout, comma-separated signature or signature-set arguments mean adjoining their union; for example, Holant($\mathcal{F}, a, b$) abbreviates Holant($\mathcal{F} \cup \{a, b\}$).

We use the unary pinnings

$$\Delta_0 = [1, 0]^\mathsf{T}, \qquad \Delta_1 = [0, 1]^\mathsf{T},$$

and the binary equality and disequality tensors

$$I = I_2 = \mathrm{EQ}_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad X = \mathrm{NEQ}_2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

We also fix the sign and skew matrices

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad Y = ZX = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

More generally, for $r \geq 1$, let

$$\mathrm{EQ}_r(x_1, \ldots, x_r) = \mathbf{1}[x_1 = \cdots = x_r].$$

For signature sets $\mathscr{E}$ and $\mathcal{F}$, the bipartite problem Holant($\mathscr{E} \mid \mathcal{F}$) is defined on a bipartite graph $H = (U, V, E)$: vertices in the left part U are assigned signatures from $\mathscr{E}$, and vertices in the right part V are assigned signatures from $\mathcal{F}$. Subdividing every ordinary edge by an EQ_2-vertex gives

$$\mathrm{Holant}(\mathcal{F}) = \mathrm{Holant}(\mathrm{EQ}_2 \mid \mathcal{F}).$$

For an integer $q \geq 1$, the counting-CSP problem $\#\mathrm{CSP}^q(\mathcal{F})$ restricts every variable degree to be divisible by q; equivalently, its variable side uses EQ_{kq}, $k \geq 1$. We use the convention that the variable list contains only variables that occur in at least one constraint. If an input encoding explicitly lists z isolated Boolean variables, they may first be deleted and the final partition value multiplied by 2^z. The notation $\#\mathrm{CSP}^q(\mathcal{F}, \mathcal{H})$ means $\#\mathrm{CSP}^q(\mathcal{F} \cup \mathcal{H})$.

2.7 Holographic transformations and the K-basis

We use the polynomial-time Turing-reduction and exact-postprocessing conventions of Section 2.2.

For tensors or matrices A and B, write

$$A \doteq B$$

when

$$A = \lambda B \qquad \text{for some } \lambda \in \overline{\mathbb{Q}}^\times.$$

In every gadget or reduction identity, this scalar is known, tracked, and restored by exact algebraic postprocessing.

For $M \in \mathrm{GL}_2(\overline{\mathbb{Q}})$ and an arity-r right-side signature f, write

$$Mf = M^{\otimes r} f, \qquad M\mathcal{F} = \{Mf : f \in \mathcal{F}\}.$$

For an arity-r left-side signature g, juxtaposition denotes the dual action $gM = (M^\mathsf{T})^{\otimes r} g$. In particular, $(\mathrm{EQ}_2)M$ has matrix $M^\mathsf{T} M$, and the holographic identity is

$$\mathrm{Holant}(\mathrm{EQ}_2 \mid \mathcal{F}) = \mathrm{Holant}((\mathrm{EQ}_2)M \mid M^{-1}\mathcal{F}), \qquad (\mathrm{EQ}_2)M = M^\mathsf{T} M.$$

This exact identity makes the transformed edge a genuine weighted binary, not a componentwise scalar.

Write $Z_X(\Omega)$ for the partition function of a grid whose internal edges carry X. The general orthogonal similitude group of X is

$$GO(X) = \{D \in \mathrm{GL}_2(\overline{\mathbb{Q}}) : D^{\mathsf{T}} X D = \lambda X \text{ for some } \lambda \in \overline{\mathbb{Q}}^{\times}\}$$

for legal global transformations preserving the native disequality edge X projectively. If $E(\Omega)$ is the internal edge set, $D\Omega$ applies D at every vertex, and $D^{\mathsf{T}} X D = \lambda X$, then

$$Z_X(D\Omega) = \lambda^{|E(\Omega)|} Z_X(\Omega). \tag{2}$$

Thus the two problems differ only by the known edge-count scalar. The matrix acts globally and is not thereby a local gadget. For $m \geq 1$, let

$$\mu_m = \{z \in \overline{\mathbb{Q}} : z^m = 1\}.$$

We use the projective linear group $\mathrm{PGL}_2(\overline{\mathbb{Q}})$ defined in Section 2.3. Thus, for a nonsingular 2-by-2 matrix T, the general projective notation $[T]$ from Section 2.3 identifies matrices differing by a nonzero algebraic scalar. The *projective order* of $[T]$ is the least $m \geq 1$ for which $T^m = \lambda I$ for some $\lambda \in \overline{\mathbb{Q}}^{\times}$, and is infinite if no such m exists.

Proposition 2.1 (Finite projective binary groups)**.** *Every finite subgroup of* $\mathrm{PGL}_2(\mathbb{C})$ *is conjugate to exactly one of the following types:*

$$C_m, \qquad D_{2m}, \qquad A_4, \qquad S_4, \qquad A_5,$$

where $m \geq 1$ *for* C_m, $m \geq 2$ *for* D_{2m}, $|C_m| = m$, $|D_{2m}| = 2m$, *and* $|A_4| = 12$, $|S_4| = 24$, $|A_5| = 60$.

Proof. This is [LM26a, Proposition 2.1]. □

Lemma 2.2 (Finite projective-group facts)**.** *The following properties will be used below.*

1. *In*
$$D_{2m} = \langle r, s \mid r^m = s^2 = 1,\ srs = r^{-1}\rangle,$$
the rotations $\langle r \rangle \cong C_m$ *form a normal index-two subgroup, and every element outside it is a reflection and an involution. If* m *is odd, the rotation group is the unique index-two subgroup. If* m *is even, the only nonidentity rotational involution is the central element* $r^{m/2}$.

2. *The group* A_4 *has no index-two subgroup, and the centralizer of an involution is* V_4.

3. *The unique index-two subgroup of* S_4 *is* A_4. *Every involution is a transposition or a double transposition; precisely the double transpositions belong to* A_4.

4. *The group* A_5 *is simple and has no index-two subgroup. The centralizer of each involution is* V_4.

5. *The projective centralizer of the distinguished involution is*
$$C_{\mathrm{PGL}_2(\mathbb{C})}([Z]) = \left\{ \left[\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}\right] : a, b \in \mathbb{C},\ ab \neq 0 \right\} \sqcup \left\{ \left[\begin{pmatrix} 0 & a \\ b & 0 \end{pmatrix}\right] : a, b \in \mathbb{C},\ ab \neq 0 \right\}.$$
Thus $[Z]$ *is the only nonidentity diagonal involution, while every nonsingular anti-diagonal class is an involution.*

Proof. This is [LM26a, Lemma 2.2]. $\square$

The conjugacy classification above is group-theoretic only. Its conjugating matrix need not lie in $GO(X)$, so it does not by itself define a legal global K-Holant normalization or provide an actual binary gadget. Later applications retain the distinguished classes $[X]$ and $[Z]$.

Fix the K-coordinate matrix

$$K = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}, \qquad K^{-1} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix}. \tag{3}$$

Since $K^{\mathsf{T}}K = X$, putting $\widehat{\mathcal{F}} = K^{-1}\mathcal{F}$ gives

$$\text{Holant}(\mathcal{F}) \equiv_{\text{T}} \text{Holant}(X \mid \widehat{\mathcal{F}}) =: \text{K-Holant}(\widehat{\mathcal{F}}). \tag{4}$$

For an original-coordinate signature f, or an original-coordinate signature set $\mathcal{F}$, a hat records passage to K-coordinates:

$$\widehat{f} = K^{-1}f, \qquad \widehat{\mathcal{F}} = K^{-1}\mathcal{F}.$$

Once we work with a signature set already in K-coordinates, we suppress these hats and denote its member signatures simply by f. Unless stated otherwise, signatures and signature sets appearing in a K-Holant problem are understood to be in K-coordinates, with native edge X. For a signature f already in K-coordinates, a tilde denotes the Walsh coordinates fixed in Equation (1); using them in a calculation does not implement a global transformation. For $s \in \overline{\mathbb{Q}}$, the identities

$$P^{-1}\left[\frac{1}{2}\begin{pmatrix} 1+s & 1-s \\ 1-s & 1+s \end{pmatrix}\right]P = \text{diag}(1, s), \qquad P^{\mathsf{T}}XP = 2Z$$

explain its use in the odd cyclic and dihedral calculations.

2.8 Gadgets, interpolation, and factor reductions

Equivalence convention. Throughout a proof, $\mathsf{H}(\mathcal{G})$ denotes the relevant Holant or K-Holant problem, and $\mathcal{G}$ denotes the current finite signature set up to polynomial-time Turing equivalence. Whenever a fixed gadget construction, polynomial interpolation, or a cited reduction supplies a finite set S such that

$$\mathsf{H}(\mathcal{G} \cup S) \leq_{\text{T}} \mathsf{H}(\mathcal{G}), \tag{5}$$

the reverse reduction is immediate by restriction. We therefore replace $\mathcal{G}$ by $\mathcal{G} \cup S$, keep the notation $\mathcal{G}$, and thereafter write $g \in \mathcal{G}$ for $g \in S$, without recording its construction history. Every such enlargement is finite. Signatures that differ by a known nonzero scalar are identified, with the scalar restored in the final value. Formal slices, linear combinations, Walsh-coordinate identities, and local changes of basis do not add signatures unless one of the stated reductions supplies them. A whole-set $GO(X)$-normalization is instead a holographic equivalence and is governed by Lemma 2.22. This is the problem-level form of the gadget and interpolation rules used by Liu and Meng [LM26a, Section 2.5] and by [MWXZ25, Section 2.3]; the factor instance is supplied by Theorem 2.14.

A *direct gadget* over $\mathcal{G}$ is a finite X-edge network of $\mathcal{G}$-vertices with ordered external ports. If it realizes h, then

$$\text{K-Holant}(\mathcal{G} \cup \{h\}) \leq_{\text{T}} \text{K-Holant}(\mathcal{G}). \tag{6}$$

A fixed gadget has constant size. A gadget realizing ρh, for $\rho \in \overline{\mathbb{Q}}^{\times}$, implements the same Turing reduction after an instance with N occurrences is rescaled by ρ^{-N}. Thus a direct gadget is one instance of the equivalence convention above. In arguments using local topology or port order, we specify a direct gadget rather than relying only on membership obtained by a Turing reduction.

Let f have arity r. For an ordered pair of distinct ports $p=(i,j)$, a binary kernel $C:\mathbb{F}_2^2\to\overline{\mathbb{Q}}$, and a residual word z indexed by $[r]\setminus\{i,j\}$ in increasing port order, let $z^{i\leftarrow a,j\leftarrow b}$ be the arity-r word obtained by inserting a,b at ports i,j. Define the two-port contraction by

$$(\partial_p^C f)(z)=\sum_{a,b\in\mathbb{F}_2} C(a,b)f(z^{i\leftarrow a,j\leftarrow b}).$$

We also write $f_{ij}^{ab}(z)=f(z^{i\leftarrow a,j\leftarrow b})$ for the residual slice. If an ordered binary B is placed on the ordered port pair $e=(u,v)$, we write B_e. The notation $B_e\mid f$ means tensor divisibility on this pair in the sense of Section 2.3. In the notation of Shao–Cai [SC20, full version, Section 2.4], their $\widehat{\partial}_{ij}\widehat{f}$ and $\partial_{ij}f$ are, respectively, our $\partial_{(i,j)}^X\widehat{f}$ and $\partial_{(i,j)}^I f$. Port order fixes the ordered pair (i,j). For $U,V\in\mathrm{Mat}_2(\overline{\mathbb{Q}})$, the Frobenius pairing is bilinear:

$$\langle U,V\rangle_{\mathrm{Frob}}=\mathrm{tr}(U^{\mathsf{T}}V)=\sum_{a,b\in\mathbb{F}_2}U(a,b)V(a,b),$$

without conjugation. Two matrices are *orthogonal* for this pairing when their Frobenius pairing is zero. Suppose a directly realized subnetwork with ordered boundary tensor B closes ports p. Its two attachment X-edges complement both boundary bits, so

$$C(a,b)=\sum_{u,v\in\mathbb{F}_2}X(a,u)B(u,v)X(v,b)=(XBX)(a,b)=B(1-a,1-b).$$

Exchanging the two ordered boundary inputs replaces B by B^{T}, and the same attachment therefore gives $XB^{\mathsf{T}}X$. Likewise, two inward-ordered binary chains L and R joined by one native X-edge have external tensor LXR^{T}. These formulas explain where transposes arise in direct gadget constructions and fix the ordered-port convention used below. Every contraction so constructed is directly realized; a contraction with arbitrary C, or a linear combination of contractions, is only a tensor calculation unless a separate reduction makes it available.

For any ordered binary signature B, its *transfer* is

$$T_B=BX.$$

Series composition through one native edge multiplies transfers. Reversing the two ports induces

$$T^{\#}=XT^{\mathsf{T}}X.$$

Thus $(ST)^{\#}=T^{\#}S^{\#}$ and $(T^{\#})^{\#}=T$. For nonsingular B, the projective order of its transfer is the order of $[BX]$. Global normalizations act on the whole problem, not by adjoining individual signatures; after applying Lemma 2.22, the transformed finite set is again denoted by the same letter.

Let $G\le\mathrm{PGL}_2(\overline{\mathbb{Q}})$ be finite, and suppose each class has a directly realized ordered-binary representative. Choose one nonsingular ordered binary B_g for each class g, normalized by $R_g=B_gX$, $[R_g]=g$, and $R_1=I$, where $1=[I]$. The matrix of every representative of class g is ρB_g for a unique nonzero scalar ρ. Since G is finite, the equivalence convention lets us adjoin these representatives once and write $B_g\in\mathcal{G}$. Displayed projective identities are read up to their known nonzero scalars; a scalar is shown explicitly whenever its zero/nonzero status or a coefficient ratio is used.

Interpolation is used only through a stated polynomial-time oracle reduction; linear combinations are not gadgets. The tensor-factor extraction result is Theorem 2.14, with its common-one-sided exception kept explicit.

2.9 Zero and nullary preprocessing

An arity-zero (or *nullary*) signature is a scalar. Let $\mathcal{F}^\circ$ be obtained from $\mathcal{F}$ by deleting all identically zero and nullary signatures. An instance containing a zero-signature vertex has value zero. Otherwise, deleting its nullary vertices gives Ω°, and

$$\mathrm{Holant}_\Omega = \left(\prod_{\substack{c \in \mathcal{F} \\ \mathrm{arity}(c)=0}} c^{m_c(\Omega)} \right) \mathrm{Holant}_{\Omega^\circ},$$

where $m_c(\Omega)$ is the multiplicity of c and the empty residual instance has value 1. Thus all structural arguments assume, without further mention, that signatures are nonzero and have positive arity; the evaluation algorithm restores the known nullary factors exactly.

All algebraic arithmetic and recognition use the exact-input conventions of Section 2.2.

2.10 Charge and Hamming layers

All charge terminology in this subsection is in K-coordinates. Recall the Hamming weight $\mathrm{wt}(x) = |\{j \in [r] : x_j = 1\}|$ from Section 2.4. For an integer $r \geq 0$, define the doubled-charge function $q_r : \mathbb{F}_2^r \to \mathbb{Z}$ by

$$q_r(x) = 2\,\mathrm{wt}(x) - r.$$

For an even integer r, we also use the integer half-charge

$$\chi_r(x) = \frac{q_r(x)}{2} = \mathrm{wt}(x) - \frac{r}{2}.$$

In the all-even analysis, gcd, divisibility, height, and interval statements use χ_r; *doubled charge* explicitly means q_r. Put

$$HW_r^{\geq} = \{x \in \mathbb{F}_2^r : q_r(x) \geq 0\}, \quad HW_r^{\leq} = \{x \in \mathbb{F}_2^r : q_r(x) \leq 0\}, \quad HW_r^{=} = \{x \in \mathbb{F}_2^r : q_r(x) = 0\},$$

and define $HW_r^{>}, HW_r^{<}$ by strict inequalities. The unsubscripted symbols use the ambient arity:

$$HW^{\geq}, \qquad HW^{\leq}, \qquad HW^{=}, \qquad HW^{>}, \qquad HW^{<}.$$

A signature has one of these properties when its support lies in the corresponding set. A signature is *single-layer* when its support lies in one Hamming layer; a signature set is single-layer when each member is so supported, with member-dependent layers allowed, and is *multilayer* otherwise. A signature set is *common one-sided* when all members are $HW^{\geq}$ or all are $HW^{\leq}$; otherwise it is *two-sided*, equivalently both charge signs occur among support words of its members. Here EO abbreviates Eulerian orientation. An arity-$2d$ *EO signature* is supported in $HW_{2d}^{=}$; this central support condition is the local balance constraint for an Eulerian orientation.

Definition 2.3 (Mountain)**.** For an integer $d \geq 1$, an arity-$2d$ signature M in K-coordinates is a mountain of height d if

$$\mathrm{supp}(M) \subseteq HW_{2d}^{=} \cup \{\mathbf{0}^{2d}, \mathbf{1}^{2d}\} \qquad \text{and} \qquad M(\mathbf{0}^{2d})M(\mathbf{1}^{2d}) \neq 0.$$

Equivalently,

$$\{\chi_{2d}(x) : M(x) \neq 0\} \subseteq \{-d, 0, d\}, \qquad \{-d, d\} \subseteq \{\chi_{2d}(x) : M(x) \neq 0\}.$$

Thus the height is both half the arity and the absolute integer half-charge of either endpoint. No restriction is placed on the central EO layer.

For an X-edge grid Ω with vertex set $V(\Omega)$, every edge contributes one 0-bit and one 1-bit. Hence each contributing assignment x, with induced local word x_v at vertex v, satisfies

$$\sum_{v\in V(\Omega)} q_{\mathrm{arity}(f_v)}(x_v) = 0.$$

Thus X-contraction preserves external total charge.

Remark 2.4 (Notation relative to prior work). We follow [MWXZ25, Sections 2.1–2.3] in writing $\widehat{\mathcal{F}} = K^{-1}\mathcal{F}$, $HW^{\geq}$, $HW^{\leq}$, $HW^{=}$, and *EO signature*. Their $\neq_2$ and $=_r$ are our X and EQ_r, respectively; subscripts in $HW_r^{\bullet}$ record arity. Thus $q_r(x) = |\{j \in [r] : x_j = 1\}| - |\{j \in [r] : x_j = 0\}|$, while the height arguments use its integer half $\chi_r = q_r/2$.

Our span is linear span; the affine span denoted $\mathrm{Span}(f)$ in [MWXZ25] is written $\mathrm{aff}(\mathrm{supp}(f))$.

2.11 Bell signatures and the Reed–Muller core

For $d, \ell, u, v \in \mathbb{F}_2$, define the ordered Bell signature

$$B_{d,\ell}(u,v) = \mathbf{1}[u \oplus v = d](-1)^{\ell u}.$$

Write $B_\gamma = B_{d,\ell}$ for $\gamma = (d,\ell)^{\mathsf{T}}$. Port order matters for $B_{1,1}$, which is negated under reversal.

With truth-table order $00, 01, 10, 11$, these four signatures are

$$\begin{aligned} B_{0,0} &= (1,0,0,1) = I, \quad & B_{0,1} &= (1,0,0,-1) = Z, \\ B_{1,0} &= (0,1,1,0) = X, \quad & B_{1,1} &= (0,1,-1,0) = Y. \end{aligned} \tag{7}$$

On an even finite port set, a *Bell product* is a nonzero scalar multiple of a tensor product of Bell signatures on the oriented pairs of a perfect matching. Reversing a pair carrying $B_{1,1}$ changes the scalar sign; on the empty port set a Bell product is a nonzero scalar. A binary proportional to a Bell signature is written $B \doteq B_\gamma$. Our real skew matrix Y differs from the conventional complex Pauli Y-matrix by a nonzero scalar.

Remark 2.5 (Bell notation in Shao–Cai). The tuple in Equation (7) is Shao–Cai's $(=_2^+, =_2^-, \neq_2^+, \neq_2^-)$ [SC20, Section 2]; $Y^{\mathsf{T}} = -Y$. Their holographic matrix named Z is our K, whereas our Z is only the diagonal sign matrix.

Index the eight ports by the points of $\mathbb{F}_2^3$ in lexicographic order. The first-order Reed–Muller code is

$$RM(1,3) = \{(\ell(t))_{t\in\mathbb{F}_2^3} \mid \ell\colon \mathbb{F}_2^3 \to \mathbb{F}_2 \text{ is affine}\}.$$

Explicitly, its elements are the truth tables of

$$\ell(t_1,t_2,t_3) = a + u_1t_1 + u_2t_2 + u_3t_3, \qquad a, u_1, u_2, u_3 \in \mathbb{F}_2,$$

with all arithmetic over $\mathbb{F}_2$. Evaluating such an ℓ on the eight inputs $t \in \mathbb{F}_2^3$, in the fixed order above, gives one binary string of length eight. The $2^4 = 16$ choices of (a, u_1, u_2, u_3) give exactly the sixteen codewords of $RM(1,3)$. For $x \in \mathbb{F}_2^8$, we write

$$\mathrm{RM}_8(x) = \mathbf{1}[x \in RM(1,3)]. \tag{8}$$

Remark 2.6 (The Shao–Cai f_8 core). With the preceding port order, RM_8 is the signature $f_8 = \widehat{f}_8$ of [SC20, full version, Section 8]. Its Bell contraction identities are established in [CFS20b, Theorems 13–14]; Lemma 8.12 proves the reduction adjoining equality to the current signature set.

2.12 Exact finite verification

The appendix states the finite calculations used in the proof and describes how the supplied programs reconstruct and check the candidates from their defining equations. All arithmetic is exact. For an integer matrix, a nonzero k-by-k minor, or rank k modulo a prime, gives rational rank at least k. For a matrix with n columns, s independent rational kernel vectors give rank at most $n-s$. We use modular arithmetic only with these stated rank bounds, not as a substitute for a characteristic-zero identity.

The six-port calculation also uses exact satisfiability checks for real polynomial systems after splitting complex variables into real and imaginary parts. The appendix records the normalization and coordinate conventions, software requirements and which conclusions depend on the solver. The correspondence between results and programs is in `certificates/theorem-to-script-map.md`.

2.13 Tractable signature sets and effective recognition

2.13.1 The standard signature classes

We include the zero signature in every class below. The affine class $\mathscr{A}$ consists of the zero signature and the nonzero arity-r signatures

$$f(x) = \lambda\, \mathbf{1}[Ax = b]\, i^{Q(x)}, \tag{9}$$

where $r, m \geq 0$ are integers, $\lambda \in \overline{\mathbb{Q}}^{\times}$, $A \in \mathrm{Mat}_{m\times r}(\mathbb{F}_2)$, $b \in \mathbb{F}_2^m$, and $Q : \mathbb{F}_2^r \to \mathbb{Z}_4$, where $\mathbb{Z}_4 = \mathbb{Z}/4\mathbb{Z}$, is represented by a multilinear polynomial over $\mathbb{Z}_4$ of degree at most two whose mixed coefficients are even. Equivalently, the phase is a product of fourth-root phases of affine linear forms interpreted as $\{0,1\}$-valued integers. Here a *mixed coefficient* is the coefficient of $x_i x_j$ for $i \neq j$, and an affine linear form over $\mathbb{F}_2$ is $a_0 + \sum_{j=1}^{r} a_j x_j$ with all coefficients in $\mathbb{F}_2$.

The product class $\mathscr{P}$ consists of zero and all signatures of the form $\lambda \prod_{j=1}^{m} h_j$, where $m \geq 0$ is an integer, $0 \neq \lambda \in \overline{\mathbb{Q}}$, and every factor h_j is a unary, binary equality, or binary disequality constraint applied to the signature variables. Empty products and repeated variables are allowed; an inconsistent product denotes the zero signature.

For every integer $q \geq 2$, put

$$\rho_q = e^{\pi i/(2q)}, \qquad T_q = \mathrm{diag}(1, \rho_q).$$

For $r \in \mathbb{Z}$, define the arity-wise affine twist $\mathscr{A}_q^r$ by

$$f \in \mathscr{A}_q^r \quad \Longleftrightarrow \quad (T_q^{-r})^{\otimes\, \mathrm{arity}(f)} f \in \mathscr{A}.$$

Since $T_q^q = \mathrm{diag}(1, i)$ preserves $\mathscr{A}$, we take $r \bmod q$; in particular, $T_2 = \mathrm{diag}(1, e^{\pi i/4})$. The local-affine class $\mathscr{L}$ consists of the signatures f such that

$$\left(\bigotimes_{j=1}^{\mathrm{arity}(f)} T_2^{\alpha_j} \right) f \in \mathscr{A} \qquad \text{for every } \alpha \in \mathrm{supp}(f).$$

Let $\mathscr{T}$ be the set of all unary and binary signatures, and let

$$\mathscr{M} = \{ f : f \text{ is a Boolean signature and } f(x) = 0 \text{ for every } x \text{ with } \mathrm{wt}(x) > 1 \}.$$

Definition 2.7 (Common transformability and K-presentation)**.** Let $\mathscr{C}$ be closed under nonzero scalars and port permutations. A finite signature set $\mathcal{F}$ is $\mathscr{C}$-transformable if one $M \in \mathrm{GL}_2(\overline{\mathbb{Q}})$ satisfies

$$M^{\mathsf{T}} M \in \mathscr{C}, \qquad M^{-1}\mathcal{F} \subseteq \mathscr{C}.$$

Equivalently, for $\widehat{\mathcal{F}} = K^{-1}\mathcal{F}$, one matrix $U = K^{-1}M$ gives the common K-presentation

$$U^{\mathsf{T}} X U \in \mathscr{C}, \qquad U^{-1}\widehat{\mathcal{F}} \subseteq \mathscr{C}.$$

In both displays the edge is included and the quantifier order is one basis for the entire signature set, not one basis per signature. We use this definition for $\mathscr{C} \in \{\mathscr{A}, \mathscr{P}, \mathscr{L}\}$ and for every fixed affine twist $\mathscr{A}_q^r$ (in particular $\mathscr{A}_2^1$).

2.13.2 The Eulerian-orientation predicate

For an arity-r signature g, define its central restriction

$$g|_{\mathrm{EO}}(x) = g(x)\mathbf{1}[q_r(x) = 0].$$

It is zero for odd r, and zero restrictions are deleted. For a signature set $\mathcal{G}$, set

$$\mathcal{G}|_{\mathrm{EO}} = \{g|_{\mathrm{EO}} : g \in \mathcal{G},\ g|_{\mathrm{EO}} \neq 0\}.$$

Let $g \neq 0$ have arity k and be supported on the single Hamming layer of weight d. Its *EO completion* is the ordered tensor

$$g_{\to\mathrm{EO}} = \begin{cases} g \otimes \Delta_0^{\otimes(2d-k)}, & 2d \geq k, \\ g \otimes \Delta_1^{\otimes(k-2d)}, & 2d < k. \end{cases} \tag{10}$$

The original ports precede the added unaries; a zeroth power is the nullary unit. The completed signature has even arity and central support, hence is EO. For a signature set of nonzero single-layer signatures, write

$$\mathcal{G}_{\to\mathrm{EO}} = \{g_{\to\mathrm{EO}} : g \in \mathcal{G}\}.$$

For arity-$2d$ EO signature g, an oriented perfect pairing $\pi = ((a_1, b_1), \ldots, (a_d, b_d))$ has unordered pairs partitioning $[2d]$. For $z \in \mathbb{F}_2^d$, let $y \in \mathbb{F}_2^{2d}$ be the unique word with $y_{a_j} = z_j$ and $y_{b_j} = 1 - z_j$; the pairing induces

$$g^{\pi}(z_1, \ldots, z_d) = g(y).$$

Definition 2.8 (The predicate $\mathsf{EOTract}$)**.** For a finite EO signature set $\mathcal{G}$, $\mathsf{EOTract}(\mathcal{G})$ holds if there exist a single sign $\varepsilon \in \{+1, -1\}$ and a single class $\mathscr{C} \in \{\mathscr{A}, \mathscr{P}\}$, both shared by the entire signature set, such that:

1. for every $g \in \mathcal{G}$ and all $\alpha, \beta, \gamma \in \mathrm{supp}(g)$, if $\delta = \alpha \oplus \beta \oplus \gamma$, then $\delta \in \mathrm{supp}(g)$ or $\varepsilon q_{\mathrm{arity}(g)}(\delta) > 0$;
2. for every $g \in \mathcal{G}$ and every oriented perfect pairing π of its ports, $g^{\pi} \in \mathscr{C}$.

We adopt the vacuous convention $\mathsf{EOTract}(\varnothing) = \text{true}$.

This is the opposite-pair restriction of [MWXZ25, GSS26]: substitute $y_{b_j} = 1 - y_{a_j}$, with the converse lifted, for a pairing with d pairs, by $\prod_{j=1}^{d} X(y_{a_j}, y_{b_j})$. Both $\mathscr{A}, \mathscr{P}$ are closed under it.

2.13.3 The tractability predicate

For an arbitrary finite signature set $\mathcal{F}$, define $\mathsf{Tract}(\mathcal{F}) := \mathsf{Tract}(\mathcal{F}^\circ)$, where $\mathcal{F}^\circ$ is defined in Section 2.9. For a preprocessed signature set $\mathcal{F}$, put $\widehat{\mathcal{F}} = K^{-1}\mathcal{F}$. Define $\mathsf{Tract}(\mathcal{F})$ to be the disjunction of the following seven alternatives.

(1) The signature set $\widehat{\mathcal{F}}$ is common one-sided, in the explicit sense that

$$[\forall \widehat{f} \in \widehat{\mathcal{F}},\ \mathrm{supp}(\widehat{f}) \subseteq HW^{\geq}_{\mathrm{arity}(\widehat{f})}] \quad \text{or} \quad [\forall \widehat{f} \in \widehat{\mathcal{F}},\ \mathrm{supp}(\widehat{f}) \subseteq HW^{\leq}_{\mathrm{arity}(\widehat{f})}],$$

and $\mathsf{EOTract}(\widehat{\mathcal{F}}|_{\mathrm{EO}})$ holds.

(2) Every $\widehat{f} \in \widehat{\mathcal{F}}$ is supported on one Hamming weight d_f. Across the signature set, some support word has negative charge and some (possibly other) support word has positive charge, and $\mathsf{EOTract}(\widehat{\mathcal{F}}_{\to\mathrm{EO}})$ holds, with the completion operation defined in Equation (10).

(3) $\mathcal{F} \subseteq \langle \mathscr{T} \rangle$.

(4) $\mathcal{F} \subseteq \langle K\mathscr{M} \rangle$, or $\mathcal{F} \subseteq \langle KX\mathscr{M} \rangle$.

(5) $\mathcal{F}$ is $\mathscr{A}$-transformable.

(6) $\mathcal{F}$ is $\mathscr{P}$-transformable.

(7) $\mathcal{F}$ is $\mathscr{L}$-transformable.

Remark 2.9. For every signature f, the matrix symbols in Item (4) act tensorwise at the arity of each signature. In particular,

$$\begin{aligned} f \in \langle K\mathscr{M} \rangle &\iff \text{every tensor-prime factor of } K^{-1}f \text{ lies in } \mathscr{M}, \\ f \in \langle KX\mathscr{M} \rangle &\iff \text{every tensor-prime factor of } XK^{-1}f \text{ lies in } \mathscr{M}. \end{aligned}$$

Lemma 2.10 (Effective algebraic transformability witnesses)**.** *Let $E \in \mathrm{Mat}_2(\overline{\mathbb{Q}})$, let $\mathcal{F}$ be a finite algebraic signature set, and let $\mathscr{C} \in \{\mathscr{A}, \mathscr{P}, \mathscr{L}\}$. Suppose there is a matrix $U \in \mathrm{GL}_2(\mathbb{C})$ such that*

$$U^{\mathsf{T}}EU \in \mathscr{C}, \qquad U^{-1}\mathcal{F} \subseteq \mathscr{C},$$

where membership in $\mathscr{C}$ allows complex values for the free scalar and unary parameters in its normal forms. Then such a witness exists in $\mathrm{GL}_2(\overline{\mathbb{Q}})$. Moreover, for exact algebraic input, the existence of a witness is decidable and an algebraic witness can be produced.

Proof. This is [LM26a, Lemma 2.15]. □

2.13.4 Dichotomies and tensor decomposition

Theorem 2.11 (Weighted Eulerian orientations [MWX25]; [GSS26, Theorems 1.1 and 2.17])**.** *For every finite algebraic EO signature set $\mathcal{G}$, if $\mathsf{EOTract}(\mathcal{G})$ holds, then* K-Holant$(\mathcal{G}) \in \mathsf{FP}$*; otherwise* K-Holant$(\mathcal{G})$ *is* $\#\mathsf{P}$*-hard. The polynomial-time side uses the LP algorithm of Guan, Shao, and Shi.*

Theorem 2.12 (Single-weight K-Holant [MWXZ25, full version, Theorem 31; conference version, Theorem 28])**.** *Let $\mathcal{G}$ be a finite preprocessed algebraic signature set in which every signature is supported on one Hamming layer. Then* K-Holant$(\mathcal{G})$ *is* $\#\mathsf{P}$*-hard unless either $\mathcal{G}$ is common one-sided with* $\mathsf{EOTract}(\mathcal{G}|_{\mathrm{EO}})$*, or both positive and negative charges occur and* $\mathsf{EOTract}(\mathcal{G}_{\to\mathrm{EO}})$*. Both exceptions are in* FP *by Theorem 2.11.*

Theorem 2.13 (Odd-arity Holant [MWXZ25, full version, Theorem 33; conference version, Theorem 30])**.** *Let $\mathcal{F}$ be a finite preprocessed algebraic signature set containing a nonzero odd-arity signature. Then* $\mathrm{Holant}(\mathcal{F}) \in \mathsf{FP}$ *if* $\mathsf{Tract}(\mathcal{F})$ *holds, and is* $\#\mathsf{P}$*-hard otherwise.*

Its first two alternatives were originally in $\mathsf{FP}^{\mathsf{NP}}$; the stated FP conclusion follows from [GSS26, Theorem 1.1 and Corollary 1.2].

Theorem 2.14 (Tensor-factor extraction [MWXZ25, full version, Theorem 32 and Remark 39; conference version, Theorem 29])**.** *Let Λ be a finite signature set and let $h = a \otimes b$ be a nonzero tensor product. Unless $\Lambda \cup \{h\}$ is common one-sided,*

$$\text{K-Holant}(\Lambda, a, b) \equiv_{\mathrm{T}} \text{K-Holant}(\Lambda, h).$$

In the common-one-sided exception, $\text{K-Holant}(\Lambda, h)$ *is in* FP *when* $\mathsf{EOTract}((\Lambda \cup \{h\})|_{\mathrm{EO}})$ *holds and is* $\#\mathsf{P}$*-hard otherwise, by charge conservation and Theorem 2.11.*

Theorem 2.15 (Quaternary eight-vertex boundary [CF23, Theorem 3.1])**.** *Let M be a quaternary signature satisfying*

$$\mathrm{supp}(M) \subseteq \{x \in \{0,1\}^4 : \mathrm{wt}(x) = 2\} \cup \{0000, 1111\}, \qquad M(0000)M(1111) \neq 0.$$

Then $\mathrm{Holant}(X \mid M)$ *is* $\#\mathsf{P}$*-hard unless there exist $\mathscr{C} \in \{\mathscr{A}, \mathscr{P}, \mathscr{L}\}$ and $U \in \mathrm{GL}_2(\overline{\mathbb{Q}})$ such that*

$$U^{\mathsf{T}} X U \in \mathscr{C}, \qquad U^{-1} M \in \mathscr{C}.$$

Every exception is in FP*. The source permits U over $\mathbb{C}$; Lemma 2.10 makes it algebraic.*

Theorem 2.16 (Even-degree counting CSP [CLX18, full version, Theorem 4.1])**.** *Let $\mathcal{G}$ be a finite algebraic Boolean signature set. Then* $\#\mathrm{CSP}^2(\mathcal{G})$ *is* $\#\mathsf{P}$*-hard unless one common alternative holds:*

$$\mathcal{G} \subseteq \mathscr{A}, \qquad \mathcal{G} \subseteq \mathscr{P}, \qquad \mathcal{G} \subseteq \mathscr{L}, \qquad \text{or} \qquad \mathcal{G} \subseteq \mathscr{A}_2^1.$$

Each exceptional case is computable in polynomial time.

The source assumes that both pinnings are freely available. In the present formulation, EQ_4 is available on the variable side of $\#\mathrm{CSP}^2$, and Corollary 3.5 supplies those pinnings. Any ordinary-Holant oracle instance arising in this use is returned to bipartite CSP form by merging adjacent even-arity equality vertices and subdividing every remaining $\mathcal{G}$–$\mathcal{G}$ edge with EQ_2. The class denoted $\mathscr{A}^\alpha$ there is our $\mathscr{A}_2^1$; the inverse-twist convention does not change the class because $T_2^2 = \mathrm{diag}(1, i)$ preserves $\mathscr{A}$.

Theorem 2.17 (Degree-multiple counting CSP [CFS20b, full version, Theorem 5.3; proof in Section 5.5])**.** *Let $q \geq 2$ be an integer and let $\mathcal{G}$ be a finite algebraic Boolean signature set. If* $\#\mathrm{CSP}^q(X, \mathcal{G})$ *is not* $\#\mathsf{P}$*-hard, then either*

$$\mathcal{G} \subseteq \mathscr{P} \qquad \text{or} \qquad \mathcal{G} \subseteq \mathscr{A}_q^r \quad \text{for one } r \in \{1, \ldots, q\}.$$

Both alternatives are computable in polynomial time.

2.13.5 Decidability and polynomial-time evaluation

Theorem 2.18 (Recognition and evaluation)**.** *For every finite algebraic signature set $\mathcal{F}$, the predicate* $\mathsf{Tract}(\mathcal{F})$*, with the preprocessing convention of Section 2.9, is decidable. If it holds, then* $\mathrm{Holant}(\mathcal{F}) \in \mathsf{FP}$.

Proof. This is [LM26a, Theorem 2.22]. □

2.14 Problem equivalence and normalization

2.14.1 Factor closure

For a preprocessed finite signature set Λ, choose an exact tensor-prime factorization of every member and let $\Pi(\Lambda)$ be the finite set of all prime factors, with their ordered port blocks. We call Λ *factor-closed* when $\Pi(\Lambda) \subseteq \Lambda$.

Lemma 2.19 (Factor closure). *If Λ is not common one-sided, then*

$$\text{K-Holant}(\Lambda) \equiv_{\mathrm{T}} \text{K-Holant}(\Lambda \cup \Pi(\Lambda)). \tag{11}$$

Accordingly, under the equivalence convention we replace Λ by this union, keep the same notation, and assume $\Pi(\Lambda) \subseteq \Lambda$. Repeated factorization terminates because every proper positive-arity factor has smaller arity than its parent.

Proof. This is [LM26a, Lemma 2.23]. □

For a factor-closed signature set Λ containing a tensor-prime signature of arity at least three, put

$$\nu(\Lambda) = \min\{\operatorname{arity}(h) : h \in \Lambda \text{ is tensor-prime and } \operatorname{arity}(h) \geq 3\}.$$

This parameter is used only after the branch in which every prime factor has arity at most two has been removed.

2.14.2 Downward inheritance and coordinate transport

Lemma 2.20 (Downward inheritance). *If $\mathcal{G} \subseteq \mathcal{G}'$ are preprocessed signature sets and $\mathsf{Tract}(K\mathcal{G}')$ holds, then $\mathsf{Tract}(K\mathcal{G})$ holds.*

Proof. This is [LM26a, Lemma 2.24]. □

Lemma 2.21 ($GO(X)$-invariance of tractability). *Let $\mathcal{G}$ be a finite preprocessed signature set and let $D \in GO(X)$. Then*

$$\mathsf{Tract}(K\mathcal{G}) \iff \mathsf{Tract}(K(D\mathcal{G})).$$

Proof. This is [LM26a, Lemma 2.25]. □

Lemma 2.22 (Coordinate transport along equivalent augmentations). *Let $\mathcal{G}_0$ be a finite preprocessed algebraic signature set, and let L_0 be a number field containing its entries, their conjugates, and every algebraic constant fixed before the construction. Consider a sequence (A_i, Λ_i), in which every Λ_i is a finite preprocessed algebraic signature set, beginning with*

$$A_0 = I, \qquad \Lambda_0 = \mathcal{G}_0,$$

in which each step is one of the following.

1. *An augmentation satisfying the following conditions: set $A_{i+1} = A_i$, choose a finite $\Lambda_{i+1} \supseteq \Lambda_i$ with*

   $$\text{K-Holant}(\Lambda_{i+1}) \leq_{\mathrm{T}} \text{K-Holant}(\Lambda_i),$$

   and require every newly adjoined positive-arity signature h to satisfy

   $$h \doteq A_i^{\otimes \operatorname{arity}(h)} h_0 \quad \textit{for a tensor } h_0 \textit{ over } L_0.$$

2. A global normalization: choose $D_i \in GO(X)$, and set

$$A_{i+1} = D_i A_i, \qquad \Lambda_{i+1} = D_i \Lambda_i,$$

recording the known nonzero scalar contributed by each transformed native edge as in Equation (2).

Then, at every index i,

$$A_i \mathcal{G}_0 \subseteq \Lambda_i, \qquad A_i^{\mathsf{T}} X A_i = \mu_i X \quad (\mu_i \in \overline{\mathbb{Q}}^{\times}), \qquad \text{K-Holant}(\Lambda_i) \equiv_{\mathrm{T}} \text{K-Holant}(\mathcal{G}_0), \tag{12}$$

and every positive-arity $h \in \Lambda_i$ satisfies

$$h \doteq A_i^{\otimes \operatorname{arity}(h)} h_0 \quad \textit{for a tensor } h_0 \textit{ over } L_0.$$

At an index where every nonzero directly realizable ordered binary is nonsingular and has finite projective transfer order, let $G(\Lambda_i)$ be the group of its projective transfer classes $[BX]$, and put

$$\overline{G}_i = [A_i]^{-1} G(\Lambda_i)[A_i].$$

Then $\overline{G}_i \le \mathrm{PGL}_2(L_0)$. On any consecutive sequence of such indices, a global normalization leaves $\overline{G}_i$ unchanged, whereas an augmentation can only enlarge it. Thus all groups compared in the finite-group argument lie in the single fixed field L_0; any generic interpolation that leaves this field is used only when it completes the complexity classification of the current case.

If $\mathsf{Tract}(K\Lambda_i)$ holds, then $\mathsf{Tract}(K\mathcal{G}_0)$ holds. Separately, if one matrix V is a common $\mathscr{C}$ K-presentation of Λ_i, where $\mathscr{C}$ is closed under nonzero scalar multiplication, then

$$W = A_i^{-1} V$$

is a common $\mathscr{C}$ K-presentation of $\mathcal{G}_0$. In particular, this applies to $\mathscr{A}, \mathscr{P}, \mathscr{L}$.

Proof. This is [LM26a, Lemma 2.26], applied inductively to the two displayed kinds of step. Restriction supplies the reverse reduction for an augmentation. A normalization conjugates the raw transfer group by $[D_i]$, which cancels after pullback by $[A_{i+1}] = [D_i A_i]$. □

2.15 Binary-equality and quaternary dichotomies

The following dichotomies are consequences of the results of Liu and Meng [LM26a].

Theorem 2.23 (Binary-equality K-Holant dichotomy)**.** *Let Λ be a finite algebraic signature set containing $I = \mathrm{EQ}_2$. If $\mathsf{Tract}(K\Lambda)$ holds, then* K-Holant$(\Lambda) \in \mathsf{FP}$*; otherwise* K-Holant$(\Lambda)$ *is* #P*-hard. The applicable side is decidable from exact algebraic input.*

Proof. [LM26a, Theorem 1.2] is stated with a reduced four-clause predicate. Since $I \in \Lambda$, its coordinate image contains $Z = K^{\otimes 2} I \in K\Lambda$; hence the binary-anchor lemma [LM26a, Lemma 2.13] identifies that reduced predicate with $\mathsf{Tract}(K\Lambda)$. The displayed formulation follows. □

An eight-vertex signature has arity four and vanishes on odd Hamming weight.

Theorem 2.24 (Tensor-prime quaternary K-Holant dichotomy)**.** *Let Λ be a finite algebraic signature set containing a tensor-prime eight-vertex signature M with $M(0000)M(1111) \neq 0$. If $\mathsf{Tract}(K\Lambda)$ holds, then* K-Holant$(\Lambda) \in \mathsf{FP}$*; otherwise* K-Holant$(\Lambda)$ *is* #P*-hard. The applicable side is decidable from exact algebraic input.*

Proof. This is [LM26a, Theorem 1.3]. $\square$

Corollary 2.25 (Eight-vertex K-Holant dichotomy)**.** *Let Γ be a finite algebraic signature set containing an eight-vertex signature M with $M(0000)M(1111) \neq 0$. If $\mathsf{Tract}(K\Gamma)$ holds, then* K-Holant$(\Gamma) \in \mathsf{FP}$*; otherwise* K-Holant(Γ) *is* $\#\mathsf{P}$*-hard.*

Proof. If $\mathsf{Tract}(K\Gamma)$ holds, Theorem 2.18 and Equation (4) give the polynomial-time conclusion. Assume $\neg\mathsf{Tract}(K\Gamma)$. If M is tensor-prime, Theorem 2.24 gives hardness.

Suppose instead that M is tensor-decomposable. Its nonzero endpoint values exclude a $1+3$ factorization: setting the one-port block to 1 and the other block to 000 would give a nonzero odd-weight entry. Thus every nontrivial factorization has type $2+2$. In such a factorization, the nonzero endpoints and the odd-weight zeros force each binary factor to be diagonal with both diagonal entries nonzero. Up to a port permutation and a nonzero scalar, one factor is therefore

$$E = \mathrm{diag}(p, q), \qquad p, q \in \overline{\mathbb{Q}}^{\times}.$$

Because M contains both charge signs, the common-one-sided exception in Lemma 2.19 is impossible. By the equivalence convention, we may therefore assume $E \in \Gamma$. Choose algebraic $a, b \neq 0$ with $a^2 p = b^2 q = 1$, and put $D = \mathrm{diag}(a, b)$. Then

$$D^{\mathsf{T}} X D = abX, \qquad DED^{\mathsf{T}} = I,$$

so $D \in GO(X)$ and $D\Gamma$ contains I. The global normalization and Lemma 2.21 preserve both complexity and the failure of Tract. Hence Theorem 2.23 gives hardness. Thus the tensor-prime and tensor-decomposable cases follow from the quaternary and binary-equality dichotomies, respectively. $\square$

3 Auxiliary Reductions

This section collects binary interpolation, reductions using one occurrence of an additional signature, and dichotomies for sets containing a nonzero unary or a higher equality. We also prove hardness when a set satisfying the stated charge and nontractability assumptions contains an EO signature outside $\mathcal{B}_{\mathrm{ARS}}$.

3.1 Binary Interpolation

Lemma 3.1 (Binary interpolation)**.** *Let Γ be a finite algebraic signature set whose supported doubled charges include both signs, let $0 \neq B \in \Gamma$ be ordered binary, and put $A = BX$.*

1. *If $B = u \otimes v$ is singular, where $u, v \neq 0$, then factor closure gives $u, v \in \Gamma$.*
2. *If A is nonsingular and has infinite projective order, there is a nonzero rank-one binary $H \in \Gamma$.*

Proof. This is [LM26a, Lemma 3.1]. $\square$

3.2 Single-occurrence Reductions and Nonzero Unaries

The single-occurrence argument used below was originally developed for an expanded version of [MWXZ25]. Since that version is not publicly available, we include the argument here for completeness.

The dichotomy for a set containing a nonzero unary will apply below; both point masses are needed only in the constructions that explicitly use them. Write

$$\Delta = \Delta_0 \otimes \Delta_1,$$

with the displayed order of its two ports. For a signature g, the superscripts $g^{=1}$ and $g^{\leq 1}$ mean that an oracle instance contains, respectively, exactly one or at most one g-vertex. They do not denote tensor powers.

Lemma 3.2 (Reduction with at most one Δ-vertex)**.** *For every fixed finite algebraic signature set* Γ*,*

$$\text{K-Holant}(\Gamma, \Delta^{\leq 1}) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

No EO, parity, charge, mountain, or tractability hypothesis is required.

Proof. This is [LM26a, Lemma 3.2]. □

The next elementary lemma isolates the only additional ingredient needed to turn one-occurrence access into unrestricted access to a new signature. It applies verbatim to ordinary Holant and to K-Holant; write H for either setting, with its native wires.

Lemma 3.3 (Single-occurrence replication)**.** *Let* Γ *be a fixed finite algebraic signature set, and let* g *and* h *be nonzero signatures of arities* r *and* s*, respectively. All connections below use the native binary edge:* EQ_2 *for ordinary Holant and* X *for K-Holant. Suppose that*

1. $\mathsf{H}(\Gamma, g^{=1}) \leq_{\mathrm{T}} \mathsf{H}(\Gamma)$*;*
2. *there is a fixed gadget* C *over* Γ *whose dangling edges are the three disjoint ordered blocks*

$$A = (a_1, \dots, a_r), \qquad P = (p_1, \dots, p_r), \qquad Y = (y_1, \dots, y_s).$$

Connecting one g*-vertex coordinatewise to* A *realizes*

$$\rho\, g(P) h(Y)$$

on P, Y*, for a known* $\rho \in \overline{\mathbb{Q}}^{\times}$*. Moreover, for consecutive copies* C_i, C_{i+1}*, the block* P_i *can be connected coordinatewise to* A_{i+1}*, with no other port identifications. Thus a chain of* k *copies, beginning with one* g*-vertex, contains exactly that one* g*-vertex and realizes*

$$\rho^k g(P_k) \prod_{i=1}^{k} h(Y_i),$$

where the blocks $Y_1, \dots, Y_k$ *are pairwise disjoint; and*

3. *there is a fixed instance* J *over* $\Gamma \cup \{g\}$*, containing exactly one* g*-vertex, whose exactly computable partition value* $Z(J)$ *is nonzero.*

Then

$$\mathsf{H}(\Gamma, h) \leq_{\mathrm{T}} \mathsf{H}(\Gamma).$$

Proof. This is [LM26a, Lemma 3.3]. An input with no h-vertex is already an instance over Γ. Let Ω be an input with $k \geq 1$ occurrences of h. Delete those vertices, retain their incident edges as k ordered boundary blocks, and attach the blocks $Y_1, \dots, Y_k$ of the displayed chain coordinatewise in the corresponding orders. Let J° be the network obtained by deleting the unique g-vertex of J, and join its boundary coordinatewise to P_k in the port order inherited from that removed vertex. Denote the resulting instance by Ω'. It contains exactly one g-vertex, and the two tensor identities give

$$Z(\Omega') = \rho^k Z(J) Z(\Omega).$$

Divide by the known nonzero scalar $\rho^k Z(J)$ and apply assumption 1. □

Lemma 3.4 (Dichotomy with a nonzero unary)**.** *Let $\mathcal{F}$ and Γ be finite algebraic signature sets, and let $u \neq 0$ be unary. If $u \in \mathcal{F}$, then* $\mathrm{Holant}(\mathcal{F})$ *is* #P*-hard unless* $\mathsf{Tract}(\mathcal{F})$ *holds. Likewise, if $u \in \Gamma$, then* K-Holant(Γ) *is* #P*-hard unless* $\mathsf{Tract}(K\Gamma)$ *holds.*

Proof. This is [LM26a, Lemma 3.4]. □

Corollary 3.5 (Higher equality supplies unaries)**.** *Let $d \geq 4$, and let Γ be a finite algebraic signature set in either the ordinary-Holant or K-Holant setting. If* $\mathrm{EQ}_d \in \Gamma$*, then $\Delta_0, \Delta_1 \in \Gamma$. Hence Lemma 3.4 applies.*

Proof. This is [LM26a, Corollary 3.6]. □

3.3 Hardness from an Eulerian-orientation Signature

For the results that invoke (REO1)–(REO2), fix a finite preprocessed algebraic set Γ of even-arity signatures in K-coordinates satisfying:

(REO1) Among the values $\chi_{\mathrm{arity}(f)}(x)$, with $f \in \Gamma$ and $x \in \mathrm{supp}(f)$, at least one is positive and at least one is negative;

(REO2) $\neg\mathsf{Tract}(K\Gamma)$.

All gadget connections in this subsection use the native edge X. Arrow-reversal symmetry (ARS) means $h(x) = \overline{h(\bar{x})}$ for every input x. An EO signature h has *projective arrow-reversal symmetry* (projective ARS) if, for some $c \in \overline{\mathbb{Q}}^{\times}$,

$$(c^{-1}h)(x) = \overline{(c^{-1}h)(\bar{x})} \qquad \text{for every } x \in \{0,1\}^{\mathrm{arity}(h)}.$$

Here $\bar{x}$ is bitwise complement and the overline on the coefficient is complex conjugation. An EO–ARS binary is both EO and ARS without rescaling. Let $\mathcal{B}_{\mathrm{ARS}}$ consist of the nonzero scalar multiples of products on arbitrary perfect matchings, with an orientation chosen for each matched pair, of binary EO–ARS signatures

$$(0, a, \overline{a}, 0), \qquad a \in \overline{\mathbb{Q}}^{\times}. \tag{13}$$

We write $\neq_4 = \delta_{0011} + \delta_{1100}$. Nonzero scalars and output-port permutations are suppressed according to the equivalence convention.

Lemma 3.6 (Substitution of fixed gadgets in a larger signature set)**.** *Let $\Lambda \subseteq \Gamma$ be finite algebraic signature sets in K-coordinates. Suppose that a finite source derivation starting from Λ adjoins a finite set S using only fixed direct X-gadgets. Then the same derivation is valid over Γ, and*

$$\text{K-Holant}(\Gamma \cup S) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

If, at some stage, a source subproblem K-Holant(Λ')*, with $\Lambda' \subseteq \Lambda \cup S$, is* #$\mathsf{P}$*-hard, then* K-Holant$(\Gamma)$ *is* #P*-hard.*

Proof. Each direct gadget is substituted independently at every occurrence of its output signature, leaving vertices labelled from $\Gamma \setminus \Lambda$ unchanged. Induction over the finite source derivation proves the displayed reduction. For the final claim, restriction and that reduction give

$$\text{K-Holant}(\Lambda') \le_{\mathrm{T}} \text{K-Holant}(\Gamma \cup S) \le_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

A problem-level Turing reduction quoted from a source is not, by itself, a direct-gadget derivation: its local signature-producing operations must first be identified. Tensor-factor extraction is not included in this statement; whenever it is used below, it is supplied separately by Theorem 2.14. □

Theorem 3.7 (Hardness from an EO signature outside $\mathcal{B}_{\mathrm{ARS}}$)**.** *Assume* (REO1)–(REO2)*. Let $h \in \Gamma$ be a nonzero positive-arity EO signature. Suppose there are no $\lambda \in \overline{\mathbb{Q}}^{\times}$, perfect matching P_h of its ports together with an orientation (i,j) of each matched pair, and coefficients $a_{ij} \in \overline{\mathbb{Q}}^{\times}$ such that*

$$h = \lambda \bigotimes_{(i,j)\in P_h} (0, a_{ij}, \overline{a_{ij}}, 0)_{ij}.$$

Then K-Holant(Γ) *is* #P*-hard.*

Lemma 3.8 (Complementary quaternaries supply a unary)**.** *Assume* (REO1)–(REO2)*, and suppose $D \in \Gamma$, where*

$$D = p\delta_s + q\delta_{\bar{s}}, \qquad s \in \{0,1\}^4, \quad \mathrm{wt}(s) = 2, \quad p, q \in \overline{\mathbb{Q}}^{\times},$$

then Γ contains a nonzero unary.

Proof. First reduce D to the unweighted complementary quaternary. After a port permutation take $s = 0101$, and connect the same last two port positions of two copies through native X-edges. The resulting ordered quaternary is

$$\begin{aligned} Q(x_1, x_2, y_1, y_2) &= \sum_{u,v\in\mathbb{F}_2} D(x_1, x_2, u, v) D(y_1, y_2, 1-u, 1-v) \\ &= pq(\delta_{0110} + \delta_{1001}). \end{aligned}$$

Thus, in output order (x_1, y_2, x_2, y_1), the output is the nonzero scalar multiple $pq \neq_4$; hence $\neq_4 \in \Gamma$.

Put $E = \neq_4$. Joining the fourth ports of $E(a_1, a_2, a_3, u)$ and $E(b_1, b_2, b_3, v)$ through $X(u, v)$ realizes, in output order $(a_1, a_2, b_3, a_3, b_1, b_2)$, the six-ary generalized disequality

$$\delta_{000111} + \delta_{111000}.$$

Attach the Δ_0-port of one Δ-vertex through X to a_1, and its Δ_1-port through X to a_3. This selects the word 111000; in remaining-port order (b_1, a_2, b_2, b_3), the output is $\delta_{0101} = \Delta \otimes \Delta$. In the notation of Lemma 3.3, delete the unique Δ-vertex to obtain a gadget with input block $A = (a_1, a_3)$ and output blocks $P = (b_1, a_2)$, $Y = (b_2, b_3)$. Connecting a Δ-vertex to A realizes $\Delta(P)\Delta(Y)$. Connect P of each copy to A of the next, in the displayed orders, with no other port identifications. A chain of k copies therefore uses exactly one Δ-vertex and realizes $\Delta(P_k) \prod_{i=1}^{k} \Delta(Y_i)$ on pairwise disjoint output blocks. Closing a single Δ-vertex by X has value one. The reductions in Lemmas 3.2 and 3.3 therefore give unrestricted access to Δ, so the equivalence convention places $\Delta \in \Gamma$.

Now use $\Delta = \Delta_0 \otimes \Delta_1$. By (REO1), the common-one-sided exception in Theorem 2.14 is impossible. Factor closure gives $\Delta_0, \Delta_1 \in \Gamma$. □

Lemma 3.9 (EO alternative)**.** *Assume* (REO1)–(REO2)*, and let* $0 \neq h \in \Gamma$ *be a positive-arity EO signature. Then either there is a* $c \in \overline{\mathbb{Q}}^{\times}$ *such that*

$$(c^{-1}h)(x) = \overline{(c^{-1}h)(\bar{x})} \qquad (x \in \{0,1\}^{\operatorname{arity}(h)}),$$

or K-Holant(Γ) *is* #P*-hard.*

Proof. We first settle the scalar field. Suppose an algebraic signature h_0 is ARS after multiplication by an arbitrary nonzero complex scalar. Then, for some $\rho \in \mathbb{C}^{\times}$,

$$h_0(x) = \rho\,\overline{h_0(\bar{x})}.$$

Choose x with $h_0(x) \neq 0$. The same identity implies $h_0(\bar{x}) \neq 0$, and hence

$$\rho = \frac{h_0(x)}{\overline{h_0(\bar{x})}} \in \overline{\mathbb{Q}}^{\times}, \qquad \rho\bar{\rho} = 1.$$

Set $c = 1 + \rho$ if $\rho \neq -1$, and set $c = i$ if $\rho = -1$. Then $c \in \overline{\mathbb{Q}}^{\times}$ and $c/\bar{c} = \rho$, so $c^{-1}h_0$ is ARS. Thus every source conclusion that ARS holds up to a nonzero scalar has the algebraic scalar required in the statement.

We next handle the binary case. Write $h = (0, a, b, 0)$. If $ab = 0$, then h is a nonzero singular binary, so Lemma 3.1 and factor extraction give a nonzero unary. By (REO2) and Lemma 3.4, this is hard. Suppose $ab \neq 0$, and put

$$A = hX = \operatorname{diag}(a, b).$$

If A has infinite projective order, Lemma 3.1 first gives a nonzero rank-one binary; applying its singular case and factor extraction again gives a unary and hence hardness. Otherwise a/b is a root of unity, so $|a| = |b|$. With $\rho = a/\bar{b}$, one has $h(x) = \rho\overline{h(\bar{x})}$ for all x, and the preceding scalar argument proves the required projective ARS conclusion.

Assume now that $\operatorname{arity}(h) \geq 4$. Let L be the number field generated over $\mathbb{Q}$ by all entries of Γ and their complex conjugates. It is a finite extension of $\mathbb{Q}$. Every output of a fixed direct X-gadget over signatures already generated from h has entries in L, and every nonzero ratio used to normalize such an output also lies in L.

Run the generating process of [MWX25, Definition 29] on the singleton source $\{h\}$, with X as its initial binary. Each self-loop and each binary path in that definition is a fixed direct X-gadget. We therefore adjoin its output to Γ by Lemma 3.6; zero outputs are discarded. If a generated nonzero binary is singular, the singular case of Lemma 3.1 gives unary factors. If it is nonsingular but its transfer has infinite projective order, the other case of that lemma gives a rank-one binary and then unary factors. Either event is hard by (REO2) and Lemma 3.4. Notice that (REO1) persists under these equivalent augmentations, so the factor-extraction exception never occurs.

Otherwise every generated nonzero binary is nonsingular and, up to a nonzero scalar, has the form

$$B_r = \begin{pmatrix} 0 & 1 \\ r & 0 \end{pmatrix}, \qquad r \in L^{\times}, \qquad B_r X = \operatorname{diag}(1, r),$$

where r is a root of unity. Let R_i be the parameter set represented by the source family $B_i(h)$ at stage i. The path clause in Definition 29 makes each nonempty R_i the subgroup generated by the root-of-unity parameters available at that stage; repeated self-loop descent [MWX25, Lemma 39] supplies nonemptiness. To make the monotonicity explicit, write the source's stage operation as Φ. Its self-loop and path construction is monotone under inclusion of its input parameter set.

Finite order and the positive-power path clause give $1 \in R_1$, so $R_0 = \{1\} \subseteq R_1$. Inductively, $R_i = \Phi(R_{i-1}) \subseteq \Phi(R_i) = R_{i+1}$. Equality at one stage forces equality thereafter by the same operation. Every R_i lies in the group of roots of unity of L, which is finite. Consequently this increasing chain stabilizes at some finite stage i_0. For a unique $k \geq 1$, writing $\mu_k = \{\omega : \omega^k = 1\}$, the stabilized source notation is therefore

$$B(h) = B_{i_0}(h) = \{B_\omega : \omega \in \mu_k\}.$$

Thus every gadget construction invoked below is finite, as required by Lemma 3.6.

If $k \geq 3$, [MWX25, Lemmas 40–42] show that h is ARS up to a nonzero scalar. The opening scalar argument completes this case.

Suppose $k = 2$, and write $Y = B_{-1} = (0, 1, -1, 0)$. Both X and Y are available in the source derivation, and Y is realized by a direct gadget. We use the proof of [MWX25, Lemma 47], not merely its displayed problem-level reduction. Its arity and support-distance inductions use only fixed X- or Y-self-loops. At arity four, the proofs of [MWX25, Lemmas 45–46], before invoking their general problem reduction, construct either an EO quaternary q for which the singleton problem is hard by [MWX25, Theorem 16], or a nonzero complementary quaternary

$$D = p\delta_s + q'\delta_{\bar{s}}, \qquad pq' \neq 0;$$

the two-copy subcase realizes $\neq_4$ itself. All these signatures arise by the fixed self-loops and two-copy gadgets just identified, so Lemma 3.6 adjoins them relative to Γ. The hard quaternary gives hardness by restriction. The complementary quaternary gives a unary by Lemma 3.8, and then hardness by Lemma 3.4. In the remaining case, the proof of Lemma 47 gives ARS up to a nonzero scalar, to which the opening argument applies.

Finally suppose $k = 1$. Follow the proof of [MWX25, Lemma 48], and let $C(h)$ denote its set of nonzero quaternary signatures obtained by successive X-self-loops. Every member of $C(h)$ is therefore available by Lemma 3.6. Moreover, every binary generated from such a signature is also generated from h. Hence its generated binary set is contained in $B(h) = \{X\}$, and Lemma 39 makes that set nonempty; it is therefore exactly $\{X\}$. If a member of $C(h)$ lies outside $\{X^{\otimes 2}, Y^{\otimes 2}\}$, the proof of Lemma 45 reaches exactly one of the hard or complementary quaternaries described in the preceding paragraph. If $Y^{\otimes 2} \in C(h)$, adjoin this signature and apply Theorem 2.14. Assumption (REO1) excludes its exceptional case, so Y is available; the proof of Lemma 47 then applies because $B(h) \subseteq \{X, Y\}$ and both binaries are available. This replaces, rather than imports, the problem-level tensor-power reduction in [MWX25, Lemma 38]. The case $C(h) = \varnothing$ is impossible for nonzero h by Lemma 39. In the remaining case $C(h) = \{X^{\otimes 2}\}$, the reorganizing argument in Lemma 48, together with [MWX25, Lemma 43], makes h a scalar multiple of a real domain-symmetric signature, hence projectively ARS. Its real and imaginary parts are used only as vectors in that argument; they are not signatures adjoined to Γ. The opening scalar argument again makes the scalar algebraic. $\square$

Lemma 3.10 (ARS alternative)**.** *Assume* (REO1)–(REO2)*, and let* $0 \neq h \in \Gamma$ *be a positive-arity EO signature for which there is a* $c \in \overline{\mathbb{Q}}^\times$ *satisfying*

$$(c^{-1}h)(x) = \overline{(c^{-1}h)(\bar{x})} \qquad (x \in \{0,1\}^{\operatorname{arity}(h)}).$$

If there are no $\lambda \in \overline{\mathbb{Q}}^\times$*, perfect matching* P_h *together with an orientation* (i, j) *of each matched pair, and* $a_{ij} \in \overline{\mathbb{Q}}^\times$ *such that*

$$h = \lambda \bigotimes_{(i,j) \in P_h} (0, a_{ij}, \overline{a_{ij}}, 0)_{ij},$$

then K-Holant(Γ) *is* #P*-hard.*

Proof. Put $h' = c^{-1}h$, which is algebraic, EO, and ARS. By unique factorization and the ARS factorization lemma [CFS20a, Lemmas 2.13–2.14], choose a tensor-prime factorization of h', up to an overall scalar, into factors $g_1, \ldots, g_m$ on disjoint port blocks, each satisfying ARS exactly. We now choose representatives without losing either algebraicity or ARS. For each j, choose $\alpha_j \in \operatorname{supp}(g_j)$, let $\alpha = \alpha_1 \cdots \alpha_m$, and define on the port block of g_j

$$q_j(z) = \frac{h'(\alpha_1, \ldots, \alpha_{j-1}, z, \alpha_{j+1}, \ldots, \alpha_m)}{h'(\alpha)}.$$

Then $q_j = g_j/g_j(\alpha_j)$. Hence every q_j is algebraic, EO, tensor-prime, and projectively ARS, and

$$h' = h'(\alpha) \bigotimes_{j=1}^{m} q_j.$$

For each q_j, the algebraic-scalar argument in the preceding lemma gives $c_j \in \overline{\mathbb{Q}}^{\times}$ such that $p_j = c_j^{-1} q_j$ satisfies ARS exactly. Thus

$$h' = h'(\alpha)\Big(\prod_{j=1}^{m} c_j\Big) \bigotimes_{j=1}^{m} p_j,$$

with an algebraic overall scalar and algebraic EO–ARS tensor-prime factors p_j. If every tensor-prime factor were binary, the prohibited displayed decomposition of h would hold. Thus there is a nonbinary tensor-prime EO–ARS factor $f \notin \mathcal{B}_{\mathrm{ARS}}$. Iterated use of Theorem 2.14 adjoins f to the current set; by (REO1), its common-one-sided exception is impossible. Whenever the cited descent subsequently extracts another factor, we repeat the displayed ratio construction and rescale by a nonzero algebraic scalar to obtain an ARS representative.

Apply the proof of [CFS20a, Lemmas 3.2 and 4.1–4.7; Theorem 4.8] to the singleton EO–ARS source generated from f, not to the possibly non-EO ambient language Γ. The signature-producing operations in that proof are local: the tensor connections in Lemma 3.2 and the mergings in Lemma 4.6 are fixed direct X-gadgets, while the selected-port attachments of the fixed binary $iY = (0, i, -i, 0)$ in Lemmas 4.5–4.6 are direct gadgets after that binary factor has been extracted. These direct gadgets can be substituted over Γ by Lemma 3.6; every factor-extraction step is supplied separately by Theorem 2.14, whose exception is again excluded by (REO1). Sections 3–4 of the cited proof use no polynomial interpolation.

Concretely, Lemma 3.2 produces a direct EO quaternary whose singleton problem is hard, directly realizes $\neq_4$, or enters its mutual-orthogonality case. Lemmas 4.1–4.4 supply the divisibility propagation used in Lemma 4.6. Lemmas 4.5–4.6 descend through signatures outside $\mathcal{B}_{\mathrm{ARS}}$, and Lemma 4.7 rules out mutual orthogonality in the base arities $4, 6, 8$. Hence the descent adjoins either a hard quaternary or $\neq_4$ by the reductions just described. The first gives hardness by restriction; the second gives a unary by Lemma 3.8 and then hardness by Lemma 3.4.

Finally, the coordinate restrictions called *pinnings* in the proof of [CFS20a, Lemma 2.14] are coordinate slices used to establish tensor-divisibility identities; they are not signatures obtained from direct pinning gadgets. The induction above starts from $h \in \Gamma$, or a factor adjoined by Theorem 2.14; a coordinate restriction of a larger tensor would not suffice, since connecting copies of the larger tensor could introduce additional terms. $\square$

Proof of Theorem 3.7. Apply Lemma 3.9. Either K-Holant(Γ) is #P-hard, or h is projectively ARS. In the latter case, $h \notin \mathcal{B}_{\mathrm{ARS}}$, so Lemma 3.10 gives hardness. $\square$

Remark 3.11 (Scope of the literature inputs). The ARS step uses [CFS20a, Theorem 4.8], not the Section 5 theorem for a wholly EO–ARS signature set. No theorem requiring the ambient set itself to be EO or EO–ARS is invoked: the source results are applied only to the indicated EO sublanguages, their direct gadgets are substituted using Lemma 3.6, and factor extraction is handled separately by Theorem 2.14.

3.4 Equalities and Generalized Equalities

Remark 3.12 (Simultaneous transformation of the edge and signature). The same matrix U in Lemma 3.13 must transform both the native edge to a nonzero multiple of EQ_2 and the quaternary signature to a nonzero multiple of EQ_4. Transforming the quaternary alone does not ensure that the transformed problem has equality edges.

Lemma 3.13 (Dichotomy from a transformed quaternary equality)**.** *Let $\mathcal{G}$ be a finite algebraic signature set containing an ordered quaternary signature g. Suppose there are*

$$U \in \mathrm{GL}_2(\overline{\mathbb{Q}}), \qquad \kappa, \lambda \in \overline{\mathbb{Q}}^{\times},$$

such that

$$(U^{-1})^{\mathsf{T}} X U^{-1} = \kappa \mathrm{EQ}_2, \qquad U^{\otimes 4} g = \lambda \mathrm{EQ}_4.$$

Then either K-Holant($\mathcal{G}$) *is* #P*-hard or* $\mathsf{Tract}(K\mathcal{G})$ *holds.*

Proof. This is [LM26a, Lemma 3.8]. ☐

Lemma 3.14 (Dichotomy with a generalized equality)**.** *Let $\mathcal{G}$ be a finite algebraic signature set containing*

$$G_q = c_0 \delta_{\mathbf{0}^q} + c_1 \delta_{\mathbf{1}^q}, \qquad c_0, c_1 \in \overline{\mathbb{Q}}^{\times}, \qquad q \geq 4 \text{ an integer.}$$

Then either K-Holant($\mathcal{G}$) *is* #P*-hard or* $\mathsf{Tract}(K\mathcal{G})$ *holds.*

Proof. This is [LM26a, Lemma 3.9]. ☐

Lemma 3.15 (Dichotomy from a two-point Walsh factor)**.** *Let Γ be a finite algebraic signature set whose supported doubled charges include both signs. Suppose a finite equivalent enlargement of Γ contains a nonzero signature with a tensor factor h of even arity $2k \geq 4$ with*

$$\widetilde{h} = (P^{-1})^{\otimes 2k} h = c_0 \delta_a + c_1 \delta_{\bar{a}}, \qquad a \in \mathbb{F}_2^{2k}, \quad c_0 c_1 \neq 0.$$

Then K-Holant(Γ) *is* #P*-hard or* $\mathsf{Tract}(K\Gamma)$ *holds. If c_1/c_0 is a root of unity, a finite equivalent enlargement contains a quaternary g satisfying the two transformation conditions of Lemma 3.13 with $U = K$. No reflection or prescribed group of binary transfers is required.*

Proof. Adjoin the specified signature and then its factor h using Equation (5) and Theorem 2.14; the supported charges of both signs exclude the common-one-sided exception. All subsequent gadgets are native K-gadgets. In Walsh coordinates an edge is $P^{\mathsf{T}} X P = 2Z$.

Choose two positions at which a has the same bit. Interchanging the names of $a, \bar{a}$, if necessary, makes that bit zero; this merely renames the two terms, and uses no bit-flip gadget. Take two copies of h, leave these two positions exposed in each copy, and join all other corresponding positions between the copies. Write J for the joined positions. The diagonal edges allow only the two equally oriented pairs of support words. Since $|J| = 2k - 2$ is even, their signs agree. Thus the four-port gadget output q has

$$\widetilde{q} = \lambda(\delta_{0000} + \rho \delta_{1111}), \quad \rho = (c_1/c_0)^2, \quad \lambda = 2^{2k-2}(-1)^{\mathrm{wt}(a|_J)} c_0^2 \neq 0.$$

Closing two ports of q gives a binary B with

$$\widetilde{B} = 2\lambda\,\mathrm{diag}(1, -\rho), \qquad P^{-1}(BX)P = 4\lambda\,\mathrm{diag}(1, \rho).$$

If ρ has infinite multiplicative order, Lemma 3.1 supplies a rank-one binary. Factor extraction supplies a nonzero unary, and Lemma 3.4 gives the conclusion.

Otherwise let ℓ be the order of ρ. A chain of $\ell - 1$ copies of B has projective Walsh transfer $\mathrm{diag}(1, \rho^{-1})$; when $\ell = 1$ use the empty chain. Attaching this chain to one port of q gives a gadget output g with $\widetilde{g} = c(\delta_{0000} + \delta_{1111})$, $c \neq 0$. Writing $E_4 = \mathbf{1}_{\mathrm{even}}$, we have

$$g = 2cE_4, \qquad K^{\otimes 4} g = 4c\mathrm{EQ}_4, \qquad (K^{-1})^{\mathsf{T}} X K^{-1} = I.$$

These are the conditions of Lemma 3.13 with $U = K$, so that lemma applies. The equivalent enlargements transfer hardness back to Γ; tractability is inherited by its subset. $\square$

4 Obtaining a Mountain

For an even-arity signature set with supported half-charges of both signs, we can adjoin either both point masses or a mountain, as defined in Definition 2.3, whose height is the gcd of the nonzero supported half-charges.

Lemma 4.1 (Producing the gcd-height signature or both point masses)**.** *Let $\mathcal{G}$ be a finite preprocessed algebraic signature set of even-arity signatures. Assume that among the supported values*

$$\mathrm{wt}(x) - \frac{\mathrm{arity}(f)}{2}, \qquad f \in \mathcal{G}, \quad f(x) \neq 0,$$

there is at least one positive value and at least one negative value. Put

$$d = \gcd\left\{ \left| \mathrm{wt}(x) - \frac{\mathrm{arity}(f)}{2} \right| : f \in \mathcal{G},\ f(x) \neq 0,\ \mathrm{wt}(x) \neq \frac{\mathrm{arity}(f)}{2} \right\}.$$

Then $d \geq 1$, every supported difference $\mathrm{wt}(x) - \mathrm{arity}(f)/2$ lies in $d\mathbb{Z}$, and at least one of the following reductions holds.

1. *There is an arity-$2d$ signature M such that*

$$\text{K-Holant}(\mathcal{G}, M) \leq_{\mathrm{T}} \text{K-Holant}(\mathcal{G}),$$

$$M(\mathbf{0}^{2d})M(\mathbf{1}^{2d}) \neq 0, \qquad M(x) = 0 \quad \textit{if } \mathrm{wt}(x) \notin \{0, d, 2d\}.$$

2. *Both point masses are accessible:*

$$\text{K-Holant}(\mathcal{G}, \Delta_0, \Delta_1) \leq_{\mathrm{T}} \text{K-Holant}(\mathcal{G}).$$

In the second outcome, if $\neg\mathsf{Tract}(K\mathcal{G})$, then Lemmas 2.20 and 3.4 and the displayed reduction imply that K-Holant$(\mathcal{G})$ *is* #P*-hard.*

Proof. Let

$$C = \{\chi_{\mathrm{arity}(f)}(x) : f \in \mathcal{G},\ f(x) \neq 0,\ \chi_{\mathrm{arity}(f)}(x) \neq 0\}$$

be the nonzero supported half-charges. The gcd assertion is immediate. The next argument expresses d and $-d$ as sums of elements of C with nonnegative multiplicities. Tensor products add half-charges, but their factor multiplicities cannot be negative. Because C contains both signs, take every pair $a, -b \in C$ with $a, b > 0$. If $g = \gcd(a, b)$, assigning multiplicity b/g to a and a/g to $-b$ gives total half-charge zero. Summing these assignments over all positive–negative pairs produces a vector $v \in \mathbb{Z}_{>0}^C$ satisfying

$$\sum_{c \in C} v_c c = 0.$$

Bézout's identity, iterated over the finite set C, supplies an integer vector $u \in \mathbb{Z}^C$ with $\sum_{c \in C} u_c c = d$ [Apo76, Section 1.3, Theorem 1.2]. The entries of u need not be nonnegative, but every entry of v is positive. Consequently, for all sufficiently large k,

$$n^+ = u + kv, \qquad n^- = -u + kv$$

are nonnegative integer vectors, while their total half-charges remain d and $-d$, respectively. Thus n^+ and n^- are the required multiplicity lists.

Choose a supported word of half-charge c for each $c \in C$. The two tensor products prescribed by n^+, n^- contain supported words of half-charges $d, -d$, and all their supported half-charges remain divisible by d. Applying the endpoint-preserving self-loop lemma [MWXZ25, full version, Lemma 40; conference version, Lemma 37] and the equivalence convention gives $P, N \in \mathcal{G}$ of arity $2d$ with $P(\mathbf{1}^{2d})N(\mathbf{0}^{2d}) \neq 0$. Self-loops preserve half-charge, so

$$\operatorname{supp}(P) \cup \operatorname{supp}(N) \subseteq HW_{2d}^{=} \cup \{\mathbf{0}^{2d}, \mathbf{1}^{2d}\}.$$

If either P or N has both endpoints, take that signature as the required M. We may therefore assume

$$P(\mathbf{0}^{2d}) = 0, \qquad N(\mathbf{1}^{2d}) = 0;$$

and their remaining entries are central.

If both central parts are nonzero, permute chosen central words to $(\mathbf{0}^d, \mathbf{1}^d)$ in P and $(\mathbf{1}^d, \mathbf{0}^d)$ in N, expose the first d ports of each signature, and join their last d ports coordinatewise. The resulting arity-$2d$ signature Q has unique endpoint extensions and

$$Q(\mathbf{0}^{2d}) = P(\mathbf{0}^d, \mathbf{1}^d)N(\mathbf{0}^{2d}), \qquad Q(\mathbf{1}^{2d}) = P(\mathbf{1}^{2d})N(\mathbf{1}^d, \mathbf{0}^d),$$

both nonzero. If only P has a central part, write $N = n\delta_{\mathbf{0}^{2d}}$ for $n \in \overline{\mathbb{Q}}^{\times}$, permute a supported central word of P to $(\mathbf{0}^d, \mathbf{1}^d)$, and join the two d-port blocks of N to the last d ports of two copies of P, leaving the first d ports of both copies exposed. Again the endpoint extensions are unique, with values $nP(\mathbf{0}^d, \mathbf{1}^d)^2$ and $nP(\mathbf{1}^{2d})^2$. (The opposite one-central case is symmetric.) Half-charge divisibility restricts either output to half-charges $-d, 0, d$; take that output as M.

If both central parts vanish, then $P = p\Delta_1^{\otimes 2d}$ and $N = n\Delta_0^{\otimes 2d}$ for $p, n \in \overline{\mathbb{Q}}^{\times}$. Since the hypothesis supplies both positive and negative supported differences, the common-one-sided exception to the tensor-factor theorem cannot occur. Factor closure therefore gives $\Delta_0, \Delta_1 \in \mathcal{G}$, proving the unary outcome. It is covered by Theorems 2.13 and 2.14. □

Remark 4.2. The factor alternative is necessary. For $\{\mathrm{EXACT}_3^4, \mathrm{EXACT}_1^4\}$, where $\mathrm{EXACT}_j^4(x) = \mathbf{1}[\mathrm{wt}(x) = j]$, conservation of integer half-charge places every fixed gadget on a single integer-half-charge layer. Thus no gadget produces a mountain, and Lemma 4.1 gives both point masses by factor extraction instead.

5 Mountain Height One

A height-one mountain is an ordered binary signature whose two diagonal entries are nonzero. We prove the dichotomy for any finite algebraic signature set containing one. We use the equivalent-signature-set convention throughout: a signature supplied by a gadget, interpolation, or an applicable factor reduction is adjoined to the current set. When direct realization matters, we specify the wiring, port order, and nonzero scalar.

A singular binary, or a nonsingular binary whose transfer has infinite projective order, is handled by Lemma 3.1 and the odd-arity theorem. Otherwise Lemma 5.3 makes the group of all directly realizable binary transfers finite, and Theorem 5.4 gives a nonsingular weighted equality or an odd cyclic or odd dihedral group. A weighted equality can be normalized to $I = \mathrm{EQ}_2$ by one global diagonal transformation in $GO(X)$, so Theorem 2.23 applies. We treat the odd-dihedral case first, since its contraction arguments are also used in the odd-cyclic case.

The distinguished binary remains the transported original mountain. A quaternary mountain satisfies Corollary 2.25. Generalized equalities and the signatures satisfying Lemma 3.13 are handled by Lemmas 3.13 and 3.14; other tensor factors are adjoined using the factor-reduction equivalence.

Theorem 5.1 (Dichotomy with a height-one mountain)**.** *Let Γ be a finite algebraic signature set containing a height-one mountain. Then*

$$\mathsf{Tract}(K\Gamma) \qquad \textit{or} \qquad \text{K-Holant}(\Gamma) \ \textit{is}\ \#\mathsf{P}\textit{-hard}.$$

The remainder proves the theorem using the current-set convention. Only the global changes of basis and the changes in $G(\Gamma)$ remain explicit.

For an ordered binary signature B, recall its transfer

$$T_B = BX.$$

Joining B and C in series through one X-edge realizes BXC, and therefore

$$T_{BXC} = T_B T_C.$$

Reversing the ports gives the map defined in Section 2.8:

$$T^{\#} := XT^{\mathsf{T}}X; \qquad T^{\#} = (\det T)ZT^{-1}Z \quad \text{if } T \text{ is nonsingular.} \tag{14}$$

Remark 5.2 (Relation to binary-signature classifications). The trichotomy above is the K-Holant form of the binary classification used by Cai–Lu–Xia [CLX09, CLX18]: singular matrices, and nonsingular matrices of infinite or finite projective order. The four ordered Bell signatures $\mathcal{B} = \{B_{0,0}, B_{0,1}, B_{1,0}, B_{1,1}\}$ represent the projective Bell classes of Shao–Cai [SC20]. Their holographic matrix, denoted Z in that paper, is our K; throughout this paper, $Z = \left(\begin{smallmatrix}1 & 0\\ 0 & -1\end{smallmatrix}\right)$, as in Equation (14). Beyond the binary classification, we need the complete gadget transfer group and uniform control of its higher-arity contractions.

Lemma 5.3 (Bounded torsion)**.** *Fix a number field L, and let G be the set of projective transfers of all directly realizable nonzero nonsingular ordered binary gadgets over a fixed signature set with entries in L. If $G \neq \varnothing$ and every element of G has finite projective order, then G is a finite subgroup of* $\mathrm{PGL}_2(L)$*.*

Proof. This is [LM26a, Lemma 5.2]. ∎

For a current signature set $\mathcal{G}$, suppose that every nonzero directly realizable ordered binary B is nonsingular and that $[BX]$ has finite order. Series composition and Lemma 5.3 then give the finite group

$$G(\mathcal{G}) = \{[BX] : B \neq 0 \text{ is an ordered binary directly realizable over } \mathcal{G}\}. \tag{15}$$

This definition includes all such binary gadgets, including those not yet adjoined to $\mathcal{G}$. We use $G(\mathcal{G})$ only under the stated nonsingularity and finite-order hypothesis; a singular binary has no projective transfer class. The decision procedure in Theorem 2.18 does not enumerate this group.

Theorem 5.4 (Weighted equality or an odd cyclic or dihedral group)**.** *Let $\mathcal{G}$ be a finite algebraic signature set containing the height-one mountain*

$$B = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad ad \neq 0,$$

and suppose every nonzero directly realizable binary over $\mathcal{G}$ is nonsingular and its transfer has finite projective order. Then one of the following holds.

1. *A nonsingular weighted equality $E = \operatorname{diag}(p, q)$, with $pq \neq 0$, belongs to the current signature set.*

2. *There exist an odd $N \geq 3$ and $D \in GO(X)$ such that the legal global normalization $\mathcal{G} \mapsto D\mathcal{G}$ puts $G(D\mathcal{G})$ in exactly one of the forms*

$$\{[R_s] : s \in \mu_N\}, \qquad \{[R_s], [ZR_s] : s \in \mu_N\}, \qquad R_s = \frac{1}{2}\begin{pmatrix} 1+s & 1-s \\ 1-s & 1+s \end{pmatrix}.$$

   Every displayed projective class has a fixed directly realizable ordered-binary representative over $D\mathcal{G}$.

Proof. Put $G = G(\mathcal{G})$ and

$$H = \langle G, [Z] \rangle.$$

All gadget entries lie in the number field generated by the entries of $\mathcal{G}$, so Lemma 5.3 makes G finite. If $[T] \in G$, port reversal and Equation (14) give $[ZT^{-1}Z] \in G$. Since inversion permutes G, it follows that $[Z]G[Z] = G$. Consequently

$$G \lhd H, \qquad [H : G] \leq 2,$$

and $H = G$ or $H = G \sqcup [Z]G$, so H is finite.

For the distinguished mountain,

$$T = BX = \begin{pmatrix} b & a \\ d & c \end{pmatrix}$$

has class $[T] \in G$. Since $a, d \neq 0$, both off-diagonal entries of T are nonzero. We first record a weighted equality. Suppose $[S] \in G$ has an anti-diagonal representative

$$S = \begin{pmatrix} 0 & p \\ q & 0 \end{pmatrix}, \qquad pq \neq 0.$$

By the definition of G, there is a directly realizable ordered binary B_S and a scalar $\rho \neq 0$ such that $B_S X = \rho S$. Hence

$$B_S = \rho S X = \rho \operatorname{diag}(p, q),$$

which is the nonsingular weighted equality in outcome 1.

We now use Proposition 2.1 and Lemma 2.2. Their use here is only group-theoretic: the distinguished class is the fixed $[Z]$, and no conjugating matrix from the abstract classification is applied to the signature set.

If H is cyclic, it is abelian, so $[T]$ commutes with $[Z]$. The projective centralizer of $[Z]$ consists of the diagonal and anti-diagonal classes. The two nonzero off-diagonal entries of T exclude the diagonal case; thus $[T] \in G$ is anti-diagonal and gives the weighted equality.

Suppose $H \cong A_4$. Since A_4 has no index-two subgroup, the index bound forces $G = H$. The centralizer of $[Z]$ in H is a V_4, so it contains an involution $u \neq 1, [Z]$. The class u belongs to G and commutes with $[Z]$. Because $[Z]$ is the only nonidentity diagonal involution, u is anti-diagonal. The same argument applies when $H \cong A_5$: simplicity gives $G = H$, and the centralizer of an involution is again V_4.

Suppose $H \cong S_4$. If $G = H$ and $[Z]$ corresponds to a transposition, choose the disjoint transposition; if $[Z]$ corresponds to a double transposition, choose either of the other two double transpositions. In both cases the chosen involution lies in G, commutes with $[Z]$, and differs from it, so it is anti-diagonal. If $G \neq H$, then the unique index-two subgroup gives $G = A_4$. Here $[Z] \notin G$ is an involution outside A_4, hence a transposition. If z' is the disjoint transposition, then $u = [Z]z'$ is a double transposition in $A_4 = G$; it commutes with $[Z]$, differs from it, and is therefore anti-diagonal.

It remains to consider

$$H = D_{2m} = \left\langle r, z : r^m = z^2 = 1,\ zrz = r^{-1} \right\rangle,$$

where the distinguished reflection is denoted by $z = [Z]$ when $[Z]$ is a reflection. If m is even and $[Z] = r^{m/2}$ is the central half-turn, then $[T] \in G$ commutes with $[Z]$; its two nonzero off-diagonal entries again make it anti-diagonal. Suppose instead that $[Z]$ is a reflection and m is even. The two involutions

$$r^{m/2}, \qquad [Z]r^{m/2}$$

commute with $[Z]$ and differ from it. If $G = H$, both belong to G. If $[H : G] = 2$, then $[Z] \notin G$, and the two displayed elements lie in opposite cosets of G; hence exactly one belongs to G. In either event G contains an involution different from $[Z]$ that commutes with $[Z]$, and this involution is anti-diagonal.

The sole unresolved case is therefore

$$H \cong D_{2N}, \qquad N \geq 3 \text{ odd}, \qquad [Z] \text{ a reflection},$$

with $G = C_N$ or $G = D_{2N}$. Indeed, for odd N the rotation subgroup is the unique index-two subgroup; thus a proper G is the rotation subgroup, while otherwise $G = H$.

Choose a rotation class $[R]$ of exact order N, and write

$$R = \begin{pmatrix} u & v \\ w & x \end{pmatrix}.$$

The dihedral relation gives $[ZRZ] = [R^{-1}]$, while

$$ZRZ = \begin{pmatrix} u & -v \\ -w & x \end{pmatrix}, \qquad R^{-1} \doteq \begin{pmatrix} x & -v \\ -w & u \end{pmatrix}.$$

The entries v, w cannot both vanish: otherwise $[R] = [R^{-1}]$, so $[R]^2 = 1$, contrary to the odd order $N \geq 3$. Since at least one of v, w is nonzero, comparison in the preceding projective equality forces the projective scalar to be 1, and hence $u = x$. Exactly one of v, w cannot vanish. In that case nonsingularity gives $u = x \neq 0$, and a projective rescaling makes R a nonidentity unipotent matrix; its powers are pairwise distinct in characteristic zero. Therefore

$$R = \begin{pmatrix} \alpha & \beta \\ \gamma & \alpha \end{pmatrix}, \qquad \beta\gamma \neq 0.$$

Choose $q \in \overline{\mathbb{Q}}^{\times}$ with $q^2 = \gamma/\beta$, and put

$$D = \operatorname{diag}(q, 1).$$

Then

$$D^{\mathsf{T}} X D = qX,$$

so $D \in GO(X)$; moreover $DZD^{-1} = Z$. With $\delta = q\beta = q^{-1}\gamma$,

$$DRD^{-1} = \begin{pmatrix} \alpha & \delta \\ \delta & \alpha \end{pmatrix}.$$

The matrix P is used only for the following analytic diagonalization:

$$P^{-1}(DRD^{-1})P = \operatorname{diag}(\alpha + \delta, \alpha - \delta).$$

Both eigenvalues are nonzero because R is nonsingular. Put

$$s_0 = \frac{\alpha - \delta}{\alpha + \delta}.$$

The exact projective order of $[R]$ implies that s_0 is a primitive N-th root of unity, and

$$DRD^{-1} = (\alpha + \delta) R_{s_0}, \qquad R_s = \frac{1}{2} \begin{pmatrix} 1+s & 1-s \\ 1-s & 1+s \end{pmatrix}.$$

Since $R_s R_t = R_{st}$, the normalized rotation classes are precisely

$$\{[R_s] : s \in \mu_N\}.$$

The diagonal matrix D fixes $[Z]$. Thus, when $G = C_N$, these are all normalized classes, and, when $G = D_{2N}$, the remaining classes are precisely

$$\{[ZR_s] : s \in \mu_N\}.$$

Finally, the normalization gives equality of complete groups, not merely an inclusion. If B_0 is a directly realizable ordered binary over $\mathcal{G}$, legal global transport gives a directly realizable boundary tensor B_0^D over $D\mathcal{G}$ with

$$B_0^D \doteq DB_0D^{\mathsf{T}}, \qquad (DB_0D^{\mathsf{T}})X = qD(B_0X)D^{-1}.$$

Hence $T_{B_0^D} \doteq DT_{B_0}D^{-1}$. Applying the same argument with $D^{-1} \in GO(X)$ gives the converse, so

$$G(D\mathcal{G}) = [D]G(\mathcal{G})[D]^{-1}.$$

Every displayed class therefore has a fixed directly realizable ordered-binary representative over $D\mathcal{G}$, with the nonzero scalar from the global transport recorded and restored. This is a single legal $GO(X)$-normalization; neither P nor an arbitrary conjugator from the finite-subgroup classification is used as a gadget or a global normalization. □

5.1 Odd cyclic and dihedral transfer groups

Definition 5.5 (Equivalent updates with strict progress)**.** Let (A, Γ) satisfy Lemma 2.22 for the initial set Γ_0 and its fixed number field L_0. An update to (A^+, Γ^+) must satisfy the same invariant and use only the augmentations and global transformations allowed there. In particular, Γ^+ is finite, preprocessed, algebraic, and factor-closed, and

$$A^+\Gamma_0 \subseteq \Gamma^+, \qquad (A^+)^{\mathsf{T}} X A^+ = \mu_{A^+} X \quad (\mu_{A^+} \neq 0),$$
$$\text{K-Holant}(\Gamma^+) \equiv_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

The transformed height-one mountain and every old tensor-prime signature remain members up to the same invertible change of basis and their recorded nonzero scalars. Every nonzero directly realizable ordered binary over both sets is nonsingular with finite projective transfer order. Put

$$H = [A]^{-1} G(\Gamma)[A], \qquad H^+ = [A^+]^{-1} G(\Gamma^+)[A^+].$$

Both groups are subgroups of $\mathrm{PGL}_2(L_0)$. When both values of ν are defined, the update is required to satisfy

$$H \subseteq H^+, \qquad \nu(\Gamma^+) \leq \nu(\Gamma), \qquad H \neq H^+ \text{ or } \nu(\Gamma^+) \neq \nu(\Gamma).$$

Thus at least one comparison is strict. If no nonbinary tensor-prime signature remains, the arity-at-most-two alternative gives tractability and no further update is made.

Lemma 5.6 (Termination of the strict updates)**.** *Fix the number field L_0 in Lemma 2.22. There is no infinite sequence*

$$(H_0, \nu_0), (H_1, \nu_1), \ldots$$

such that every H_i is a finite subgroup of $\mathrm{PGL}_2(L_0)$, every ν_i is an integer at least three, and every update satisfies

$$H_i \subseteq H_{i+1}, \qquad \nu_{i+1} \leq \nu_i,$$

with at least one of these two inequalities strict. In the height-one construction these quantities are

$$H_i = [A_i]^{-1} G(\Gamma_i)[A_i], \qquad \nu_i = \min\{\text{arity}(f) : f \in \Gamma_i \textit{ is tensor-prime and } \text{arity}(f) \geq 3\}.$$

If the second set is empty, the arity-at-most-two case is tractable and no further strict update is made. This assertion concerns only the transitions represented by the displayed pairs; operations preserving both H_i and ν_i are handled by the finite recursive extension of Lemma 5.8.

Proof. Write $\delta = [L_0 : \mathbb{Q}]$. If $[T] \in \mathrm{PGL}_2(L_0)$ has projective order m, the ratio of the two eigenvalues of T is a primitive m-th root of unity ζ_m. The eigenvalues lie in an extension of L_0 of degree at most two, whence

$$\varphi(m) = [\mathbb{Q}(\zeta_m) : \mathbb{Q}] \leq [L_0(\zeta_m) : \mathbb{Q}] \leq 2\delta.$$

There are therefore only finitely many possible projective element orders; let M_0 be their maximum. By Proposition 2.1, every finite subgroup of $\mathrm{PGL}_2(L_0)$ has order at most

$$C_0 = \max\{2M_0, 60\}.$$

For a sequence in the statement, define

$$\Phi_i = (C_0 - |H_i|) + (\nu_i - 3) \in \mathbb{Z}_{\geq 0}.$$

The inclusions and inequalities make Φ_i nonincreasing, and the strict condition decreases it by at least one at every update. Hence the sequence has at most Φ_0 updates. In the application, Lemma 2.22 places every pulled-back group in $\mathrm{PGL}_2(L_0)$, equivalent enrichment makes the groups nested, and Definition 5.5 makes ν nonincreasing. Factor closure between updates is finite because every nontrivial tensor split replaces a signature by positive-arity factors of strictly smaller arity; it is not part of the displayed measure. $\square$

For either group with odd N, the proof will establish that the single matrix P fixed in Equation (1) satisfies

$$(P^{-1})^{\otimes \operatorname{arity}(f)} f \in \mathscr{P} \quad (f \in \Lambda), \qquad P^{\mathsf{T}} X P = 2Z \in \mathscr{P}. \tag{16}$$

These are the simultaneous signature and edge conditions for a common K-presentation in $\mathscr{P}$; the same matrix must handle every member of Λ.

Theorem 5.7 (The odd-dihedral case). *Let $(A, \mathcal{G})$ be a stage of the sequence in Lemma 2.22 at which $\mathcal{G}$ is finite, preprocessed, algebraic, and factor-closed and contains the transported height-one mountain. Suppose that, after the normalization of Theorem 5.4, every nonzero directly realizable binary is nonsingular with finite projective transfer order, and $G(\mathcal{G})$ is*

$$D_{2N} = \{[R_s], [ZR_s] : s \in \mu_N\}$$

for an odd $N \geq 3$. Then at least one of the following holds.

1. K-Holant$(\mathcal{G})$ *is* #P*-hard or* $\mathsf{Tract}(K\mathcal{G})$ *holds.*
2. *There is a later stage $(A^+, \mathcal{G}^+)$ of the same sequence at which $\mathcal{G}^+$ is finite, preprocessed, algebraic, and factor-closed and still contains the correspondingly transported height-one mountain, such that*

$$\text{K-Holant}(\mathcal{G}^+) \equiv_{\mathrm{T}} \text{K-Holant}(\mathcal{G}),$$

every nonzero directly realizable ordered binary over $\mathcal{G}^+$ is nonsingular with finite projective transfer order, and either

$$[A]^{-1} G(\mathcal{G})[A] \subsetneq [A^+]^{-1} G(\mathcal{G}^+)[A^+]$$

or, with both values defined, $\nu(\mathcal{G}^+) < \nu(\mathcal{G})$. In either case the full update conditions of Definition 5.5 hold.

3. *There is an equivalent finite preprocessed algebraic factor-closed set $\mathcal{G}^{\mathrm{sat}} \supseteq \mathcal{G}$, in the same coordinates and with the same matrix A, for which the fixed Walsh matrix P satisfies*

$$(P^{-1})^{\otimes \operatorname{arity}(f)} f \in \mathscr{P} \quad (f \in \mathcal{G}^{\mathrm{sat}}), \qquad P^{\mathsf{T}} X P = 2Z \in \mathscr{P}.$$

Hence $K\mathcal{G}$ is $\mathscr{P}$-transformable.

Fix a signature set Γ for which $G(\Gamma)$ is defined, a subgroup $H \leq G(\Gamma)$, and one directly realizable ordered binary B_h with $[B_h X] = h$ for each $h \in H$. We will refer to the following contraction condition for $f \in \Gamma$:

(C) For every ordered pair of distinct ports of f, every chosen B_h, and both kernel orientations XB_hX and $XB_h^{\mathsf{T}}X$, the contraction is zero, or all its positive-arity tensor-prime factors are nonsingular binaries B with $[BX] \in H$.

Any nullary factors are absorbed into the recorded scalar. A zero contraction is allowed and is not assigned a projective transfer class. The conditions are in the current K-coordinates; after a Walsh change of coordinates the group membership is conjugated by P as well. Applications below take $H = G(\Gamma)$, with the complete group from Equation (15). Before using condition (C), any factor that gives hardness, tractability, or an update in Definition 5.5 is handled separately. This condition does not assert that arbitrary gadget outputs have been adjoined.

Lemma 5.8 (Finite recursive contraction and factor extraction)**.** *Let (A, Γ) be a stage of a sequence satisfying Lemma 2.22, over its fixed number field L_0. Suppose that Γ is finite, preprocessed, algebraic, and factor-closed, contains the transported height-one mountain, and every nonzero directly realizable ordered binary over Γ is nonsingular and has finite projective transfer order. Let $H = G(\Gamma)$, and fix one directly realizable ordered-binary representative of each class in this finite group.*

Start with every nonbinary tensor-prime factor currently in Γ. For each retained factor f, process every choice of an ordered pair of distinct ports, a fixed representative of a class in H, and either kernel orientation. Adjoin the resulting nonzero two-port contraction and, through Lemma 2.19, all of its positive-arity tensor-prime factors; apply the same procedure to each newly retained nonbinary factor. This finite recursive procedure reaches one of the following outcomes:

1. K-Holant(Γ) *is* #P*-hard or* $\mathsf{Tract}(K\Gamma)$ *holds;*

2. *there are a matrix A^+ and a finite preprocessed algebraic factor-closed set Γ^+, continuing the same coordinate-transport sequence and still containing the correspondingly transported height-one mountain, with* K-Holant$(\Gamma^+) \equiv_{\mathrm{T}}$ K-Holant(Γ)*, such that every nonzero directly realizable ordered binary over Γ^+ is nonsingular with finite projective transfer order, and either*

$$[A]^{-1}G(\Gamma)[A] \subsetneq [A^+]^{-1}G(\Gamma^+)[A^+]$$

or, with both values defined, $\nu(\Gamma^+) < \nu(\Gamma)$. The update satisfies Definition 5.5, including the non-strict comparisons for both quantities;

3. *there is a finite preprocessed algebraic factor-closed set $\Gamma^{\mathrm{sat}} \supseteq \Gamma$, with*

$$\text{K-Holant}(\Gamma^{\mathrm{sat}}) \equiv_{\mathrm{T}} \text{K-Holant}(\Gamma), \qquad G(\Gamma^{\mathrm{sat}}) = H,$$

for which the coordinate matrix remains A. Every indicated two-port contraction has been processed, every nonbinary tensor-prime factor supplied by factor closure is retained in Γ^{sat}, and every binary factor it produces is nonsingular with transfer class in H. A nonbinary factor produced from a parent f has arity at most $\operatorname{arity}(f) - 2$.

The third outcome may preserve ν; it is exhaustive closure with the same displayed group and coordinate matrix. It does not require a strict change in either quantity.

Proof. Begin with the finitely many nonbinary tensor-prime factors in Γ. For each retained factor, process its finite set of ordered-port, representative, and orientation choices. For each choice, adjoin the resulting contraction and all positive-arity tensor-prime factors under Equations (5) and (11). The transported mountain rules out the common-one-sided factor exception; if a factor theorem instead proves hardness or tractability, this is outcome 1.

An odd or unary factor is handled by the odd-arity or unary theorem. A singular binary, or a nonsingular binary whose transfer has infinite projective order, is handled by Lemma 3.1 and those same theorems. After each augmentation, consider all directly realizable nonzero ordered binaries. If any is singular or its transfer has infinite projective order, the preceding argument gives outcome 1. Otherwise the group is defined, and an element outside H strictly enlarges the pulled-back group, and a new nonbinary tensor-prime factor of arity below $\nu(\Gamma)$ strictly decreases ν. These give outcome 2, with both comparisons preserved as in Definition 5.5. If neither occurs, retain all the binary factors and repeat the procedure on each newly retained nonbinary tensor-prime factor.

There are finitely many initial factors and finitely many choices at each step. Every newly produced nonbinary factor has arity at most the arity of its parent minus two. The recursive depth is therefore bounded, so only finitely many factors and contraction choices occur. If neither of the first two outcomes occurs, processing all these choices gives outcome 3. Every augmentation retains the transformed mountain and satisfies the fixed-field and scalar conditions of Lemma 2.22. $\square$

5.2 Contractions for the odd cyclic and dihedral groups

The next two lemmas are shared by the odd cyclic and odd dihedral cases. For a signature $f \in \mathcal{G}$, recall that a tilde denotes only its Walsh coordinates

$$\widetilde{f} = (P^{-1})^{\otimes \operatorname{arity}(f)} f.$$

The matrix P is never inserted into a gadget: every contraction below is a native K-gadget, and every pencil identity is only an analysis of those contractions; no displayed linear combination is asserted to be realizable.

5.2.1 Contraction kernels and the four-coefficient relation

Fix odd $N \geq 3$. With R_s as in Theorem 5.4, the actual rotation and reflection classes are

$$[R_s], \qquad [ZR_s], \qquad s \in \mu_N. \tag{17}$$

They have Walsh forms

$$P^{-1} R_s P = D_s := \operatorname{diag}(1, s), \qquad P^{-1} Z P = X, \qquad F_s := X D_s.$$

Fix a direct binary-chain representative of every projective class in Equation (17). Contracting two ports then gives, up to the recorded nonzero representative scalar, the Walsh kernels

$$\mathsf{K}_s^+ = \begin{pmatrix} 1 & 0 \\ 0 & -s \end{pmatrix}, \qquad \mathsf{K}_s^- = \begin{pmatrix} 0 & 1 \\ -s & 0 \end{pmatrix}, \qquad s \in \mu_N. \tag{18}$$

The plus kernels occur in both groups; the minus kernels occur only in the dihedral group. Under port reversal, D_s is fixed, whereas F_s is sent projectively to $F_{s^{-1}}$. Attaching binary representatives to the ports adds no kernels because

$$\tau_Z(T) := Z T^{\mathsf{T}} Z$$

preserves $\{D_s : s \in \mu_N\}$ and $\{D_s, F_s : s \in \mu_N\}$. Hence Equation (18) lists every oriented kernel needed below.

Although this subsection otherwise has N odd, the next pencil fact is valid for every integer $N \geq 3$ and will also be used in the even-order cyclic case. For such an N and $a, b, c, d \in \overline{\mathbb{Q}}$, let $\mathcal{R}_N(a, b, c, d)$ mean that, for some $\alpha \in \mu_N$, one of the following three systems holds:

$$c = -\alpha a, \qquad d = -\alpha b, \qquad \text{(proportional)}, \tag{19}$$

$$b = c = 0, \qquad d = \alpha a, \qquad \text{(diagonal monomial)}, \tag{20}$$

$$a = d = 0, \qquad c = \alpha b, \qquad \text{(anti-diagonal monomial)}. \tag{21}$$

The three systems are not required to be disjoint; in particular, their union contains the all-zero quadruple.

Lemma 5.9 (Shared cyclic phase pencil, including zeros)**.** *For every integer $N \geq 3$, the relation $\mathcal{R}_N(a,b,c,d)$ is equivalent to the following condition: for every $s \in \mu_N$,*

$$\operatorname{diag}(a - sb, c - sd)Z \quad \textit{is zero or is projectively } D_\eta \textit{ for some } \eta \in \mu_N.$$

Proof. This is [LM26a, Lemma 7.1]. □

5.2.2 Support parity

Lemma 5.10 (Walsh parity from contractions)**.** *Let $N \geq 3$ be odd, let $H = C_N$ or D_{2N}, and let $f \neq 0$ have positive even arity. At arity two, suppose the Walsh transfer of f belongs to H. At arity at least four, suppose that every positive-arity prime factor of every nonzero contraction of f using an ordered pair of distinct ports, a chosen representative of H, and either kernel orientation is supported on one Walsh parity class; in the cyclic case suppose that class is even. Then $\widetilde{f}$ is supported on one parity class, which is even in the cyclic case.*

Proof. At arity two, a Walsh binary is, up to a nonzero scalar, a diagonal or anti-diagonal transfer times Z, and hence has one parity. Only the diagonal alternative occurs in the cyclic group. At larger arity, the tensor product of the prime factors of a nonzero contraction has one parity; in the cyclic case it is even.

At arity four, each K_s^+-contraction is zero, diagonal, or anti-diagonal, and every entry is affine in s. Among three phases, either two contractions vanish, annihilating the pencil, or two nonzero contractions have the same support parity, which forces the opposite entries to vanish identically. In the only remaining pattern—one zero, one diagonal, and one anti-diagonal contraction—each opposite pair acquires two roots, with the same conclusion. Hence every equal-bit slice has one parity. Two supported words of opposite parity have a coordinate pair equal in both unless their four ordered coordinate pairs exhaust $\mathbb{F}_2^2$; in the latter case both words have weight two and therefore already have the same parity, a contradiction. In the former case deleting the repeated pair contradicts slice parity. In the cyclic case, odd Walsh parity would realize a reflection and is therefore excluded.

For arity at least six, fix distinct ports i, j. The actual plus contraction is

$$\widetilde{h}_s(y) = \lambda_s\big(\widetilde{f}(00, y) - s\widetilde{f}(11, y)\big), \qquad \lambda_s \neq 0.$$

A zero contraction makes all others proportional. Otherwise two phases have the same contraction parity. Every coordinate in the opposite parity then vanishes at two distinct phases, so its affine coefficient vanishes identically; the two nonzero phases also recover both endpoint slices with the common parity.

In the cyclic case, delete two equal coordinates of any supported word and choose a phase other than the at most one value cancelling its residual coefficient; the resulting contraction is even. In the dihedral case, two opposite-parity supported words at arity at least six have two coordinates with the same ordered pair. At most one phase cancels each residual coefficient, so $|\mu_N| \geq 3$ leaves a phase preserving both words. Deleting the chosen pair preserves their opposite parities, contradicting the common slice parity. □

For every f covered by Lemma 5.10, Walsh duality gives the equivalent complement relation in the native K-coordinates,

$$f(\bar{x}) = f(x) \quad \text{or} \quad f(\bar{x}) = -f(x),$$

with only the plus sign in the cyclic subcase.

5.3 The odd-dihedral case

5.3.1 Support and coefficients on even blocks

For a finite port set V and $C \subseteq V$, let $\mathbf{1}_C \in \mathbb{F}_2^V$ denote the incidence vector of C. For a partition π of V into nonempty even blocks, define the subspace

$$L_\pi = \operatorname{span}_{\mathbb{F}_2}\{\mathbf{1}_C : C \in \pi\}.$$

Its cosets are precisely the sets obtained by toggling each block as one unit.

The class $\mathcal{P}_N(V)$ consists of the zero tensor and all tensors g for which there exist such a partition π, an offset $a \in \mathbb{F}_2^V$, a scalar $c \in \overline{\mathbb{Q}} \setminus \{0\}$, and block ratios $r_C \in (-1)^{|C|/2}\mu_N$ such that

$$\begin{aligned} \operatorname{supp}(g) &= a + L_\pi, \\ g\left(a + \sum_{C\in\pi} t_C \mathbf{1}_C\right) &= c \prod_{C\in\pi} r_C^{t_C} \qquad ((t_C)_{C\in\pi} \in \mathbb{F}_2^\pi). \end{aligned}$$

For $C \subseteq V$, write $a|_C$ for the restriction of a to C, and $\overline{a|_C}$ for its bitwise complement. Equivalently, with the tensor factors placed on their indicated port sets,

$$g = c \bigotimes_{C\in\pi} (\delta_{a|_C} + r_C \delta_{\overline{a|_C}}).$$

Thus, in the variables (t_C), the coefficient function is the product $c \prod_C r_C^{t_C}$. A factor on one block has the displayed two-point form and the stated ratio restriction. For nonzero g, its support determines π: the inclusion-minimal nonzero vectors of L_π are precisely the block indicators. We omit V when clear. The class $\mathcal{P}_N$ differs from the standard product class $\mathscr{P}$ of Section 2.13.1; its even-block and coefficient-ratio restrictions remain part of its definition.

Lemma 5.11 (Closure of $\mathcal{P}_N$ under contractions)**.** *The class $\mathcal{P}_N$ is closed under nonzero scalar multiplication, port permutation, tensor product on disjoint port sets, and, up to the zero tensor, contraction of two ports against any kernel K_s^+, K_s^-, or its transpose from Equation* (18)*, where $s \in \mu_N$. Consequently, if every positive-arity prime factor of a contraction has Walsh tensor in $\mathcal{P}_N$, then the contraction itself has Walsh tensor in $\mathcal{P}_N$.*

Proof. For a tensor product, take the disjoint union of the two even-block partitions and concatenate their offsets and block ratios. Scalars and port permutations merely change the leading coefficient or relabel that data.

Consider a two-port contraction of a nonzero tensor in $\mathcal{P}_N$. If the ports lie in the same block C, the kernel with the wrong xor gives zero. For the supported xor, delete the two ports. When ports remain in C, its new block ratio is

$$-sr_C \quad \text{or} \quad -r_C/s,$$

and therefore belongs to $(-1)^{(|C|-2)/2}\mu_N$. If $|C| = 2$, the block instead contributes a scalar, possibly zero, which is absorbed into the leading coefficient.

If the ports lie in distinct blocks C, D, the kernel affinely identifies their two block variables. After complementing one residual offset when needed, replace those blocks by

$$(C \setminus \{i\}) \,\dot\cup\, (D \setminus \{j\}).$$

Its ratio is one of
$$-s^{\pm 1} r_C r_D^{\pm 1} \in (-1)^{(|C|+|D|-2)/2} \mu_N.$$
All other blocks are unchanged. Transposing K_s^- only supplies a nonzero scalar and replaces s by s^{-1}, so it is covered as well. Thus every nonzero contraction is again in $\mathcal{P}_N$. The last assertion follows from prime factorization; nullary factors are absorbed into the scalar. $\square$

5.3.2 Four-port support and coefficients

Fix $N \geq 3$ and a quaternary signature q in the current set, and put $\widetilde{q} = (P^{-1})^{\otimes 4} q$. When $\widetilde{q}$ is supported on the even Hamming-weight positions in Walsh coordinates, write

$$p = \widetilde{q}_{0000}, \qquad t = \widetilde{q}_{1111}, \qquad x_{ab} = \widetilde{q}_{\mathbf{1}_{\{a,b\}}}. \tag{22}$$

For every ordering (i, j, k, l) of $[4]$, the equal and mixed phase-contraction matrices are, respectively,

$$H_{ij}^{+}(s) = \begin{pmatrix} p - s x_{ij} & 0 \\ 0 & x_{kl} - st \end{pmatrix},$$
$$H_{ij}^{-}(s) = \begin{pmatrix} 0 & x_{jl} - s x_{il} \\ x_{jk} - s x_{ik} & 0 \end{pmatrix}.$$

Thus condition (C) for the cyclic group and Lemma 5.9 give $\mathcal{R}_N(p, x_{ij}, x_{kl}, t)$. For the dihedral group, the mixed contraction also gives $\mathcal{R}_N(x_{jk}, x_{ik}, x_{jl}, x_{il})$.

When $\widetilde{q}$ is supported on the odd Hamming-weight positions in Walsh coordinates, write

$$u_i = \widetilde{q}(e_i), \qquad v_i = \widetilde{q}(\mathbf{1}^4 \oplus e_i).$$

For every such port ordering, condition (C) for the dihedral group gives

$$\mathcal{R}_N(u_k, v_l, u_l, v_k), \qquad \mathcal{R}_N(u_j, u_i, v_i, v_j).$$

Lemma 5.12 (Shared cyclic quaternary family)**.** *Let $N \geq 3$, let Λ be a finite algebraic signature set in K-coordinates, and, for every $s \in \mu_N$, fix an actual ordered binary representative C_s over Λ, in a specified orientation, whose Walsh matrix is*

$$\widetilde{C}_s = \kappa_s D_s Z, \qquad \kappa_s \in \overline{\mathbb{Q}}^{\times}.$$

Let $q \in \Lambda$ be quaternary, and suppose $\widetilde{q} = (P^{-1})^{\otimes 4} q$ is supported on the even Hamming-weight positions in Walsh coordinates. Define p, t, x_{ab} by Equation (22). Assume that, for every complementary pair, both orders of its two residual ports, and every $s \in \mu_N$, contraction with C_s gives either the zero contraction or a nonsingular binary contraction whose Walsh transfer class belongs to

$$C_N = \{[D_\eta] : \eta \in \mu_N\}.$$

Organize the eight coordinates into four complementary pairs. Put $B_0 = (p, t)$ and $\mathfrak{I} = \{\{1,2\}, \{1,3\}, \{1,4\}\}$. For $I = \{a, b\} \in \mathfrak{I}$, set $x_I = x_{ab}$, put $I^c = [4] \setminus I$, and define the other three complementary pairs $B_I = (x_I, x_{I^c})$, $I \in \mathfrak{I}$. Here $B_I \neq (0, 0)$ initially means only that at least one coefficient is nonzero; the conclusion below makes both coefficients nonzero. Then exactly one of the following holds:

1. If $p, t \neq 0$, then $\lambda = t/p \in \mu_N$, and each pair B_I is zero or

$$B_I = (-p\lambda\alpha_I^{-1}, -p\alpha_I), \qquad \alpha_I \in \mu_N. \tag{23}$$

2. If $p = t = 0$, each pair B_I is zero or

$$B_I = \rho_I(1, \alpha_I), \qquad \rho_I \neq 0, \quad \alpha_I \in \mu_N.$$

In particular, exactly one of p, t cannot be nonzero. Conversely, every even Walsh table in these alternatives satisfies the stated contraction condition under every phase contraction.

Proof. This is [LM26a, Lemma 7.2]. □

Proposition 5.13 (Dihedral quaternary classification)**.** *Let $N \geq 3$ be odd, and let $q \in \Gamma$ be a nonzero quaternary satisfying condition* (C) *with $H = G(\Gamma) = D_{2N}$ in the normalized coordinates above. Put $\widetilde{q} = (P^{-1})^{\otimes 4}q$. Then $\widetilde{q} \in \mathcal{P}_N([4])$. More precisely, $\widetilde{q}$ has one Walsh parity and exactly one or two complementary pairs whose coefficient pairs are nonzero. After choosing an orientation, there is a scalar $\rho \in \overline{\mathbb{Q}} \setminus \{0\}$ such that, for one nonzero complementary pair $\{x, \bar{x}\}$, its nonzero coefficients in the order $(\widetilde{q}_x, \widetilde{q}_{\bar{x}})$ are*

$$\rho(1, \delta), \qquad \delta \in \mu_N, \tag{24}$$

or, for two nonzero complementary pairs $\{x, \bar{x}\}$ and $\{y, \bar{y}\}$, its nonzero coefficients in the order $(\widetilde{q}_x, \widetilde{q}_{\bar{x}}, \widetilde{q}_y, \widetilde{q}_{\bar{y}})$ are

$$\rho(1, \delta, -\gamma, -\delta\gamma^{-1}), \qquad \delta, \gamma \in \mu_N. \tag{25}$$

Conversely, every Walsh table in Equations (24) *and* (25) *satisfies the stated contraction condition under all phase and reflection contractions.*

Proof. Equal kernels give a μ_N-ratio on every nonzero complementary pair. If $p, t \neq 0$, Equation (23) puts all nonzero central coefficients in $-p\mu_N$. For two nonzero central coefficient pairs, the 2-by-2 coefficient matrix of the reflected two-port contraction has all four entries nonzero, and its only possible proportional alternative in Lemma 5.9 requires a ratio in $-\mu_N$, not μ_N. The cosets are disjoint for odd N.

With zero endpoints, three nonzero central coefficient pairs would require three pairwise ratios in $-\mu_N$, although the product of two lies in μ_N. Thus at most two complementary pairs have nonzero coefficients. Applying the same argument to both orientations of the reflected two-port kernels K_s^- gives the identical conclusion on the odd Hamming-weight positions.

Substitution gives Equations (24) and (25) and proves their converse, including zeros. One block uses $\pi = \{[4]\}$; two blocks form a $2+2$ partition coset with ratios $-\gamma, -\delta\gamma^{-1} \in -\mu_N$, whose product is δ. Thus both parities lie in $\mathcal{P}_N([4])$, with odd parity absorbed by the coset offset.

A one-block tensor is a weighted GE. Up to flips and permutation, a two-block support is $\{0000, 1111, 1100, 0011\}$, and Equation (25) gives

$$\widetilde{q}_x\widetilde{q}_{\bar{x}} = \widetilde{q}_y\widetilde{q}_{\bar{y}},$$

for the two complementary pairs $\{x, \bar{x}\}$ and $\{y, \bar{y}\}$ in the support. Hence it is a product of two nonsingular weighted equality/disequality binaries. Thus $\widetilde{q} \in \mathscr{P}$, as required by Equation (16). □

5.3.3 Products of binaries and the six-port case

Write $\mathcal{B}_N(V) \subseteq \mathcal{P}_N(V)$ for zero together with the tensors whose partition consists entirely of pairs. Every nonzero member is a product of nonsingular Walsh equality/disequality binaries, with ratio in $-\mu_N$ on each pair. For the resulting perfect matching M, its support is $a+L_M$, is complement-closed, and has cardinality $2^{|V|/2}$. The contraction formulas in Lemma 5.11 show that $\mathcal{B}_N$ is itself closed under all residual-group contractions: joining different pairs makes one pair, and closing a pair gives a scalar.

Theorem 5.14 (Six-port factorization from contractions)**.** *Let $N \geq 3$ be odd. If every nonzero oriented phase or reflection contraction of an arity-six signature f has Walsh tensor in $\mathcal{B}_N([4])$, then $f = 0$ or f is a tensor product of three nonsingular binaries.*

Proof. Suppose $f \neq 0$, and write $F = \widetilde{f}$. By Lemma 5.10, F has one parity. Fix a port pair and fix $x_i \oplus x_j = \xi$, with the two corresponding slices A, B, so the normalized actual contractions are $q_s = A - sB$. Put $U = \operatorname{supp}(A) \cup \operatorname{supp}(B)$. Each nonzero coordinate remains nonzero at at least $N-1$ phases, while every nonzero contraction has four support points. Thus

$$(N-1)|U| \leq 4N.$$

Moreover $U = \bigcup_s \operatorname{supp}(q_s)$ is complement-closed. For $N \geq 5$ it is therefore empty or consists of at most two complement pairs. For $N = 3$, the same conclusion is exactly Theorem A.3(i). Part (ii) of that theorem now applies for every odd N: the support of F is a partition coset of type (6), $(4+2)$, or $(2+2+2)$, or belongs to the orbit of

$$S_\star = \{000000, 000011, 001100, 110011, 111100, 111111\}. \tag{26}$$

Here and below, bit relabelings describe support coordinates only. They permute the available xor kernels projectively, with possible phase inversion, so all contraction arguments cover the full indicated orbits.

A (6) support has a supported-xor contraction with exactly two nonzero entries, each with a unique lift. This contradicts the four-point support required by $\mathcal{B}_N([4])$. On a $(4+2)$ support, use block variables and denote the four nonzero coefficients by c_{uv}. Closing the binary block gives, after normalizing its orientation,

$$(c_{00} - sc_{01})\delta_{0000} + (c_{10} - sc_{11})\delta_{1111}.$$

At most two phases are forbidden by cancellation; since $N \geq 3$, some contraction again has exactly two nonzero entries. No tensor factorization of the original four coefficients was assumed. On $S_\star$, contract ports $3, 5$ with K_s^+. Only the two endpoints have equal bits there, giving

$$F_{000000}\delta_{0000} - sF_{111111}\delta_{1111},$$

again with two nonzero entries and unique lifts.

Only the $(2+2+2)$ support remains. Its eight coefficients c_{uvw} are nonzero. Closing any one of its binary blocks leaves a 2-by-2 table whose determinant vanishes at every $s \in \mu_N$: a nonzero contraction has the same four-point matching support, and a zero contraction also has zero determinant. For example,

$$\det\big(c_{uv0} - sc_{uv1}\big)_{u,v\in\mathbb{F}_2} = 0.$$

This polynomial has degree at most two and at least three roots. Its constant and quadratic coefficients show that the two endpoint faces have rank one. Doing this for all three axes makes every axis face minor zero. Since all eight entries are nonzero, the ratios along each axis are independent of the other two bits, so $c_{uvw} = cr_1^u r_2^v r_3^w$. This is the required three-binary factorization, which pulls back through the analytic Walsh change of coordinates. $\square$

5.3.4 Matching pencils

Lemma 5.15 (Three matching products on a line)**.** *Let three pairwise nonproportional tensor products of nonsingular diagonal or anti-diagonal binaries on the same m ports satisfy a linear relation with all coefficients nonzero. Either their supports coincide and the three tensors lie in one Segre ruling, or their matchings differ only on four ports, where the three different pairings occur. In the latter case all binary factors outside those four ports are common up to nonzero scalars.*

Proof. Let the supports be S_0, S_1, S_2, matching cosets of dimension $m/2$. No point can belong to just one support. If two supports coincide, the third is contained in them and has the same cardinality, so all coincide. On their common block coordinates, every matrix flattening minor has degree at most two along the line. Three rank-one points make these minors vanish identically. A line of simple tensors varies in at most one factor: if two factors of two of its points were nonproportional, flattening across one of them would give rank two at a generic third point. This is one Segre ruling.

Otherwise each S_i is the union of its intersections with the other two supports. Two proper affine subspaces cover an affine space only when they are disjoint parallel hyperplanes: each has at most half its points, so both must attain that bound. Thus the pairwise intersections are hyperplanes and the triple intersection is empty. The union of two perfect matchings is a disjoint union of alternating even cycles, counting a common edge as a two-cycle. If there are $k = m/2$ edges and b components, their common direction code has dimension b. Here $b = k-1$, so exactly one component is a four-cycle and all other edges coincide. The offsets on those common edges agree because the supports intersect. After removing their free binary coordinates, S_2 is the four-point symmetric difference of the other two local supports; its matching is the third pairing on the same four ports. Finally, on each pairwise intersection only two tensors contribute to the dependence. Fixing one four-port assignment with nonzero coefficient leaves the full external matching cube, on which their products are proportional. Uniqueness of the factors on disjoint pairs gives the asserted common factors. □

Proposition 5.16 (Phase-independent matching and common factors)**.** *Let $N \geq 3$ be odd, and let f have even arity $n \geq 8$. Suppose every nonzero result of one or more successive two-port contractions using the kernels in Equation* (18) *has Walsh tensor in $\mathcal{B}_N$. Then each two-port pencil can be written as $q_s = H \otimes h_s$, where H is a fixed product of nonsingular binaries on disjoint pairs and h_s is a binary tensor on one fixed remaining port pair. Zero members and proportional pencils are included.*

Proof. Write a pencil as $q_s = A - sB$. A zero member makes all other members proportional, which settles that case. If the pencil is not projectively constant, any three distinct phases give a linear dependence with all three coefficients nonzero, so Lemma 5.15 applies to these members of $\mathcal{B}_N$. In its common-support case, the quadratic flattening minors vanish identically, giving the fixed matching ruling for the whole pencil.

In the second case of Lemma 5.15, the matchings differ on four ports and all other binary factors are common. Close every such common pair, on the original gadget, with its supported-xor kernel at phase 1. For a binary ratio $r \in -\mu_N$, this gives a nonzero multiple of $1 - r$, nonzero because N is odd. These operations multiply the three contractions by the same nonzero scalar. At least one pair is closed since $n \geq 8$. Before closing the original pencil ports, the resulting six-port tensor is a nonzero result of contractions of f, hence is itself a binary matching product. Its contractions cannot have three different matchings: closing ports of a fixed matching deletes an edge or merges two edges, independently of the phase. This contradiction proves that the matching is independent of phase. □

5.3.5 A common binary factor suffices

Lemma 5.17 (A binary factor from contractions)**.** *Under the hypotheses of Proposition 5.16, every nonzero f is tensor-decomposable.*

Proof. Put $F = \widetilde{f}$, fix ports i, j, and for each $\xi \in \{0, 1\}$ consider the restriction to $x_i \oplus x_j = \xi$, viewed as an n-ary tensor that is zero outside that constraint. Its two slices are $A = F_{ij}^{0,\xi}$, $B = F_{ij}^{1,1-\xi}$. Exactly one of A, B cannot be nonzero. Indeed, if $A \neq 0$ and $B = 0$, take a supported word of A, close any two external ports with its supported xor, and choose a phase avoiding cancellation of the projected word. The resulting nonzero contraction has complement-closed support. Its complementary word would require an assignment in B, a contradiction. The case $A = 0$, $B \neq 0$ is identical.

For each nonzero restriction, Proposition 5.16 writes the pencil as $H \otimes h_s$, where H is a fixed product of external binaries and

$$h_s = (a - sb)\delta_e + (c - sd)\delta_{\bar{e}}$$

on one external pair. Projectively constant pencils are included by choosing any pair. Every nonzero h_s has ratio in $-\mu_N$, so Lemma 5.9 applies. Restoring the ports i, j, its proportional alternative gives a binary block $\{i, j\}$; either monomial alternative gives one four-port GE block containing i, j. All other blocks are nonsingular binaries. Both slices A, B are nonzero, so the restored block has both coefficients nonzero. This is a factorization in Walsh coordinates of the restriction; it does not assert that either this restriction or either slice is realizable by a gadget.

If only one of the two restrictions is nonzero, this factorization of F already has an external binary factor, since $n \geq 8$. Otherwise denote the restrictions for $\xi = 0, 1$ by U, V. We use the following observation repeatedly. Close an external pair with an xor supported by both restrictions, and choose a phase at which both resulting contractions are nonzero. Such a phase exists: a contraction of either displayed block product with a supported xor has unique lifts, except when an entire binary block is closed, in which case only one phase can cancel its scalar. The nonzero contraction G is a binary matching product with both ij xors present. Its two ij-xor restrictions have:

(O1) the same four-port block containing i, j;

(O2) the same external matching and the same offset on each external pair;

(O3) proportional full binary factors on each external pair, including equality of their ordered coefficient ratios.

Indeed, i, j belong to different matching edges in G. Either xor restriction joins those edges into one four-port block and leaves every other binary factor unchanged. In particular, this observation identifies coefficients as well as supports.

First, neither U nor V can have $\{i, j\}$ as a binary block. If U did, close ports in two distinct external binary blocks of U. Both xors occur in U, so one is also supported by V. The observation would require a four-port block through i, j, whereas this contraction leaves their binary block unchanged.

Write the four-port blocks of U, V as $K = \{i, j, a, b\}$, $L = \{i, j, c, d\}$. Close a, b with its fixed xor in U. That restriction now has $\{i, j\}$ as a binary block, so the observation implies that V does not support this xor. Therefore a, b lie in the same block of V, with the opposite xor. Either $L = K$, or $\{a, b\}$ is an external binary block of V. By symmetry, in the latter case $\{c, d\}$ is an external binary block of U, and the two pairs are disjoint. Choose a port e outside $\{i, j, a, b, c, d\}$, possible since $n \geq 8$, and close a, e. Both restrictions have unique lifts for either xor. The four-port block through i, j in the contracted U contains b, whereas that in the contracted V remains $\{i, j, c, d\}$. This contradicts the observation. Thus $K = L$.

Let M_U, M_V be the external matchings on $E = [n] \setminus K$. For $e \in E$, closing a, e makes the four-port block through i, j equal to $\{i, j, b, M_U(e)\}$ in one restriction and $\{i, j, b, M_V(e)\}$ in the other. The observation gives $M_U(e) = M_V(e)$, so the external matchings coincide. There are at least two external pairs. For any one of them, R, choose e on another pair and close a, e. The pair R remains external. The observation now identifies its full binary factor in the two restrictions, hence in U, V, since the contraction avoids R. Therefore $U + V = F$ has this common factor. Invertible Walsh coordinates preserve tensor factorization, proving the lemma. □

Proof of Theorem 5.7. Apply Lemma 5.8. Its first two outcomes are the first two conclusions of the theorem. In outcome 3, work in the equivalent set $\Gamma = \mathcal{G}^{\mathrm{sat}}$, with the same matrix and group. If a nonbinary prime remains, choose one, f, of minimum arity n. Odd and unary factors have already been handled by their respective complexity theorems. Every positive-arity prime factor of a two-port contraction of f has smaller arity, and hence is binary by minimality. Its transfer is in D_{2N}, so every nonzero two-port contraction is in $\mathcal{B}_N$ in Walsh coordinates. Closure of $\mathcal{B}_N$ gives the same assertion for all further contractions involving at least one port pair.

If $n = 4$, Proposition 5.13 gives either a four-port generalized equality, to which Lemma 3.15 applies, or a product of two binaries, contradicting primeness. If $n = 6$, Theorem 5.14 contradicts primeness. If $n \geq 8$, Lemma 5.17 gives the same contradiction. Thus either hardness or tractability has been proved, or every tensor-prime in Γ is binary. In the latter case one Walsh basis puts all signatures in $\mathscr{P}$, and $P^{\mathsf{T}} X P = 2Z \in \mathscr{P}$, giving outcome 3 and the stated common K-presentation. □

5.4 The odd-cyclic case

Theorem 5.18 (The odd-cyclic case)**.** *Let $(A, \mathcal{G})$ be a stage of the sequence in Lemma 2.22 at which $\mathcal{G}$ is finite, preprocessed, algebraic, and factor-closed and contains the transported height-one mountain. Let $N \geq 3$ be odd, and suppose every nonzero directly realizable binary is nonsingular with finite projective transfer order, and*

$$P^{-1} G(\mathcal{G}) P = C_N = \{[\mathrm{diag}(1, s)] : s \in \mu_N\}.$$

Then at least one of the following holds.

1. K-Holant$(\mathcal{G})$ *is* #P*-hard or* Tract$(K\mathcal{G})$ *holds.*
2. *There is a later stage $(A^+, \mathcal{G}^+)$ of the same sequence at which $\mathcal{G}^+$ is finite, preprocessed, algebraic, and factor-closed and still contains the correspondingly transported height-one mountain, such that*

$$\text{K-Holant}(\mathcal{G}^+) \equiv_{\mathrm{T}} \text{K-Holant}(\mathcal{G}),$$

every nonzero directly realizable ordered binary over $\mathcal{G}^+$ is nonsingular with finite projective transfer order, and either

$$[A]^{-1} G(\mathcal{G}) [A] \subsetneq [A^+]^{-1} G(\mathcal{G}^+) [A^+]$$

or, with both values defined, $\nu(\mathcal{G}^+) < \nu(\mathcal{G})$. In either case the full update conditions of Definition 5.5 hold.

3. *There is an equivalent finite preprocessed algebraic factor-closed set $\mathcal{G}^{\mathrm{sat}} \supseteq \mathcal{G}$, in the same coordinates and with the same matrix A, for which the fixed Walsh matrix P satisfies*

$$(P^{-1})^{\otimes \operatorname{arity}(f)} f \in \mathscr{P} \quad (f \in \mathcal{G}^{\mathrm{sat}}), \qquad P^{\mathsf{T}} X P = 2Z \in \mathscr{P}.$$

Hence $K\mathcal{G}$ is $\mathscr{P}$-transformable.

5.4.1 Phase-independent matching

Throughout this odd-cyclic proof, $\widetilde{f}$ denotes the fixed Walsh coordinates, and $P^{-1}G(\Gamma)P = C_N = \{[D_s] : s \in \mu_N\}$, with one direct representative fixed for each class and $N \geq 3$ odd.

For the conditional analysis below, work in outcome 3 of Lemma 5.8, retaining its finite, preprocessed, algebraic, factor-closed set Γ, the matrix A, and the transported height-one mountain. Odd-arity tensor-prime signatures have already been handled by Theorem 2.13. Choose a nonbinary tensor-prime $f \in \Gamma$ of minimum arity n, which is even. Attaching the chosen binary representative to distinct ports i, j and closing through the native edges gives, up to a known nonzero representative scalar, the Walsh kernel

$$\mathsf{K}_s = \operatorname{diag}(1, -s), \qquad s \in \mu_N.$$

After dividing out that recorded scalar, the contraction has Walsh coefficients

$$C_{ij,s}(y) = A(y) - sB(y), \qquad A(y) = \widetilde{f}(00, y), \quad B(y) = \widetilde{f}(11, y). \tag{27}$$

By Lemma 5.8(3), each indicated nonzero contraction and all its positive-arity tensor-prime factors belong to Γ. A nonbinary factor would have arity at most $n-2$, contradicting minimality. Every binary factor is nonsingular and its Walsh transfer class belongs to C_N. Thus condition (C) holds for f; each nonzero contraction is a product of such binaries on a perfect matching. A transfer D_t has Walsh tensor

$$D_t Z = \operatorname{diag}(1, -t).$$

For a perfect matching M of the residual ports, put

$$E_M = \{y \in \mathbb{F}_2^{[n]\setminus\{i,j\}} : y_a = y_b \text{ for every } \{a, b\} \in M\}. \tag{28}$$

Every nonzero contraction has support E_M, with a nowhere-zero rank-one table on the matching variables. By Lemma 5.10, $\widetilde{f}$ has even Walsh parity.

Lemma 5.19 (Phase-independent matching)**.** *Let $N \geq 3$ be odd, and let $f \in \Gamma$ be a nonbinary tensor-prime signature of minimum arity under the assumptions of this subsection, including condition* (C) *with $P^{-1}G(\Gamma)P = C_N$. Define $C_{ij,s}$ and E_M by Equations* (27) *and* (28)*. For each pair i, j, either every contraction vanishes, or $C_{ij,s} = 0$ for at most one phase $s \in \mu_N$ and there is a fixed perfect matching M_{ij} such that*

$$C_{ij,s} \neq 0 \quad \Longrightarrow \quad \operatorname{supp}(C_{ij,s}) = E_{M_{ij}} \qquad (s \in \mu_N).$$

This includes the boundary case $N = 3$.

Proof. If $A - sB$ and $A - tB$ vanish at distinct phases, then $A = B = 0$, so the first alternative holds. Otherwise at most one phase gives a zero contraction. Any two distinct nonzero contractions satisfy

$$\operatorname{supp}(A) \cup \operatorname{supp}(B) = \operatorname{supp}(A - sB) \cup \operatorname{supp}(A - tB), \tag{29}$$

because a scalar affine polynomial cannot vanish at two phases unless both coefficients vanish.

If no contraction vanishes, choose three phases. By Equation (29), the third matching code lies in the union of the first two. A vector space contained in $U \cup V$ lies in one of them (otherwise sum vectors from its two exclusive parts), and all matching codes have the same dimension. Thus the third code equals one of the first two. Two phases with one code and a third with another would make a word in the latter difference vanish at two phases, contradicting the affine-pencil identity. Hence all codes coincide.

If $A - s_0 B = 0$, then every other contraction is $(s_0 - s)B$, and the conclusion is immediate. $\square$

5.4.2 Pairs with every phase contraction zero

Define

$$Z_f^{\mathrm{cyc}} = \{\{i,j\} : C_{ij,s} = 0 \text{ for every } s \in \mu_N\}.$$

Two phases in Equation (27) show that an edge is exactly the support equation

$$\widetilde{f}(x) \neq 0 \quad \Longrightarrow \quad x_i + x_j = 1.$$

It follows that every connected component of Z_f^{cyc} is complete bipartite. An odd path forces unequal endpoint bits and hence supplies the corresponding edge; an even path forces equal bits and cannot also be a zero edge.

Lemma 5.20 (Every port pair has a nonzero phase contraction above arity four)**.** *Under the assumptions of Lemma 5.19, if* $\operatorname{arity}(f) = n \geq 6$*, then*

$$Z_f^{\mathrm{cyc}} = \varnothing.$$

Proof. Each color class has size at most two: if one contained p, q, r and the other b, a nonzero contraction on the nonzero pair p, q would have an all-zero residual coefficient while retaining the forced disequality $r \neq b$, impossible for a nonsingular matching product.

An isolated $\{1\} \sqcup \{2\}$ is impossible. For every $k \geq 3$, the all-zero word of a nonzero $(1,k)$-contraction has the unique lift $x_1 = x_k = 1$, whence

$$\widetilde{f}(e_1 + e_k) \neq 0 \qquad (k = 3, \ldots, n).$$

A nonzero $(2,3)$-contraction then has each $e_1 + e_k$, $k \geq 4$, by the unique lift $x_2 = x_3 = 0$. Its support matching would contain every $\{1,k\}$, impossible.

Thus only $2+1$ and $2+2$ remain. Label the size-two color $\{1,3\}$ and the other $\{2\}$ or $\{2,4\}$. In a $(1,3)$ contraction the forced equations give the unique lift

$$x_1 = x_3 = 1 - x_2.$$

Put

$$h(x_2, \mathbf{y}) = \widetilde{f}(1 - x_2, x_2, 1 - x_2, \mathbf{y}),$$

where $\mathbf{y}$ denotes the remaining coordinates; in the $2+2$ case, also impose the support equation $x_4 = x_2$. Then

$$C_{13,s}(x_2, \mathbf{y}) = u_s(x_2) h(x_2, \mathbf{y}), \qquad u_s(0) = -s, \quad u_s(1) = 1.$$

Every contraction is nonzero and equals a nowhere-zero unary u_s times h. Dividing analytically by u_s only modifies the matching factor incident with port 2, so h factors over that matching. Reversing the unique lift gives a four-port GE factor on $\{1,2,3,4\}$ in the $2+2$ case, or on $\{1,2,3,k\}$ in the $2+1$ case, tensor the remaining binaries. For $n \geq 6$ this contradicts tensor-primeness. $\square$

5.4.3 Weight-two coefficients and the arity bound

Lemma 5.21 (Rank-one diagonal restrictions)**.** *Let $r \geq 3$, and let $h : \mathbb{F}_2^r \to \mathbb{C}^\times$ have full support and $h(0) = 1$. Suppose that, for every $i \neq j$, the restriction $h|_{t_i=t_j}$, viewed as a function of the common bit and the other $r-2$ bits, has tensor rank one. If $r \geq 4$, then*

$$h(t) = \prod_{i=1}^{r} \rho_i^{t_i} \qquad (\rho_i \neq 0).$$

If $r = 3$, there are $\rho_1, \rho_2, \rho_3, \kappa \neq 0$ such that

$$h(e_i) = \rho_i, \qquad h(e_i + e_j) = \kappa\rho_i\rho_j, \qquad h(111) = \kappa\rho_1\rho_2\rho_3. \tag{30}$$

Proof. This is [LM26a, Lemma 7.8]. □

Theorem 5.22 (The minimum nonbinary arity is four in the cyclic case)**.** *Let Γ be a finite preprocessed algebraic factor-closed signature set containing the transported height-one mountain. Suppose every nonzero directly realizable binary is nonsingular and has finite projective transfer order, and, in the fixed Walsh coordinates,*

$$P^{-1}G(\Gamma)P = C_N = \{[\mathrm{diag}(1, s)] : s \in \mu_N\}, \qquad N \geq 3 \textit{ odd.}$$

Let $f \in \Gamma$ be an even-arity nonbinary tensor-prime signature of minimum arity among the nonbinary tensor-prime signatures in Γ. For every ordered pair of distinct ports of f, every fixed direct representative of a class in $G(\Gamma)$, and either kernel orientation, assume that the resulting two-port contraction is zero or all of its positive-arity tensor-prime factors are nonsingular binaries B satisfying

$$[P^{-1}(BX)P] \in C_N.$$

Then $\mathrm{arity}(f) \leq 4$*, and hence* $\mathrm{arity}(f) = 4$.

Proof. The contraction hypothesis is condition (C) for the stated group, so the two preceding lemmas apply. Assume that f has arity $n \geq 6$. By Lemma 5.20, every pair has a nonzero phase contraction. For $S \subseteq [n]$, write $a_S = \widetilde{f}(\mathbf{1}_S)$. We abbreviate $a_{i_1 \cdots i_k} = a_{\{i_1,\ldots,i_k\}}$.

Step 1: the constant coefficient is nonzero. If $a_\varnothing = 0$, the all-zero coefficient $C_{ij,s}(0) = -sa_{ij}$ of a nonzero matching contraction forces $a_{ij} \neq 0$ for every pair. Hence all contractions are nonzero. For disjoint pairs ij, pq,

$$C_{ij,s}(e_p + e_q) = a_{pq} - sa_{ijpq}$$

is nonzero at some phase. The common matching from Lemma 5.19 would therefore contain every edge on its at least four residual ports, impossible. Normalize henceforth $a_\varnothing = 1$, and put

$$W_f = \{\{p, q\} : a_{pq} \neq 0\}.$$

Step 2: a nonempty W_f factors f. For deleted i, j and a disjoint edge $pq \in W_f$, the coefficients $1 - sa_{ij}$ and $a_{pq} - sa_{ijpq}$ exclude at most one phase each. Thus some contraction contains $e_p + e_q$, and its common support matching contains pq. Deleting two ports outside any two adjacent W_f-edges would put both edges in one matching, so W_f is a matching.

Moreover, every $pq \in W_f$ is a global equality constraint. Indeed, if some supported z had $z_p \neq z_q$, choose equal coordinates i, j outside p, q, possible by even parity and $n \geq 6$. The residual word violates the pq-edge in every nonzero (i,j)-contraction and therefore vanishes at all phases. Two phases of Equation (27) then make both lifts vanish, including z. Consequently

$$pq \in W_f \quad \Longrightarrow \quad \operatorname{supp}(\widetilde{f}) \subseteq \{z : z_p = z_q\}. \tag{31}$$

Fix $E_1 = \{u,v\} \in W_f$ and $R = [n] \setminus E_1$. In an (u,k)-contraction, the all-zero word has the unique lift $x_u = x_k = x_v = 0$; hence, for $p, q \in R \setminus \{k\}$, its coefficient at $e_p + e_q$ is exactly a_{pq}. Its matching therefore consists of the W_f-edges on $R \setminus \{k\}$ and one edge incident with v. Since that last edge cannot cover all $n-2 \geq 4$ ports, the induced matching on R is nonempty. Choose k on one of its edges, with partner ℓ. In the (u,k)-contraction, ℓ has no incident W_f-edge left, so it must be the partner of v. Any vertex unmatched in $W_f|_R$ would also have to be matched to v, which is impossible. Thus $W_f = \{E_1, \ldots, E_r\}$ is perfect, where $r = n/2$.

Equation (31) identifies $\widetilde{f}$ with a coefficient function $h(t_1, \ldots, t_r)$ on the common bits of the E_k's. Contracting one port from E_i and one from E_j imposes $t_i = t_j$, and the remaining partners determine the deleted bits uniquely. The contraction is nonzero since its all-zero coefficient is one. Its support is contained in the matching code obtained by merging E_i, E_j. Both this code and the contraction's matching support have 2^{r-1} points, so they coincide. Removing the nowhere-zero unary phase on the merged variable therefore gives a full-support rank-one diagonal slice of h. Every assignment in $\mathbb{F}_2^r$ has two equal coordinates, so these restrictions cover $\mathbb{F}_2^r$ and h has full support. After $h(0) = 1$, Lemma 5.21 factors h when $r \geq 4$. If $r = 3$, use Equation (30). Let $H_i(s)$ be the two-bit coefficient function obtained by closing E_i at phase s. For $\{i,j,k\} = [3]$,

$$\det H_i(s) = \rho_j \rho_k (\kappa - 1)(1 - \kappa s^2 \rho_i^2).$$

If $\kappa \neq 1$, the contraction is nonzero and hence rank one for every phase, whereas $s \mapsto s^2$ permutes the odd-order group μ_N; the last factor cannot vanish at two phases. Thus $\kappa = 1$, and the six-port table also factors. Both cases contradict tensor-primeness.

Step 3: an empty W_f is impossible. Fix i, j. Every contraction now has constant coefficient one and is nonzero. Choose two edges e, e' of its common residual matching. Full matching support gives $a_{\{i,j\}\cup e} a_{\{i,j\}\cup e'} \neq 0$, while its two-bit rank-one identity is

$$a_{e \cup e'} - s a_{\{i,j\} \cup e \cup e'} = s^2 a_{\{i,j\}\cup e} a_{\{i,j\}\cup e'} \qquad (s \in \mu_N).$$

This nonzero quadratic cannot vanish at the $N \geq 3$ distinct phases. The contradiction completes the proof. $\square$

5.4.4 The quaternary cyclic case

We now treat arity four without assuming that the quaternary is a mountain. The following lemma gives the output obtained by attaching binary representatives to the four ports.

Lemma 5.23 (Attaching cyclic binary representatives to four ports)**.** *Let $N \geq 3$, let Λ be a finite algebraic signature set in K-coordinates, and, for every $s \in \mu_N$, fix an actual ordered binary representative C_s over Λ, in a specified orientation, whose Walsh matrix is*

$$\widetilde{C}_s = \kappa_s D_s Z, \qquad \kappa_s \in \overline{\mathbb{Q}}^\times.$$

Let $q \in \Lambda$ be quaternary. For $r_1, r_2, r_3, r_4 \in \mu_N$, attach C_{r_i} to port i of q through one native X-edge and leave the other port of each representative exposed. After division by the recorded nonzero scalar $2^4 \prod_{i=1}^4 \kappa_{r_i}$, the resulting actual output h satisfies

$$\widetilde{h}_y = \widetilde{q}_y \prod_{i=1}^4 r_i^{y_i}.$$

Consider the map

$$(r_1, r_2, r_3, r_4) \longmapsto (R, r_{12}, r_{13}, r_{14}), \qquad R = \prod_{i=1}^{4} r_i, \quad r_I = \prod_{i \in I} r_i,$$

Write $(\mu_N)^2 := \{z^2 : z \in \mu_N\}$*. A target* $(R, u, v, w) \in \mu_N^4$ *occurs exactly when*

$$\frac{uvw}{R} \in (\mu_N)^2. \tag{32}$$

For odd N*, this map is bijective.*

Proof. This is [LM26a, Lemma 7.3]. □

Theorem 5.24 (The quaternary odd-cyclic case)**.** *Let* Γ *be a finite preprocessed algebraic factor-closed signature set containing the transported height-one mountain. Suppose every nonzero directly realizable ordered binary is nonsingular and has finite projective transfer order, and, in the fixed Walsh coordinates,*

$$P^{-1}G(\Gamma)P = C_N = \{[\mathrm{diag}(1, s)] : s \in \mu_N\}, \qquad N \geq 3 \text{ odd}.$$

Let $0 \neq q \in \Gamma$ *be quaternary. For every ordered pair of distinct ports of* q*, every fixed direct representative of a class in* $G(\Gamma)$*, and either kernel orientation, assume that the resulting two-port contraction is zero or all of its positive-arity tensor-prime factors are nonsingular binaries* B *with*

$$[P^{-1}(BX)P] \in C_N.$$

Then at least one of the following holds.

1. K-Holant(Γ) *is* #P*-hard or* Tract($K\Gamma$) *holds.*
2. *Up to a permutation of its ports,*

$$q \doteq a \otimes b,$$

where a, b *are nonzero positive-arity signatures of arity strictly below four, and*

$$\text{K-Holant}(\Gamma, a, b) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

Proof. The stated contractions satisfy condition (C) for the cyclic group. In the Walsh tensor $\widetilde{q}$, put

$$p = \widetilde{q}_{0000}, \qquad t = \widetilde{q}_{1111}, \qquad x_I = \widetilde{q}_{\mathbf{1}_I} \quad (|I| = 2).$$

For each such I, put $B_I = (x_I, x_{I^c})$, so B_{I^c} has the same two coefficients in reverse order. By Lemma 5.10, $\widetilde{q}$ is even; the shared family Lemma 5.12 excludes exactly one nonzero Walsh endpoint. Suppose first that $p, t \neq 0$, and put $\lambda = t/p$. For each $B_I \neq (0,0)$, it gives $(x_I, x_{I^c}) = (-p\lambda\alpha_I^{-1}, -p\alpha_I)$. For each zero pair choose $\alpha_I \in \mu_N$ arbitrarily, using one index from each complementary pair as in Lemma 5.12. By Lemma 5.23, prescribe $R = \lambda^{-1}$ and $r_I = \alpha_I\lambda^{-1}$. Let h be the normalized representative of the gadget output, obtained by division by the recorded nonzero gadget scalar. Its Walsh coefficient array $\widetilde{h}$ has endpoints (p, p), value $(-p, -p)$ on each complementary pair with $B_I \neq (0,0)$, and zero on every pair with $B_I = (0,0)$.

Here the eight even-parity Walsh positions form the endpoint pair and three central pairs $\{\mathbf{1}_I, \mathbf{1}_{I^c}\}$, where $|I| = 2$. In the inverse Walsh formula of Equation (1), a pair $\{y, \bar{y}\}$ with common coefficient b contributes

$$b\big((-1)^{x \cdot y} + (-1)^{x \cdot \bar{y}}\big).$$

This vanishes when $\mathrm{wt}(x)$ is odd. At either endpoint all signs are positive. At $x = \mathbf{1}_J$, $|J| = 2$, the endpoint pair and the pair $\{\mathbf{1}_J, \mathbf{1}_{J^c}\}$ have positive signs, and the other two central pairs have negative signs. Thus, if exactly k of the three pairs B_I are nonzero,

$$h(\mathbf{0}^4) = h(\mathbf{1}^4) = 2p(1-k),$$
$$h(\mathbf{1}_J) = \begin{cases} 2p(k-1), & \text{if } B_J \neq (0,0), \\ 2p(k+1), & \text{if } B_J = (0,0). \end{cases}$$

This computes the K-coordinate entries of the same normalized output; Fourier inversion is not an additional gadget operation. Thus $k = 1$ gives a nonzero multiple of $X \otimes X$, up to a port permutation, whose factorization pulls back through the invertible binary attachments just constructed. For $k = 0$, the original Walsh tensor is already a higher-arity GE, so Lemma 3.15 gives outcome 1. For $k = 2, 3$, the output in native coordinates is even and has two nonzero endpoints. Adjoin this actual quaternary and apply Corollary 2.25. Its conclusion transfers to Γ by the gadget reduction and downward inheritance, again giving outcome 1.

It remains that $p = t = 0$. If just one complementary pair has nonzero coefficients, Lemma 3.15 gives outcome 1 directly; no reflection is needed. Otherwise label the three complementary pairs $\{\mathbf{1}_{I_\ell}, \mathbf{1}_{I_\ell^c}\}$, $\ell = 1, 2, 3$. Write each nonzero pair B_I as $\rho_I(1, \alpha_I)$, choose the unique $\beta_I \in \mu_N$ with $\beta_I^2 = \alpha_I$; for a zero pair choose $\beta_I \in \mu_N$ arbitrarily. Apply Lemma 5.23 with $R = 1, r_I = \beta_I$. Again divide the actual output by its recorded nonzero gadget scalar and denote the normalized tensor by h. Its nonzero Walsh coefficient pairs become (σ_I, σ_I), where $\sigma_I = \rho_I \beta_I$; set $\sigma_I = 0$ when $B_I = (0,0)$, and write $\sigma_\ell = \sigma_{I_\ell}$. These attachments again give an invertible operation on each port. Put $c = \sigma_1 + \sigma_2 + \sigma_3$. The same sign calculation, now with zero Walsh endpoints, gives

$$h(\mathbf{0}^4) = h(\mathbf{1}^4) = 2c,$$
$$h(\mathbf{1}_{I_\ell}) = h(\mathbf{1}_{I_\ell^c}) = 2(2\sigma_\ell - c) \qquad (\ell = 1, 2, 3),$$

and all odd-weight entries vanish. If $c \neq 0$, this is a complement-symmetric mountain; after division by $2c$, its endpoints are one and its three central coefficient-pair values are

$$r_\ell = 2\sigma_\ell/c - 1, \qquad 1 + r_1 + r_2 + r_3 = 1 + 2c/c - 3 = 0.$$

Up to the recorded nonzero scalar, this mountain is the output of invertible binary attachments to q, so it may be adjoined by Equation (5). It is endpoint-nondegenerate, and Corollary 2.25 makes the enlarged problem either tractable or #P-hard. Downward inheritance in the first case and the gadget reduction in the second transfer that conclusion to Γ, giving outcome 1.

Assume $c = 0$. Exactly one nonzero pair is impossible. Two nonzero Walsh pairs have values $\sigma, -\sigma$; after permuting ports, $\widetilde{h}$ factors as

$$\sigma \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}_{14} \otimes \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}_{23},$$

so the factor-reduction equivalence applies. Finally suppose that all three σ_i are nonzero. Choose distinct i, j, k with $\{i, j, k\} = \{1, 2, 3\}$ and $\sigma_j^2 + \sigma_k^2 \neq 0$, which exists since otherwise subtraction would give $2\sigma_i^2 = 0$. Put pair i at $12 \mid 34$, and join ports $3, 4$ of one copy to ports $1, 2$ of another by native edges. The actual output is

$$\frac{1}{4}\widetilde{G}_i = \begin{pmatrix} \sigma_i^2 & 0 & 0 & 0 \\ 0 & -(\sigma_j^2 + \sigma_k^2) & -2\sigma_j\sigma_k & 0 \\ 0 & -2\sigma_j\sigma_k & -(\sigma_j^2 + \sigma_k^2) & 0 \\ 0 & 0 & 0 & \sigma_i^2 \end{pmatrix}.$$

If every phase contraction of G_i were zero or had transfer class in C_N in Walsh coordinates, its equal nonzero endpoints and Lemma 5.12 would give $\lambda = 1$; complement symmetry and odd N then force every nonzero central coefficient to be $-\sigma_i^2$. In particular,

$$\sigma_j^2 + \sigma_k^2 = \sigma_i^2 = (\sigma_j + \sigma_k)^2,$$

because $\sigma_i = -(\sigma_j + \sigma_k)$, contradicting $\sigma_j \sigma_k \neq 0$. Hence an actual phase contraction of G_i is a nonzero binary which is singular or has Walsh transfer class outside C_N. A singular binary gives unary factors and the odd-arity theorem; a nonsingular binary of infinite projective transfer order is handled by Lemma 3.1. Otherwise its transfer has finite projective order, contradicting the definition of $G(\Gamma)$. Every hardness or tractability conclusion above gives outcome 1 through the cited theorem and the equivalent-enlargement reduction. In a factor case, the attached binaries act invertibly, so the displayed factorization pulls back to q; the transported mountain excludes the common-one-sided exception, and the factor-reduction equivalence gives the reduction in outcome 2. □

Completion of Theorem 5.18. Apply Lemma 5.8. Its first two outcomes are the first two conclusions of the theorem. In outcome 3, the listed two-port contractions and all their positive-arity tensor-prime factors are retained in the equivalent set supplied by outcome 3. Other fixed gadget outputs and any required factors are adjoined separately by Equation (5) and Theorem 2.14, as in the quaternary closure proof. If a nonbinary prime remains, choose one of minimum arity. The argument preceding Lemma 5.19 proves condition (C) for this signature. Theorem 5.22 bounds its arity by four, and Theorem 5.24 then gives hardness, tractability, or a tensor factorization. The last possibility contradicts the choice of a tensor-prime signature, so outcome 1 holds.

In the remaining case every tensor-prime signature is binary. By Lemma 5.8(3), each has transfer class in $G(\Gamma)$ and Walsh tensor proportional to $D_r Z$. Hence Equation (16) holds, which proves the third conclusion of Theorem 5.18. □

Theorem 5.25 (Hardness, tractability, or a strict update)**.** *Let Γ_0 be the initial signature set, and let (A, Γ) be a stage of the sequence in Lemma 2.22; in particular,*

$$A\Gamma_0 \subseteq \Gamma, \qquad A^{\mathsf{T}} X A = \mu_A X \quad (\mu_A \neq 0), \qquad \text{K-Holant}(\Gamma) \equiv_{\mathrm{T}} \text{K-Holant}(\Gamma_0).$$

Suppose Γ is finite, preprocessed, algebraic, and factor-closed, and contains the height-one mountain

$$B = \begin{pmatrix} c_0 & u \\ v & c_1 \end{pmatrix}, \qquad c_0 c_1 \neq 0.$$

Then one of the following occurs:

1. K-Holant(Γ) *is* #P*-hard, or* Tract$(K\Gamma)$ *holds;*
2. *there are $A^+ \in GO(X)$ and a finite preprocessed algebraic factor-closed set Γ^+, continuing the same coordinate-transport sequence, such that*

$$A^+\Gamma_0 \subseteq \Gamma^+, \qquad (A^+)^{\mathsf{T}} X A^+ = \mu_{A^+} X \quad (\mu_{A^+} \neq 0), \qquad \text{K-Holant}(\Gamma^+) \equiv_{\mathrm{T}} \text{K-Holant}(\Gamma),$$

and, for $D = A^+ A^{-1}$, the transported binary DBD^{T} remains in Γ^+ up to its known nonzero scalar. In this alternative every nonzero directly realizable ordered binary over both Γ and Γ^+ is nonsingular with finite projective transfer order, so $G(\Gamma)$ and $G(\Gamma^+)$ are defined. Moreover, either

$$[A]^{-1} G(\Gamma) [A] \subsetneq [A^+]^{-1} G(\Gamma^+) [A^+],$$

or both values of ν *are defined and*

$$\nu(\Gamma^+) < \nu(\Gamma).$$

Both quantities satisfy the non-strict comparisons and fixed-field conditions of Definition 5.5.

Proof. Maintain factor closure using the factor-reduction equivalence. Each nontrivial factorization produces positive-arity factors of strictly smaller arity, so this process terminates. An odd-arity member is handled by the odd-signature theorem; a rank-one binary supplies unaries and hence the same conclusion. If no nonbinary tensor-prime signature remains, the arity-at-most-two alternative gives tractability.

All equivalent augmentations and global transformations satisfy Lemma 2.22. Whenever the groups are defined, their pullbacks are nested. The old tensor-prime signatures, up to one common invertible transformation, remain in the set, so ν cannot increase. A strict group inclusion or decrease of ν gives outcome 2. An augmentation preserving both quantities is retained with the same matrix and group; Lemma 5.8 proves separately that the required recursive contraction and factor extraction is finite.

If B is singular, or if BX is nonsingular of infinite projective order, Lemma 3.1 supplies the rank-one binary used in the preceding paragraph. The same argument applies to every other directly realizable binary that is singular or whose nonsingular transfer has infinite projective order. We may therefore assume that every nonzero directly realizable binary is nonsingular and of finite projective transfer order. Then Lemma 5.3 and Theorem 5.4 give either a nonsingular weighted equality in the current set or, after one legal $GO(X)$-normalization, an odd cyclic or odd dihedral residual form. In the equality case, write $E = \operatorname{diag}(p, q) \in \Gamma$, where $pq \neq 0$, and choose a diagonal $D \in GO(X)$ with $DED^{\mathsf{T}} \doteq I$. The resulting legal global normalization satisfies the hypotheses of Theorem 2.23. Its hard conclusion composes back to Γ, while its tractable conclusion pulls back through the problem equivalence and Lemma 2.22; hence this is outcome 1.

Apply Theorems 5.7 and 5.18 in the two residual forms. Their first conclusions give outcome 1, and their updates satisfy outcome 2 by Definition 5.5. Their remaining conclusions give Equation (16), which Lemma 2.22 transports back to $\mathsf{Tract}(K\Gamma)$. This again gives outcome 1 and completes the case analysis. □

Proof of Theorem 5.1. Maintain (A, Γ) and the invariant of Lemma 2.22 over the fixed initial number field, and apply Theorem 5.25. Its first outcome is the required hardness or tractability conclusion. The second outcome gives an update satisfying Definition 5.5; by Lemma 5.6, only finitely many such updates are possible. Between these updates, factor closure terminates by strict arity descent, and the recursive contractions and factor extraction terminate by Lemma 5.8. Repeated application therefore reaches the first outcome. Finally, Lemma 2.22 transports the conclusion back to the initial signature set. □

6 Mountain Height Two

A height-two mountain is a quaternary signature with $M(0000)M(1111) \neq 0$ that vanishes on odd Hamming weight. It is an eight-vertex signature with those nonzero endpoints. If it is tensor-prime, Theorem 2.24 applies. Otherwise the parity and endpoint conditions force a $2 + 2$ product of nonsingular weighted equalities; the factor-reduction equivalence places such an equality in the current signature set. A legal diagonal $GO(X)$-normalization then sends it to I. Then Theorem 2.23 closes that case. Both arguments are included in Corollary 2.25.

Theorem 6.1 (Dichotomy with a height-two mountain)**.** *Let Γ be a finite algebraic signature set containing a height-two mountain. Then*

$$\mathsf{Tract}(K\Gamma) \qquad \textit{or} \qquad \text{K-Holant}(\Gamma) \textit{ is } \#\mathsf{P}\textit{-hard.}$$

Proof. The mountain is an eight-vertex signature with two nonzero endpoints, so apply Corollary 2.25. Hardness already reduces to K-Holant(Γ), while tractability is the predicate required here. ☐

7 Mountain Height Three

Theorem 7.1 (Height-three mountain dichotomy)**.** *Let Γ be a finite algebraic set of even-arity signatures such that $\chi_r(x) \in 3\mathbb{Z}$ for every arity-r member $f \in \Gamma$ and every $x \in \operatorname{supp}(f)$. Suppose that Γ contains a mountain*

$$M = a\delta_{\mathbf{0}^6} + b\delta_{\mathbf{1}^6} + Q, \qquad ab \neq 0, \qquad Q \textit{ EO.}$$

Then $\mathsf{Tract}(K\Gamma)$ *holds or* K-Holant(Γ) *is* $\#\mathsf{P}$*-hard.*

The aim is to adjoin a quaternary EO signature outside the binary matching-product class. Then Theorem 3.7 gives the conclusion for the entire signature set. A two-copy construction handles a matching-product central layer; the other central layers are covered by the six-port loop and two-copy calculations below.

7.1 Matching Products and Two-Copy Constructions

Let V be an ordered set of even cardinality. For distinct $i, j \in V$, write $\partial_{ij} f$ for the contraction with a native X-edge

$$(\partial_{ij} f)(x_{V\setminus\{i,j\}}) = f(0, 1, x_{V\setminus\{i,j\}}) + f(1, 0, x_{V\setminus\{i,j\}}).$$

We also write ∂_e when $e = \{i, j\}$. Define the matching-product class

$$\mathcal{B}_\times(V) = \left\{ \lambda \bigotimes_{\{u,v\}\in P} \begin{pmatrix} 0 & p_{uv} \\ q_{uv} & 0 \end{pmatrix}_{u,v} : \begin{array}{l} P \text{ is a perfect matching of } V, \\ \lambda \prod_{\{u,v\}\in P} p_{uv} q_{uv} \neq 0 \end{array} \right\},$$

and put $\mathcal{B}_\times^0 = \mathcal{B}_\times \cup \{0\}$. Each factor has a specified orientation; reversing its ports exchanges its two nonzero entries. In particular, $\mathcal{B}_{\mathrm{ARS}} \subseteq \mathcal{B}_\times$. An *$X$-matching product* is a nonzero scalar multiple of $\bigotimes_{e\in P} X_e$. Such binary factors are tensor-prime, so uniqueness of tensor-prime factorization identifies both their port pairs and their projective values. We abbreviate this uniqueness as UPF.

Lemma 7.2 (A matching-product central layer yields a hard quaternary)**.** *Let $d \geq 3$, and let Γ be a finite algebraic signature set containing*

$$M = a\delta_{\mathbf{0}^{2d}} + b\delta_{\mathbf{1}^{2d}} + Q, \qquad ab \neq 0,$$

where, for a perfect matching $P = \{\{u_1, v_1\}, \ldots, \{u_d, v_d\}\}$ of the ordered ports,

$$Q = \lambda \bigotimes_{k=1}^{d} \begin{pmatrix} 0 & p_k \\ q_k & 0 \end{pmatrix}_{u_k, v_k}, \qquad \lambda \prod_{k=1}^{d} p_k q_k \neq 0.$$

There is a fixed two-copy gadget over $\{M\}$, using only native X-edges, whose ordered boundary signature h is EO and has the six weight-two entries

$$(h_{0011}, h_{0101}, h_{0110}, h_{1001}, h_{1010}, h_{1100}) = (ab, c_1, c_2, c_3, c_4, ab), \qquad c_1c_2c_3c_4 \neq 0. \tag{WT.18a}$$

In particular,

$$\text{K-Holant}(\Gamma, h) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

The signature h is not a tensor product of two nonsingular anti-diagonal binaries. Moreover, its supported triples include

$$1100 \oplus 1010 \oplus 0110 = 0000, \qquad 0011 \oplus 0101 \oplus 1001 = 1111.$$

The first identity rules out the positive sign and the second rules out the negative sign in the support clause of $\mathsf{EOTract}$*. Hence* $\neg\mathsf{EOTract}(\{h\})$*, and Theorem 2.11 makes* K-Holant$(\{h\})$ #P*-hard. Since*

$$\text{K-Holant}(\{h\}) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma, h) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma),$$

the problem K-Holant(Γ) *is* #P*-hard.*

Proof. Write the matching of Q as $P = \{\{w_0, w_1\}, \ldots, \{w_{2d-2}, w_{2d-1}\}\}$. On two copies of the mountain, define a port bijection p by fixing w_0, w_{2d-1} and transposing

$$(w_1, w_2), (w_3, w_4), \ldots, (w_{2d-3}, w_{2d-2}).$$

Leave (w_0, w_1) external on the first copy and (w_0, w_2) external on the second, in the external order $(w_0^{(1)}, w_1^{(1)}, w_0^{(2)}, w_2^{(2)})$. For every $r \notin \{0, 1\}$, join the first-copy port w_r to the second-copy port $p(w_r)$ by one native X-edge. With both copies central, the matching edges and wires form two alternating components: the edge between the first external pair and a path of length $4d-3$ between the second. Both lengths are odd, so each weight-one/weight-one boundary pattern has a unique extension and a nonzero coefficient c_i.

If both copies use endpoints, the native wires force the two constant words to be complementary. This gives exactly the other two EO patterns, with unique extensions of value ab. A mixed endpoint–central extension is impossible: for $d \geq 3$, some central matching edge has both ports wired to the constant copy and is forced equal. Thus each nonzero coefficient in *(WT.18a)* has exactly one nonzero extension.

The six-word support excludes membership in $\mathcal{B}_\times$. Moreover

$$1100 \oplus 1010 \oplus 0110 = 0000, \qquad 0011 \oplus 0101 \oplus 1001 = 1111.$$

These two supported triples rule out the two common signs in the support clause of $\mathsf{EOTract}$. The weighted EO dichotomy therefore makes the output hard. By the equivalence convention, the output may be adjoined to the signature set containing the mountain. $\square$

7.2 Six-Port Loop and Two-Copy Contractions

The finite calculation for both families of contractions is recorded in the appendix; the reduction uses only its stated conclusion.

Definition 7.3 (Loop and two-copy contractions of a six-port mountain)**.** Let $U = [6]$, and let

$$M = a\delta_{\mathbf{0}^6} + b\delta_{\mathbf{1}^6} + Q, \qquad Q = \sum_{S \in \binom{U}{3}} q_S \delta_{\mathbf{1}_S}, \qquad ab \neq 0.$$

For each unordered pair $\{i,j\} \in \binom{U}{2}$, let $L_{ij} = \partial_{ij} M$ be the four-port EO signature obtained by one native X-self-loop, with the residual ports in increasing order. For $R \in \binom{U \setminus \{i,j\}}{2}$,

$$L_{ij}(\mathbf{1}_R) = q_{R \cup \{i\}} + q_{R \cup \{j\}}. \tag{Q6.3}$$

These are the fifteen one-copy contractions.

For a two-copy contraction, choose unordered two-element external sets $A, B \in \binom{U}{2}$, give them their canonical increasing orders

$$A = (a_0, a_1), \qquad B = (b_0, b_1),$$

and retain the first/second labels of the two copies of M. Put $I_A = U \setminus A$ and $I_B = U \setminus B$, and choose a bijection $\pi : I_A \to I_B$. Leave A external on the first copy and B external on the second copy, in the order (a_0, a_1, b_0, b_1), and join every $r \in I_A$ to $\pi(r) \in I_B$ by a native X-edge. For $x, y \in \{0,1\}^2$, put

$$A_x = \{a_r : x_r = 1\}, \qquad B_y = \{b_r : y_r = 1\}.$$

The resulting four-port EO signature $C_{A,B,\pi}$ has, when $\mathrm{wt}(x) + \mathrm{wt}(y) = 2$,

$$C_{A,B,\pi}(x,y) = ab\, \mathbf{1}[(x,y) \in \{(00,11),(11,00)\}] + \sum_{\substack{T \subseteq I_A \\ |T| = 3 - \mathrm{wt}(x)}} q_{A_x \cup T}\, q_{B_y \cup (I_B \setminus \pi(T))}. \tag{Q6.4}$$

There are $15^2 \cdot 4! = 5400$ canonically labelled two-copy contractions. The right-hand side of (Q6.4) is the coefficient of the complete direct contraction; no individual summand is treated as a separately realizable signature.

Lemma 7.4 (Support forced by six-port contractions)**.** *Let M, L_{ij}, and $C_{A,B,\pi}$ be as in Definition 7.3. Suppose that every one of the fifteen signatures L_{ij} and every one of the 5400 canonically labelled signatures $C_{A,B,\pi}$ belongs to $\mathcal{B}^0_\times([4])$ on its four ordered external ports. Then*

$$\mathrm{supp}(Q) \subseteq \{S, U \setminus S\} \qquad \textit{for some } S \in \binom{U}{3}. \tag{Q6.5}$$

This holds for arbitrary complex values of the q_S and every $ab \neq 0$.

Proof. With $t = ab$, the signatures L_{ij} and $C_{A,B,\pi}$ in Equations (Q6.3) and (Q6.4) are precisely the loop and labelled two-copy arrays covered by Theorem A.6. Its conclusion is *(Q6.5)*. ☐

Theorem 7.5 (A non-matching quaternary output from a six-port mountain)**.** *Let Γ be a finite algebraic signature set containing*

$$M = a\delta_{\mathbf{0}^6} + b\delta_{\mathbf{1}^6} + Q, \qquad ab \neq 0, \qquad Q \textit{ EO}.$$

There is a nonzero quaternary EO signature $h \notin \mathcal{B}_\times$ such that

$$\text{K-Holant}(\Gamma, h) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma).$$

The signature h can be chosen as one of the fifteen one-copy contractions L_{ij}, one of the 5400 labelled two-copy contractions $C_{A,B,\pi}$ from Definition 7.3, or the explicit two-copy output of Lemma 7.2.

Proof. We select the signature from one contraction L_{ij} or one contraction $C_{A,B,\pi}$ when $Q \notin \mathcal{B}_\times$, and from the explicit two-copy output of Lemma 7.2 otherwise. If $Q = 0$, a contraction $C_{A,B,\pi}$ has exactly the two endpoint–endpoint coefficients ab, hence support two rather than four; adjoin this contraction to Γ.

Suppose $Q \neq 0$. If a loop or two-copy contraction lies outside $\mathcal{B}^0_\times$, adjoin it to Γ. Otherwise Lemma 7.4 gives $\mathrm{supp}(Q) \subseteq \{S, U \setminus S\}$. Choose $i \in S$ and $j \notin S$. The two possible nonzero coefficients of $\partial_{ij} M$ have the unique extensions $S, U \setminus S$; hence this contraction is nonzero with support at most two, contradicting its membership in $\mathcal{B}^0_\times$.

In particular, if

$$Q = \lambda \bigotimes_{r=0}^{2} \begin{pmatrix} 0 & p_r \\ q_r & 0 \end{pmatrix}_{w_{2r}, w_{2r+1}}, \qquad \lambda p_0 q_0 p_1 q_1 p_2 q_2 \neq 0,$$

the case $d = 3$ of Lemma 7.2 gives the six-word output with one nonzero extension per supported word *(WT.18a)*, which belongs to Γ by Equation (5). □

Proof of Theorem 7.1. Apply zero/nullary preprocessing and assume $\neg\mathsf{Tract}(K\Gamma)$. By Theorem 7.5, a finite equivalent enlargement contains a nonzero quaternary EO signature $h \notin \mathcal{B}_\times$. It has charge zero, so the enlargement preserves the charge hypothesis; it retains the two-sided mountain and remains nontractable by Lemma 2.20. Since $\mathcal{B}_{\mathrm{ARS}} \subseteq \mathcal{B}_\times$, Theorem 3.7 proves hardness, which transfers to the original Γ. □

8 Mountain Height at Least Four

Theorem 8.1 (Mountain dichotomy at every height at least four)**.** *Let $d \geq 4$, and let Γ be a finite algebraic set of even-arity signatures such that*

$$\chi_r(x) \in d\mathbb{Z} \qquad (f \in \Gamma,\ \mathrm{arity}(f) = r,\ x \in \mathrm{supp}(f)).$$

Suppose that Γ contains a mountain

$$M = a\delta_{\mathbf{0}^{2d}} + b\delta_{\mathbf{1}^{2d}} + Q, \qquad ab \neq 0, \qquad Q \text{ EO}.$$

Then $\mathsf{Tract}(K\Gamma)$ holds or K-Holant(Γ) *is* #P*-hard.*

The common step attaches extracted binary signatures to selected ports so that every self-loop contraction becomes an X-matching product. A single scalar then makes the central tensor satisfy ARS. For $d \geq 5$, the matching-product reconstruction of Shao and Cai makes Lemma 7.2 applicable. We prove this general case first. At height four, a classification of the matching family gives the Reed–Muller central layer. We then normalize the endpoint coefficients and obtain equality by a Turing reduction.

Remark 8.2 (Relation to the real-valued Holant argument)*.* The port normalization below follows the contraction and binary-attachment strategy of [SC20, full version, Lemma 8.1]. Their transformed ambient language satisfies ARS globally. Here the ambient language is complex-valued: charge conservation makes the self-loop outputs EO, and Theorem 3.7 forces their binary ARS factorization unless hardness has already been proved. ARS alone implies neither EO support nor this factorization.

We therefore prove the relative normalization and phase recovery in Lemmas 8.5 and 8.6, before applying the algebraic reconstruction restated in Lemma 8.7. The central layer is not adjoined as a signature. Hardness comes from our actual mountain gadgets and, at height four, the relative Bell reduction; it does not require extending the global-ARS hardness arguments to the whole language.

8.1 Common Normalization and Phase Recovery

For the proof of Theorem 8.1, apply zero/nullary preprocessing and assume $\neg\mathsf{Tract}(K\Gamma)$. If $Q = 0$, Lemma 3.14 already gives hardness. Hence it suffices to consider $Q \neq 0$. Every native self-loop on the complete mountain kills its endpoints:

$$\partial_{ij} M = \partial_{ij} Q. \tag{HN.1}$$

More generally, charge conservation makes every native gadget output of positive arity less than $2d$ EO. By Theorem 3.7, any such nonzero output outside $\mathcal{B}_{\mathrm{ARS}}$ gives hardness. Otherwise its factors are nonsingular binaries

$$B(z) = \begin{pmatrix} 0 & z \\ \bar{z} & 0 \end{pmatrix}, \qquad z \in \overline{\mathbb{Q}}^{\times},$$

up to nonzero scalars. These factors may be adjoined using Theorem 2.14: the retained mountain is two-sided, excluding the common-one-sided exception. We absorb each finite equivalent enlargement into Γ by Equation (5). The extracted binaries have charge zero, and Lemma 2.20 preserves nontractability.

Attaching the first port of an ordered $B(z)$ to a mountain port through X acts on that port by

$$T_z = B(z)^{\mathsf{T}} X = \mathrm{diag}(\bar{z}, z), \qquad T_z B(z) = |z|^2 X. \tag{HN.2}$$

This invertible diagonal map preserves support and both nonzero endpoints. For consistently oriented factors,

$$B(z)^{\mathsf{T}} X B(w) \doteq X \iff w/z \in \mathbb{R}^{\times} \iff B(w) \doteq B(z). \tag{HN.3}$$

Indeed the two off-diagonal entries are $\bar{z}w$ and $z\bar{w}$. In this case $T_z^{\mathsf{T}} X T_w$ is also a nonzero scalar multiple of X. All these operations use the extracted ordered binaries on the complete mountain. The conjugate entry in $B(z)$ is already part of that binary; no conjugate mountain is required.

Lemma 8.3 (Vanishing pair contractions)**.** *Let h have arity at least three. If $\partial_{ij} h = 0$ for every pair $i \neq j$, then $h(x) = 0$ whenever $0 < \mathrm{wt}(x) < \mathrm{arity}(h)$. In particular, an EO signature with all pair contractions zero is zero.*

Proof. This is [CFS20a, Lemma 2.7]; its three-equation argument is purely linear and uses neither ARS nor nonvanishing contractions. ☐

Lemma 8.4 (Propagation of X-matching products)**.** *Let F be EO on an even port set V, with $|V| \geq 8$, and suppose every nonzero $\partial_e F$ lies in $\mathcal{B}_{\mathrm{ARS}}$. Either of the following conditions implies that every $\partial_e F$ is an X-matching product.*

1. *Every $\partial_e F$ with $e \cap \{1, 2\} = \varnothing$ is an X-matching product.*
2. *Both $\partial_{12} F$ and $\partial_{ij} F$ are X-matching products, $\{1, 2\} \cap \{i, j\} = \varnothing$, and the matching of $\partial_{ij} F$ pairs $1, 2$ with distinct mates u, v.*

The remaining pair contractions are allowed to be zero in the hypotheses.

Proof. Contracting any pair in an X-matching product gives another nonzero X-matching product: it either removes a matching edge with scalar 2, or splices two edges with scalar 1. Thus, whenever a deleted pair e is disjoint from a pair r for which $\partial_r F$ is an X-matching product,

$$0 \neq \partial_e \partial_r F = \partial_r \partial_e F. \tag{HN.4}$$

In particular $\partial_e F \neq 0$. A factor whose ports avoid r survives in this double contraction, so UPF makes it projectively X.

For condition 1, only contractions with $e \cap \{1,2\} \neq \varnothing$ remain. A pair r disjoint from e, with $\partial_r F$ an X-matching product, first proves nonvanishing. For each binary factor on ports a, b, choose a pair r avoiding $\{1,2\} \cup e \cup \{a,b\}$. This excludes at most five ports, so such an r exists and $\partial_r F$ is an X-matching product. The factor survives in (HN.4), proving condition 1.

For condition 2, put $W = V \setminus \{1,2,i,j\}$, where $u, v \in W$ and $|W| \geq 4$. For $e \in \binom{W}{2} \setminus \{\{u,v\}\}$, nonvanishing follows using the 12-contraction. The ports $1,2$ cannot be paired in $\partial_e F$: that pair would survive deletion of ij, whereas deleting e from the X-matching product $\partial_{ij} F$ pairs $1,2$ only if $e = \{u,v\}$. Consequently write

$$\partial_e F = C(x_1, x_r) D(x_2, x_s) H.$$

The 12-double contraction makes H an X-product and the splice of C, D projectively X. By (HN.3), $C \doteq D$ in these orientations. Both factors cannot be affected by deletion of ij, since this would splice $1,2$ together, contradicting the other contraction order. A surviving factor is projectively X by comparison with the X-matching product $\partial_{ij} F$. Hence $C, D \doteq X$, and the whole contraction is an X-matching product.

For $e = \{u,v\}$, choose two distinct ports $r, s \in W \setminus \{u,v\}$. The contractions at the three disjoint pairs $12, ij, rs$ are now X-matching products. The uv-contraction is nonzero by the 12-comparison. Each of its binary factors avoids at least one of these three pairs, and hence is projectively X by (HN.4). Thus every $\partial_e F$ with $e \subseteq W$ is an X-matching product.

Next let $e = \{i,k\}$ or $\{j,k\}$, with $k \in W$. It is nonzero by the 12-comparison. For a factor on a, b, use $r = 12$ if $\{a,b\}$ avoids 12; otherwise choose

$$r \in \binom{W \setminus (\{k\} \cup \{a,b\})}{2}.$$

At most two ports of W are excluded, so the choice exists. Here r is disjoint from both e and the factor's ports, and $\partial_r F$ is an X-matching product. The factor remains in (HN.4), so it is proportional to X. Along with the already established X-matching product $\partial_{ij} F$, this covers all deleted pairs avoiding 12. Condition 1 completes the proof. $\square$

Lemma 8.5 (Relative port normalization at every arity at least eight)**.** *Under the hypotheses of Theorem 8.1, assume* $\neg\mathsf{Tract}(K\Gamma)$ *and* $Q \neq 0$. *Then either* K-Holant(Γ) *is* #P*-hard, or a finite equivalent enlargement of* Γ*, still supported on half-charges in* $d\mathbb{Z}$*, contains an actual mountain*

$$M^\circ = A\delta_{\mathbf{0}^V} + B\delta_{\mathbf{1}^V} + F, \qquad |V| = 2d, \quad AB \neq 0, \quad 0 \neq F \ \mathit{EO},$$

such that

$$\partial_e F = \lambda_e \bigotimes_{t \in P_e} X_t, \qquad \lambda_e \in \overline{\mathbb{Q}}^\times \qquad \left(e \in \binom{V}{2}\right). \tag{HN.5}$$

Here P_e *is a perfect matching of* $V \setminus e$. *No nonvanishing assumption is imposed on the original contractions of* Q.

Proof. At every stage retain the complete mountain, adjoin the finitely many contractions or factors used, and apply Theorem 3.7 again. The charge and nontractability invariants described above remain valid. Thus every current nonzero contraction is an ARS matching product, or hardness has already been proved.

By Lemma 8.3, some contraction of Q is nonzero; relabel it as $\partial_{12}Q$. Extract its $d-1$ binary factors and extend one end of each, on the complete mountain, by a matching copy as in (HN.2). The resulting mountain M_0 has nonzero endpoints and central layer F_0, with $\partial_{12}F_0$ an X-matching product. For every pair e avoiding 12, $\partial_{12}\partial_e F_0 = \partial_e\partial_{12}F_0 \neq 0$. Hence all these contractions are nonzero.

Suppose first that $1, 2$ form a matching edge in every such contraction. Write $\partial_e F_0 = C_e(x_1, x_2)H_e$. The 12-comparison shows that the contraction of C_e is nonzero and H_e is an X-product. If e, f are disjoint pairs avoiding 12, their double contraction is nonzero: in the e-order, deletion of f acts only on the X-product H_e. The factor on 12 survives in both orders, so $C_e \doteq C_f$. The graph of disjoint pairs on $V \setminus \{1, 2\}$ is connected. Indeed, intersecting distinct pairs exclude only three of its at least six ports, leaving a pair disjoint from both. Thus all C_e are projectively one ordered factor C. Extract C from one actual contraction and extend port 1 of M_0 by a matching copy. Every contraction whose deleted pair avoids 12 is now an X-matching product. After reapplying Theorem 3.7, condition 1 of Lemma 8.4 makes every pair contraction an X-matching product.

Otherwise choose a pair ij avoiding 12 whose contraction has distinct mates u, v for $1, 2$, and write

$$\partial_{ij}F_0 = C(x_1, x_u)D(x_2, x_v)H.$$

The 12-comparison makes H an X-product and gives $C \doteq D$ by (HN.3). Extract both factors from this actual contraction and extend ports $1, 2$ of M_0 by matching copies. The new ij-contraction is an X-matching product. The effective 12-kernel is $T_C^{\mathsf{T}} X T_D \doteq X$, so the 12-contraction remains an X-matching product as well. The diagonal maps preserve the distinct mates of ports $1, 2$. Reapply Theorem 3.7 to the current contractions, then use condition 2 of Lemma 8.4.

All extensions are invertible and preserve the central support and the nonzero endpoints. Their number depends only on the fixed source signature. Factor extraction and direct gadget substitution give the claimed finite equivalence. Only contractions of the central layer are adjoined; the central layer itself is used only in the algebraic argument. □

Lemma 8.6 (A common phase for EO contractions)**.** *Let F be algebraic and EO on an even port set V, with $|V| \geq 6$, and suppose* (HN.5) *holds with every $\lambda_e \neq 0$. There is $c \in \overline{\mathbb{Q}}^{\times}$ such that $G = c^{-1}F$ is real-valued and invariant under bitwise complement, and every coefficient λ_e/c in* (HN.5) *is positive real. In particular G satisfies exact ARS.*

Proof. For disjoint pairs e, f, the two contraction orders give

$$\partial_f\partial_e F = m_{ef}\lambda_e X_{P'}, \qquad \partial_e\partial_f F = m_{fe}\lambda_f X_{P''}, \qquad m_{ef}, m_{fe} \in \{1, 2\}.$$

Here $X_P = \bigotimes_{t\in P} X_t$ has only $0, 1$ entries. Both tensors are nonzero, so commutation identifies their supports and gives

$$\lambda_e/\lambda_f = m_{fe}/m_{ef} \in \{1, 2, 1/2\}.$$

The graph of disjoint pairs on V is connected: two intersecting pairs have a common disjoint pair since $|V| \geq 6$. Fix one pair e_0 and take $c = \lambda_{e_0}$. Every contraction of $G = c^{-1}F$ is therefore real and complement-invariant, with a positive coefficient.

Let $\overline{G}$ denote coefficientwise conjugation and $G^{\mathrm{rev}}(x) = G(\bar{x})$. Contraction commutes with both operations. Hence every pair contraction of $G - \overline{G}$ and of $G - G^{\mathrm{rev}}$ vanishes. Both differences are EO, so Lemma 8.3 makes them zero. This is an algebraic conclusion and does not assert that G is a realizable signature. □

We use precisely the following algebraic part of the real-valued classification.

Lemma 8.7 (Shao–Cai matching-product reconstruction [SC20, full version, Lemma 2.11]). *Let G be an EO tensor of even arity $n \geq 8$ satisfying exact ARS. Suppose*

$$\partial_e G = \mu_e \bigotimes_{t \in P_e} X_t, \qquad \mu_e \in \mathbb{R}^\times \qquad \left(e \in \binom{[n]}{2}\right),$$

where each P_e is a perfect matching of the residual ports. If $n \geq 10$, then G is an X-matching product. The same conclusion holds for $n = 8$ if $P_A \cap P_B \neq \varnothing$ for some $A \neq B$ with $A \cap B \neq \varnothing$; in that case the X-factor on that shared edge is also a factor of G.

8.2 Height at Least Five

Proof of Theorem 8.1 for $d \geq 5$. As above, assume $\neg\mathsf{Tract}(K\Gamma)$ and $Q \neq 0$. Apply Lemma 8.5. Outside its hard outcome there is an actual mountain $M^\circ = A\delta_{\mathbf{0}^{2d}} + B\delta_{\mathbf{1}^{2d}} + F$ with $AB \neq 0$ and every pair contraction of F an X-matching product. By Lemma 8.6, $G = c^{-1}F$ is EO with exact ARS and nonzero real X-product contractions. Since $2d \geq 10$, Lemma 8.7 makes G, and hence F, an X-matching product. Apply Lemma 7.2 to the complete M°. Its two-copy output is hard, and the finite equivalence transfers hardness to the original Γ. $\square$

8.3 Height Four

The eight-port exception is the central-layer Reed–Muller tensor

$$R_8^{\mathrm{cen}}(x) = \mathbf{1}\left[x = (a + u \cdot t)_{t \in \mathbb{F}_2^3} \text{ for some } a \in \mathbb{F}_2,\ 0 \neq u \in \mathbb{F}_2^3\right]. \tag{WT.20}$$

This tensor has fourteen supported words and is not a matching product, but each contraction is an X-product. For deleted ports p, q, the surviving pairs are $\{t, t + p + q\}$: each choice of one endpoint from every remaining pair has a unique extension with $u \cdot (p + q) = 1$. A matching edge can therefore belong to both P_A and P_B, for distinct A, B, only when $A \cap B = \varnothing$, so the eight-port overlap clause of Lemma 8.7 does not apply.

For an eight-port set V and a bijection $\iota : V \to \mathbb{F}_2^3$, write

$$(\iota_* F)((x_t)_{t \in \mathbb{F}_2^3}) = F((x_{\iota(v)})_{v \in V}).$$

All canonical Reed–Muller statements below use this explicit port relabelling.

8.3.1 Translation Form of the Matching Family

Lemma 8.8 (A shared edge forces a matching product). *Let F be an arity-eight EO tensor satisfying* (HN.5). *If $P_A \cap P_B \neq \varnothing$ for some $A \neq B$ with $A \cap B \neq \varnothing$, then F is an X-matching product. Consequently, any finite algebraic signature set containing an actual mountain $M^\circ = A\delta_{\mathbf{0}^8} + B\delta_{\mathbf{1}^8} + F$, with $AB \neq 0$, has a #P-hard K-Holant problem.*

Proof. Apply Lemma 8.6 and then the eight-port clause of Lemma 8.7. This makes F an X-matching product. The complete mountain now satisfies Lemma 7.2. $\square$

Lemma 8.9 (Translation form of the matching family). *Let F be an arity-eight EO tensor on a port set V, and suppose that for every $u \neq v$ one has*

$$\partial_{uv} F = \lambda_{uv} \bigotimes_{e \in P_{uv}} X_e, \qquad \lambda_{uv} \neq 0,$$

where P_{uv} is a perfect matching of $V \setminus \{u, v\}$. Suppose $P_A \cap P_B = \varnothing$ whenever $A \neq B$ and $A \cap B \neq \varnothing$. Then there is a bijection $\iota : V \to \mathbb{F}_2^3$ such that, after relabelling each port v by $\iota(v)$, for $p \neq q$ and $h = p + q$ one has

$$P_{pq} = \{\{t, t+h\} : t \in \mathbb{F}_2^3 \setminus \{p, q\}\}, \tag{Q8.14}$$

where duplicate unordered pairs are suppressed.

Proof. We apply uniqueness of tensor-prime factorization and a finite matching classification. For disjoint deleted pairs e, r, let $\kappa_r(P_e)$ be the two-edge matching left after contracting r in the X-product on P_e: delete r when $r \in P_e$, and otherwise splice the mates of the two ends of r and retain the untouched edge. Both operations are nonzero for X factors. Commutativity of the two loops and unique factorization give the exact constraint

$$\kappa_r(P_e) = \kappa_e(P_r) \qquad (e \cap r = \varnothing). \tag{Q8.14a}$$

The condition on $P_A \cap P_B$ and *(Q8.14a)* are exactly the hypotheses of Theorem A.9, which supplies the required bijection and *(Q8.14)*. Thus the identity obtained by commuting two direct contractions forces the matching family in *(Q8.14)*; an arbitrary partition of the complete graph into perfect matchings would not suffice. □

Lemma 8.10 (The matching family determines the central $RM(1,3)$ coefficients)**.** *Let F be an arity-eight EO tensor on a port set V such that*

$$\partial_{uv} F = \lambda_{uv} \bigotimes_{e \in P_{uv}} X_e, \qquad \lambda_{uv} \neq 0,$$

for every $u \neq v$, where P_{uv} is a perfect matching of $V \setminus \{u, v\}$. Suppose $P_A \cap P_B = \varnothing$ whenever $A \neq B$ and $A \cap B \neq \varnothing$. Choose a bijection $\iota : V \to \mathbb{F}_2^3$ supplied by Lemma 8.9, put $F^\iota = \iota_* F$, and write

$$c_S = F^\iota(\mathbf{1}_S), \qquad S \in \binom{\mathbb{F}_2^3}{4}.$$

Then

$$\operatorname{supp}(F^\iota) = \left\{ \mathbf{1}_S : S \in \binom{\mathbb{F}_2^3}{4},\ \bigoplus_{s \in S} s = 0 \right\}, \tag{Q8.17}$$

the set of the fourteen affine planes. Moreover all fourteen nonzero coefficients in (Q8.17) *are equal.*

Proof. Relabel the ports through ι, continue to denote the resulting tensor by $F^\iota = \iota_* F$, and write $V = \mathbb{F}_2^3$. For $S \in \binom{V}{4}$, $p \in S$, and $q \notin S$, put

$$S' = S \triangle \{p, q\}, \qquad T = S \setminus \{p\}, \qquad \sigma(S) = \bigoplus_{s \in S} s, \qquad h = p + q.$$

The coefficient of the actual contraction at $\mathbf{1}_T$ is the complete two-extension sum

$$(\partial_{pq} F^\iota)(\mathbf{1}_T) = c_S + c_{S'}.$$

The word $\mathbf{1}_T$ selects one endpoint from each edge of the matching *(Q8.14)* if and only if

$$\sigma(S) \in \{0, h\}. \tag{Q8.19}$$

Indeed, in $V/\langle h \rangle$ the fibre $\{p, q\}$ contains one point of S, while the other occupancies are $1, 1, 1$ or a permutation of $2, 1, 0$. Their quotient xor is respectively zero or the sum of the two odd fibres, proving *(Q8.19)*. Coefficients at these words equal λ_{pq} and all others vanish. Thus the c_S and nonzero λ_{pq} satisfy the swap rule in Theorem A.11. That theorem eliminates all coefficients on four-sets that are not affine planes and identifies the fourteen plane coefficients with one $c \neq 0$, proving *(Q8.17)*. $\square$

Equivalently, the central-layer Reed–Muller tensor from *(WT.20)* is

$$R_8^{\mathrm{cen}}(x) = \mathbf{1}\left[\mathrm{wt}(x) = 4, \bigoplus_{t:x_t=1} t = 0\right].$$

It consists of the fourteen nonconstant words of $RM(1,3)$. Consequently the full core defined in Equation (8) is

$$\mathrm{RM}_8(x) = \delta_{\mathbf{0}^8}(x) + \delta_{\mathbf{1}^8}(x) + R_8^{\mathrm{cen}}(x) = \mathbf{1}[x \in RM(1,3)].$$

Thus RM_8 is the full sixteen-word Reed–Muller signature.

8.3.2 Normalizing the Endpoint Coefficients

The preceding lemma has made all central coefficients equal. One information-set contraction detects the remaining endpoint invariant.

Lemma 8.11 (Normalizing the endpoint coefficients of RM_8)**.** *Let*

$$M^\circ = A\delta_{\mathbf{0}^8} + B\delta_{\mathbf{1}^8} + cR_8^{\mathrm{cen}}, \qquad ABc \neq 0.$$

Then either the two-copy native-X construction below yields a nonzero arity-six EO signature outside $\mathcal{B}_\times$, or

$$AB = c^2 \tag{Q8.25}$$

and, for any algebraic s with $s^8 = c/A$,

$$D(s)^{\otimes 8} M^\circ = c\mathrm{RM}_8, \qquad D(s) = \mathrm{diag}(s, s^{-1}), \qquad D(s)^\mathsf{T} X D(s) = X. \tag{Q8.26}$$

The first outcome uses two source vertices, four information-set edges, and one final X-closure.

Proof. Index the ports by $V = \mathbb{F}_2^3$ and take the complementary affine information sets

$$I_0 = \{000, 100, 010, 001\}, \qquad I_1 = V \setminus I_0.$$

Join corresponding I_0-ports of two copies of M° by four native X-edges. An affine word is determined by its values on either information set. Hence the internal equations force the two $RM(1,3)$ words to be complementary on all eight coordinates. The remaining ports form four opposite pairs indexed by I_1, and their coefficient function on the four independent pair-choice bits is

$$H(y) = \begin{cases} AB, & y = 0^4 \text{ or } 1^4, \\ c^2, & \text{otherwise.} \end{cases}$$

Close one of these four pairs by a native X-edge. On the three remaining matching bits the two antipodal coefficients are $AB + c^2$, whereas the other six coefficients are $2c^2$. If $AB = -c^2$, the output is nonzero but its support omits two of the eight choices of one endpoint from each remaining pair. Otherwise that support uniquely determines its three opposite pairs, and a rank-one coefficient table would have all 2×2 minors vanish after fixing any one of the three pair-choice bits. Choosing a minor through an antipodal word instead gives determinant

$$\det \begin{pmatrix} AB + c^2 & 2c^2 \\ 2c^2 & 2c^2 \end{pmatrix} = 2c^2(AB - c^2).$$

Thus $AB \neq c^2$ gives the first outcome, and the remaining case is (Q8.25).

Finally, $D(s)^{\otimes 8}$ multiplies a weight-k coefficient by s^{8-2k}. It fixes the central layer, sends A to $As^8 = c$, and sends B to $Bs^{-8} = AB/c = c$. This proves (Q8.26); direct multiplication gives the displayed $GO(X)$ identity. □

8.3.3 Obtaining Equality from the Reed–Muller Signature

We use the Bell contraction identity and the occurrence reduction of [SC20, full version, Definition 8.7 and Lemma 8.9]. The proof below supplies the ambient-relative step for our complex-valued high-charge language; it does not assume global ARS. Recall that RM_8 is the full sixteen-word $RM(1,3)$ signature, and put $\mathcal{B} = \{I, X, Y, Z\}$, with the ordered matrices of Equation (7).

Lemma 8.12 (The Reed–Muller Bell reduction under high charge)**.** *Let* Γ *be a finite algebraic set of even-arity signatures such that*

$$\chi_r(x) \in 4\mathbb{Z} \qquad (f \in \Gamma,\ \operatorname{arity}(f) = r,\ x \in \operatorname{supp}(f)).$$

Suppose Γ *contains* $c\mathrm{RM}_8$ *for a known* $c \neq 0$*. Then*

$$\text{K-Holant}(\Gamma, \mathcal{B}) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma). \tag{Q8.31}$$

Consequently, if $\neg\mathsf{Tract}(K\Gamma)$*, then* K-Holant$(\Gamma)$ *is* #P*-hard.*

Proof. Throughout, boundary tensors refer to direct native-X gadgets over the indicated signature set, with ordered boundary ports. The gadgets may be disconnected; components without boundary contribute scalars.

Occurrence reduction. For each ordered $b \in \mathcal{B}$, attaching b to two ports of $c\mathrm{RM}_8$ leaves, up to a known nonzero scalar and a fixed ordering of the six remaining ports, $b^{\otimes 3}$. This is the identity used in [SC20, full version, Lemma 8.9]. It remains valid in native X-coordinates, since $XbX = \pm b$. Thus any three b-vertices can be replaced by one b-vertex and one core vertex, reducing their number by two. Repeating this operation leaves at most two occurrences, with linear growth in the instance size and a tracked nonzero scalar. This is a tensor identity, so arbitrary complex labels in the remaining instance do not affect the reduction.

For $b = Y$ or I, a self-loop on a b-vertex has value zero, while two such vertices sharing an edge collapse to a native X-wire, up to their recorded orientation signs. If they share both edges, they form a component of known scalar value. We may therefore assume that the remaining b-vertices are disjoint. Removing $k \leq 2$ of them leaves a direct gadget with $2k$ ordered boundary ports.

Adjoining Y*.* The following identity holds on all sixteen four-bit assignments:

$$Y_{12} \otimes Y_{34} = X_{14} \otimes X_{23} - X_{13} \otimes X_{24}. \tag{Q8.32}$$

Hence two disjoint Y-occurrences can be evaluated by the difference of two native-X closures of the residual gadget. The effective kernel of a Y-vertex attached through two native edges is $XYX = -Y$, so the two signs cancel here.

It remains to evaluate one occurrence. Write $\langle A,B\rangle = \sum_x A(x)B(x)$ for the bilinear coefficient pairing. If every binary boundary tensor H over Γ satisfies $\langle H,Y\rangle = 0$, every such instance has value zero. Otherwise fix one direct binary gadget G over Γ with

$$\kappa = \langle G,Y\rangle \neq 0.$$

Its size and the nonzero algebraic number κ depend only on the fixed language Γ, not on the input instance; its truth table is computed by finite exact summation. For the residual binary H of an input, apply (Q8.32) to the disconnected gadget $H \otimes G$. Two oracle closures compute

$$\langle H \otimes G, Y \otimes Y\rangle = \kappa\langle H,Y\rangle.$$

Dividing by the known κ, and restoring the sign from $XYX = -Y$, gives the required one-occurrence value. The case distinction is a statement about the fixed language: in the second case one fixed witness gadget is part of the reduction, and in the first case the zero rule applies to every input. Together with occurrence reduction, this proves

$$\text{K-Holant}(\Gamma, Y) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma). \tag{Q8.33}$$

Adjoining I. Put $\Lambda = \Gamma \cup \{Y\}$. The added signature has charge zero, so all supported half-charges in Λ still lie in $4\mathbb{Z}$. Charge conservation across native edges therefore makes every binary boundary tensor over Λ anti-diagonal and every quaternary boundary tensor EO. No factorization or ARS assumption on these tensors is required.

Use occurrence reduction for I. One disjoint occurrence evaluates to zero, since $XIX = I$ is orthogonal to every anti-diagonal binary for the bilinear pairing. For two disjoint occurrences, let H be the four-port residual tensor, with the removed vertices on 12 and 34. On its EO support,

$$(I_{12} \otimes I_{34})|_{HW_4^=} = \left.\frac{X_{13} \otimes X_{24} + X_{14} \otimes X_{23} - X_{12} \otimes X_{34}}{2}\right|_{HW_4^=}. \tag{Q8.34}$$

Thus three native-X closures, with respective coefficients $1/2, 1/2, -1/2$, evaluate the desired contraction. These are K-Holant(Λ)-oracle queries, valid for arbitrary complex coefficients of H. We have proved

$$\text{K-Holant}(\Gamma, Y, I) \leq_{\mathrm{T}} \text{K-Holant}(\Gamma, Y). \tag{Q8.35}$$

Finally $YXI = Z$ supplies the remaining Bell signature by a direct gadget, and an X-vertex merely subdivides a native wire. This proves (Q8.31).

If $\neg\mathsf{Tract}(K\Gamma)$, restriction gives $\neg\mathsf{Tract}(K(\Gamma \cup \{Y, I\}))$. The binary-equality K-Holant dichotomy Theorem 2.23, namely [LM26a, Theorem 1.2 and Lemma 2.13], makes the enlarged problem hard. The two displayed reductions transfer this hardness to K-Holant(Γ). □

Proof of Theorem 8.1 for $d = 4$. Assume $\neg\mathsf{Tract}(K\Gamma)$ and $Q \neq 0$, as in Section 8.1. Apply Lemma 8.5 and suppose hardness has not already been proved. If $P_A \cap P_B \neq \varnothing$ for some $A \neq B$ with $A \cap B \neq \varnothing$, hardness follows from Lemma 8.8. Otherwise Lemmas 8.9 and 8.10 give a bijection $\iota : V \to \mathbb{F}_2^3$ under which the actual mountain is

$$M^\circ = A\delta_{\mathbf{0}^8} + B\delta_{\mathbf{1}^8} + cR_8^{\mathrm{cen}}, \qquad ABc \neq 0.$$

The first outcome of Lemma 8.11 is an actual nonzero EO signature outside $\mathcal{B}_\times$, so Theorem 3.7 gives hardness. Otherwise (Q8.26) supplies $D(s) = \mathrm{diag}(s, s^{-1}) \in GO(X)$ with

$$D(s)^{\otimes 8} M^\circ = c\mathrm{RM}_8.$$

Apply $D(s)$ to the entire current signature set and again call it Γ. This whole-set equivalence preserves support and charge, and Lemma 2.21 preserves nontractability. Now $c\mathrm{RM}_8 \in \Gamma$, so Lemma 8.12 completes the proof. □

9 Synthesis and Proof of the Dichotomy

9.1 Combining the Height Cases

The height analyses now give a dichotomy for the entire signature set; every tractable conclusion concerns one common presentation of that set.

Theorem 9.1 (Dichotomy for sets containing a mountain)**.** *Let $\mathcal{G}$ be a finite algebraic set of even-arity signatures containing a mountain of height d. If $d \geq 3$, assume that $\chi_r(x) \in d\mathbb{Z}$ for every $f \in \mathcal{G}$ of arity r and every $x \in \operatorname{supp}(f)$, as for the gcd-height mountain produced by Lemma 4.1. Then*

$$\mathsf{Tract}(K\mathcal{G}) \qquad \textit{or} \qquad \text{K-Holant}(\mathcal{G}) \textit{ is } \#\mathsf{P}\textit{-hard.}$$

Proof. Apply Theorem 5.1 for $d = 1$, Theorem 6.1 for $d = 2$, Theorem 7.1 for $d = 3$, and Theorem 8.1 for $d \geq 4$. These cases exhaust all heights. □

9.2 Proof of the Main Theorem

Recall that Theorem 1.1 gives an FP-versus-$\#\mathsf{P}$-hard dichotomy characterized by the tractability predicate Tract defined in Section 2.13. We now complete its proof by combining the preceding results to establish the corresponding dichotomy in K-coordinates. For a signature set $\mathcal{G}$ in these coordinates, write $\mathsf{Tract}_K(\mathcal{G}) := \mathsf{Tract}(K\mathcal{G})$.

Theorem 9.2 (Structural classification in K-coordinates)**.** *Let $\mathcal{G}$ be a finite algebraic signature set. If $\mathsf{Tract}_K(\mathcal{G})$ holds, then* K-Holant$(\mathcal{G}) \in \mathsf{FP}$*; otherwise* K-Holant$(\mathcal{G})$ *is $\#\mathsf{P}$-hard. The applicable side is decidable from exact algebraic input.*

Proof. The tractable implication and its effective recognition follow from Theorem 2.18. For the converse, assume that K-Holant$(\mathcal{G})$ is not $\#\mathsf{P}$-hard; we prove $\mathsf{Tract}_K(\mathcal{G})$. Apply zero and nullary preprocessing and let Λ be the resulting positive-arity signature set. We use the equivalence convention throughout; Lemma 2.22 transports both common presentations and complexity conclusions back to $\mathcal{G}$.

Odd arity is closed first by Theorem 2.13. In the all-even case, test common one-sided support before factor extraction; this is alternative (1) by charge conservation and Theorem 2.11. If every member is single-layer, the remaining two-sided case is alternative (2) by Theorem 2.12. Otherwise choose two signature–support-point pairs with

$$f_\pm \in \Lambda, \qquad x_\pm \in \operatorname{supp}(f_\pm),$$
$$\chi_{\operatorname{arity}(f_+)}(x_+) > 0 > \chi_{\operatorname{arity}(f_-)}(x_-).$$

The two signs need not occur in the same signature.

Apply Lemma 2.19 and continue to write Λ for the factor-closed set. Both selected signature–support-point pairs remain present. An odd-arity factor returns to Theorem 2.13; if no higher-arity prime remains, alternative (3) holds.

We are left with a two-sided, multilayer all-even signature set containing a higher-arity prime. Let d be the gcd of its nonzero supported half-charges, computed on the current set Λ after factor closure. Apply Lemma 4.1 to Λ; the two selected pairs verify its sign hypothesis. In its pin outcome, take $\Delta_0, \Delta_1 \in \Lambda$ by the equivalence convention and apply the odd-arity theorem. Otherwise that lemma places in Λ a mountain of the charge gcd d, with

$$\chi_{\operatorname{arity}(f)}(x) \in d\mathbb{Z} \qquad (f \in \Lambda, x \in \operatorname{supp}(f)).$$

Thus Theorem 9.1 applies. Its hard outcome is excluded, so $\mathsf{Tract}_K(\Lambda)$, and hence $\mathsf{Tract}_K(\mathcal{G})$, holds.

For completeness, if the proof gives a common K-presentation W, the ordinary-coordinate basis is $M = KW$. Since $K^\mathsf{T} K = X$,

$$M^{-1}(K\mathcal{G}) = W^{-1}\mathcal{G}, \qquad M^\mathsf{T} M = W^\mathsf{T} X W,$$

so one basis handles both the edge and the entire signature set. □

Proof of Theorem 1.1. Decidability and polynomial-time evaluation in every tractable case follow from Theorem 2.18. For hardness, put $\mathcal{G} = K^{-1}\mathcal{F}$. By Equation (4),

$$\text{Holant}(\mathcal{F}) \equiv_\text{T} \text{K-Holant}(\mathcal{G}).$$

By the definition of Tract_K,

$$\mathsf{Tract}_K(\mathcal{G}) = \mathsf{Tract}(K\mathcal{G}) = \mathsf{Tract}(\mathcal{F}).$$

If $\neg\mathsf{Tract}(\mathcal{F})$, then $\neg\mathsf{Tract}_K(\mathcal{G})$, so Theorem 9.2 implies that K-Holant($\mathcal{G}$) is #P-hard. The displayed problem equivalence transfers this hardness to Holant($\mathcal{F}$). All reductions are polynomial-time Turing reductions, as required. □

10 Conclusion

We have proved an FP-versus-#P-hard dichotomy for Boolean Holant on general graphs with arbitrary finite sets of algebraic complex-valued signatures, under polynomial-time Turing reductions. The classification uses the same seven conditions as the odd-arity dichotomy, and this criterion is decidable from exact algebraic input.

For the all-even cases, charge conservation and tensor-factor extraction either lead to a known classification or produce a mountain signature. The analysis by height, together with the existing odd-arity, Eulerian-orientation, binary-equality, and quaternary dichotomies, yields the full result. In each holographic tractable case, one basis works simultaneously for the edge and every signature in the set.

A Exact Finite Certificates

This appendix gives the exact finite calculations used in the odd-dihedral order-three and high-charge six- and eight-port arguments. Each detailed proposition is followed by the consequence used in the main proof. We retain the exhaustive counts, equations, verification details, and the distinction between the finite calculations and the reductions proved in the text.

The accompanying programs reconstruct and check the finite objects from the definitions below. Reproduction instructions and the correspondence between manuscript statements and programs are given under Code Availability. The algebraic reductions and direct-gadget arguments remain in the text.

Throughout this appendix, ranks and linear systems are computed exactly over the indicated characteristic-zero field. A projective vector is normalized by making its first nonzero coordinate one, except when a primitive integer representative with first nonzero coordinate positive is explicitly requested. For a nonzero rational projective vector, this integer representative is obtained by clearing denominators, dividing all coordinates by their greatest common divisor, and choosing the sign so that the first nonzero coordinate is positive. No numerical tolerance or random sampling is used.

As a supplementary check, the script `verify_high_charge_q8_normalized.py` verifies the rank of each relevant 560-by-98 integer matrix over $\mathbb{Q}$: a nonzero vector in its kernel gives rank at most 97, while rank 97 modulo the prime 1,000,000,007 exhibits a nonzero integer minor and gives rank at least 97. For its integer relation matrix, an explicit kernel and a nonzero 13-by-13 minor of absolute value 16 give rank 13; rank 9 modulo two then forces the nonzero Smith invariant factors to be nine 1's and four 2's. These are exact rank and divisibility arguments over the integers, not heuristic finite-field specializations.

A.1 Odd-dihedral and high-charge certificates

A.1.1 The order-three odd-dihedral six-port certificate

Fix $\omega \in \overline{\mathbb{Q}}$ with $\omega^2 + \omega + 1 = 0$, and identify a residue $r \in \mathbb{F}_3$ with its representative $0, 1, 2$ whenever it appears as an exponent ω^r. For $\epsilon \in \mathbb{F}_2$, let

$$\mathcal{B}_\epsilon = \{\{x, \bar{x}\} : x \in \mathbb{F}_2^4,\ \mathrm{wt}(x) \equiv \epsilon \pmod 2\}$$

be the four complement blocks of parity ϵ. Choose an orientation $(x_p, \bar{x}_p)$, $0 \le p \le 3$, of these four blocks, and for $x \in \mathbb{F}_2^4$ let $e_x \in \mathbb{C}^{\mathbb{F}_2^4}$ be the coordinate vector given by $e_x(y) = \mathbf{1}[y = x]$. For $r, t \in \mathbb{F}_3$, put

$$V_{p,r} = \omega^r e_{x_p} + \omega^{2r} e_{\bar{x}_p}, \qquad E_{pq}(r,t) = V_{p,r} - V_{q,t}. \tag{33}$$

Define the finite projective set

$$\mathcal{Q}_3^\epsilon = \{[V_{p,r}] : 0 \le p \le 3,\ r \in \mathbb{F}_3\} \cup \{[E_{pq}(r,t)] : 0 \le p < q \le 3,\ r, t \in \mathbb{F}_3\}.$$

For a six-port tensor F, an unordered pair $i < j$, and $a, b \in \mathbb{F}_2$, write F_{ij}^{ab} for the quaternary slice obtained by setting the two selected ports to a, b, with the residual ports kept in increasing order. For $\xi \in \mathbb{F}_2$ and $s \in \mu_3$, define

$$C_{ij,\xi,s}(F) = F_{ij}^{0,\xi} - sF_{ij}^{1,1-\xi}. \tag{34}$$

For $\eta \in \mathbb{F}_2$, if $\mathrm{supp}(F) \subseteq \{x \in \mathbb{F}_2^6 : \mathrm{wt}(x) \equiv \eta \pmod 2\}$, then this contraction has residual parity $\eta \oplus \xi$.

For a support $S \subseteq \mathbb{F}_2^6$, define

$$U_{ij}^\xi(S) = \{x_{\widehat{\{i,j\}}} : x \in S,\ x_i \oplus x_j = \xi\}.$$

The (i, j, ξ) *xor sector* consists of the assignments satisfying $x_i \oplus x_j = \xi$. We will use the condition that $U_{ij}^\xi(S)$ is empty or is the union of one or two complete blocks from a single parity class $\mathcal{B}_\epsilon$. In the table below, the middle column lists the numbers of complement blocks supporting the three projective points; for example, 1/1/2 means that two points each use one block and the third uses two. For three distinct blocks p, q, u, the points $[E_{pq}(r,t)]$, $[E_{pu}(r,v)]$, and $[E_{qu}(t,v)]$ are collinear because $E_{pq}(r,t) - E_{pu}(r,v) + E_{qu}(t,v) = 0$.

We call a pencil $A - sB$ *projectively nonconstant* when A, B are linearly independent. Otherwise it is zero or has constant projective class wherever it is nonzero. Consider labelled pencils $q_s = A - sB$, $s \in \mu_3$, modulo one common nonzero scalar, such that all three q_s are nonzero and their three distinct projective classes are points of $\mathcal{Q}_3^\epsilon$ assigned bijectively to the phase labels. Only pencils satisfying these conditions are counted below.

Proposition A.1 (Order-three odd-dihedral six-port classification)**.** *The following statements hold.*

(i) *Each set* $\mathcal{Q}_3^\epsilon$ *has* 66 *points. Its projective lines containing at least three of its points are exactly*

number	*supported blocks*	*geometry*
4	1/1/1	*one fixed complement block*
36	2/2/2	*one fixed Segre ruling*
54	1/1/2	*a secant between two blocks*
108	2/2/2	*the displayed relation on three blocks*

Every such line contains exactly three points of $\mathcal{Q}_3^\epsilon$*. Hence there are* 202 *such lines and* $202 \cdot 6 = 1212$ *labelled pencils satisfying the stated conditions in each parity.*

(ii) *Let* $\eta \in \mathbb{F}_2$*, let* $\operatorname{supp}(F) \subseteq \{x \in \mathbb{F}_2^6 : \operatorname{wt}(x) \equiv \eta \pmod 2\}$*, and suppose that every contraction in Equation* (34) *is zero or has projective class in the appropriate set* $\mathcal{Q}_3^{\eta\oplus\xi}$*. For no fixed* $i < j$ *and* ξ *does the projectively nonconstant family* $s \mapsto C_{ij,\xi,s}(F)$ *lie on a line from the third or fourth row of part* (i)*. Consequently, for every such* F*, each projectively nonconstant family* $s \mapsto C_{ij,\xi,s}(F)$ *either is supported on one complement block or lies in a fixed Segre ruling.*

(iii) *Exactly* 873 *supports* $S \subseteq \mathbb{F}_2^6$ *make every* $U_{ij}^\xi(S)$ *empty or a union of one or two complete complement blocks of a single parity. They are the empty support; the* 392 *distinct cosets, counted without duplication,*

$$a + L_\pi, \qquad L_\pi = \operatorname{span}_{\mathbb{F}_2}\{\mathbf{1}_C : C \in \pi\},$$

where $a \in \mathbb{F}_2^6$ *and* π *is a partition of the six ports into even blocks; and the* 480*-element bit-flip/port-permutation orbit of*

$$S_\star = \{000000, 000011, 001100, 110011, 111100, 111111\}.$$

The distinct affine supports occur in the respective numbers

$$32, \qquad 240, \qquad 120$$

for partition types (6)*,* $(4+2)$*, and* $(2+2+2)$*, where a partition type is the multiset of its block sizes. Every nonempty support in the classification has one parity.*

Proof. The vectors in Equation (33) give twelve one-block and fifty-four two-block projective points. Normalize every nonzero vector at its first nonzero coordinate. For every one of the $\binom{66}{2}$ point pairs, take the exact reduced-row-echelon form of its two-dimensional span and group equal row spaces. The groups containing at least three points of $\mathcal{Q}_3^\epsilon$ all have size three and give the table in part (i).

The table also has a direct block interpretation. The three points on one fixed block give four one-block lines. On a fixed pair of blocks, the nine two-block points form a 3-by-3 Segre grid in the transformed phase coordinates $u = r + t$, $v = r - t$ in $\mathbb{F}_3$. Fixing u or v, not the original r or t in Equation (33), gives its two rulings and hence $6 \cdot 2 \cdot 3 = 36$ lines. Every $E_{pq}(r,t)$ lies on the secant through $V_{p,r}$ and $V_{q,t}$, giving $6 \cdot 3^2 = 54$ secants. For three distinct blocks,

$$E_{pq}(r,t) - E_{pu}(r,v) + E_{qu}(t,v) = 0$$

gives $\binom{4}{3}3^3 = 108$ lines on three blocks. Conversely, in a dependent triple every block occurring in one vector must occur in another. According as the triple uses one, two, or three blocks, it is one of the displayed types; four blocks would leave a block occurring only once.

Assigning the phases $1, \omega, \omega^2$ to the three points on one of these 202 lines gives a unique affine pencil up to one common nonzero scalar. Indeed, after the unique dependence is scaled, solve

$$q_1 = A - B, \qquad q_\omega = A - \omega B$$

and verify $q_{\omega^2} = A - \omega^2 B$. All six phase assignments are distinct. This proves the last assertion of part (i).

For part (ii), after a bit flip to even parity, fix $x_0 = 0000$, $x_1 = 0011$, $x_2 = 0101$, and $x_3 = 0110$. For $t = (t_1, \ldots, t_4) \in \mathbb{F}_3^4$, define $(R_t q)(x) = \omega^{\sum_{j=1}^4 t_j x_j} q(x)$. Under R_t, the four oriented block labels change by Ht, where

$$H = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}, \qquad \det H = -16 \neq 0 \pmod 3.$$

The phase permutations are generated by $s \mapsto \alpha s$ and $s \mapsto \alpha/s$, with $\alpha \in \mu_3$; these maps form $D_6 \cong S_3$ and realize all six permutations of μ_3. Together with port permutations and bit flips, these actions reduce the two-block secant and three-block cases to the following representative pencils for a fixed restriction $x_i \oplus x_j = \xi$. A semicolon separates the four oriented complement blocks, and 0 denotes the pair $(0, 0)$:

$$A_{\text{sec}} = (1, 1; \omega, \omega; 0; 0), \qquad B_{\text{sec}} = (\omega^2, \omega^2; \omega, \omega; 0; 0),$$
$$A_{\text{tri}} = (1, 1; \omega, \omega; \omega^2, \omega^2; 0), \qquad B_{\text{tri}} = (\omega, \omega; 1, 1; \omega^2, \omega^2; 0).$$

Each displayed pair (A_*, B_*) denotes the pencil $q_s = A_* - sB_*$, $s \in \mu_3$. Since H is invertible over $\mathbb{F}_3$, the maps R_t shift the block labels independently. Port permutations and bit flips act transitively on the block choices and the two total parities. Together with the phase permutations, this proves that every orbit of the two-block secant and three-block cases under these operations is represented.

Fix the selected pair $i < j$ and sector ξ. It remains to exclude six-port tensors F satisfying the hypotheses of part (ii) and realizing either selected representative pencil. Put $\zeta = \xi \oplus 1$. Under the restriction $x_i \oplus x_j = \zeta$, a projectively nonconstant pencil is one of the 1212 labelled pencils from part (i). Its slices $C = F_{ij}^{0,\zeta}$ and $D = F_{ij}^{1,1-\zeta}$ are fixed up to a common scalar $\lambda \in \overline{\mathbb{Q}}^\times$. For each of the fourteen port pairs other than the selected pair, and for each xor sector and phase, substitution into Equation (34) gives a quaternary contraction $q_0 + \lambda q_1$, where the computed vectors $q_0, q_1 \in \overline{\mathbb{Q}}^{\mathbb{F}_2^4}$ are independent of λ. For a fixed representative $v \in \overline{\mathbb{Q}}^{\mathbb{F}_2^4} \setminus \{0\}$ of a point of the appropriate set $\mathcal{Q}_3^\epsilon$, the set of λ for which this vector is zero or belongs to $\operatorname{span}\{v\}$ is all of the field, empty, or one point; exact two-column row reduction determines which case occurs. Intersecting the resulting finite unions over all remaining pairs, both xor sectors, and all three phases leaves no nonzero λ for either representative.

For an opposite pencil with constant projective class, write $C = cv$ and $D = dv$, where $v \in \overline{\mathbb{Q}}^{\mathbb{F}_2^4} \setminus \{0\}$ represents one of the 66 points of the appropriate $\mathcal{Q}_3^\epsilon$ and $(c, d) \in \overline{\mathbb{Q}}^2 \setminus \{(0, 0)\}$. For each remaining contraction, substitution gives $q_0 + cq_1 + dq_2$, where the computed $q_0, q_1, q_2 \in \overline{\mathbb{Q}}^{\mathbb{F}_2^4}$ are independent of c, d. Its preimage over all target points of $\mathcal{Q}_3^\epsilon$ is a finite union of affine linear systems in (c, d). Exact RREF, followed by intersection over all contractions in the hypothesis of part (ii), leaves no nonzero pair (c, d). The identically zero opposite sector is checked separately and also fails. All matrices have entries in $\mathbb{Q}(\omega)$; their ranks and consistency are unchanged after extension to $\overline{\mathbb{Q}}$. Reversing the two ports when $\xi = 0$ changes nothing, whereas for $\xi = 1$ reversal at s gives a nonzero scalar multiple of the contraction at s^{-1}, so unordered port pairs exhaust the ordered conditions.

For part (iii), fix one port pair and one xor sector. A union satisfying the stated condition is empty, one of the eight one-block unions across the two parities, or one of the twelve two-block unions. Each residual word in a nonempty union occurs only in $F_{ij}^{0,\xi}$, only in $F_{ij}^{1,1-\xi}$, or in both. Hence the exact finite superset of local candidates has

$$1 + 8 \cdot 3^2 + 12 \cdot 3^4 = 1045$$

members. Combine the two disjoint xor sectors on the first port pair, giving 1045^2 candidates. Retain a candidate precisely when, on each of the other fourteen pairs, each of the two unions is empty or is the union of one or two complete complement blocks of its own parity. Encode a support by a 64-bit integer whose bit k corresponds to the k-th word in lexicographic order on $\{0,1\}^6$, starting at $k = 0$. Exact intersection of these encoded support lists leaves 873 distinct supports.

Independently generate every partition into nonempty even blocks and every coset of its block-indicator space. This gives

$$32 + 240 + 120 = 392$$

distinct supports. Generating all bit flips and port permutations of $S_\star$ gives 480 further supports. These two families and the empty support are disjoint, and their union is exactly the 873-support intersection. Since every block indicator has even weight and $S_\star$ has even parity, every nonempty support has one parity.

The exact calculation constructs the sets $\mathcal{Q}_3^\epsilon$, their lines containing at least three points, labelled pencils, affine preimages, local support lists, partition cosets, and exceptional orbit from the displayed definitions. It loads no precomputed data or stored accepted-case list and uses no random choice, floating-point comparison, finite-field specialization, or external solver. $\square$

Remark A.2 (Scope of the odd-dihedral certificate). Proposition A.1 is a normalized finite statement. The algebraic quaternary classification, the construction of the realizable kernels in Equation (18), and the reduction to Equation (34) remain in the main proof. The certificate does not prove the $N \geq 5$ support bound, the exclusions of contractions supported on two points, the coefficient factorization for support that selects one endpoint of each matching edge, the arbitrary-arity common-factor argument, or the reduction for the entire signature set.

Theorem A.3 (Six-port support classification)**.** *The order-three classification has the following consequences.*

(i) *Let F be a six-port tensor of one parity, and suppose that every contraction $C_{ij,\xi,s}(F)$, for every port pair, xor, and $s \in \mu_3$, is zero or belongs to $\mathcal{B}_3([4])$. Then, for every fixed pair and xor, with $A = F_{ij}^{0,\xi}$, $B = F_{ij}^{1,1-\xi}$,* $\operatorname{supp}(A) \cup \operatorname{supp}(B)$ *is empty or is the union of at most two complement pairs in its residual parity.*

(ii) *Let F be a six-port tensor of one parity. If the support union in part* (i) *is empty or consists of at most two complement pairs for every pair and xor, then* $\operatorname{supp}(F)$ *is empty, a partition coset of type* (6), $(4+2)$, *or* $(2+2+2)$, *or a bit-flip/port-permutation image of $S_\star$ in Equation* (26).

Proof. For part (i), every nonzero contraction over all port pairs and xors has projective class in the corresponding set $\mathcal{Q}_3^\epsilon$. By Proposition A.1(ii), a projectively nonconstant pencil lies on a one-block line or one fixed Segre ruling. Its support union consequently contains at most two complete complement pairs. Coordinatewise,

$$\operatorname{supp}(A) \cup \operatorname{supp}(B) = \bigcup_{s \in \mu_3} \operatorname{supp}(A - sB),$$

since a nonzero affine coefficient has at most one root. This also handles zero and projectively constant pencils. For part (ii), the hypothesis says precisely that every local union $U^{\xi}_{ij}(\mathrm{supp}(F))$ is empty or consists of one or two complete complement blocks of a single parity, as required in the enumeration of Proposition A.1(iii). That enumeration gives the asserted alternatives. Neither statement asserts coefficient factorization merely from the support shape; the required coefficient argument and the exclusions of contractions supported on two points are proved in Theorem 5.14. $\square$

A.1.2 High-charge six-port loop and two-copy arrays

Let

$$\mathcal{W}_4 = (0011, 0101, 0110, 1001, 1010, 1100).$$

The arrays below use this order. The Python loop generator lists its entries in reverse order; both (Q6.6) and (Q6.7) are invariant under $z_k \mapsto z_{5-k}$. The two-copy generator uses $\mathcal{W}_4$ directly. Let $U = [6]$, let $t \in \mathbb{C}^\times$, and assign arbitrary complex coefficients q_S to the twenty triples $S \in \binom{U}{3}$. For each deleted pair $i < j$, order the residual ports as $v_0 < v_1 < v_2 < v_3$; for $w \in \mathcal{W}_4$, write $w = (w_0, w_1, w_2, w_3)$ and put $R_w = \{v_r : w_r = 1\}$, and define the six EO entries

$$L_{ij}(w) = q_{R_w \cup \{i\}} + q_{R_w \cup \{j\}}, \qquad w \in \mathcal{W}_4. \tag{35}$$

For unordered two-element external sets $A, B \in \binom{U}{2}$, use the canonical increasing orders $A = (a_0, a_1)$, $B = (b_0, b_1)$ and the external order (a_0, a_1, b_0, b_1). The first and second copies retain their first/second copy labels. Put $I_A = U \setminus A$, $I_B = U \setminus B$, and let $\pi : I_A \to I_B$ be a bijection. Define the six EO entries, indexed by the concatenated words $xy \in \mathcal{W}_4$, where $x = (x_0, x_1)$ and $y = (y_0, y_1)$, by

$$C_{A,B,\pi}(x, y) = t\,\mathbf{1}[(x, y) \in \{(00, 11), (11, 00)\}] + \sum_{\substack{T \subseteq I_A \\ |T| = 3 - \mathrm{wt}(x)}} q_{A_x \cup T}\, q_{B_y \cup (I_B \setminus \pi(T))}, \tag{36}$$

where $A_x = \{a_r : x_r = 1\}$, and similarly for B_y.

For any quaternary EO array $z = (z_0, z_1, z_2, z_3, z_4, z_5)$ indexed in the order $\mathcal{W}_4$, put $P_0 = z_0 z_5$, $P_1 = z_1 z_4$, and $P_2 = z_2 z_3$.

Proposition A.4 (High-charge six-port coefficient classification)**.** *Suppose every one of the fifteen arrays in Equation* (35) *and all* $15^2 \cdot 4! = 5400$ *labelled arrays in Equation* (36) *satisfy*

$$z_0 = 0 \iff z_5 = 0, \qquad z_1 = 0 \iff z_4 = 0, \qquad z_2 = 0 \iff z_3 = 0, \tag{Q6.6}$$

$$P_0 P_1 P_2 = 0, \qquad P_i P_j (P_i^2 - P_j^2) = 0 \quad (0 \le i < j \le 2). \tag{Q6.7}$$

Then

$$\{R \in \tbinom{U}{3} : q_R \neq 0\} \subseteq \{S, U \setminus S\} \qquad \textit{for some } S \in \binom{U}{3}.$$

This holds for arbitrary complex q_S *and every* $t \neq 0$.

Proof. For each of the fifteen loop arrays, Equations (Q6.6) and (Q6.7) force at least one of its three complementary EO-entry pairs to vanish. Selecting such a pair gives, for a suitable $S \in \binom{U}{3}$ containing exactly one of i, j, two signed equations

$$q_S = -q_{S \triangle \{i,j\}}, \qquad q_{U \setminus S} = -q_{U \setminus (S \triangle \{i,j\})}.$$

Thus every coefficient vector lies in at least one of exactly $3^{15} = 14\,348\,907$ branches defined by signed linear equations. Signed union–find deduplicates them to 1868 linear subspaces; modulo the S_6 port action and simultaneous complementation, they form exactly 24 orbits. For the C++ enumeration, number ports from 0 to 5 and order triple coordinates by increasing $\sum_{k\in S} 2^k$. In each nonzero component C of the signed equations, write its coordinates as $q_i = \varepsilon_i u_C$, where $\varepsilon_i \in \{1, -1\}$, choose $\varepsilon_{\min C} = 1$, and use $\min C$ to name the free parameter u_C. After this sign normalization, choose the lexicographically least encoding under port permutations and simultaneous complementation, taking exactly one representative per orbit.

For each family, substitute independent complex parameters into the displayed polynomials Equations (35) and (36) and impose only Equations (Q6.6) and (Q6.7). Split every complex parameter into real and imaginary parts. A selected nonzero parameter may be normalized to 1, because all equations are homogeneous under $(q, t) \mapsto (cq, c^2 t)$ for $c \in \mathbb{C}^\times$. Exact satisfiability checks for the resulting real polynomial systems find no solution with two distinct nonzero coefficient pairs $\{q_S, q_{U\setminus S}\}$.

The resulting exhaustive calculation is

$$14{,}348{,}907 \text{ branches} \longrightarrow 1868 \text{ signed spaces} \longrightarrow 24 \text{ orbits.}$$

All generated expressions have integer coefficients, and all reductions performed before the solver use exact rational polynomial arithmetic. The all-zero orbit is immediate. The exact real-polynomial queries for the remaining twenty-three nonzero orbit families were evaluated with Z3 4.14.0. The answers are `UNSAT`, meaning that no satisfying assignment exists. The calculation uses 61 queries: 51 for projective charts of sparse orbits 1–19 and 10 for the implications used in dense orbits 20–23. For sparse orbits, the script checks that parameter sets omitted by the chosen normalizations are supported on at most one complementary block and hence cannot meet the counterexample condition. The dense branches explicitly reduce to a previously checked sparse family, the zero vector, or one complementary block. No parameter sample replaces a quantified polynomial query.

The six-port wrapper requires Z3 4.14.0 and records the compiler and solver versions and source-file hashes. Before any solver query, the Python stage reindexes the C++ output and checks that both stages give exactly the same orbit representatives. The programs regenerate the orbit representatives and exact real-polynomial queries on each run. The suite option `-q6-audit-dir` saves the generated inputs, parameter encoding, branch labels, and solver responses for inspection; the package README gives the reproduction commands.

The runs neither request nor check solver proof objects. The `UNSAT` answers therefore rely on the correctness of Z3; exporting or rerunning the inputs does not remove this dependence. Subject to this dependence, the exhaustive calculation establishes the proposition over arbitrary complex coefficients. □

Remark A.5 (Scope of the six-port certificate). Proposition A.4 starts from the coefficient arrays Equations (35) and (36). It does not prove that these arrays are direct loop or two-copy gadgets, does not make any summand of Equation (36) separately realizable, and does not establish a reduction for the entire signature set. Those arguments remain in Section 7.2 and Theorem 7.5.

Theorem A.6 (Consequence for loop and two-copy arrays). *Let $U = [6]$, let $t \in \mathbb{C}^\times$, and let $(q_S)_{S\in\binom{U}{3}}$ be arbitrary complex coefficients. Form the fifteen loop arrays in Equation* (35) *and all* 5400 *labelled two-copy arrays in Equation* (36)*. If every one of these arrays belongs to $\mathcal{B}^0_\times([4])$, then*

$$\{R \in \tbinom{U}{3} : q_R \neq 0\} \subseteq \{S, U \setminus S\} \qquad \textit{for some } S \in \binom{U}{3}.$$

Proof. Every member of $\mathcal{B}^0_\times([4])$ satisfies Equations (Q6.6) and (Q6.7); hence Proposition A.4 applies and yields the assertion. □

Remark A.7 (Why the two-copy six-port arrays are necessary). Order each triple as $S = \{s_0 < s_1 < s_2\}$, and order the twenty triples by the lexicographic order of (s_0, s_1, s_2). In that order, the full-support coefficient vector

$$(1, -1, -1, 1, i, -i, -1, 1, -i, -i, -i, -i, -1, 1, -i, i, -1, 1, 1, -1)$$

does not define a member of $\mathcal{B}_\times([6])$, although all fifteen arrays Equation (35) lie in $\mathcal{B}^0_\times([4])$. For this same vector, one labelled array Equation (36) is

$$(t, 2, -2, -2, 2, t),$$

where $t \neq 0$; its six nonzero entries put it outside that class. Thus the one-loop conditions alone cannot imply the conclusion of Theorem A.6.

A.1.3 Eight-port matching families

Let $V = \{0, 1, \ldots, 7\}$, and for distinct labels a, b write ab for the unordered pair $\{a, b\}$. For every pair $e \in \binom{V}{2}$, let P_e be a perfect matching of $V \setminus e$. For disjoint $e, r \in \binom{V}{2}$, write $r = \{u, v\}$ and define $\kappa_r(P_e)$ as follows. If $r \in P_e$, delete that edge. Otherwise let u', v' be the mates of u, v, remove uu' and vv', add $u'v'$, and retain the unique untouched edge. In both cases $\kappa_r(P_e)$ is a two-edge matching on $V \setminus (e \cup r)$.

Proposition A.8 (Eight-port matching classification)**.** *Suppose*

$$\kappa_r(P_e) = \kappa_e(P_r) \qquad (e \cap r = \varnothing) \tag{37}$$

for every pair of disjoint deleted pairs. After relabeling so that $P_{01} = \{23, 45, 67\}$, exactly nine labelled families satisfy these constraints. Seven are induced by a global perfect matching M on eight ports, meaning $P_e = \kappa_e(M)$ under the same deletion-and-splicing rule. Each of these seven families has $P_A \cap P_B \neq \varnothing$ for some distinct pairs A, B with $A \cap B \neq \varnothing$. With labels $0, \ldots, 7$ interpreted as three-bit vectors, one remaining family is given by the matchings below and the other is obtained by swapping $6, 7$, where the two-character string ab abbreviates the unordered pair $\{a, b\}$:

h	$\{\{t, t+h\} : t \in \mathbb{F}_2^3\}$
001	$01, 23, 45, 67$
010	$02, 13, 46, 57$
011	$03, 12, 47, 56$
100	$04, 15, 26, 37$
101	$05, 14, 27, 36$
110	$06, 17, 24, 35$
111	$07, 16, 25, 34.$

(Q8.14b)

For a deleted pair p, q, removing the edge pq from the row $h = p + q$ gives P_{pq}.

Proof. Once P_{01} is fixed, each of the other twenty-seven port pairs has the fifteen perfect matchings on its six residual ports as candidates. An exhaustive backtracking search imposes Equation (37) whenever both sides become known. The search has 236 search nodes and exactly nine leaves. Direct inspection of each leaf gives the stated $7 + 2$ split and the table in Equation (Q8.14b).

The scripts `verify_high_charge_q8_atlas.py` and `verify_high_charge_q8_normalized.py` reconstruct the same nine leaves from the displayed definitions, and the latter also checks the

236-node count and the linear systems for the original tensor's seventy weight-four coefficients and twenty-eight contraction scales. The two programs reuse the same matching-generation logic and are therefore not independent implementations. This proposition is only the finite matching classification. The direct double-contraction identity that yields Equation (37), and the reduction that excludes families with $P_A \cap P_B \neq \varnothing$ for distinct A, B satisfying $A \cap B \neq \varnothing$, remain in the main proof. $\square$

Theorem A.9 (Translation form of the matching family)**.** *Let V be an eight-element port set and, for every $e \in \binom{V}{2}$, let P_e be a perfect matching of $V \setminus e$. Suppose*

$$\kappa_r(P_e) = \kappa_e(P_r) \qquad (e \cap r = \varnothing).$$

If $P_A \cap P_B = \varnothing$ whenever $A \neq B$ and $A \cap B \neq \varnothing$, then there is a bijection $\iota : V \to \mathbb{F}_2^3$ such that, after relabelling, for $p \neq q$ and $h = p + q$,

$$P_{pq} = \{\{t, t+h\} : t \in \mathbb{F}_2^3 \setminus \{p, q\}\},$$

where duplicate unordered pairs are suppressed.

Proof. Relabel V so that $P_{01} = \{23, 45, 67\}$. By Proposition A.8, seven of the nine labelled families violate the condition on $P_A \cap P_B$. The remaining two differ only by a relabelling and give the displayed matching family. $\square$

Proposition A.10 (Coefficients on affine planes)**.** *Identify $V = \mathbb{F}_2^3$, with the labels $0, \ldots, 7$ interpreted as their three-bit binary expansions, and for $S \in \binom{V}{4}$ put $\sigma(S) = \bigoplus_{s \in S} s$. Let $c_S \in \mathbb{C}$, and let $\lambda_{pq} \in \mathbb{C}^\times$ for every unordered pair $\{p, q\} \in \binom{V}{2}$. Suppose that, whenever $p \in S$, $q \notin S$, $h = p + q$, and $S' = S \triangle \{p, q\}$,*

$$c_S + c_{S'} = \begin{cases} \lambda_{pq}, & \sigma(S) \in \{0, h\}, \\ 0, & \sigma(S) \notin \{0, h\}. \end{cases} \tag{Q8.19a}$$

Then the four-sets with nonzero coefficients are exactly the fourteen affine planes of V, and there is one $c \in \mathbb{C}^\times$ such that

$$c_S = \begin{cases} c, & \sigma(S) = 0, \\ 0, & \sigma(S) \neq 0. \end{cases} \tag{Q8.19b}$$

Call a four-subset a nonplane *when it is not an affine plane. Let G be the graph on the fifty-six nonplanes, joining S to $S \triangle \{p, q\}$ whenever $p \in S$, $q \notin S$, and the latter four-set is also a nonplane. Let H be the bipartite incidence graph between planes and unordered port pairs, with an edge between S and $\{p, q\}$ exactly when $|S \cap \{p, q\}| = 1$. Then G is connected and contains an odd cycle, and H is connected.*

Proof. A four-set has xor zero exactly when it is an affine plane: xor zero makes the fourth point the affine sum of the other three, and the seven nonzero linear functionals on $\mathbb{F}_2^3$ each have two affine level-set planes. Thus there are fourteen such sets and fifty-six nonplanes. For a nonplane S, the exchanged set $S' = S \triangle \{p, q\}$ is again a nonplane exactly when $\sigma(S) \notin \{0, p + q\}$. Lift $t \in V$ to $\widehat{t} = (1, t) \in \mathbb{F}_2^4$, and let M be the represented matroid on ground set V with representing vectors $\widehat{t}$. A four-set is a nonplane exactly when its four lifted points form a basis. Thus G is the basis-exchange graph of the rank-four binary matroid M, whose vertices are bases and whose edges exchange one element while preserving a basis. This graph is connected.

In binary labelling it contains the triangle below, where, for example, 0124 abbreviates the four-set $\{0, 1, 2, 4\}$:

$$\texttt{0124} — \texttt{0125} — \texttt{0245} — \texttt{0124}. \tag{Q8.20}$$

The three xor sums are $7, 6, 3$, while the corresponding swap directions are $1, 5, 4$; hence all three edges belong to G. The second line of Equation (Q8.19a) gives, on every edge of G,

$$c_{S'} = -c_S.$$

Alternating around Equation (Q8.20) and using characteristic zero forces the three coefficients on the triangle to vanish. Connectivity then forces $c_S = 0$ for every nonplane.

Let S now be a plane and let $p \in S, q \notin S$. Then $S' = S \triangle \{p, q\}$ has xor $p + q \neq 0$, so it is a nonplane. The first line of Equation (Q8.19a) therefore gives

$$c_S = \lambda_{pq} \neq 0. \tag{38}$$

It remains only to show that these values agree. Regard a plane as incident with every unordered pair having one endpoint in the plane and one outside it. Two nonparallel affine planes meet in two points and their union has six points; a pair from their intersection to the complement of their union is incident with both planes in H. Two distinct parallel planes are complementary and hence have the same incident port pairs. Thus H is connected. Every pair $\{p, q\}$ is incident with a plane: choose a linear functional ℓ with $\ell(p + q) = 1$ and take $\{x : \ell(x) = \ell(p)\}$. Equation (38) makes all fourteen plane coefficients and all incident λ_{pq}'s equal to one common nonzero c, proving Equation (Q8.19b).

The script `verify_high_charge_q8_atlas.py` also checks the $14 + 56$ split, the connectivity of G, the displayed odd cycle, and the connectivity of H. These are supplementary checks of the argument above. □

Theorem A.11 (Affine-plane coefficients)**.** *Identify $V = \mathbb{F}_2^3$. Let $c_S \in \mathbb{C}$ for $S \in \binom{V}{4}$, and let $\lambda_{pq} \in \mathbb{C}^\times$ for every unordered pair $\{p, q\} \in \binom{V}{2}$. Suppose that, whenever $p \in S$, $q \notin S$, $h = p + q$, and $S' = S \triangle \{p, q\}$,*

$$c_S + c_{S'} = \begin{cases} \lambda_{pq}, & \bigoplus_{s \in S} s \in \{0, h\}, \\ 0, & \bigoplus_{s \in S} s \notin \{0, h\}. \end{cases}$$

Then there is one $c \in \mathbb{C}^\times$ such that

$$c_S = \begin{cases} c, & \bigoplus_{s \in S} s = 0, \\ 0, & \bigoplus_{s \in S} s \neq 0. \end{cases}$$

Equivalently, the support consists exactly of the fourteen affine planes of $\mathbb{F}_2^3$, and all supported coefficients are equal.

Proof. This follows from Proposition A.10. □

References

[Apo76] Tom M. Apostol. *Introduction to Analytic Number Theory.* Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1976.

[Bac17] Miriam Backens. A new Holant dichotomy inspired by quantum computation. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:14, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[Bac21] Miriam Backens. A full dichotomy for Holant^c, inspired by quantum computation. *SIAM Journal on Computing*, 50(6):1739–1799, 2021.

[Bac25] Miriam Backens. Erratum: A full dichotomy for Holant^c, inspired by quantum computation. *SIAM Journal on Computing*, 54(3):814–818, 2025.

[BDG$^+$09] Andrei Bulatov, Martin Dyer, Leslie Ann Goldberg, Markus Jalsenius, and David Richerby. The complexity of weighted Boolean #CSP with mixed signs. *Theoretical Computer Science*, 410(38–40):3949–3961, 2009.

[BDG$^+$12] Andrei Bulatov, Martin Dyer, Leslie Ann Goldberg, Markus Jalsenius, Mark Jerrum, and David Richerby. The complexity of weighted and unweighted #CSP. *Journal of Computer and System Sciences*, 78(2):681–688, 2012.

[BDR19] Cornelius Brand, Holger Dell, and Marc Roth. Fine-grained dichotomies for the Tutte plane and Boolean #CSP. *Algorithmica*, 81(2):541–556, 2019.

[BG05] Andrei Bulatov and Martin Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2–3):148–186, 2005.

[Bul13] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *Journal of the ACM*, 60(5):34:1–34:41, 2013.

[CC17] Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. *Journal of the ACM*, 64(3):19:1–19:39, 2017.

[CC19] Jin-Yi Cai and Xi Chen. A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights. *Computational Complexity*, 28:345–408, 2019.

[CCL13] Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM Journal on Computing*, 42(3):924–1029, 2013.

[CCL16] Jin-Yi Cai, Xi Chen, and Pinyan Lu. Nonnegative weighted #CSP: An effective complexity dichotomy. *SIAM Journal on Computing*, 45(6):2177–2198, 2016.

[CF22] Jin-Yi Cai and Zhiguo Fu. Holographic algorithm with matchgates is universal for planar #CSP over Boolean domain. *SIAM Journal on Computing*, 51(2):STOC17–50–STOC17–151, 2022.

[CF23] Jin-Yi Cai and Zhiguo Fu. Complexity classification of the eight-vertex model. *Information and Computation*, 293:105064, 2023.

[CFGW22] Jin-Yi Cai, Zhiguo Fu, Heng Guo, and Tyson Williams. FKT is not universal—a planar Holant dichotomy for symmetric constraints. *Theory of Computing Systems*, 66(1):143–308, 2022.

[CFS20a] Jin-Yi Cai, Zhiguo Fu, and Shuai Shao. Beyond #CSP: A dichotomy for counting weighted Eulerian orientations with ARS. *Information and Computation*, 275:104589, 2020.

[CFS20b] Jin-Yi Cai, Zhiguo Fu, and Shuai Shao. From Holant to quantum entanglement and back. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. Full version: arXiv:2004.05706.

[CFST26] Jin-Yi Cai, Austen Fan, Shuai Shao, and Zhuxiao Tang. New planar algorithms and a full complexity classification of the eight-vertex model. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*, pages 967–978. Association for Computing Machinery, 2026.

[CFX18] Jin-Yi Cai, Zhiguo Fu, and Mingji Xia. Complexity classification of the six-vertex model. *Information and Computation*, 259:130–141, 2018.

[CGW16a] Jin-Yi Cai, Heng Guo, and Tyson Williams. A complete dichotomy rises from the capture of vanishing signatures. *SIAM Journal on Computing*, 45(5):1671–1728, 2016.

[CGW16b] Jin-Yi Cai, Heng Guo, and Tyson Williams. The complexity of counting edge colorings and a dichotomy for some higher domain Holant problems. *Research in the Mathematical Sciences*, 3(1):18, 2016.

[CH96] Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996.

[CHL12] Jin-Yi Cai, Sangxia Huang, and Pinyan Lu. From Holant to #CSP and back: Dichotomy for Holant^c problems. *Algorithmica*, 64:511–533, 2012.

[CI25] Jin-Yi Cai and Jin Soo Ihm. Holant* dichotomy on domain size 3: A geometric perspective. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, volume 334 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 148:1–148:18, Dagstuhl, Germany, 2025. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[CLX09] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009)*, pages 715–724, 2009.

[CLX11] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for Holant* problems of Boolean domain. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)*, pages 1714–1728. Society for Industrial and Applied Mathematics, 2011.

[CLX13] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for Holant problems with a function on domain size 3. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013)*, pages 1278–1295. Society for Industrial and Applied Mathematics, 2013.

[CLX14] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. The complexity of complex weighted Boolean #CSP. *Journal of Computer and System Sciences*, 80(1):217–236, 2014.

[CLX17] 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.

[CLX18] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for real Holant^c problems. In *Proceedings of the Twenty-Ninth Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2018)*, pages 1802–1821. SIAM, 2018. Full version: arXiv:1702.02693.

[CLX26] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for Holant* problems with one ternary function on domain size 3. *Information and Computation*, 309:105421, 2026. Journal version of the SODA 2013 paper.

[CM23] Jin-Yi Cai and Ashwin Maran. The complexity of counting planar graph homomorphisms of domain size 3. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1285–1297. Association for Computing Machinery, 2023.

[CM24] Jin-Yi Cai and Ashwin Maran. Polynomial and analytic methods for classifying complexity of planar graph homomorphisms. arXiv:2412.17122v1, 2024.

[CMY26] Jin-Yi Cai, Ashwin Maran, and Ben Young. Planar graph homomorphisms: A dichotomy and a barrier from quantum groups. arXiv:2601.23198v1, 2026.

[CY26] Jin-Yi Cai and Ben Young. Vanishing signatures, orbit closure, and the converse of the Holant theorem. In Shubhangi Saraf, editor, *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:20, Dagstuhl, Germany, 2026. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[DG00] Martin Dyer and Catherine Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3–4):260–289, 2000.

[DGJ09] Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. The complexity of weighted Boolean #CSP. *SIAM Journal on Computing*, 38(5):1970–1986, 2009.

[DGP07] Martin Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *Journal of the ACM*, 54(6):27:1–27:23, 2007.

[DR13] Martin Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM Journal on Computing*, 42(3):1245–1274, 2013.

[GGJT10] Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM Journal on Computing*, 39(7):3336–3402, 2010.

[GSS26] Jincheng Guan, Shuai Shao, and Ke Shi. An LP algorithm for counting Eulerian orientations through the lens of quasi-polymorphism. arXiv:2607.27961v1, 2026.

[GW20] Heng Guo and Tyson Williams. The complexity of planar Boolean #CSP with complex weights. *Journal of Computer and System Sciences*, 107:1–27, 2020.

[HL16] Sangxia Huang and Pinyan Lu. A dichotomy for real weighted Holant problems. *Computational Complexity*, 25:255–304, 2016.

[LFC25] Yin Liu, Austen Z. Fan, and Jin-Yi Cai. Restricted Holant dichotomy on domain sizes 3 and 4. *Theoretical Computer Science*, 1023:114931, 2025.

[Liu24] Ying Liu. Exponential time complexity of the complex weighted Boolean #CSP. In Weili Wu and Guangmo Tong, editors, *Computing and Combinatorics (COCOON 2023)*, volume 14422 of *Lecture Notes in Computer Science*, pages 83–96. Springer, 2024.

[LM26a] Chenghua Liu and Boning Meng. A dichotomy for complex Boolean Holant with binary disequality. arXiv:2609.00219v1, 2026. Version 1, submitted 31 August 2026.

[LM26b] Chenghua Liu and Boning Meng. From block orthogonality to decidability in complex-weighted counting CSP. arXiv:2608.14845v1, 2026. Version 1, submitted 14 August 2026.

[LW18] Jiabao Lin and Hanpin Wang. The complexity of Boolean Holant problems with nonnegative weights. *SIAM Journal on Computing*, 47(3):798–828, 2018.

[MP26] Boning Meng and Yicheng Pan. Dichotomies for #CSP on graphs that forbid a clique as a minor. In Philip Bille, Seth Pettie, and Sabine Storandt, editors, *34th Annual European Symposium on Algorithms (ESA 2026)*, volume 388 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:17, Dagstuhl, Germany, 2026. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[MWX25] Boning Meng, Juqiu Wang, and Mingji Xia. The $\mathrm{FP}^{\mathrm{NP}}$ versus #P dichotomy for #EO. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pages 1795–1806. Association for Computing Machinery, 2025. Full version: arXiv:2502.02012.

[MWXZ25] Boning Meng, Juqiu Wang, Mingji Xia, and Jiayi Zheng. From an odd arity signature to a Holant dichotomy. In Srikanth Srinivasan, editor, *40th Computational Complexity Conference (CCC 2025)*, volume 339 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:20, Dagstuhl, Germany, 2025. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. Full version: arXiv:2502.05597.

[SC20] Shuai Shao and Jin-Yi Cai. A dichotomy for real Boolean Holant problems. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1091–1102. IEEE, 2020. Full version: arXiv:2005.07906.

[TW24] Dimitrios M. Thilikos and Sebastian Wiederrecht. Killing a vortex. *Journal of the ACM*, 71(4):27:1–27:56, 2024.

[Val08] Leslie G. Valiant. Holographic algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008.

[Xia26] Mingji Xia. The framework to unify all complexity dichotomy theorems for Boolean tensor networks. arXiv:2603.09417v3, 2026. Version 3, revised 19 July 2026.

[YHF24] Peng Yang, Yuan Huang, and Zhiguo Fu. The computational complexity of Holant problems on 3-regular graphs. *Theoretical Computer Science*, 982:114256, 2024.