A Dichotomy for Planar Graph Homomorphisms with Nonnegative Weights

Chenghua Liu[1] and Boning Meng[2]

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

Abstract

We prove a complete complexity dichotomy for planar graph homomorphism counting with any fixed symmetric nonnegative matrix of arbitrary finite order. The tractable matrices are precisely direct sums of zero blocks and rank-one or bipartite rank-two blocks tensored with zero-field Boolean Ising interactions; all other cases are #P-hard. We also characterize exactly which fixed positive vertex weights preserve tractability. The results hold for algebraic weights and extend to fixed real weights in a prescribed exact representation. Our proof recovers tensor structure even when planar gadgets cannot distinguish colors, combining exact interpolation with matrix-logarithm expansions and entropy maximization. Maximizing logarithmic support produces distance kernels, while an entropy-based continuation argument extends their positive definiteness. The resulting geometry forces Cartesian products of cliques, from which counting-hardness arguments isolate the Boolean tensor factors. Together, these arguments establish the necessity of the stated tractable forms.

Acknowledgments

The authors used OpenAI's ChatGPT in preparing this manuscript, including for language editing and LaTeX preparation. ChatGPT also contributed to the exploratory development of most of the arguments. All claims and proofs were independently checked and finalized by the authors, who take full responsibility for the content.

Contents

1 Introduction

For a fixed symmetric nonnegative matrix M of order q, the partition function

$$Z_M(G) = \sum_{\sigma:V(G)\to[q]} \prod_{uv\in E(G)} M_{\sigma(u),\sigma(v)}$$

counts weighted homomorphisms from G to a fixed target. When M is a zero–one adjacency matrix, this is ordinary graph homomorphism counting. More generally, its colors are spin states and its entries are pair interactions. We ask which fixed interactions admit exact polynomial-time evaluation on planar input graphs.

Planarity changes this question substantially. The zero-field Ising partition function is computable by Pfaffian methods [Kas61, Kas63, TF61], although its nontrivial positive interactions are hard on general graphs [BG05]. Tensor products of Ising interactions give larger tractable state spaces. Do these constructions account for every tractable nonnegative interaction, regardless of the number of states? How does the answer change when each state also carries a fixed positive vertex weight?

Counting complexity and earlier planar classifications. On unrestricted input graphs, the classification developed through a sequence of increasingly general results. Dyer and Greenhill classified unweighted targets [DG00, DG04]; Bulatov and Grohe extended the dichotomy to nonnegative weights [BG05]. Goldberg, Grohe, Jerrum, and Thurley treated algebraic real weights of mixed signs [GGJT10], and Cai, Chen, and Lu completed the classification for algebraic complex weights [CCL13]. The planar problem asks which additional interactions become tractable when the input graphs have no crossings.

Graph homomorphism counting is also a special case of counting constraint satisfaction, or #CSP, with one symmetric binary constraint. Creignou and Hermann established the Boolean unweighted dichotomy [CH96], and Bulatov extended it to arbitrary finite domains [Bul08]. Dyer and Richerby gave an effective criterion for this classification [DR13]. The weighted theory then reached nonnegative algebraic constraints through Cai, Chen, and Lu [CCL16], and complex algebraic constraints through Cai and Chen [CC17]. These results classify unrestricted instances, leaving the extra tractability of planar instances to be understood.

Valiant's holographic algorithms [Val08] supplied a way to obtain such planar algorithms from matchgates and the FKT method. Cai, Lu, and Xia introduced the Holant framework, which expresses counting problems through local constraints on edge assignments [CLX09]. In work first presented in 2010, they established a planar Boolean #CSP classification for symmetric real-valued constraints and showed that holographic matchgate algorithms capture precisely the additional tractable cases [CLX17]. Subsequent classifications extended this picture to complex weights and nonsymmetric Boolean constraints [GW20, CF16]. For the broader class of symmetric Boolean planar Holant problems, however, further tractable families occur beyond the FKT method [CFGW22].

For graph homomorphisms specifically, Cai and Maran classify all real symmetric three-state matrices [CM23], and all nonnegative symmetric four-state matrices of full rank [CM24]. We recall their complete statements in Section 2.7. Their polynomial and analytic methods supply important interpolation tools. Cai, Maran and Young obtain arbitrary-domain dichotomies under diagonal separation: each pair of colors must be distinguishable on the diagonal by a planar edge gadget [CMY26]. Their connection to quantum automorphisms also exposes the obstruction to treating such separation as an automatic preliminary step.

For example, the tensor product of $\begin{pmatrix} 1 & 1/2 \\ 1/2 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 1/3 \\ 1/3 & 1 \end{pmatrix}$ has rank four and a transitive group of coordinate flips. Consequently, every planar edge gadget has equal diagonal entries on all four colors. Yet its partition function is a product of two tractable Ising partition functions. Our classification accommodates these symmetries directly, without requiring diagonal separation.

Our main results: a complete structural classification. Our two main theorems give a complete answer to these questions. For every fixed symmetric nonnegative interaction matrix, with or without fixed positive vertex weights, they identify explicit structural conditions for exact polynomial-time evaluation on planar graphs. Every case outside the stated classes is #P-hard: exact evaluation is at least as hard as counting the satisfying assignments of a Boolean formula. The size of the state space and the interactions are fixed; the polynomial running time is measured in the size of the input graph.

Theorem 1.1 completes the planar counterpart of the general-graph nonnegative dichotomy of Bulatov and Grohe [BG05]. It closes the classification gap left by the small-domain results of Cai and Maran [CM23, CM24] and the color-separation results of Cai, Maran and Young [CMY26]. In particular, it covers the symmetric interactions for which planar gadgets cannot distinguish colors, identified as a remaining obstacle in [CMY26, Section 7]. The classification now applies to every finite number of colors, including singular interactions and forbidden color pairs.

The tractable forms share a simple principle: their partition functions reduce to elementary sums together with independent zero-field Ising computations on the same input graph. To state this structure precisely, for $\rho > 0$, put

$$W(\rho) = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \qquad T(\boldsymbol{\rho}) = \bigotimes_{r=1}^{d} W(\rho_r). \tag{1.1}$$

The interaction $W(\rho)$ assigns weight 1 to equal states and weight ρ to different states. A color of $T(\boldsymbol{\rho})$ is a tuple of d bits, and the interaction is the product of the interactions of the individual bits. An empty tensor product, $d = 0$, is the one-by-one matrix $[1]$. Here $\otimes$ is the matrix Kronecker product. A Boolean-domain matrix has size 2×2, with rows and columns indexed by $\{0, 1\}$; its entries need not be zero or one. The support graph of a nonnegative symmetric matrix M has a possible loop or edge ij precisely when $M_{ij} > 0$. Its connected components index the diagonal blocks used below. Polynomial-time computability refers to exact finite encodings of the output; FP denotes the corresponding class of polynomial-time computable functions. Section 2 specifies these conventions and the computation model.

Write $\mathrm{Sym}_q(S)$ for the symmetric $q \times q$ matrices with entries in S, and $A \cong B$ for equality after a simultaneous permutation of rows and columns.

Theorem 1.1 (Structural dichotomy)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ be fixed with algebraic entries. Its planar graph homomorphism problem* Pl-GH(M) *is in* FP *if $M \cong \bigoplus_j M_j$, where each support component M_j has one of the following forms:*

$$[0], \tag{1.2}$$

$$(aa^{\mathsf{T}}) \otimes T(\boldsymbol{\rho}), \qquad a \in \mathbb{R}^k_{>0}, \tag{1.3}$$

$$\begin{pmatrix} 0 & ab^{\mathsf{T}} \\ ba^{\mathsf{T}} & 0 \end{pmatrix} \otimes T(\boldsymbol{\rho}), \qquad a \in \mathbb{R}^k_{>0},\ b \in \mathbb{R}^{\ell}_{>0}. \tag{1.4}$$

Here $k, \ell \geq 1$, $d \geq 0$, and $\rho_r > 0$ for every $r \in [d]$; the integer dimensions and parameters may differ between components. In every other case the problem is #P-hard under polynomial-time Turing reductions.

The formulas describe how every easy case is assembled. A rank-one interaction $a_i a_j$ separates into one factor for each endpoint of an edge, so its color choices can be summed independently at the input vertices. In (1.3), the remaining bit coordinates are independent Ising models, and their partition functions multiply. The bipartite form (1.4) adds only a choice of which input side maps to which target side; a connected nonbipartite input contributes zero. Direct sums mean that each connected input component chooses one target support component. The zero singleton contributes only to isolated input vertices. Proposition 2.6 gives the resulting evaluation formulas.

The force of the theorem is that these constructions exhaust the possibilities: every fixed nonnegative interaction outside this list remains hard even when all inputs are planar. Thus the additional tractability supplied by planarity is accounted for entirely by the zero-field Ising factors. The factor $W(1)$ can be absorbed into the rank-one part, and the two bipartite sides may have different sizes. Section 10 proves the necessity of the classification.

Our second main result: fixed positive vertex weights. Theorem 1.3 gives a complete classification also when each color has its own fixed preference at every input vertex, as in a spin system with a single-site field. It determines exactly when those preferences can be incorporated into the elementary vertex sums while leaving independent zero-field Ising factors. A fixed vector $w > 0$ is used at every input vertex; its entries may differ between colors. Define

$$Z_{M,w}(G) = \sum_{\sigma:V(G)\to[q]} \prod_{v\in V(G)} w_{\sigma(v)} \prod_{uv\in E(G)} M_{\sigma(u),\sigma(v)}, \qquad Z_M = Z_{M,\mathbf{1}},$$

and denote its planar evaluation problem by Pl-GH(M, w). The classification uses the total weight of each class of identical rows of the original matrix.

Definition 1.2 (Identical-row quotient)**.** For a symmetric matrix M and a positive vector w, partition $[q]$ into the classes $K_1, \dots, K_r$ of exactly identical rows of M. For representatives $i \in K_a$, $j \in K_b$, set

$$R_{ab} = M_{ij}, \qquad \bar{w}_a = \sum_{i\in K_a} w_i.$$

Symmetry makes this well-defined. We call $(R, \bar{w})$ the identical-row quotient of (M, w).

As in Corollary 3.8, direct summation within the classes gives, for every graph,

$$Z_{M,w}(G) = Z_{R,\bar{w}}(G). \tag{1.5}$$

In particular all zero rows, if any, form a single zero class.

We use $W(\rho)$ and $T(\boldsymbol{\rho})$ from (1.1), and write $\mathbf{1}_{2^d} = (1, \dots, 1)^\mathsf{T}$. An empty tensor product is $[1]$.

Theorem 1.3 (One fixed positive vertex-weight vector)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ and $w \in \mathbb{R}^q_{>0}$ have fixed algebraic entries, and let $(R, \bar{w})$ be their quotient in Definition 1.2. For a support-component vertex set C of R, write $R_C = R[C]$ and $\bar{w}_C = \bar{w}|_C$. The problem* Pl-GH(M, w) *is in* FP *if each pair $(R_C, \bar{w}_C)$, after a simultaneous permutation of its colors and weights, has one of the following forms:*

(i) $R_C = [0]$, *with an arbitrary positive scalar weight;*

(ii) *for positive vectors $a, \lambda \in \mathbb{R}^k_{>0}$,*

$$R_C = (aa^\mathsf{T}) \otimes T(\boldsymbol{\rho}), \qquad \bar{w}_C = \lambda \otimes \mathbf{1}_{2^d}; \tag{1.6}$$

(iii) *for $a, \lambda \in \mathbb{R}^k_{>0}$ and $b, \eta \in \mathbb{R}^\ell_{>0}$,*

$$R_C = \begin{pmatrix} 0 & ab^\mathsf{T} \\ ba^\mathsf{T} & 0 \end{pmatrix} \otimes T(\boldsymbol{\rho}), \qquad \bar{w}_C = \begin{pmatrix} \lambda \\ \eta \end{pmatrix} \otimes \mathbf{1}_{2^d}. \tag{1.7}$$

In each component, $k, \ell \geq 1$ when present, $d \geq 0$, and $\rho_j > 0$, $\rho_j \neq 1$; the parameters may differ between components. If these conditions fail, Pl-GH(M, w) *is* #P*-hard. All reductions are polynomial-time Turing reductions in the exact algebraic bit model.*

In (1.6), the condition is $\bar{w}_{(j,x)} = \lambda_j$ for $x \in \{0,1\}^d$: the weight may depend on the amplitude index j, but not on an Ising coordinate. Here λ_j is an aggregated vertex weight, not an eigenvalue. In (1.7) the same condition holds separately on each side, with vectors λ and η that need not agree. The class sum, rather than constancy of the original w_i, is what matters. This identifies exactly which positive vertex fields are compatible with the tractable matrix forms. As R has distinct rows, amplitudes within each of a, b are pairwise distinct; $W(1)$ factors disappear under quotienting. For $d = 0$, $\mathbf{1}_{2^d} = [1]$ and the tensor condition on each side is vacuous. Section 11 proves the theorem; Corollary 11.2 allows zero vertex weights by deleting their colors before quotienting.

For example, let $M = J_q$ be the all-ones matrix and let $w > 0$ be arbitrary. All rows merge into one class, giving

$$R = [1], \qquad \bar{w} = \left[\sum_{i=1}^{q} w_i\right], \qquad Z_{J_q,w}(G) = \left(\sum_{i=1}^{q} w_i\right)^{|V(G)|}.$$

This is case (ii) with $a = [1]$, $\lambda = \bar{w}$ and $d = 0$. There are no nondegenerate Ising coordinates, so the theorem imposes no equality requirement on the original vertex weights.

Consequences for graph homomorphisms and spin systems. The structural forms give several concrete classifications, whose proofs are collected in Section 12.

(i) *Unweighted target graphs.* For every fixed graph H, with loops permitted, counting homomorphisms from planar graphs to H is polynomial-time computable precisely in the familiar cases: every component of H is a fully looped clique, a loopless complete bipartite graph, or a loopless isolated vertex. All other cases are #P-complete (Corollary 12.1). These are exactly the tractable targets in the general-graph dichotomy of Dyer and Greenhill [DG00, DG04]. Thus restricting the input to planar graphs introduces no new tractable unweighted target graphs. The extra tractability available to nonnegative weighted targets comes from their interaction values.

(ii) *Planar bipartite inputs and two sublattices.* Doubling the color domain converts evaluation on planar bipartite inputs into our standard planar problem. This yields a complete classification also for rectangular interactions between two prescribed sublattices, with a fixed positive weight vector on each side (Corollary 12.2). Bipartite input restrictions can genuinely enlarge the tractable class, as Remark 12.3 illustrates.

(iii) *Clock models.* For every fixed $q \geq 2$ and nonzero real coupling K, the standard clock interaction $M_{ij} = \exp(K \cos(2\pi(i-j)/q))$ is tractable on planar graphs exactly for $q = 2, 4$; every other case is #P-hard. A fixed positive vertex-weight vector preserves these tractable cases exactly when it is constant (Corollary 12.5). To our knowledge, no previous exact-complexity classification covers this family for every fixed q. The cyclic symmetry of the zero-field model prevents diagonal separation by edge gadgets. This consequence therefore addresses a natural family beyond that hypothesis, including state counts studied by stochastic estimation and numerical methods [AM14, Kim17].

(iv) *Coupled Ising systems and single-site potentials.* For interactions that couple products of Ising spins across an edge, tractability is equivalent to linear independence, over $\mathbb{F}_2$, of the active character labels. With positive single-site weights, the total weight on each resulting character fiber must also agree (Corollary 12.6). This gives an explicit criterion for generalized Ashkin–Teller systems [GW81]. For the spin-one Blume–Capel model, every nonzero pair coupling remains hard with any finite fixed crystal and magnetic fields (Corollary 12.7). Finally, the same main classifications hold on graphs of any fixed bounded orientable genus (Corollary 12.8).

These applications use the usual statistical-mechanical convention $M_{ij} = e^{-\beta E_{ij}}$, $w_i = e^{-\beta U_i}$: E specifies pair energies and U a fixed single-site potential, applied at every vertex. Zero edge weights impose hard exclusions. This includes state-dependent fields in the Potts model [EM11], crystal fields in the Blume–Capel model [BP18], and coupled Ising layers [GW81]. Identical states contribute their combined Boltzmann weight, so the quotient in Theorem 1.3 has a direct physical interpretation. Our conclusions concern exact evaluation on arbitrary planar input graphs. Approximation and sampling have a different complexity landscape; low-temperature algorithms, for example, are available for several spin systems under the hypotheses of algorithmic Pirogov–Sinai theory [HPR18].

Technical contribution: logarithms, interpolation and entropy. The algorithmic ingredients of the tractable cases are established ones: elementary sums and products of zero-field Ising computations (Proposition 2.6). The main task is proving their necessity without a rank or separation assumption. Our approach uses the matrix logarithm to extract discrete geometry from the family of interactions available by planar reductions. If A is positive definite, the expansion $A^t = I + t \log A + O(t^2)$ exposes its logarithmic support. In Lemma 4.2, orders of vanishing at $t = 0$ and exact multiplicative identities transfer this information to distance kernels. The effective interpolation in Lemma 3.10 turns these identities into finite exact queries to planar gadgets, with polynomial size bounds.

We maximize the off-diagonal support of $\log A$ within one family closed under the required gadget and interpolation operations. The resulting graph R controls the kernel $E_R(x) = (x^{d_R(i,j)})_{i,j}$. Local expansions restrict shortest paths and clique neighborhoods (Lemmas 4.3 and 4.4), but the global geometric argument requires positivity throughout $0 < x < 1$. Lemma 4.5 supplies this extension by maximizing matrix entropy over positive semidefinite completions with prescribed entries. The first-order condition forces zeros in the matrix logarithm outside the prescribed support. Strict concavity and an analytic implicit-function argument identify the completion with $E_R(x)$ and extend its positive definiteness from a neighborhood of zero to the entire interval. Lemma 4.6 then uses the expansion at $x = 1$ to realize graph distance as squared Euclidean distance; together with the local constraints, this forces a Cartesian product of cliques. The proof of Theorem 6.1 eliminates non-Boolean factors by counting hardness.

Matrix functions and their expansions are classical [Hig08]; logarithmic sparsity and entropy geometry are studied in [Pav24, PST24]. Polynomial and analytic interpolation already play a central role in planar counting [CM24]. The contribution here is their combination in exact counting reductions: maximum logarithmic support, entropy-based continuation, and Cartesian geometry yield the common tensor coordinates of Proposition 4.8, without separating colors. Integer Gram arguments yield the zero–one classification (Theorem 7.1), and normalization handles strictly positive matrices, including singular ones (Theorem 8.1). Finite weighted moments then determine the permitted vertex weights (Sections 11.2 and 11.3). For rectangular blocks, Proposition 9.2 uses a Fourier degree argument and a parallel-edge norm identity to align the coordinates on the two sides.

Organization and computation model. Section 2 fixes notation and the exact algebraic computation model, recalls known counting results, and gives the algorithms for the tractable forms. Section 3 establishes the planar interpolation and vertex-weight reductions. Section 4 develops maximum logarithmic support, entropy-based continuation, and Cartesian geometry.

Section 5 proves hardness for Boolean tensor products with an unequal-diagonal factor, and Section 6 completes the positive definite classification.

Section 7 derives the zero–one dichotomy through integer Gram matrices. Section 8 treats strictly positive matrices and proportional rows, while Section 9 handles positive rectangular blocks and aligns their two sets of tensor coordinates. Section 10 assembles the support and block classifications to prove Theorem 1.1. Section 11 proves the fixed positive vertex-weight classification, Theorem 1.3. Section 12 gives the consequences for graph homomorphisms, bipartite inputs, spin systems, and bounded genus.

Appendix A extends both main classifications to arbitrary fixed real weights in the prescribed field representation of [CMY26, Appendix A], the model used for real Boltzmann parameters in the applications. It also proves signed support and magnitude reductions; a full signed-edge dichotomy remains open.

2 Preliminaries

2.1 Counting problems and exact computation

For a positive integer q, write $[q] = \{1, \dots, q\}$. All input graphs are finite undirected multigraphs, including graphs with loops. Edge products count parallel edges with their multiplicities, and a loop contributes twice to the degree of its incident vertex. An isolated vertex has degree zero. We write $c(G)$ for the number of connected components, including isolated vertices, and use the value 1 for an empty product.

A target matrix and its domain size are fixed; only the input graph varies. For a symmetric matrix M and positive vertex weights $w = (w_1, \dots, w_q)$, define

$$Z_M(G) = \sum_{\sigma:V(G)\to[q]} \prod_{uv\in E(G)} M_{\sigma(u),\sigma(v)}, \tag{2.1}$$

$$Z_{M,w}(G) = \sum_{\sigma:V(G)\to[q]} \prod_{v\in V(G)} w_{\sigma(v)} \prod_{uv\in E(G)} M_{\sigma(u),\sigma(v)}. \tag{2.2}$$

The unweighted vertex convention is $w = \mathbf{1}$, so $Z_M = Z_{M,\mathbf{1}}$. In particular,

$$Z_{M,w}(\varnothing) = 1, \qquad Z_{M,w}(K_1) = \sum_{i=1}^{q} w_i,$$

where K_1 is a loopless isolated input vertex.

The vector w is fixed and is the same at every input vertex; this notation does not make arbitrary unary labels part of the input. Positive vertex weights arise when identical colors are merged, when edge matrices are normalized, and when unary gadgets are attached. Their removal in Lemma 3.7 is a one-way reduction: under the stated no-zero-row and no-proportional-row hypotheses, $\text{Pl-GH}(M) \le_T \text{Pl-GH}(M, w)$. It transfers unweighted hardness to the weighted problem, but does not transfer unweighted tractability or assert a Turing equivalence. The separate weighted classification is proved in Section 11.

Unless explicitly stated otherwise, the counting problems in the main text have fixed algebraic weights in the domains

$$M \in \mathrm{Sym}_q(\mathbb{R}), \qquad w \in \mathbb{R}^q_{>0}.$$

Thus signed edges do not supply signed background vertex weights. For algebraic weights, exact arithmetic uses a fixed number field containing all source weights and the finitely many auxiliary constants. The extension to arbitrary fixed real weights is treated separately in Appendix A.

Following [CMY26], Pl-GH(M) denotes exact evaluation of Z_M on planar graphs; Pl-GH(M, w) denotes the weighted variant. For a zero–one matrix M, $Z_M(G)$ is the number of graph homomorphisms from G to the graph represented by M.

All reductions are polynomial-time Turing reductions, written $\leq_T$. We call two problems A, B polynomial-time Turing equivalent if $A \leq_T B$ and $B \leq_T A$, and write $A \equiv_T B$ for this relation. Membership in FP means polynomial-time computation of the specified finite encoding of the output. Hardness and completeness for #P refer to these reductions. The class #P consists of functions counting the accepting computations of nondeterministic polynomial-time machines. Algebraic weighted partition functions need not be integers; their hard cases are therefore stated as #P-hard. Zero–one homomorphism counting is integer-valued and belongs to #P.

Initially all fixed weights are real algebraic numbers. Computation takes place in a fixed number field containing them, with exact arithmetic and bit complexity. Fixed gadgets, fixed auxiliary algebraic constants and other data depending only on the fixed counting problem may be included in a reduction. Appendix A specifies the separate representation and exact arithmetic for arbitrary fixed real weights. The two algebraic main theorems and their proofs use only the number-field model above.

2.2 Planar gadgets and joint simulation

For a single unweighted matrix M, an edge gadget is a finite graph $\mathbf{K}$ with two distinct labelled vertices ℓ_1, ℓ_2, called its *terminals*. Its *signature* is the matrix

$$\mathbf{K}(M)_{ij} = \sum_{\substack{\tau: V(\mathbf{K}) \to [q] \\ \tau(\ell_1)=i,\ \tau(\ell_2)=j}} \prod_{uv \in E(\mathbf{K})} M_{\tau(u),\tau(v)}.$$

It is a planar edge gadget if it has a planar embedding with both labelled vertices on the outer face. We use a signature as a symmetric binary constraint only when the signature is symmetric. The collection of these actual symmetric signatures is denoted by Pl-$\mathfrak{E}(M)$, as in [CMY26, Definition 1]. Replacing each edge of a planar graph by $\mathbf{K}$ gives

$$Z_M(\mathbf{K}G) = Z_{\mathbf{K}(M)}(G).$$

Thus an actual gadget gives a reduction. A converse is not asserted: a matrix obtained by Turing interpolation need not belong to Pl-$\mathfrak{E}(M)$. Without the planarity and outer-face conditions, the corresponding collection of symmetric signatures is denoted by $\mathfrak{E}(M)$.

We also need mixed instances. A finite constraint language $\mathcal{F}$ consists of fixed binary matrices and fixed unary functions on a specified domain. An instance Ω consists of a planar graph G, a binary label $C_e \in \mathcal{F}$ on each edge, and a finite multiset $\mathcal{U}_v$ of unary labels at each vertex. For fixed background vertex weights $w > 0$, its partition function is

$$Z_{\mathcal{F},w}(\Omega) = \sum_{\sigma: V(G) \to [q]} \prod_{e=uv \in E(G)} (C_e)_{\sigma(u),\sigma(v)} \prod_{v \in V(G)} \left(w_{\sigma(v)} \prod_{u \in \mathcal{U}_v} u_{\sigma(v)} \right).$$

The unit-background convention is $w = \mathbf{1}$, and Pl-GH($\mathcal{F}$) := Pl-GH($\mathcal{F}, \mathbf{1}$). When a weighted source satisfies Lemma 3.7, this unit-background model is obtained by adding a D^{-1}-loop at each vertex, where $D = \mathrm{diag}(w)$. Adding a unary u at an existing vertex v means

$$\mathcal{U}_v \longleftarrow \mathcal{U}_v \uplus \{u\}, \qquad \text{with additional factor } u(\sigma(v)).$$

This changes the labels at v and leaves $V(G)$ unchanged; any gadget realizing that factor is specified separately. An ordinary Pl-GH(M, w)-instance has only the edge type M and no additional unary labels. Edge matrices are symmetric unless their two endpoint domains are explicitly distinguished. When domains are distinguished, each vertex v has a prescribed colour set X_v, and every assignment satisfies $\sigma(v) \in X_v$. The signature of a two-terminal gadget is indexed by $X_{\ell_1} \times X_{\ell_2}$; its internal vertices may have other prescribed domains. In particular, an X–Y constraint is a rectangular matrix in $\mathbb{R}^{X \times Y}$. These prescribed domains are part of the instance model; they do not supply an unrestricted pinning oracle.

On the domain $[q]$, a weighted mixed gadget $\mathbf{K}$ with terminal set T and coloring $\eta : T \to [q]$ has conditional signature

$$\sum_{\substack{\tau:V(\mathbf{K})\to[q]\\ \tau|_T=\eta}} \prod_{e=uv\in E(\mathbf{K})} (C_e)_{\tau(u),\tau(v)} \prod_{v\in V(\mathbf{K})\setminus T} w_{\tau(v)} \prod_{v\in V(\mathbf{K})} \prod_{u\in\mathcal{U}_v} u_{\tau(v)}.$$

With prescribed endpoint domains, restrict τ and η to those domains. Background weights are omitted precisely at the terminals; explicit terminal unaries are retained. A rooted graph is a pair (F, r_F) with $r_F \in V(F)$; its terminal set is $\{r_F\}$. For $|T| = 1$, the displayed signature is the rooted conditional function. Gluing therefore counts each background weight once. A *port* in a vertex replacement corresponds to an incident half-edge; a loop has two such half-edges. Replacements preserve the cyclic order of ports in a planar embedding.

Definition 2.1 (Joint availability)**.** An auxiliary finite set of constraints $\mathcal{A}$ is *jointly available* from a fixed source language $\mathcal{F}$ if, with the same fixed background weights and endpoint domains,

$$\text{Pl-GH}(\mathcal{F} \cup \mathcal{A}, w) \leq_T \text{Pl-GH}(\mathcal{F}, w).$$

Here both problems permit arbitrary planar mixed instances. The number of constraint types is fixed; their occurrence counts are input data. For a family with a rational parameter x, uniform availability means cost $\text{poly}(|\Omega| + \text{bit}(x))$, where $|\Omega|$ is the encoded instance size and $\text{bit}(x)$ is the binary length of its reduced numerator and denominator.

The words *available* and *simulation* always have this joint meaning. They permit the auxiliary and source constraints to coexist in the same instance. A single homogeneous reduction $\text{Pl-GH}(N) \leq_T \text{Pl-GH}(M)$ alone does not establish this property. An interpolated signature is said to be simulated, while an actual gadget is said to realize its signature. Gadget substitutions and interpolation retain the specified colour coordinates. Basis changes, pinning, and linear combinations require their own justification.

A known nonzero scalar is removed from values by

$$Z_{cM}(G) = c^{|E(G)|} Z_M(G) \qquad (c \neq 0).$$

In a mixed instance with s marked occurrences, the factor is c^s; exact multiplication or division performs this normalization.

Series and parallel composition of planar edge gadgets, and attachments at a root, preserve planarity. In the unweighted vertex convention, an n-edge path realizes M^n, and n parallel edges realize $M^{\circ n}$, for positive integers n. These are the stretching and thickening gadgets $\mathbf{S}_n, \mathbf{T}_n$ of [CMY26]. With $D = \text{diag}(w)$, a path of integer length $n \geq 1$ has signature

$$M(DM)^{n-1}; \qquad n = 2: \quad MDM.$$

If $D_M = \text{diag}(M_{11}, \ldots, M_{qq})$, adding n loops at each endpoint realizes $D_M^n M D_M^n$. This is the loop gadget $\mathbf{L}_n$. For entrywise positive M, the notation $\mathcal{T}_M(t)$ in that paper means $M^{\circ t}$; for a symmetric positive definite matrix M, its $\mathcal{S}_M(t)$ means M^t. Noninteger powers and interpolated projectors are handled by reductions, not by noninteger-sized graphs.

2.3 Matrices, supports and graph products

Write $\mathrm{Mat}_{p,q}(S)$ for the $p \times q$ matrices with entries in S, abbreviate $\mathrm{Mat}_{q,q}(S)$ to $\mathrm{Mat}_q(S)$, and write $\mathrm{Sym}_q(S)$ for the symmetric ones. Following [CMY26], $\mathrm{Sym}_q^{\mathrm{F}}(S)$ denotes the full-rank members of $\mathrm{Sym}_q(S)$, and $\mathrm{Sym}_q^{\mathrm{pd}}(S)$ denotes its positive definite members. Throughout this paper, positive definiteness and positive semidefiniteness are understood for real symmetric matrices. We abbreviate these terms as PD and PSD, respectively. For such a matrix A, we write $A \succeq 0$ when $x^{\mathsf{T}}Ax \geq 0$ for every nonzero real vector x, and $A \succ 0$ when $x^{\mathsf{T}}Ax > 0$ for every nonzero real vector x. These conditions are distinct from entrywise nonnegativity and strict positivity, denoted by $A \geq 0$ and $A > 0$, respectively. The matrices I_q, J_q are the identity and all-ones matrices; $\mathbf{1}$ is the all-ones column vector of the required dimension. For a subset X of the domain, $\mathbf{1}_X$ is its indicator vector: its i-th entry is one if $i \in X$, and zero otherwise. For row and column index sets S, T, $A[S,T] = (A_{ij})_{i \in S, j \in T}$, and $A[S] = A[S,S]$ when both index sets are S. An asterisk retains every index: $A[S,*] = A[S,[q]]$ for a matrix with q columns. We also write $A_{S,T}$ for $A[S,T]$, $A_i = A[i,*]$ for a row, and $w|_S = (w_i)_{i \in S}$ for a restricted vector. When a subscript C labels a support component, $A_C = A[C]$ and $w_C = w|_C$. The transpose is A^{T}, the trace is $\operatorname{Tr} A = \sum_i A_{ii}$, and $\operatorname{diag}(v)$ is the diagonal matrix with diagonal v.

Let $\operatorname{sgn} : \mathbb{R} \to \{-1, 0, 1\}$ be the sign function, defined by

$$\operatorname{sgn}(x) = \begin{cases} 1, & x > 0, \\ 0, & x = 0, \\ -1, & x < 0. \end{cases}$$

For a real symmetric matrix M, define

$$P(M)_{ij} = \mathbf{1}_{\{M_{ij} \neq 0\}}, \qquad |M|_{ij} = |M_{ij}|, \qquad S(M)_{ij} = \operatorname{sgn}(M_{ij}).$$

Thus $S(M)$ records the sign of each entry of M. The support graph Γ_M has vertex set $[q]$ and an edge ij, including a possible loop, precisely when $M_{ij} \neq 0$. Thus its zero–one adjacency matrix is $P(M)$; for a nonnegative matrix this is the existing notation $\mathbf{1}_{M>0}$. A support component of M means the principal block indexed by a connected component of Γ_M. We call M connected or bipartite when Γ_M has that property. A loop in Γ_M contributes once to a matrix row sum; this differs from the input-graph degree convention.

We call a support component *basic* if it is a fully looped complete graph, a loopless complete bipartite graph with two nonempty sides, or a loopless isolated vertex. For the vertex set C of a support component, put $N(i) = \{j : M_{ij} \neq 0\}$. The equivalent row-support condition is

$$\forall i, j \in C : \quad N(i) = N(j) \text{ or } N(i) \cap N(j) = \varnothing.$$

For the converse, choose a vertex with nonempty neighborhood U. Any two vertices in U have a common neighbor, so they have the same neighborhood, say V. Similarly, all vertices of V have neighborhood U. The component is therefore $U \cup V$. If $U \cap V \neq \varnothing$, then $U = V$ and the component is a fully looped clique; otherwise it is complete bipartite. A zero row gives the isolated case.

When a support quotient is used, let C be the vertex set of a chosen support component and let $A = P(M)[C]$ be the principal submatrix with row and column indices in C. Declare two vertices equivalent when their rows of A are identical, and write the classes as $I_1, \ldots, I_r$. Choose a representative $r_a \in I_a$ from each class and set

$$B_{ab} = A_{r_a r_b} \qquad (a, b \in [r]).$$

Symmetry and equality of the rows make this independent of the representatives. The matrix B retains one row and column from each class; it does not add rows or columns. For $i \in I_a$,

$$\sum_{j\in C} A_{ij} = \sum_{b=1}^{r} B_{ab}|I_b|, \qquad \sum_{b=1}^{r} B_{ab} \text{ is the unweighted row sum of } B.$$

Thus the original row sums count class sizes, while the row sums of the quotient do not. Regularity must be checked separately.

For example, start with $K_{2,3}$ on parts $\{u_1, u_2\}$ and $\{v_1, v_2, v_3\}$, add the edge v_1v_2, and add a loop at v_3. In the displayed vertex order, the adjacency matrix and its identical-row classes are

$$A = \begin{pmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 \end{pmatrix}, \qquad \begin{array}{ll} I_1 = \{u_1, u_2\}, & I_2 = \{v_1\}, \\ I_3 = \{v_2\}, & I_4 = \{v_3\}. \end{array}$$

Every row sum of A is 3: each u_i sees the three v's, each of v_1, v_2 sees both u's and the other one, and v_3 sees both u's and its loop. With representatives u_1, v_1, v_2, v_3, the quotient is

$$B = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}, \qquad B\begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ 2 \\ 2 \\ 2 \end{pmatrix}, \qquad B\begin{pmatrix} 2 \\ 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ 3 \\ 3 \\ 3 \end{pmatrix}.$$

A loop in the graph represented by A contributes one to these matrix row sums, as stipulated above. This example shows that merging identical rows can destroy regularity even though class-size-weighted row sums are preserved.

This construction is the twin quotient of $P(M)[C]$: its classes are defined by equal support rows. For a numerical matrix, an identical-row quotient instead uses equality of its actual entries. Merging or expanding support twins preserves the three basic forms, so A is basic if and only if B is basic. For counting with the zero–one matrix A, the sizes $|I_a|$ must be retained as vertex weights on B, by Corollary 3.8; the unweighted partition functions of A and B need not agree.

The Hadamard product is $(A \circ B)_{ij} = A_{ij}B_{ij}$; $A^{\circ n}$ is an entrywise power, whereas A^n is an ordinary matrix power. For noninteger entrywise powers the positive real branch is used on positive entries, and any treatment of zeros is specified separately. The Kronecker product satisfies $(A \otimes B)_{(i,a),(j,b)} = A_{ij}B_{ab}$. Direct sum, denoted by $A \oplus B$, means block diagonal sum. A Kronecker sum means an expression such as $A \otimes I + I \otimes B$, rather than a direct sum. For matrices of equal order, $A \cong B$ means equality after one simultaneous row and column permutation. Vector norms are Euclidean. For matrices we use the Frobenius inner product and norm

$$\langle A, B\rangle_F = \mathrm{Tr}(A^\mathsf{T} B), \qquad \|A\|_F^2 = \sum_{i,j} |A_{ij}|^2.$$

The Schur product theorem gives

$$A, B \succeq 0 \Rightarrow A \circ B \succeq 0, \qquad A, B \succ 0 \Rightarrow A \circ B \succ 0.$$

For real vectors $v_1, \ldots, v_r$, their Gram matrix $G = (\langle v_i, v_j\rangle)_{i,j=1}^r$ always satisfies $G \succeq 0$. Moreover,

$$G \succ 0 \iff v_1, \ldots, v_r \text{ are linearly independent.}$$

For a connected nonnegative symmetric matrix with positive diagonal, Perron–Frobenius theory implies that all sufficiently large integer powers are entrywise positive. For an entrywise positive $n \times n$ matrix M, define

$$r(M) := \max\{|\lambda| : \lambda \text{ is an eigenvalue of } M\}.$$

The Perron–Frobenius theorem states that $r(M) > 0$ is a simple eigenvalue, strictly larger in modulus than every other eigenvalue, and that

$$\exists v \in \mathbb{R}^n_{>0} : \quad Mv = r(M)v, \qquad \ker(M - r(M)I) = \operatorname{span}\{v\}.$$

Such a positive eigenvector is called a *Perron vector*. Every positive eigenvector of M has eigenvalue $r(M)$ and equals cv for some $c > 0$; imposing $\sum_i v_i = 1$ makes it unique. We use these conclusions under the stated positivity hypotheses; see [HJ13].

For a connected simple graph R, $d_R(i,j)$ is graph distance. Write $V(R), E(R)$ for its vertex and edge sets and $N_R(i)$ for the set of vertices adjacent to i. We write $i \sim_R j$ if $ij \in E(R)$, and omit the subscript when the graph is clear. The Cartesian product $R_1 \square R_2$ has vertex set $V(R_1) \times V(R_2)$ and

$$(i,a) \sim (j,b) \iff \big(i = j,\ a \sim_{R_2} b\big) \text{ or } \big(a = b,\ i \sim_{R_1} j\big).$$

Write K_s for the complete simple graph on s vertices. A graph used as the parameter of Pl-GH$(\,\cdot\,)$ denotes its adjacency matrix. For an integer $d \geq 1$, $R^{\square d}$ denotes the Cartesian product of d copies of R. In particular, the Boolean cube is

$$Q_d = K_2^{\square d} = \underbrace{K_2 \square \cdots \square K_2}_{d \text{ factors}}.$$

Its vertices are indexed by $\{0,1\}^d$, with two vertices adjacent exactly when they differ in one coordinate. For example, Q_2 is the four-cycle $00, 01, 11, 10, 00$. Its graph distance is the Hamming distance $d_H(x,y) = |\{r : x_r \neq y_r\}|$. By convention, $Q_0 = K_1$, the graph with one vertex and no edges. A clique is a set of pairwise adjacent vertices; it is maximal if it is not properly contained in another clique.

2.4 Matrix functions and analytic facts

Section 4 uses matrix logarithms to extract a graph, entropy to extend a distance kernel across a real interval, and a commutator coefficient to separate tensor coordinates. The facts below supply the analytic operations in those arguments.

For a PD matrix $H = \sum_a \lambda_a P_a$, the $\lambda_a > 0$ are its distinct eigenvalues and P_a projects orthogonally onto the λ_a-eigenspace. In particular, $P_a^{\mathsf{T}} = P_a$, $P_a P_b = \mathbf{1}_{\{a=b\}} P_a$, and $\sum_a P_a = I$. Functional calculus gives

$$\log H = \sum_a (\log \lambda_a) P_a, \qquad H^t = \sum_a \lambda_a^t P_a = \exp(t \log H).$$

Here $\log H$ is the matrix logarithm, never an entrywise logarithm. It is real symmetric and depends analytically on H in the PD cone. The matrix exponential is $\exp C = \sum_{n \geq 0} C^n / n!$. It maps real symmetric matrices bijectively onto PD matrices, with inverse the matrix logarithm. For $A \in \mathbb{R}^{m \times m}$ and $B \in \mathbb{R}^{n \times n}$, the Kronecker product multiplication rule gives

$$(A \otimes I_n)(I_m \otimes B) = A \otimes B = (I_m \otimes B)(A \otimes I_n).$$

Thus these two summands commute; no condition $AB = BA$ is required (the matrices A, B may even have different orders). For commuting matrices X, Y, the exponential series gives $e^{X+Y} = e^X e^Y$. The same series gives $e^{A\otimes I_n} = e^A \otimes I_n$ and $e^{I_m \otimes B} = I_m \otimes e^B$, so

$$\begin{aligned}\exp(A \otimes I_n + I_m \otimes B) &= \exp(A \otimes I_n)\exp(I_m \otimes B) \\ &= (e^A \otimes I_n)(I_m \otimes e^B) \\ &= e^A \otimes e^B.\end{aligned}$$

For a PSD matrix, the orthogonal projector onto its range is the sum of the spectral projectors for its positive eigenvalues. A principal matrix function and an entrywise matrix function are different operations.

For a differentiable matrix function f, its Fréchet derivative Df_C is the linear map satisfying

$$f(C + V) = f(C) + Df_C[V] + o(\|V\|_F) \quad (V \to 0).$$

For symmetric C, V,

$$D\exp_C[V] = \int_0^1 e^{(1-t)C} V e^{tC}\, dt, \qquad \operatorname{Tr}(V D\exp_C[V]) > 0 \quad \text{if } V \neq 0.$$

On the PD cone put

$$\Phi(H) = \operatorname{Tr}(H \log H - H).$$

This function is strictly convex and $D\Phi_H[V] = \operatorname{Tr}((\log H)V)$. The von Neumann entropy functional is

$$-\operatorname{Tr}(H \log H) = -\sum_j \lambda_j(H) \log \lambda_j(H), \qquad -\Phi(H) = -\operatorname{Tr}(H \log H) + \operatorname{Tr} H.$$

Thus minimizing Φ on a fixed-trace set maximizes this entropy. For $\operatorname{Tr} H = q > 0$, the normalized matrix $\rho = H/q$ has entropy $-\operatorname{Tr}(\rho \log \rho) = -\operatorname{Tr}(H \log H)/q + \log q$, with the same optimizers. The spectral formulas extend continuously to PSD matrices by $0 \log 0 = 0$. Matrix functional calculus and derivatives are treated in [Hig08]. For the entropy functional, the derivative follows from $\frac{d}{d\lambda}(\lambda \log \lambda - \lambda) = \log \lambda$; the positive Hessian below gives strict convexity. Lemma 4.5 proves existence, interiority and continuation for the constrained optimizer used here. The derivative of the logarithm is

$$D\log_H[V] = \int_0^\infty (H + sI)^{-1} V (H + sI)^{-1}\, ds.$$

Consequently the Hessian of Φ, meaning its second derivative as a quadratic form, satisfies $\langle V, D\log_H[V]\rangle_F > 0$ for nonzero symmetric V.

A real analytic function agrees locally with a convergent power series. The identity theorem says that two such functions on a connected open set that agree on a nonempty open subset agree everywhere; on an interval, an interior accumulation point of zeros also suffices. The analytic implicit function theorem states that, for an analytic map $F : U \to \mathbb{R}^n$ on an open set $U \subseteq \mathbb{R}^n$,

$$F(c_0) = y_0, \quad \det DF_{c_0} \neq 0 \quad \Longrightarrow \quad F(c(y)) = y$$

has a unique local analytic solution with $c(y_0) = c_0$. We use this with the finite-dimensional space of symmetric matrices having a specified support.

For matrices, the commutator is $[A, B] = AB - BA$. The local logarithm of a product of matrix exponentials has a convergent Baker–Campbell–Hausdorff expansion. The coefficient needed in Proposition 4.8 is

$$[t^2 u] \log(e^{tA/2} e^{uB} e^{tA/2}) = -\frac{1}{24}[A, [A, B]].$$

Here $[t^iu^j]f$ denotes the coefficient of t^iu^j in the Taylor series of f at $(0,0)$. Expanding the exponentials and $\log(I+V) = V - V^2/2 + V^3/3 + O(\|V\|_F^4)$ verifies this finite-degree identity directly. The notations $O(t^k)$ and $o(t^k)$ refer to the limit under discussion, entrywise for fixed-size matrices.

2.5 Algebraic functions and bit complexity

The tools in this subsection have two roles in the reductions. Algebraic functions and their branches constrain identities that hold as a parameter varies. Field norms then bound the number of exceptional samples, while traces and logarithmic heights control output representations and exact bit complexity. The eigenvalue interpolation in Section 5 uses all of these steps; the examples below indicate how they fit together.

A number field K is a finite extension of $\mathbb{Q}$. In a fixed basis, its elements have rational coordinates, so exact equality and field arithmetic reduce to rational arithmetic. A splitting field of a polynomial over K is the smallest extension of K containing all its roots, inside a chosen algebraic closure. For distinct field elements $\lambda_1, \ldots, \lambda_s$, the Vandermonde matrix $(\lambda_j^{i-1})_{i,j=1}^s$ has determinant $\prod_{i<j}(\lambda_j - \lambda_i) \neq 0$. Thus evaluations at distinct points recover a polynomial of degree at most $s-1$, and values of $\sum_j a_j \lambda_j^h$ for $h = 0, \ldots, s-1$ recover its coefficients a_j. Write $K[x]$ for the polynomials in a symbolic parameter x with coefficients in K. Its rational-function field is

$$K(x) = \{p(x)/q(x) : p, q \in K[x],\ q \not\equiv 0\}.$$

For example, $(1+x)/(1-x)$ belongs to $K(x)$, but not to $K[x]$. Equality here is equality of rational functions, rather than agreement at a single point. Substitution $x = r$, away from the poles, turns the symbolic expression into a numerical value.

An algebraic function over $K(x)$ satisfies a nonzero polynomial equation whose coefficients lie in $K(x)$. For example, adjoining a root s of $s^2 = 1 + x^2$ gives the quadratic extension $L = K(x)(s)$, with basis $1, s$ over $K(x)$. Symbolically, one reduces powers using $s^2 = 1 + x^2$. Numerically, at a real sample r, the two possible values are $\pm\sqrt{1+r^2}$. A branch is a locally single-valued analytic choice of the algebraic function away from its singularities. Near $x = 0$, for example,

$$s_+(x) = \sqrt{1+x^2}, \quad s_+(0) = 1, \qquad s_-(x) = -s_+(x), \quad s_-(0) = -1.$$

The choice must remain analytic as x varies locally; it is not an arbitrary new choice of root at every point. The order $\operatorname{ord}_{x=a} f$ of a nonzero meromorphic branch is the exponent k in $f(x) = (x-a)^k g(x)$, with $g(a) \neq 0$. Orders add in products; negative orders denote poles.

For analytic continuation we allow the parameter to move in the complex plane. Continuing a branch along a loop can return a different branch; this permutation is called monodromy. In the basic example $y^2 = x$, the loop $x = e^{i\theta}$, $0 \leq \theta \leq 2\pi$, lifts from $y = 1$ to $y = e^{i\theta/2}$, which ends at -1. Thus going once around the branch point changes the sign of the square root. For $s^2 = 1+x^2$, the corresponding branch points are i and $-i$, the two simple zeros of $1+x^2$. More generally, continuation around a simple zero of a polynomial changes the sign of its square root. If the other radicands are holomorphic and nonzero there, a sufficiently small loop leaves their local square-root branches unchanged. Positive square roots at a real sample specify the starting branches. In Section 5, such a sign change exchanges two eigenvalue branches and inverts their ratio, thereby constraining multiplicative function identities. Complex continuation is used to prove those identities; the oracle queries still use the prescribed real sample parameters. For $f_1, \ldots, f_b \in K(x)$, the extension $K(x)(\sqrt{f_1(x)}, \ldots, \sqrt{f_b(x)})$ has degree at most 2^b; no independence of the radicands is implicit.

For a finite extension L/K, its degree $[L : K]$ is its dimension over K. Relative trace and norm use the multiplication map:

$$m_z : L \to L, \quad y \mapsto zy, \qquad \mathrm{Tr}_{L/K}(z) = \mathrm{Tr}(m_z), \quad \mathrm{Norm}_{L/K}(z) = \det(m_z).$$

In characteristic zero they are also the sum and product over the K-embeddings into an algebraic closure. The images under these embeddings are called algebraic conjugates. An element is integral over $K[x]$ when it satisfies a monic polynomial with coefficients in $K[x]$. Its norm in a finite extension of $K(x)$ is a rational function integral over $K[x]$. Such a rational function, written with coprime numerator and denominator, cannot have a nonconstant denominator, so this norm belongs to $K[x]$. We also use unique factorization for polynomials over a field: an irreducible polynomial dividing a product divides a factor. Two polynomials are associates if they differ by a nonzero constant. This field norm is not a Euclidean or Frobenius norm. For the example $L = K(x)(s)$, $s^2 = 1 + x^2$, the multiplication map for $z = a + bs$, with $a, b \in K(x)$, has matrix

$$[m_{a+bs}]_{(1,s)} = \begin{pmatrix} a & b(1+x^2) \\ b & a \end{pmatrix}.$$

Consequently

$$\mathrm{Tr}_{L/K(x)}(a + bs) = 2a, \qquad \mathrm{Norm}_{L/K(x)}(a + bs) = a^2 - b^2(1 + x^2).$$

These are the sum and product of the conjugates $a + bs$ and $a - bs$. If $a, b \in K[x]$, this norm is a polynomial; arbitrary rational-function coefficients need not give a polynomial norm. For instance, the nonzero integral element $s - 1$ has norm $-x^2$, so a zero of either of its finite branches can occur only at $x = 0$. In Section 5, the same idea is applied to a nonzero difference of formal spectral products that is integral over $K[x]$: its nonzero polynomial norm bounds the number of sample points at which the products can accidentally coincide.

Trace has a different use. For a finite extension E/F, if a computed $z \in E$ is already known to lie in the base field F, then

$$[E : F]^{-1} \,\mathrm{Tr}_{E/F}(z) = z.$$

This identity converts its representation back to F. It does not hold for an arbitrary extension element. In the quadratic example, half the trace of $a + bs$ is a, which equals $a + bs$ only when $b = 0$.

The absolute logarithmic height of an algebraic number α with primitive integer minimal polynomial $a_d \prod_{i=1}^{d}(X - \alpha_i)$, $a_d > 0$, is

$$h(\alpha) = \frac{1}{d}\left(\log a_d + \sum_{i=1}^{d} \log \max\{1, |\alpha_i|\}\right).$$

Here primitive means that the integer coefficients have greatest common divisor one; this polynomial is an integer multiple of the monic minimal polynomial over $\mathbb{Q}$. For a reduced rational number p/q, with $q > 0$, this gives $h(p/q) = \log \max\{|p|, q\}$: height measures the size of the exact numerator and denominator, up to constant factors in their binary lengths. In bounded algebraic degree, polynomial logarithmic-height bounds likewise give polynomial bounds on the sizes of exact algebraic encodings. Fixed-size matrices over a fixed number field have eigenvalues in an extension of bounded degree. With polynomial input logarithmic heights, their eigenvalues and spectral products of polynomial length have polynomial logarithmic-height bounds. The same holds for the unique solutions of polynomial-size nonsingular exact linear systems over a common bounded-degree number field, with polynomial coefficient logarithmic heights. These assertions require a fixed base field or uniformly bounded extension degrees; matrix size alone is insufficient if entry degrees are allowed to grow.

In a parameter-dependent interpolation, each sample may require its own extension field. The compositum is the smallest field containing all of these fields inside a fixed algebraic closure. For example,

$$\mathbb{Q}(\sqrt{2})\mathbb{Q}(\sqrt{3}) = \mathbb{Q}(\sqrt{2}, \sqrt{3})$$

has degree four over $\mathbb{Q}$, although each individual field has degree two. The compositum is not their set-theoretic union: it must also contain sums and products of elements from both fields. A uniform degree bound on the individual sample fields does not give a uniform degree bound on their compositum as the number of samples grows.

We avoid that compositum. First perform the inner computation in the bounded-degree extension for its own sample. When the recovered result is known to belong to the fixed base field, use relative trace to convert its representation back to that field. Then perform the outer interpolation entirely over the base field; see Remark 3.4 and Section 5. Every uniform parameter claim also requires effectively computable entries of polynomial bit length in the stated fixed or bounded-degree fields. A power with a fixed rational exponent may use a fixed extension; allowing an arbitrary exponent as part of the input does not automatically satisfy these bounds. The effective interpolation arguments below establish the specific degree, sample-count and logarithmic-height bounds they require.

For $m \in \mathbb{R}^s_{>0}$, its multiplicative relations form

$$\Lambda(m) = \left\{ z \in \mathbb{Z}^s : \prod_{\alpha=1}^{s} m_\alpha^{z_\alpha} = 1 \right\}.$$

This subgroup of $\mathbb{Z}^s$ has a finite integer basis. Its orthogonal complement is taken in $\mathbb{R}^s$ with the usual dot product; because it is defined by integer linear equations, it has a basis of rational, and hence after rescaling integer, vectors.

2.6 Fourier coordinates on the Boolean cube

In Section 9, Fourier coordinates diagonalize known cube kernels. Varying their eigenvalues distinguishes sets of different cardinalities; an entrywise-square gadget then determines how the two colour domains correspond.

For $S \subseteq [d]$ and $x \in \{0,1\}^d$, the character $\chi_S(x) = (-1)^{\sum_{r\in S} x_r}$ has degree $|S|$. The normalized Walsh matrix

$$F_{x,S} = 2^{-d/2}\chi_S(x)$$

is orthogonal. The degree-h Fourier subspace is the span of its columns indexed by $|S| = h$. The symmetric difference of subsets is denoted by $S \triangle T$; characters multiply by $\chi_S\chi_T = \chi_{S\triangle T}$. For vectors f, g indexed by the cube, write $\widehat{f} = F^\mathsf{T} f$, and similarly for other vectors. In these orthonormal coordinates, entrywise multiplication gives

$$\widehat{f \circ g}_S = 2^{-d/2} \sum_{T\subseteq[d]} \widehat{f}_T \widehat{g}_{S\triangle T}.$$

Orthogonal changes to these coordinates preserve Euclidean and Frobenius norms.

The coordinate flip X_r sends a color x to $x + e_r$, where e_r is the r-th unit vector and addition is modulo two. It has eigenvalue $(-1)^{\mathbf{1}_{r\in S}}$ on χ_S. For scalars a_S indexed by $S \subseteq [d]$, write $\operatorname{diag}_{S\subseteq[d]}(a_S)$ for the $2^d \times 2^d$ diagonal matrix whose (S,S)-entry is a_S, with row and column indices ordered as the columns of F. All off-diagonal entries are zero. We abbreviate this notation to $\operatorname{diag}_S(a_S)$ when the index set is clear. For example, when $d = 2$ and the subset order is $\varnothing, \{1\}, \{2\}, \{1, 2\}$,

$$\operatorname{diag}_{S\subseteq[2]}(z^{|S|}) = \operatorname{diag}(1, z, z, z^2).$$

For the factors in (1.1),

$$F^{\mathsf{T}} T(\boldsymbol{\rho}) F = \operatorname{diag}_{S\subseteq[d]} \left(\prod_{r=1}^{d} (1 + (-1)^{\mathbf{1}_{r\in S}} \rho_r) \right).$$

The Boolean noise operator with parameter z is the operator having eigenvalue $z^{|S|}$ on χ_S, namely

$$F \operatorname{diag}_S(z^{|S|}) F^{\mathsf{T}} = 2^{-d} \bigotimes_{r=1}^{d} \begin{pmatrix} 1+z & 1-z \\ 1-z & 1+z \end{pmatrix}.$$

These are Fourier identities. Whether a given matrix in these identities is jointly available from a particular source is a separate question settled by the reductions.

2.7 Known counting results and algorithms

A matrix *over the Boolean domain* is a 2×2 matrix; it need not have zero–one entries. A strictly positive symmetric Boolean-domain matrix is called zero-field here when its diagonal entries are equal. Up to a positive scalar, it is $W(\rho)$ from (1.1).

Proposition 2.2 (Boolean and Potts counting problems)**.** *The following facts for fixed matrices will be used.*

- *(i) If a strictly positive, invertible, symmetric matrix over the Boolean domain has unequal diagonal entries, its planar homomorphism problem is #P-hard. Equal diagonal entries give a polynomial-time zero-field Ising algorithm.*
- *(ii) For every integer $s \geq 3$, the positive Potts matrix $I_s + J_s$ has a #P-hard planar homomorphism problem.*

The first statement is [CMY26, Theorem 22], which invokes the Boolean classification of [GW20]. The rank-one case is excluded from its hardness hypothesis and is evaluated by elementary vertex sums. For the second, let $T_G(x, y)$ be the Tutte polynomial, with the subset expansion

$$T_G(x, y) = \sum_{A\subseteq E(G)} (x-1)^{c(V(G),A)-c(G)} (y-1)^{|A|-|V(G)|+c(V(G),A)}.$$

Here $c(V(G), A)$ counts components of the spanning subgraph, including its isolated vertices. The random-cluster expansion gives

$$Z_{I_s+J_s}(G) = s^{c(G)} T_G(s+1, 2).$$

The planar Tutte classification [Ver05] is tractable on $(x-1)(y-1) \in \{1, 2\}$ and at $(1, 1), (-1, -1), (\omega, \omega^2), (\omega^2, \omega)$, where $\omega = e^{2\pi \mathrm{i}/3}$, and is hard elsewhere at algebraic points. At $(x, y) = (s+1, 2)$,

$$(x-1)(y-1) = s \geq 3,$$

so this point is outside the stated exceptions, giving part (ii).

Complete three-state and full-rank four-state classifications. For reference, the Boolean real symmetric matrix $B = \left(\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right)$ is planar tractable exactly in the following algebraic cases:

$$ac = b^2, \qquad b = 0, \qquad a = c, \qquad \text{or} \quad ac = -b^2 \text{ and } a = -c. \tag{2.3}$$

Outside these cases it is #P-hard [CM23, Theorem 2]. For nonnegative entries, the last case adds no matrices beyond the first three.

Theorem 2.3 (Three states, [CM23])**.** *Let M be a fixed real symmetric 3×3 matrix. Its planar homomorphism problem is polynomial-time computable if, up to a simultaneous permutation of rows and columns, one of the following holds:*

(i) $\operatorname{rank} M \leq 1$*;*

(ii) $M = \left(\begin{smallmatrix} a & b & 0 \\ b & c & 0 \\ 0 & 0 & t \end{smallmatrix}\right)$*, where the 2×2 block satisfies at least one condition in (2.3), and t is arbitrary;*

(iii) $M = \left(\begin{smallmatrix} 0 & 0 & a \\ 0 & 0 & b \\ a & b & 0 \end{smallmatrix}\right)$*, where a, b are arbitrary real numbers.*

In every other case it is #P-hard under polynomial-time Turing reductions.

This is the classification in [CM23, Theorem 43]. The rank-one alternative includes both signs of a real symmetric rank-one matrix: an overall fixed nonzero scalar multiplies the partition function by its $|E(G)|$-th power.

Theorem 2.4 (Four states of full rank, [CM24])**.** *Let M be a fixed nonnegative real symmetric 4×4 matrix of rank four. Its planar homomorphism problem is polynomial-time computable precisely when, up to simultaneous permutation,*

(i) $M = A \oplus B$*, where A, B are smaller nonempty symmetric nonnegative matrices whose planar homomorphism problems are polynomial-time computable; or*

(ii) $M = A \otimes B$*, where A, B are symmetric nonnegative 2×2 matrices satisfying (2.3).*

Every other such M gives a #P-hard problem under polynomial-time Turing reductions.

This is [CM24, Theorem 117]. The direct sums have sizes $1 + 3$ or $2 + 2$, so Theorem 2.3 and (2.3) make the criterion explicit. In the strictly positive case the criterion reduces to a positive scalar multiple of $W(\rho_1) \otimes W(\rho_2)$ with $\rho_1, \rho_2 > 0$ and $\rho_1, \rho_2 \neq 1$ [CM24, Theorem 112]. The full-rank assumption is part of this earlier four-state theorem; Theorem 1.1 removes it and permits every finite domain size. The cited real-weight statements use fixed exact representations; for the body of this paper their algebraic specializations suffice, and Appendix A discusses the real-weight model.

Matchgate representation and product evaluation. For this discussion, a matchgate is a planar graph with complex edge weights and distinct external vertices listed in counterclockwise order on its outer face. Its signature at a bit string is the weighted perfect-matching sum after deleting the listed vertices whose bits are one. Define the equality signature $\mathrm{EQ}_{n,q} : [q]^n \to \{0, 1\}$ to be one exactly when all arguments agree. Listing its values gives a row vector of length q^n. For $U \in \mathbb{C}^{q\times 2^\ell}$, its transformed signature is $\mathrm{EQ}_{n,q} U^{\otimes n}$, a function of $n\ell$ bits. The external bits are ordered by the successive blocks in this tensor expression. Fu proves that, for $n, q \geq 3$, this function cannot be a matchgate signature when $\operatorname{rank} U = q$ [Fu17, Theorem 1.2].

The structural forms are evaluated as products of zero-field Ising partition functions. The rank-one and bipartite rank-two factors are evaluated directly; each remaining Boolean factor is handled by its own planar Pfaffian calculation. No complexity conclusion here is inferred merely from the failure of one particular matchgate representation.

Proposition 2.5 (Diagonal separation, [CMY26]). *Let X be a fixed nonnegative symmetric matrix.*

(i) If $X > 0$ and $X_{ii} \neq X_{jj}$ for all $i \neq j$, then

$$\operatorname{rank} X = 1 \implies \text{Pl-GH}(X) \in \mathsf{FP}, \qquad \operatorname{rank} X > 1 \implies \text{Pl-GH}(X) \textit{ is } \#\text{P-hard}.$$

(ii) Suppose that for each $i \neq j$ there is a planar edge gadget K_{ij}, with conditional signature F_{ij}, such that $(F_{ij})_{ii} \neq (F_{ij})_{jj}$. If every support component of X has rank zero, rank one, or the bipartite rank-two form $\left(\begin{smallmatrix} 0 & V \\ V^{\mathsf{T}} & 0 \end{smallmatrix}\right)$ with $\operatorname{rank} V = 1$, then $\text{Pl-GH}(X) \in \mathsf{FP}$. In every other case, $\text{Pl-GH}(X)$ is $\#$P-hard.

These are Theorems 32 and 34 of [CMY26]. The terminals of each edge gadget are cofacial, as required for planar edge replacement. They provide optional shorter hardness proofs; the structural proof below does not assume that such separating gadgets exist.

2.8 Algorithms for the tractable forms

We collect the algorithmic statements and their evaluation formulas here. All target dimensions and weights are fixed; the input is a graph, with loops and parallel edges allowed. The Boolean classification and the two diagonal-separation dichotomies from [CMY26, Theorems 22, 32 and 34] are stated in Propositions 2.2 and 2.5, with their respective hypotheses. The algorithms below require no diagonal-separation condition.

Rank-one and bipartite rank-two matrices. The real symmetric case of [BG05, Theorem 6(2)] states the following: for a fixed real symmetric matrix A, if each support component has rank at most one, or is bipartite and has rank at most two, then $G \mapsto Z_A(G)$ is polynomial-time computable on arbitrary finite graphs. The theorem in that source also permits polynomial weights; here only fixed real weights are used. In the nonnegative case, the nonzero connected matrices in this statement are precisely

$$aa^{\mathsf{T}}, \qquad A_{ab^{\mathsf{T}}} := \begin{pmatrix} 0 & ab^{\mathsf{T}} \\ ba^{\mathsf{T}} & 0 \end{pmatrix}, \qquad a \in \mathbb{R}^k_{>0}, \quad b \in \mathbb{R}^\ell_{>0}.$$

The formulas below give their algorithms directly in our exact output model.

Zero-field planar Ising. For every fixed $\rho > 0$, the problem $\text{Pl-GH}(W(\rho))$ is in FP. For $\rho \neq 1$ this is the equal-diagonal case of [CMY26, Theorem 22], stated in Proposition 2.2(i); $W(1) = J_2$ gives $Z_{W(1)}(G) = 2^{|V(G)|}$. The Fisher–Kasteleyn–Temperley (FKT) method evaluates weighted perfect matchings of a planar graph by a Pfaffian [Kas61, TF61, Kas63]. For a skew-symmetric matrix of even order, its Pfaffian is the signed sum over pairings of indices and its square is the determinant. A planar Pfaffian orientation makes this algebraic sum give the matching partition function with its prescribed edge weights, up to an efficiently computable overall sign. A reference perfect matching determines this sign whenever a matching exists; otherwise the partition function is zero. The zero-field Ising reduction is the conversion of the even-subgraph expansion into the weighted perfect-matching sum of a decorated planar graph. Fisher's correspondence gives this conversion; see [BTR20, Section 2.1.3, pp. 7–8], where the weights and the partition-function identity are explicit. For a loopless G, the expansion in our normalization is

$$Z_{W(\rho)}(G) = 2^{|V(G)|} \left(\frac{1+\rho}{2}\right)^{|E(G)|} \sum_{\substack{A \subseteq E(G) \\ \deg_A(v) \text{ even for all } v}} \left(\frac{1-\rho}{1+\rho}\right)^{|A|}.$$

The correspondence is an identity of polynomials in the even-subgraph edge weights. Hence it also holds for their negative values when $\rho > 1$; the signed matching sum is still evaluated by the Pfaffian orientation above. Isolated vertices contribute 2 each. This gives the zero-field algorithm of Proposition 2.2(i) for every $\rho > 0$. Loops in an Ising input contribute their constant diagonal weight and may be removed; parallel edges cause no difficulty.

Direct sums. For fixed real symmetric matrices $M_1, \dots, M_r$ and $M \cong \bigoplus_{j=1}^r M_j$, [CMY26, Theorem 58] states both implications

$$(\forall j,\ \text{Pl-GH}(M_j) \in \mathsf{FP}) \implies \text{Pl-GH}(M) \in \mathsf{FP},$$
$$(\exists j,\ \text{Pl-GH}(M_j) \text{ is \#P-hard}) \implies \text{Pl-GH}(M) \text{ is \#P-hard}.$$

For the first implication, the component formula in the next proof is the complete algorithm. Only this implication is needed for composition of the tractable forms.

Proposition 2.6 (Composition of known tractable cases)**.** *Let M be a fixed real symmetric algebraic matrix such that, after a simultaneous row and column permutation, $M = \bigoplus_{j=1}^r M_j$, with each M_j of one of the forms*

$$[0], \qquad (aa^\mathsf{T}) \otimes T(\boldsymbol{\rho}), \qquad \begin{pmatrix} 0 & ab^\mathsf{T} \\ ba^\mathsf{T} & 0 \end{pmatrix} \otimes T(\boldsymbol{\rho}).$$

For each nonzero block, $a \in \mathbb{R}^k_{>0}$, $b \in \mathbb{R}^\ell_{>0}$ when present, $k, \ell \geq 1$, and $T(\boldsymbol{\rho}) = \bigotimes_{h=1}^d W(\rho_h)$ with fixed $d \geq 0$ and $\rho_h > 0$. The dimensions and parameters may differ between blocks; $d = 0$ and $\rho_h = 1$ are allowed. Then Pl-GH$(M) \in \mathsf{FP}$ *in the algebraic bit model.*

Proof. For a rank-one block, summation separates at the input vertices:

$$Z_{aa^\mathsf{T}}(G) = \prod_{v \in V(G)} \left(\sum_{i=1}^k a_i^{\deg_G(v)} \right). \tag{2.4}$$

Each loop counts twice in $\deg_G(v)$, matching the contribution a_i^2 of its edge; an isolated vertex contributes k.

For a connected bipartite input with sides L, R and at least one edge, define

$$\Phi_{a,b}(G; L, R) := \prod_{u \in L} \left(\sum_{i=1}^k a_i^{\deg_G(u)} \right) \prod_{v \in R} \left(\sum_{j=1}^\ell b_j^{\deg_G(v)} \right).$$

Its assignments into the two sides of A_{ab^T} have two possible orientations. For any connected input G, this gives

$$Z_{A_{ab^\mathsf{T}}}(G) = \begin{cases} \Phi_{a,b}(G; L, R) + \Phi_{a,b}(G; R, L), & G \text{ bipartite, } |E(G)| > 0, \\ 0, & G \text{ nonbipartite,} \\ k + \ell, & G \text{ a single isolated vertex.} \end{cases} \tag{2.5}$$

The two orientation sums may differ. An input self-loop makes G nonbipartite. The zero block $[0]$ contributes 1 on a single isolated vertex and 0 on a connected input with an edge.

For fixed $A \in \mathrm{Sym}_p(\mathbb{R})$ and $B \in \mathrm{Sym}_q(\mathbb{R})$, write a tensor color as (i, j). Separating the assignments $\sigma : V(G) \to [p]$ and $\tau : V(G) \to [q]$ gives

$$\begin{aligned} Z_{A \otimes B}(G) &= \sum_{\sigma, \tau} \prod_{uv \in E(G)} A_{\sigma(u),\sigma(v)} B_{\tau(u),\tau(v)} \\ &= Z_A(G) Z_B(G). \end{aligned} \tag{2.6}$$

Thus a nonzero displayed block is evaluated by (2.4) or (2.5), multiplied by $\prod_{h=1}^{d} Z_{W(\rho_h)}(G)$. Each Ising factor uses the planar algorithm above, for all $\rho_h > 0$, including $\rho_h > 1$. When $d = 0$, this product is 1.

Let $G_1, \dots, G_m$ be the connected components of G. Every nonzero assignment on G_α stays in one target block. Summing its choices of block, then multiplying over input components, gives

$$Z_M(G) = \prod_{\alpha=1}^{m} \left(\sum_{j=1}^{r} Z_{M_j}(G_\alpha) \right). \tag{2.7}$$

The empty input has value 1. A singleton input contributes the sum of the block sizes, which is the order of M.

These formulas use a fixed number of Ising computations per component and polynomially many arithmetic operations. The displayed parameters can be chosen in a fixed finite algebraic extension of the field generated by the entries of M. Indeed, in a positive block indexed by (i, z), normalize

$$a_i = \sqrt{M_{(i,0),(i,0)}}, \qquad \rho_h = \frac{M_{(1,0),(1,e_h)}}{M_{(1,0),(1,0)}}.$$

Here 0 is the zero Boolean vector and e_h has its sole 1 in coordinate h. For a bipartite block, write its positive rectangular part as V and choose

$$a_1 = 1, \qquad b_j = V_{(1,0),(j,0)}, \qquad a_i = \frac{V_{(i,0),(1,0)}}{b_1}, \qquad \rho_h = \frac{V_{(1,0),(1,e_h)}}{V_{(1,0),(1,0)}}.$$

Thus only fixed square roots and ratios are needed. The factors and degree powers have polynomial bit length in the fixed number field. Exact arithmetic and conversion back to the prescribed output field are as in Section 2.5 and Remark 3.4. This proves the algorithm independently of the later structural necessity arguments. $\square$

3 Planar reductions

We use the counting model of Section 2.1 and joint availability from Definition 2.1. Product interpolation replaces entries while preserving equal-length product identities; spectral interpolation does the same for eigenvalues. Root restrictions and weight removal then connect these mixed simulations to ordinary counting problems. All fixed matrix entries, background vertex weights and auxiliary constraint values in the complexity statements of this section are fixed real algebraic numbers. The number of constraint types is fixed; their number of occurrences in an input is unrestricted. Other jointly available constraints may remain in an instance during each interpolation below. Background vertex weights are fixed and positive, and remain unchanged unless their removal is explicitly proved. Purely linear-algebraic statements below hold over $\mathbb{R}$ as stated.

3.1 Product and spectral interpolation

Lemma 3.1 (Joint product interpolation). *Let $C, C' \in \mathrm{Sym}_q(\mathbb{R})$ be fixed algebraic matrices, with C jointly available from a fixed source language $\mathcal{F}$ with fixed positive algebraic background vector w. Put*

$$E = \{(i, j) \in [q]^2 : C_{ij} \neq 0\}, \qquad C_{(i,j)} = C_{ij},$$

and, for every integer $s \geq 1$, define

$$\Lambda_s(C) = \left\{ \prod_{a=1}^{s} C_{e_a} : (e_1, \dots, e_s) \in E^s \right\}.$$

Assume:

(i) $\forall i, j \in [q] : \; C_{ij} = 0 \Rightarrow C'_{ij} = 0$*;*

(ii)

$$\forall s \in \mathbb{Z}_{\geq 1}, \quad \exists\, \theta_s : \Lambda_s(C) \longrightarrow \mathbb{R},$$
$$\forall (e_1, \ldots, e_s) \in E^s, \quad \theta_s\left(\prod_{a=1}^{s} C_{e_a}\right) = \prod_{a=1}^{s} C'_{e_a}.$$

Then C' is jointly available. The same assertion holds for real algebraic unary functions. In particular, positive unary functions may be raised to any fixed rational power.

Proof. Consider an instance with s marked occurrences to be replaced. If $s = 0$, no replacement is needed. If $s > 0$ and $C = 0$, condition (i) makes the desired value zero. Otherwise let $c_1, \ldots, c_t$ be the distinct nonzero entries of C. Replace each marked occurrence by h parallel copies of C, where $h \geq 1$. The resulting value has the form

$$F(h) = \sum_{\lambda \in \Lambda_s(C)} A_\lambda \lambda^h, \qquad \Lambda_s(C) = \left\{ \prod_{\alpha=1}^{t} c_\alpha^{n_\alpha} : n_\alpha \geq 0, \; \sum_\alpha n_\alpha = s \right\}.$$

Equal numerical products are merged. The coefficients A_λ contain all other edge constraints, attached unaries, and background vertex weights; they may have either sign. An assignment using a zero entry of C contributes zero for every $h \geq 1$ and, by (i), also contributes zero to the desired replacement.

Since t is fixed, the number $N = |\Lambda_s(C)|$ satisfies

$$N \leq \binom{s+t-1}{t-1}.$$

List the distinct nonzero bases as $\lambda_1, \ldots, \lambda_N$. The values $F(1), \ldots, F(N)$ determine all coefficients because

$$\det(\lambda_j^h)_{1 \leq h, j \leq N} = \left(\prod_{j=1}^{N} \lambda_j\right) \prod_{1 \leq i < j \leq N} (\lambda_j - \lambda_i) \neq 0.$$

Condition (ii) gives the required value directly:

$$\sum_{\lambda \in \Lambda_s(C)} A_\lambda \theta_s(\lambda).$$

New zero positions are allowed, since $\theta_s(\lambda)$ may be zero. Value representatives can be chosen consistently because

$$C_e = C_f \neq 0 \quad \Longrightarrow \quad C'_e = \theta_1(C_e) = \theta_1(C_f) = C'_f.$$

There are polynomially many bases and queries. The largest query replaces the marked edges by at most sN edges. Products of fixed algebraic numbers with exponents $O(sN)$ have polynomial bit length in a fixed number field. Exact equality testing, merging products, and solving the displayed linear system therefore have polynomial bit complexity. Parallel replacement is planar also for an input loop and introduces no new vertices, so no background vertex weight changes.

For a unary function, attach h copies at every marked occurrence and use the same calculation. For a positive unary, replacing a nonzero product λ by λ^r, with fixed rational r, is unambiguous and uses only a fixed algebraic extension.

For a fixed finite collection $C'_1, \dots, C'_k$, each satisfying (i) and (ii) relative to C, put $\mathcal{G}_0 = \mathcal{F} \cup \{C\}$ and $\mathcal{G}_j = \mathcal{G}_0 \cup \{C'_1, \dots, C'_j\}$ for $1 \le j \le k$. Applying the preceding argument with source language $\mathcal{G}_{j-1}$ gives

$$\text{Pl-GH}(\mathcal{G}_j, w) \le_T \text{Pl-GH}(\mathcal{G}_{j-1}, w) \qquad (1 \le j \le k).$$

The other types remain fixed constraints at each stage. Since k is fixed, the polynomial bounds compose. If C itself is obtained by a previous joint simulation, compose that simulation with these queries. This proves joint availability from the original source. $\square$

Corollary 3.2 (Support, magnitude, sign, and extremal positions)**.** *For a fixed real symmetric algebraic matrix M, the matrices $P(M)$, $|M|$, and $S(M)$ are jointly available from M. The entrywise square $M^{\circ 2}$ is realized by two parallel M-edges. With any fixed positive algebraic background vertex weights w,*

$$\text{Pl-GH}(P(M), w) \le_T \text{Pl-GH}(M^{\circ 2}, w) \le_T \text{Pl-GH}(M, w),$$
$$\text{Pl-GH}(P(M), w) \le_T \text{Pl-GH}(|M|, w) \le_T \text{Pl-GH}(M, w).$$

Moreover,

$$\text{Pl-GH}(M, w) \equiv_T \text{Pl-GH}(\{|M|, S(M)\}, w).$$

For $M \ne 0$, put

$$a = \min_{M_{ij} \ne 0} |M_{ij}|, \qquad b = \max_{i,j} |M_{ij}|.$$

The matrices $\mathbf{1}_{\{|M|=a\}}$ and $\mathbf{1}_{\{|M|=b\}}$ are jointly available as well. For $M > 0$, these select its numerically smallest and largest entries.

Proof. For $s \ge 1$, $M_{e_a} \ne 0$ for every $a \in [s]$, and $\lambda = \prod_{a=1}^{s} M_{e_a} \in \Lambda_s(M)$, set

$$\theta_{P,s}(\lambda) = 1 = \prod_{a=1}^{s} P(M)_{e_a},$$
$$\theta_{|M|,s}(\lambda) = |\lambda| = \prod_{a=1}^{s} |M|_{e_a},$$
$$\theta_{S,s}(\lambda) = \text{sgn}(\lambda) = \prod_{a=1}^{s} S(M)_{e_a}.$$

The zero positions satisfy

$$M_{ij} = 0 \quad \Longrightarrow \quad P(M)_{ij} = |M|_{ij} = S(M)_{ij} = 0.$$

Lemma 3.1, iterated over these three replacement matrices, gives their joint availability; when $M = 0$, all three are zero as well.

Two parallel edges, including two parallel copies of a loop, realize $M^{\circ 2}$ without adding vertices. This matrix is nonnegative and has support $P(M)$, so its support is jointly available by the same lemma. Similarly, $|M|$ has support $P(M)$. These observations prove both reduction chains and preserve every specified background vertex weight. In particular, isolated vertices retain their original weight sum.

The preceding joint availability gives the mixed problem from the M-source. In the reverse direction, replace each M-edge by an $|M|$-edge and an $S(M)$-edge in parallel, using $M = |M| \circ S(M)$. This proves the mixed equivalence.

For $s \ge 1$, $c \in \{a, b\}$, and $x_r = |M_{e_r}| > 0$, put $\lambda = \prod_{r=1}^{s} x_r \in \Lambda_s(|M|)$. Then

$$a \le x_r \le b \quad (r \in [s]) \quad \Longrightarrow \quad (\lambda = c^s \iff \forall r \in [s],\ x_r = c).$$

Consequently the indicator matrices satisfy condition (ii) with

$$\theta_{c,s}(\lambda) = \mathbf{1}_{\{\lambda=c^s\}} = \prod_{r=1}^{s} \mathbf{1}_{\{x_r=c\}}, \qquad \lambda \in \Lambda_s(|M|).$$

Since $c > 0$, they retain the zero positions of $|M|$. Apply Lemma 3.1 to the already available matrix $|M|$, successively for $c = a, b$, to obtain both masks jointly. $\square$

The problem on $\{|M|, S(M)\}$ in this equivalence is a mixed constraint problem. Separate polynomial-time algorithms for the two homogeneous problems Pl-GH$(|M|, w)$ and Pl-GH$(S(M), w)$ do not by themselves supply an algorithm for this mixed language. Hardness of Pl-GH$(|M|, w)$ transfers to Pl-GH(M, w), whereas tractability of Pl-GH$(|M|, w)$ alone does not settle Pl-GH(M, w). The extremal assertion concerns absolute values and excludes zero when taking the minimum; it does not select an arbitrary internal magnitude level or the numerically least or greatest signed entry.

The next lemma uses unit background vertex weights, with the convention Pl-GH$(\mathcal{F})$ = Pl-GH$(\mathcal{F}, \mathbf{1})$ from Section 2.2.

Lemma 3.3 (Spectral interpolation)**.** *Let $\mathcal{F}$ be a fixed source language and $C \in \mathrm{Sym}_q(\mathbb{R})$ a fixed algebraic matrix such that*

$$C \succeq 0, \qquad \text{Pl-GH}(\mathcal{F} \cup \{C\}) \le_T \text{Pl-GH}(\mathcal{F}).$$

Write

$$C = U \operatorname{diag}(\lambda_1, \ldots, \lambda_q) U^\mathsf{T}, \qquad U^\mathsf{T} U = I_q, \qquad \lambda_i \ge 0,$$

and define its orthogonal range projector by

$$\Pi_C = U \operatorname{diag}(\mathbf{1}_{\{\lambda_1>0\}}, \ldots, \mathbf{1}_{\{\lambda_q>0\}}) U^\mathsf{T}.$$

Thus, for $\operatorname{im} C = \{Cx : x \in \mathbb{R}^q\}$,

$$\Pi_C x = \begin{cases} x, & x \in \operatorname{im} C, \\ 0, & x \in \ker C. \end{cases}$$

Then

$$\text{Pl-GH}(\mathcal{F} \cup \{C, \Pi_C\}) \le_T \text{Pl-GH}(\mathcal{F}).$$

If $C \succ 0$, then, for every fixed $r \in \mathbb{Q}$,

$$C^r = U \operatorname{diag}(\lambda_1^r, \ldots, \lambda_q^r) U^\mathsf{T},$$

and

$$\text{Pl-GH}(\mathcal{F} \cup \{C, C^r\}) \le_T \text{Pl-GH}(\mathcal{F}).$$

Proof. Let $\Omega(T)$ denote a mixed instance with s designated edges carrying the matrix T; all other constraints belong to $\mathcal{F} \cup \{C\}$. If $s = 0$, the assumed reduction already applies. If $C = 0$, then $\Pi_C = 0$, and an instance with a designated Π_C-edge has value zero. Hence assume $s \ge 1$ and $C \ne 0$. For its distinct positive eigenvalues, write

$$C = \sum_{a=1}^{t} \alpha_a P_a, \qquad 0 < \alpha_1 < \cdots < \alpha_t, \qquad \Pi_C = \sum_{a=1}^{t} P_a.$$

Replacing each designated edge by a C-path of length $h \geq 1$ gives an instance Ω_h with

$$C^h = \sum_{a=1}^{t} \alpha_a^h P_a,$$
$$F(h) := Z(\Omega_h) = Z(\Omega(C^h)) = \sum_{\mu \in \Lambda_s} B_\mu \mu^h,$$

where

$$\Lambda_s = \left\{ \prod_{a=1}^{t} \alpha_a^{n_a} : n_a \in \mathbb{Z}_{\geq 0},\ \sum_{a=1}^{t} n_a = s \right\}, \qquad N = |\Lambda_s| \leq \binom{s+t-1}{t-1}.$$

All unchanged constraints and the spectral projectors occur in the coefficients B_μ. List the distinct products as $\mu_1, \ldots, \mu_N$. Since

$$\det(\mu_j^h)_{1 \leq h,j \leq N} = \left(\prod_{j=1}^{N} \mu_j \right) \prod_{1 \leq i < j \leq N} (\mu_j - \mu_i) \neq 0,$$

the values $F(1), \ldots, F(N)$ recover every B_μ, and

$$Z(\Omega(\Pi_C)) = \sum_{\mu \in \Lambda_s} B_\mu.$$

If $C \succ 0$, then

$$C^r = \sum_{a=1}^{t} \alpha_a^r P_a, \qquad \prod_{a=1}^{t} (\alpha_a^r)^{n_a} = \left(\prod_{a=1}^{t} \alpha_a^{n_a} \right)^r,$$

and consequently

$$Z(\Omega(C^r)) = \sum_{\mu \in \Lambda_s} B_\mu \mu^r.$$

The path substitutions preserve planarity, including at input loops, and the largest query has size $O(|\Omega| + sN)$. For fixed C, r, the fixed-number-field bit bounds of Lemma 3.1 apply in a fixed algebraic extension containing the eigenvalues, projectors and required powers. Composing with the assumed joint simulation of C completes both reductions to Pl-GH($\mathcal{F}$). □

Remark 3.4 (Exact fields in parameter-dependent interpolation). A fixed-size algebraic matrix whose entries lie in a fixed number field (or in extensions of uniformly bounded degree over it), and have polynomial bit length, has eigenvalues of bounded algebraic degree and polynomial logarithmic height. Its spectral products and interpolation solutions therefore have polynomial logarithmic height for every fixed matrix size. When the matrix varies with a rational parameter, the splitting fields at different sample points need not have a bounded-degree compositum.

We will never form that compositum. If a recovered answer z_x at sample x belongs to a fixed base field K, first express it in its individual bounded-degree extension K_x/K, and then replace its representation by

$$\frac{1}{[K_x : K]} \operatorname{Tr}_{K_x/K}(z_x).$$

This equals z_x and is computed using a multiplication matrix of bounded dimension. Only then is an outer interpolation performed over K. For the same reason an algorithm using a fixed auxiliary algebraic extension returns its final answer in the prescribed output field. Whenever a parameter-dependent argument below invokes an extension, its base-field descent is stated explicitly.

3.2 Root restrictions and positive vertex weights

Lemma 3.5 (Restriction at one root)**.** *Fix $A \in \mathrm{Sym}_q(\mathbb{R})$, $w \in \mathbb{R}^q_{>0}$, and $X \subseteq [q]$, with all entries of A, w algebraic. There exist fixed rooted planar graphs $(F_1, r_1), \ldots, (F_s, r_s)$ and fixed coefficients $c_1, \ldots, c_s \in \mathbb{Q}(A, w)$, depending only on (A, w, X), such that every connected planar input (G, r), with $r \in V(G)$, satisfies*

$$\sum_{\substack{\sigma: V(G) \to [q] \\ \sigma(r) \in X}} \prod_{v \in V(G)} w_{\sigma(v)} \prod_{uv \in E(G)} A_{\sigma(u), \sigma(v)} = \sum_{j=1}^{s} c_j Z_{A,w}(G \vee F_j).$$

Here $G \vee F_j$ identifies r with r_j. In particular, for every connected component of Γ_A with vertex set $C \subseteq [q]$,

$$\text{Pl-GH}(A[C], w|_C) \leq_T \text{Pl-GH}(A, w).$$

If each support-component problem is in FP*, the whole problem is in* FP*.*

Proof. For a rooted planar graph (F, r_F), let $f_F(i)$ be its conditional partition function with root color i, omitting the background weight of the root itself. Put $D = \mathrm{diag}(w)$ and $U = \mathrm{span}\{f_F\} \subseteq \mathbb{R}^q$. For vectors u, v, use the positive inner product $\langle u, v \rangle_D = u^\mathsf{T} D v$. Choose rooted graphs $F_1, \ldots, F_s$ whose conditional functions form a basis of U. The D-orthogonal projection of $\mathbf{1}_X$ onto U has an expansion

$$\mathrm{proj}_{U,D}\, \mathbf{1}_X = \sum_{j=1}^{s} c_j f_{F_j}.$$

Since $f_G \in U$ for the rooted input G,

$$\sum_{i \in X} w_i f_G(i) = \sum_{j=1}^{s} c_j \langle f_{F_j}, f_G \rangle_D = \sum_{j=1}^{s} c_j Z_{A,w}(G \vee F_j),$$

where $G \vee F_j$ identifies r with r_j. Each root can be placed on the outer face, so this gluing is planar. The root background weight is counted once on both sides. The basis graphs and coefficients depend only on (A, w, X). To compute the coefficients in the prescribed field $F = \mathbb{Q}(A, w)$, set

$$H_{ab} = f_{F_a}^\mathsf{T} D f_{F_b}, \qquad g_a = f_{F_a}^\mathsf{T} D \mathbf{1}_X.$$

All entries of H, g lie in F, and the basis property and positivity of D make H positive definite. Thus $(c_1, \ldots, c_s)^\mathsf{T} = H^{-1} g \in F^s$. The entries and coefficients lie in a fixed number field. Section 2.5 therefore gives the polynomial bit bounds for this fixed linear combination of oracle answers, proving the claimed polynomial-time Turing reduction.

For the vertex set C of a connected component of Γ_A and a connected rooted input (G, r), following input paths gives the termwise implication

$$\sigma(r) \in C, \qquad \prod_{uv \in E(G)} A_{\sigma(u), \sigma(v)} \neq 0 \quad \Longrightarrow \quad \sigma(V(G)) \subseteq C.$$

The restricted value is thus $Z_{A[C], w|_C}(G)$, also for signed entries. It equals $\sum_{i \in C} w_i$ on an isolated vertex. For the connected components $G_1, \ldots, G_k$ of the input,

$$Z_{A[C], w|_C}(G) = \prod_{a=1}^{k} Z_{A[C], w|_C}(G_a).$$

This extends the component reduction to arbitrary planar inputs. If each Pl-GH$(A[C], w|_C)$ belongs to FP, then

$$Z_{A,w}(G) = \prod_{a=1}^{k} \left(\sum_{C} Z_{A[C],w|_C}(G_a) \right)$$

gives a polynomial-time algorithm for Pl-GH(A, w), since the number of components of Γ_A is fixed. Here the sum ranges over their vertex sets C, and the empty product is 1. The positive inner product in this proof comes from $w > 0$, not from nonnegativity of the edge entries. This is a single-root calculation and does not grant simultaneous pinning at an arbitrary number of vertices. $\square$

Lemma 3.6 (Linear independence of tensor powers)**.** *Let $v_1, \dots, v_r \in \mathbb{R}^n$ be nonzero, with no two proportional over $\mathbb{R}$. For $p = \max(1, r-1)$, the tensors $v_i^{\otimes p}$ are linearly independent. Consequently, if $B \in \mathbb{R}^{r \times n}$ has nonzero, pairwise nonproportional rows and D is a positive $n \times n$ diagonal matrix, then $(BDB^{\mathsf{T}})^{\circ p}$ is positive definite.*

Proof. For $r \geq 2$, nonproportionality gives linear forms satisfying

$$\ell_{ij}(v_j) = 0, \qquad \ell_{ij}(v_i) \neq 0 \quad (i \neq j).$$

Set $p = r - 1$ and $Q_i(v) = \prod_{j \neq i} \ell_{ij}(v)$. Its associated linear functional on tensor powers satisfies

$$\widetilde{Q}_i(v_j^{\otimes p}) = Q_i(v_j) = 0 \quad (j \neq i), \qquad \widetilde{Q}_i(v_i^{\otimes p}) = Q_i(v_i) \neq 0.$$

Applying $\widetilde{Q}_i$ to $\sum_j a_j v_j^{\otimes p} = 0$ gives $a_i = 0$ for every i. For $r = 1$, use $p = 1$ and $v_1 \neq 0$.

The rows u_i of $BD^{1/2}$ remain nonzero and nonproportional, since $D^{1/2}$ is invertible. Their tensor Gram matrix is

$$\left(\langle u_i^{\otimes p}, u_j^{\otimes p} \rangle\right)_{i,j} = \left(\langle u_i, u_j \rangle^p\right)_{i,j} = (BDB^{\mathsf{T}})^{\circ p} \succ 0.$$

No entrywise nonnegativity of B is used; opposite rows are excluded by nonproportionality over $\mathbb{R}$. $\square$

Lemma 3.7 (Removal of positive vertex weights)**.** *Let $B \in \mathrm{Sym}_q(\mathbb{R})$ have fixed algebraic entries, no zero row and no two rows proportional over $\mathbb{R}$. For every fixed algebraic vector $w = (w_1, \dots, w_q) \in \mathbb{R}^q_{>0}$,*

$$\text{Pl-GH}(B) \leq_T \text{Pl-GH}(B, w).$$

Put $D = \mathrm{diag}(w)$. The matrix D^{-1} is jointly available from (B, w), and a D^{-1}-loop at a vertex gives $w_i(D^{-1})_{ii} = 1$ for its color $i \in [q]$. For every fixed $r \in \mathbb{Q}$, define

$$D^r = \mathrm{diag}(w_1^r, \dots, w_q^r),$$

and

$$u_r : [q] \longrightarrow \mathbb{R}_{>0}, \quad u_r(i) = (D^r)_{ii} = w_i^r.$$

Then

$$\text{Pl-GH}(\{B, D^r, u_r\}, w) \leq_T \text{Pl-GH}(B, w).$$

For every fixed finite constraint language $\mathcal{F}$ containing B,

$$\text{Pl-GH}(\mathcal{F}, w) \leq_T \text{Pl-GH}(B, w) \quad \Longrightarrow \quad \text{Pl-GH}(\mathcal{F}, \mathbf{1}) \leq_T \text{Pl-GH}(B, w).$$

Proof. Set $D = \operatorname{diag}(w)$ and $p = \max(1, q-1)$. A two-edge path in the weighted source has signature BDB. Taking p parallel copies gives the symmetric matrix

$$K = (BDB)^{\circ p} \succ 0$$

by Lemma 3.6. This is a fixed planar edge gadget. Its entries may be negative, but its eigenvalues after the following congruence are positive:

$$T = D^{1/2} K D^{1/2} \succ 0.$$

A chain of $h \geq 1$ copies of K, counting the background weights at its internal vertices, has signature

$$N_h = K(DK)^{h-1} = D^{-1/2} T^h D^{-1/2}.$$

Write $T = \sum_{a=1}^{t} \lambda_a P_a$ using its distinct positive eigenvalues. On s marked edges, replacing each one by this chain yields an expansion

$$F(h) = \sum_{\lambda \in \Lambda_s} A_\lambda \lambda^h, \qquad \Lambda_s = \left\{ \prod_{a=1}^{t} \lambda_a^{n_a} : n_a \geq 0,\ \sum_a n_a = s \right\}.$$

The exterior $D^{-1/2}$ factors and the spectral projectors are part of the coefficients, independent of h. All other edge and unary constraints remain in those coefficients as well. After merging equal bases there are at most $\binom{s+t-1}{t-1}$ of them. For $s > 0$, consecutive positive values of h give the nonsingular Vandermonde system used in Lemma 3.1. Recovering its coefficients and summing them evaluates this expansion at $h = 0$. Since $\sum_a P_a = I$, this jointly simulates

$$N_0 = D^{-1/2} I D^{-1/2} = D^{-1}.$$

For $s = 0$ there is nothing to simulate. All chain substitutions are planar, including substitutions on a loop, and have polynomial size. The eigenvalues and projectors lie in a fixed algebraic extension. The same product-count and fixed-field bit bounds as above give polynomial bit complexity. This argument explicitly uses the weighted chain formula; it does not apply the unit-background path formula to a weighted source.

Let $\mathcal{F}$ be a fixed finite language containing B with $\text{Pl-GH}(\mathcal{F}, w) \leq_T \text{Pl-GH}(B, w)$. For an $\mathcal{F}$-instance Ω, let Ω° add one D^{-1}-loop at each original vertex, retaining all original edges and explicit unary factors. For every assignment,

$$\prod_{v \in V(G)} w_{\sigma(v)} (D^{-1})_{\sigma(v),\sigma(v)} = 1.$$

Consequently,

$$Z_{\mathcal{F} \cup \{D^{-1}\}, w}(\Omega^\circ) = Z_{\mathcal{F}, \mathbf{1}}(\Omega),$$

and the joint simulation just proved gives

$$\begin{aligned} \text{Pl-GH}(\mathcal{F}, \mathbf{1}) &\leq_T \text{Pl-GH}(\mathcal{F} \cup \{D^{-1}\}, w) \\ &\leq_T \text{Pl-GH}(\mathcal{F}, w) \leq_T \text{Pl-GH}(B, w). \end{aligned}$$

The loop additions are planar and apply also to isolated vertices. Taking $\mathcal{F} = \{B\}$ proves the unweighted reduction.

For fixed $r \in \mathbb{Q}$, the nonzero entries of D^{-1} are its positive diagonal entries, and

$$\prod_{\ell=1}^{s} (D^r)_{i_\ell i_\ell} = \left(\prod_{\ell=1}^{s} (D^{-1})_{i_\ell i_\ell} \right)^{-r} \qquad (i_1, \ldots, i_s \in [q]).$$

The off-diagonal zero entries are preserved. Hence the binary version of Lemma 3.1, with $\theta_s(\lambda) = \lambda^{-r}$, supplies D^r jointly with the retained constraints. A D^r-loop at a vertex has the additional factor

$$u_r(i) = (D^r)_{ii} = w_i^r, \qquad i \in [q],$$

so it simulates one occurrence of the unary function u_r. This proves the remaining assertion for each fixed exponent and, by finite iteration, for every fixed finite collection of exponents.

If a matrix $A \in \mathrm{Sym}_q(\mathbb{R})$ has zero rows, first restrict it to the index set of its nonzero rows before using the lemma: on a connected input containing an edge, assignments using a zero-row color contribute zero; on an isolated vertex their weight sum is known. The case $A = 0_{q\times q}$ is handled below. □

Corollary 3.8 (Actual twins and zero–one matrices)**.** *Let A be a fixed real symmetric algebraic matrix with identical-row classes $I_1, \ldots, I_s$, and let w be fixed positive algebraic vertex weights. Retain one row and column from each class to obtain R, including its zero-row class if present, and put $\bar{w}_a = \sum_{i\in I_a} w_i$. Then, for every input,*

$$Z_{A,w}(G) = Z_{R,\bar{w}}(G).$$

If $A \neq 0$, let B be obtained by deleting the zero row and matching column of R, if present. If the rows of B are pairwise nonproportional over $\mathbb{R}$, then

$$\text{Pl-GH}(B) \leq_T \text{Pl-GH}(A, w).$$

In particular, for a zero–one $A \neq 0$, deleting zero rows and merging twins always gives Pl-GH$(B) \leq_T$ Pl-GH(A)*, and the same hardness lifting holds with arbitrary fixed positive algebraic vertex weights on A.*

Proof. Let $\pi : [q] \to [s]$ send each actual color to its twin class. Symmetry and equality of rows give $A_{ij} = R_{\pi(i),\pi(j)}$. For a fixed quotient coloring $\tau : V(G) \to [s]$,

$$\sum_{\sigma:\, \pi\circ\sigma=\tau} \prod_{v\in V(G)} w_{\sigma(v)} = \prod_{v\in V(G)} \sum_{i\in I_{\tau(v)}} w_i = \prod_{v\in V(G)} \bar{w}_{\tau(v)}.$$

Multiplying by the class-dependent edge product and summing over τ proves the identity, including loops, isolated vertices, and the empty input.

Write w_B for the restriction of $\bar{w}$ to the colors of B. On a connected input containing an edge, an assignment using a zero-row color has zero edge product. The displayed exact identity therefore evaluates Z_{B,w_B} using the (A, w) oracle on such inputs. For an isolated input vertex the value for (B, w_B) is the known sum of its weights. Multiplication over input components gives Pl-GH$(B, w_B) \leq_T$ Pl-GH(A, w). The vector w_B is positive. Under the stated nonproportionality hypothesis, Lemma 3.7 gives Pl-GH$(B) \leq_T$ Pl-GH(B, w_B), proving the claimed reduction.

For a zero–one quotient, distinct nonzero rows cannot be proportional: any proportionality constant must equal 1 at a nonzero entry, and would make the rows identical. Deleting all-zero columns does not alter this fact. Thus the additional hypothesis is automatic in the zero–one case, proving the original hardness-lifting statement and its positively weighted version. □

The exact quotient identity retains the positive class weights; it does not by itself remove them. For a real quotient, opposite rows or other nonunit proportional rows can remain after identical rows have been merged, so the extra hypothesis in the weight-removal step is essential to that argument. Moreover, identical rows in $P(A)$ need not be identical rows in A.

A support twin quotient therefore does not license merging the corresponding numerical real colors.

For the zero matrix $A = 0_{q\times q}$,

$$Z_{A,w}(G) = \begin{cases} 0, & E(G) \neq \varnothing, \\ \left(\sum_{i=1}^{q} w_i\right)^{|V(G)|}, & E(G) = \varnothing. \end{cases}$$

Thus $w = \mathbf{1}$ gives $q^{|V(G)|}$ on edgeless inputs, and the empty input has value 1.

3.3 Function identities and finite oracle reductions

A functional identity records equality for every parameter value. The next two lemmas find a polynomial-size sample that separates all nonidentical products of the length needed by the input, then use the Vandermonde interpolation of Section 2.5. This converts an analytic invariant into a finite oracle reduction.

Lemma 3.9 (Zeros of exponential sums)**.** *If $0 < \mu_1 < \cdots < \mu_N$ and the real coefficients c_i are not all zero, then $\sum_{i=1}^{N} c_i \mu_i^t$ has at most $N-1$ distinct real zeros.*

Proof. Discard zero coefficients and induct on their number N. The case $N = 1$ is immediate. Dividing by $\mu_1^t > 0$ preserves zeros and gives

$$g(t) = c_1 + \sum_{i=2}^{N} c_i (\mu_i/\mu_1)^t, \qquad g'(t) = \sum_{i=2}^{N} c_i \log(\mu_i/\mu_1)(\mu_i/\mu_1)^t.$$

By induction g' has at most $N-2$ distinct zeros, so Rolle's theorem gives at most $N-1$ distinct zeros for g. □

Lemma 3.10 (Effective transfer of product identities)**.** *Let $\mathcal{F}$ be a fixed finite algebraic constraint language on $[q]$, with unit background vertex weights, and let $K \subseteq \mathbb{R}$ be a fixed number field containing its weights. Fix $B \in \mathrm{Sym}_q(K)$ and a matrix family $F : \mathbb{R} \to \mathrm{Sym}_q(\mathbb{R})$. For $I = \{(i,j) : 1 \le i \le j \le q\}$, put $r = |I| = q(q+1)/2$, and write $F_e(t) = F_{ij}(t)$ and $B_e = B_{ij}$ for $e = (i,j) \in I$. For each integer $m \ge 1$, define*

$$\begin{aligned} \mathcal{A}_m &= \left\{ \alpha \in \mathbb{Z}_{\ge 0}^{I} : \sum_{e\in I} \alpha_e = m \right\}, \\ p_\alpha(t) &= \prod_{\substack{e\in I\\ \alpha_e>0}} F_e(t)^{\alpha_e}, \qquad\qquad b_\alpha = \prod_{\substack{e\in I\\ \alpha_e>0}} B_e^{\alpha_e}. \end{aligned}$$

The common length m counts entry factors, not polynomial degree. Assume the following four conditions.

(i) The fixed source family has one of the two forms

$$\begin{array}{lll} \text{(P)} & F(t) \in \mathrm{Sym}_q(K[t]), & \max_{i,j} \deg_t F_{ij} \le d, \\ \text{(S)} & F(t) = M^t, & M \in \mathrm{Sym}_q(K), \quad M \succ 0, \end{array}$$

where the degree bound d in (P) *is fixed.*

(ii) For every positive integer n,

$$\text{Pl-GH}(\mathcal{F} \cup \{F(n)\}) \le_T \text{Pl-GH}(\mathcal{F}), \tag{3.1}$$

by one simulation algorithm with cost $\mathrm{poly}(|\Omega| + n)$ *on an instance Ω and parameter n.*

(iii) The integer samples are eventually entrywise strictly positive:

$$\exists n_0 \in \mathbb{Z}_{\geq 1} \quad \forall n \in \mathbb{Z}_{\geq n_0}: \qquad F(n) > 0. \tag{3.2}$$

(iv) Equal-length product identities preserve the target product:

$$\forall m \geq 1\ \forall \alpha, \beta \in \mathcal{A}_m, \qquad p_\alpha \equiv p_\beta \implies b_\alpha = b_\beta. \tag{3.3}$$

Here $\equiv$ means equality as functions of the source parameter t.

Condition (iv) ensures that every group merged by a source product identity has one well-defined target product. Then

$$\text{Pl-GH}(\mathcal{F} \cup \{B\}) \leq_T \text{Pl-GH}(\mathcal{F}). \tag{3.4}$$

The replacement matrix B may contain zero or negative entries.

For a target family $B(x)$, $x \in X \subseteq \mathbb{Q}$, keep the source family $F(t)$, the field K, and the simulation algorithm fixed. The reduction is uniform in $|\Omega| + \text{bit}(x)$ under these additional conditions:

(a) For every $x \in X$, condition (iv) holds with $b_\alpha(x) = \prod_{e:\,\alpha_e > 0} B_e(x)^{\alpha_e}$.

(b) There are effectively represented real number fields $K_x \supseteq K$ and a fixed constant c such that

$$B(x) \in \text{Sym}_q(K_x), \qquad [K_x : K] \leq c.$$

Their presentations are computable in time polynomial in $\text{bit}(x)$, and exact field arithmetic has polynomial bit cost.

(c) The entries satisfy

$$\text{bit}(B(x)) + \text{time}(B(x)) \leq \text{poly}(\text{bit}(x)),$$

where $\text{time}(B(x))$ is the time to compute all entries in the prescribed field presentation.

Thus x parametrizes the target, while t parametrizes the fixed source family. Output-field conversions are as in Remark 3.4; descent is used when the recovered value belongs to the requested smaller field.

Proof. Let Ω contain m marked occurrences of B. If $m = 0$, there is nothing to replace. Otherwise put

$$P_m = (m+1)^r, \qquad |\mathcal{A}_m| = \binom{m+r-1}{r-1} \leq P_m.$$

Zero bounds. In case (S), write $M = \sum_{a=1}^{s} \lambda_a P_a$ with distinct eigenvalues $0 < \lambda_1 < \cdots < \lambda_s$. Set

$$D_m = \begin{cases} dm, & \text{(P)}, \\ \dbinom{m+s-1}{s-1} - 1, & \text{(S)}. \end{cases}$$

For every $\alpha, \beta \in \mathcal{A}_m$,

$$p_\alpha \not\equiv p_\beta \quad \implies \quad \left|\{t \in \mathbb{R} : p_\alpha(t) = p_\beta(t)\}\right| \leq D_m. \tag{3.5}$$

In case (P), this follows from $\deg(p_\alpha - p_\beta) \leq dm$. In case (S), each difference expands as

$$p_\alpha(t) - p_\beta(t) = \sum_{\xi \in \Xi_m} c_\xi \xi^t, \qquad \Xi_m = \left\{ \prod_{a=1}^{s} \lambda_a^{\nu_a} : \nu \in \mathbb{Z}_{\geq 0}^s,\ \sum_a \nu_a = m \right\}.$$

After merging equal bases, there are at most $\binom{m+s-1}{s-1}$ terms. Lemma 3.9 gives (3.5).

A separating integer for the current length. Fix an integer n_0 as in (3.2). By (3.5),

$$p_\alpha \equiv p_\beta \quad \Longleftrightarrow \quad p_\alpha(n_0+j) = p_\beta(n_0+j) \quad \text{for every } 0 \le j \le D_m. \tag{3.6}$$

Exact evaluation at these integers therefore determines the equivalence relation $\alpha \sim \beta \Longleftrightarrow p_\alpha \equiv p_\beta$. The exceptional set

$$\mathcal{E}_m = \{n \in \mathbb{Z}_{\ge n_0} : \text{some } \alpha \not\sim \beta \text{ satisfy } p_\alpha(n) = p_\beta(n)\}$$

has size at most $\binom{P_m}{2} D_m$. Consequently, exact comparison finds an integer

$$n_m \in \{n_0, \dots, n_0 + \tbinom{P_m}{2} D_m\} \setminus \mathcal{E}_m.$$

For this integer and all $\alpha, \beta \in \mathcal{A}_m$,

$$p_\alpha(n_m) = p_\beta(n_m) \quad \Longleftrightarrow \quad p_\alpha \equiv p_\beta \quad \Longrightarrow \quad b_\alpha = b_\beta. \tag{3.7}$$

The sample n_m is chosen for the current product length m.

Parallel interpolation. Apply the length-m interpolation argument in the proof of Lemma 3.1, with $C = F(n_m)$. For the current input, (3.7) is exactly the length-m compatibility needed by that Vandermonde calculation. Write

$$\mathcal{A}_m/{\sim} = \{\mathcal{C}_1, \dots, \mathcal{C}_N\}, \qquad N \le P_m,$$

choose $\alpha^{(\ell)} \in \mathcal{C}_\ell$, and put

$$\mu_\ell = p_{\alpha^{(\ell)}}(n_m) > 0, \qquad \eta_\ell = b_{\alpha^{(\ell)}}.$$

The μ_ℓ are pairwise distinct, and each η_ℓ is independent of its representative. Replace every marked edge by $h \ge 1$ parallel $F(n_m)$-edges to obtain Ω_h. Grouping assignments into these classes gives

$$y_h := Z(\Omega_h) = \sum_{\ell=1}^{N} a_\ell \mu_\ell^h, \qquad Z(\Omega) = \sum_{\ell=1}^{N} a_\ell \eta_\ell,$$

where the a_ℓ retain all unchanged constraints. Thus, for the nonsingular Vandermonde matrix $V = (\mu_\ell^h)_{1 \le h, \ell \le N}$,

$$Z(\Omega) = (\eta_1, \dots, \eta_N) V^{-1} (y_1, \dots, y_N)^{\mathsf{T}}. \tag{3.8}$$

The bounds $P_m, D_m, n_m, N = m^{O(1)}$ show that all samples and parallel-copy counts are polynomial in the input size, and $|\Omega_h| = O(|\Omega| + mh)$ for $1 \le h \le N$. In case (P), integer samples are exact polynomial evaluations over K; in case (S), they are ordinary integer powers $M^n \in \mathrm{Sym}_q(K)$. Their entries, the products, and the interpolation systems have polynomial bit length and are computed in the fixed field K. The spectral expansion is used to bound zeros, not to approximate a sample or its logarithm. Parallel substitutions preserve planarity, including at loops. Applying (3.1) to the N queries therefore proves (3.4).

For $B = B(x)$, the chosen n_m, the bases μ_ℓ, and the coefficients a_ℓ are unchanged. Only the replacement products become

$$\eta_\ell(x) = \prod_{\substack{e \in I \\ \alpha_e^{(\ell)} > 0}} B_e(x)^{\alpha_e^{(\ell)}}.$$

The stated field and entry bounds make (3.8) computable in $\mathrm{poly}(|\Omega| + \mathrm{bit}(x))$ time. Apply the required output conversion within the individual field K_x, as in Remark 3.4; no compositum of fields for different parameter values is formed. $\square$

The following application illustrates how an order of vanishing can verify the product-identity hypothesis of Lemma 3.10. It is not needed as a separate premise in the later classification proof.

Lemma 3.11 (Full logarithmic support is hard)**.** *Let $M \in \mathrm{Sym}_q^{\mathrm{pd}}(\mathbb{R}_{\geq 0})$ have connected support, with $q \geq 3$. If every off-diagonal entry of* $\log M$ *is nonzero, then* Pl-GH(M) *is* #P*-hard. The conclusion also applies when M is jointly available in a larger jointly available constraint set.*

Proof. We apply Lemma 3.10 with source M, $F(t) = M^t = e^{tL}$, $L = \log M$, and $B = I_q + J_q$. Take a fixed number field K containing the entries of M. The four conditions are verified as follows.

(i) The assumption $M \succ 0$ gives source type (S).

(ii) For every $n \geq 1$, an n-edge path has signature M^n. Replacing all marked edges in Ω uses $O(n|\Omega|)$ edges, so the same construction gives the required uniform simulation. If M is already jointly simulated from another fixed source, composition with that simulation preserves a polynomial bound.

(iii) Positive definiteness gives $M_{ii} > 0$ for every i. Its connected nonnegative support contains a positive-weight path of length at most $q-1$ between any two colors. Diagonal steps extend this path to every length $n \geq q-1$, so $M^n > 0$ for all such n. Thus $n_0 = q-1$ suffices.

(iv) The exponential series gives

$$F(t) = I_q + tL + \frac{t^2}{2}L^2 + O(t^3), \qquad F_{ij}(t) = \begin{cases} 1 + tL_{ii} + O(t^2), & i = j, \\ tL_{ij} + O(t^2), & i \neq j. \end{cases}$$

For $m \geq 1$ and positions $(i_a, j_a) \in [q]^2$, $a \in [m]$, set

$$p(t) = \prod_{a=1}^{m} F_{i_a j_a}(t), \qquad S = \{a \in [m] : i_a \neq j_a\}, \qquad k = |S|.$$

Every off-diagonal L_{ij} is nonzero, so $c = \prod_{a \in S} L_{i_a j_a} \neq 0$, with $c = 1$ if S is empty. Consequently,

$$\begin{aligned} p(t) &= \prod_{a \notin S} (1 + O(t)) \prod_{a \in S} (tL_{i_a j_a} + O(t^2)) \\ &= t^k (c + O(t)). \end{aligned}$$

Thus

$$\mathrm{ord}_{t=0}\, p = k, \qquad \prod_{a=1}^{m} B_{i_a j_a} = 2^{m-k} = 2^{\,m - \mathrm{ord}_{t=0} p}. \tag{3.9}$$

For two products of the same length m, it follows that

$$\begin{aligned} \prod_{a=1}^{m} F_{i_a j_a}(t) \equiv \prod_{a=1}^{m} F_{u_a v_a}(t) & \\ \implies \quad & |\{a : i_a \neq j_a\}| = |\{a : u_a \neq v_a\}| \\ \implies \quad & \prod_{a=1}^{m} B_{i_a j_a} = \prod_{a=1}^{m} B_{u_a v_a}. \end{aligned}$$

This is condition (iv). The leading coefficient need only be nonzero; positivity of the entries of L is not required.

Lemma 3.10 therefore gives

$$\text{Pl-GH}(B) \leq_T \text{Pl-GH}(M), \qquad B = I_q + J_q.$$

The same substitutions retain any other fixed jointly available constraints. Proposition 2.2(ii), with $q \geq 3$, therefore gives hardness. The expansion near $t = 0$ verifies a functional identity condition; every oracle query uses a positive integer path length and no numerically approximated logarithm. ☐

4 Maximum logarithmic support and Cartesian products

We study symmetric matrices on one colour set through the largest possible off-diagonal support of their matrix logarithms. The resulting graph is a Cartesian product of complete graphs; when all factors are K_2, its coordinates also factor every nonnegative positive definite matrix in the family.

Fix a finite source language $\mathcal{F}$, with the unit-background convention of Section 2.2, a finite colour set X, and $q = |X|$. Let $\mathcal{A} \subseteq \{H \in \mathbb{R}^{X \times X} : H = H^\mathsf{T}\}$ satisfy

$$\forall \mathcal{B} \subseteq \mathcal{A}, \quad |\mathcal{B}| < \infty : \qquad \text{Pl-GH}(\mathcal{F} \cup \mathcal{B}) \leq_T \text{Pl-GH}(\mathcal{F}).$$

Thus each fixed finite $\mathcal{B}$ has a reduction permitting all its constraints in the same instance, with unrestricted occurrence counts. The reduction and its polynomial bound may depend on $\mathcal{B}$; uniformity in a variable rational parameter is established separately when needed.

Assume that $\mathcal{A}$ is closed under planar gadget compositions whose signature is symmetric on $X \times X$, and under the simulations of Lemmas 3.1, 3.3, and 3.10 whenever their hypotheses hold. Prescribed colour domains are retained as in Section 2.2. This closure ensures that each matrix constructed below remains in the same family used for maximization.

These assumptions hold for the family obtained by finite applications of those operations to jointly available constraints. Indeed, for a fixed finite $\mathcal{B}$, list its construction dependencies as $C_1, \ldots, C_m$ and put $\mathcal{G}_0 = \mathcal{F}$, $\mathcal{G}_j = \mathcal{G}_{j-1} \cup \{C_j\}$. Each gadget substitution or interpolation retains the other labels, so

$$\text{Pl-GH}(\mathcal{F} \cup \mathcal{B}) \leq_T \text{Pl-GH}(\mathcal{G}_m) \leq_T \cdots \leq_T \text{Pl-GH}(\mathcal{G}_0) = \text{Pl-GH}(\mathcal{F}).$$

The number m depends only on the fixed constructions, hence this finite composition is polynomial-time.

Section 4.1 constructs distance kernels and derives local restrictions on shortest paths and neighbourhoods. Section 4.2 extends positive definiteness to $0 < x < 1$ and proves the Cartesian product structure. Section 4.3 assumes that this product is a Boolean cube and obtains a common tensor decomposition. Factors K_s with $s \geq 3$ are excluded later by the hardness argument in Section 6.

4.1 Choosing a maximum support

For $H \succ 0$, define the simple graph

$$R_H = (X, E(R_H)), \qquad E(R_H) = \{\{i,j\} : i,j \in X,\ i \neq j,\ (\log H)_{ij} \neq 0\}.$$

Here $\log H$ is the matrix logarithm. Thus R_H has no loops, whereas Γ_H records the nonzero entries of H, including its diagonal. Assume

$$\mathcal{C} = \{H \in \mathcal{A} : H \geq 0,\ H \succ 0,\ \Gamma_H \text{ connected}\} \neq \varnothing,$$

and choose

$$M \in \arg\max_{H \in \mathcal{C}} |E(R_H)|, \qquad L := \log M, \qquad R := R_M.$$

The maximum is attained because

$$\varnothing \neq \{|E(R_H)| : H \in \mathcal{C}\} \subseteq \{0, 1, \dots, \tbinom{q}{2}\}.$$

The matrix family itself can be infinite.[1] This is an existence choice for the fixed source: M and its joint simulation are fixed before the input graph is given. The objective is the number of edges, rather than maximality under inclusion.

Lemma 4.1 (Positive entries of the logarithm)**.** *The maximizing matrix M can be chosen so that $L_{ij} > 0$ for every $ij \in E(R)$. The graph R is connected.*

Proof. The graph R is connected: otherwise $L = \log M$, and hence $M = e^L$, would be block diagonal. For a positive even integer s and small positive rational t, put $K_s(t) = (M^t)^{\circ s}$. Lemma 3.3 makes M^t jointly available, even when it has some negative entries. The Schur product theorem shows that $K_s(t)$ is PD, and the even power makes it nonnegative. For $ij \in E(R)$,

$$(\log K_s(t))_{ij} = t^s L_{ij}^s + O(t^{s+1}).$$

Indeed the off-diagonal entries of $K_s(t)$ have order at least s, whereas its diagonal perturbation from I has order at least one. Consequently every edge of R persists, with a positive entry of $\log K_s(t)$ for small t. These edges keep the support connected; the maximum edge count forbids an additional logarithmic edge. Choose one such rational parameter and replace M by the resulting matrix, retaining $L = \log M$ and $R = R_M$. $\square$

We fix this choice of M for the rest of the section.

Lemma 4.2 (Simulating $E_R(x)$)**.** *For each positive rational x, the matrix*

$$E_R(x)_{ij} = x^{d_R(i,j)}$$

belongs to $\mathcal{A}$, where d_R is shortest-path distance. The simulation is uniform in the bit length of x. More generally, let $H \succ 0$ be any fixed algebraic matrix with connected R_H and $(\log H)_{ij} > 0$ on $E(R_H)$. Then $E_{R_H}(x)$ is jointly available from H, uniformly in $\mathrm{bit}(x)$*. If $H \in \mathcal{A}$, closure also gives $E_{R_H}(x) \in \mathcal{A}$. For the maximizing graph R, for all sufficiently small positive real x, it is PD and*

$$(\log E_R(x))_{ij} = 0 \qquad (ij \notin E(R),\ i \neq j). \tag{4.1}$$

Proof. Apply Lemma 3.10 with source M, source family $F(t) = M^t = e^{tL}$, and target $B(x) = E_R(x)$, $x \in \mathbb{Q}_{>0}$. Fix a number field K containing the entries of M. Lemma 4.1 gives $L_{ij} > 0$ on every edge of R.

(i) Since $M \succ 0$, the source family has type (S).

(ii) An n-edge path with unit internal vertex weights has signature M^n. Replacing the marked edges of an instance Ω uses $O(n|\Omega|)$ edges. Composing with the fixed joint simulation of M from the original source gives the required uniform bound.

[1] In fact, if $q \geq 2$, any $H \in \mathcal{C}$ has pairwise distinct powers $H^n \in \mathcal{C}$, $n \geq 1$. They are nonnegative and PD, and $(H^n)_{ij} \geq H_{ii}^{n-1} H_{ij} > 0$ on each support edge, so connectedness is retained. Connectedness also gives $H \neq I$; a positive eigenvalue different from one distinguishes its powers.

(iii) Choose $a \geq 0$ such that $L + aI \geq 0$. Connectedness of R gives, for each i, j, a positive term in the exponential series, so

$$M^t = e^{-at} \sum_{h \geq 0} \frac{t^h (L + aI)^h}{h!} > 0 \qquad (t > 0).$$

Thus $n_0 = 1$ suffices.

(iv) For $d = d_R(i, j)$, expansion at $t = 0$ gives

$$(M^t)_{ij} = \frac{(L^d)_{ij}}{d!} t^d + O(t^{d+1}), \qquad (L^d)_{ij} > 0, \qquad \operatorname{ord}_{t=0}(M^t)_{ij} = d.$$

This includes $i = j$, for which $d = 0$ and $L^0 = I$. Fewer than d off-diagonal steps cannot join i to j; at length d, the contributing products are precisely the positive shortest-path products. Orders of vanishing add under multiplication. Consequently,

$$\prod_{a=1}^{m} (M^t)_{i_a j_a} \equiv \prod_{a=1}^{m} (M^t)_{k_a \ell_a} \implies \sum_{a=1}^{m} d_R(i_a, j_a) = \sum_{a=1}^{m} d_R(k_a, \ell_a).$$

The corresponding target products are therefore equal, because

$$\prod_{a=1}^{m} B(x)_{i_a j_a} = x^{\sum_{a=1}^{m} d_R(i_a, j_a)}.$$

The parameter conditions of the lemma also hold:

(a) The last argument applies to every $x \in \mathbb{Q}_{>0}$.

(b) All target entries are rational, so take $K_x = K$.

(c) The exponents satisfy $0 \leq d_R(i, j) \leq q - 1$; hence all entries $x^{d_R(i,j)}$ are computable with bit length and running time polynomial in $\operatorname{bit}(x)$.

More explicitly, in the notation of (3.8), the recovered coefficients a_ℓ are retained, and their target multipliers are

$$\eta_\ell(x) = x^{\sum_{e=(i,j) \in I} \alpha_e^{(\ell)} d_R(i,j)}, \qquad Z(\Omega) = \sum_{\ell=1}^{N} a_\ell \eta_\ell(x).$$

Only these multipliers depend on x; the integer source samples and their interpolated coefficients do not. This proves joint simulation uniformly in $|\Omega| + \operatorname{bit}(x)$, and closure gives $E_R(x) \in \mathcal{A}$.

The same checks with $H, \log H, R_H$ use only $H \succ 0$, connectedness of R_H, and positive logarithmic edge entries. They prove the general simulation assertion; if $H \in \mathcal{A}$, closure again gives $E_{R_H}(x) \in \mathcal{A}$.

Since $E_R(0) = I$, the distance kernel is PD near zero, and $(\log E_R(x))_{ij} = x + O(x^2)$ for $ij \in E(R)$. For sufficiently small positive rational x, put $H_x = E_R(x)$; then $H_x \in \mathcal{C}$, since it belongs to $\mathcal{A}$, is PD, and has strictly positive entries, so $|E(R_{H_x})| \leq |E(R)|$. The preceding expansion gives $E(R) \subseteq E(R_{H_x})$, hence $E(R_{H_x}) = E(R)$. The logarithmic entries are analytic near zero, so continuity and density of the rational parameters give (4.1). $\square$

Lemma 4.3 (Weighted shortest paths)**.** *Let $N \in \mathcal{A}$, $N \succ 0$, and $K = \log N$ satisfy*

$$R_N = R, \qquad K_{uv} > 0 \quad (uv \in E(R)).$$

For distinct $i, j \in X$, let $h = d_R(i, j)$ and let $\mathcal{P}_{ij}$ be the set of shortest i-j paths in R. Then

$$|\mathcal{P}_{ij}| = h!, \qquad \prod_{uv \in P} K_{uv} = \prod_{uv \in Q} K_{uv} \quad (P, Q \in \mathcal{P}_{ij}).$$

In particular, this applies to $N = M$ and $K = L$.

Proof. Put $a_P = \prod_{uv\in P} K_{uv} > 0$. For each positive even integer s, the matrix $H_s(t) = (N^t)^{\circ s}$ belongs to $\mathcal{A}$ for positive rational t, by spectral interpolation and parallel composition. It is nonnegative and PD by the Schur product theorem. On every edge of R, both its entries and its logarithmic entries have leading term $t^s K_{uv}^s > 0$. Thus, for small rational $t > 0$, it belongs to $\mathcal{C}$ and retains every edge of R in its logarithmic support. The maximum edge count forces $\log H_s(t)$ to vanish at all nonedges, hence analytically near zero.

At distance h, comparison of the t^{sh}-coefficients in $H_s(t) = \exp(\log H_s(t))$ gives

$$\left(\frac{\sum_{P\in\mathcal{P}_{ij}} a_P}{h!}\right)^s = \frac{\sum_{P\in\mathcal{P}_{ij}} a_P^s}{h!}. \tag{4.2}$$

Only h off-diagonal factors along a shortest path contribute at this order; diagonal factors increase the order. Put $A = \sum_P a_P/h!$ and $a_* = \max_P a_P$. Taking s-th roots in (4.2) along the positive even integers gives

$$A = \lim_{\substack{s\to\infty\\ 2|s}} \left(\frac{\sum_P a_P^s}{h!}\right)^{1/s} = a_*.$$

Substituting $A = a_*$ back into (4.2) yields $\sum_P (a_P/a_*)^s = h!$ for every positive even s. Since $0 < a_P/a_* \le 1$ and the sum is finite,

$$h! = \lim_{\substack{s\to\infty\\ 2|s}} \sum_P (a_P/a_*)^s = |\{P \in \mathcal{P}_{ij} : a_P = a_*\}|.$$

Also $\sum_P a_P = h! a_*$, so

$$\sum_{P: a_P < a_*} a_P = 0.$$

Every a_P is positive; hence there are no smaller terms. □

Lemma 4.4 (Nonadjacent common neighbours)**.** *Every pair of vertices at distance two in R has exactly two common neighbours, and these are nonadjacent. Every neighbourhood in R is a disjoint union of cliques.*

Proof. The number of common neighbours is two by Lemma 4.3. For $E = E_R(x)$, use the planar gadget in Figure 1: its terminals ℓ_1, ℓ_2 are joined by the paths $\ell_1 p_1 \ell_2$ and $\ell_1 p_2 \ell_2$, with an additional edge $p_1 p_2$. Every edge carries E. Fixing terminal colours i, j and summing over internal colours α, β gives

$$\Omega(E)_{ij} = \sum_{\alpha,\beta\in X} E_{i\alpha} E_{\alpha j} E_{i\beta} E_{\beta j} E_{\alpha\beta}.$$

Both terminals lie on the outer face, so $\Omega(E) \in \mathcal{A}$ for positive rational x. Moreover, $\Omega(E) > 0$ for $x > 0$ and $\Omega(E) \to I$ as $x \to 0$. It is therefore PD, and belongs to $\mathcal{C}$, for sufficiently small positive rational x.

For any terminal colours a, b, define

$$g_{ab}(z) := d_R(a, z) + d_R(z, b) \ge d_R(a, b).$$

Substitution of $E_{ab} = x^{d_R(a,b)}$ gives the contribution

$$E_{a\alpha} E_{\alpha b} E_{a\beta} E_{\beta b} E_{\alpha\beta} = x^{g_{ab}(\alpha)+g_{ab}(\beta)+d_R(\alpha,\beta)}.$$

Thus each ordered pair (α, β) contributes one monomial; in particular, $\Omega(E)_{ab} = O(x^{2d_R(a,b)})$ for $a \ne b$.

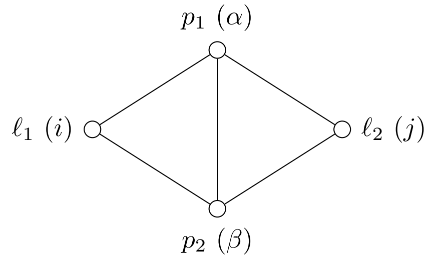

Figure 1: The Wheatstone gadget defining $\Omega(E)$. Every edge carries E; parentheses indicate the assigned colours.

Now fix $d_R(i,j) = 2$, let u, v be the two common neighbours, and put

$$S = \{i, j, u, v\}, \qquad \varepsilon = \mathbf{1}_{\{uv \in E(R)\}}.$$

Then

$$g_{ij}(z) = 2 \iff z \in S, \qquad g_{ij}(z) \geq 3 \quad (z \notin S).$$

Degree four in x requires $\alpha = \beta \in S$, giving four terms. Degree five requires an ordered adjacent pair in S: the edges iu, iv, ju, jv give eight terms, and uv, if present, gives two more. If an internal colour is outside S, coincident colours give degree at least $3+3 = 6$, and distinct colours give degree at least $3+2+1 = 6$. Therefore

$$\Omega(E)_{ij} = 4x^4 + (8 + 2\varepsilon)x^5 + O(x^6).$$

For an edge $ab \in E(R)$, we have $g_{ab}(z) = 1 \iff z \in \{a, b\}$. The internal assignments $(a,a), (b,b)$ contribute $2x^2$, and $(a,b), (b,a)$ contribute $2x^3$. Every other assignment has degree at least four: equal internal colours give $2g_{ab}(\alpha) \geq 4$, and distinct colours give at least $2+1+1 = 4$. Hence

$$\Omega(E)_{ab} = 2x^2 + 2x^3 + O(x^4) \qquad (ab \in E(R)).$$

On the diagonal, $g_{aa}(z) = 2d_R(a,z)$, so

$$\begin{aligned}\Omega(E)_{aa} &= 1 + 2\sum_{z \neq a} x^{3d_R(a,z)} \\ &\quad + \sum_{\substack{\alpha \neq a \\ \beta \neq a}} x^{2d_R(a,\alpha) + 2d_R(a,\beta) + d_R(\alpha,\beta)} = 1 + O(x^3).\end{aligned}$$

The last sum is $O(x^4)$.

Set $Q = \Omega(E) - I$. In $(Q^2)_{ij} = \sum_k Q_{ik}Q_{kj}$, only $k = u, v$ contribute below degree six. Each contributes

$$(2x^2 + 2x^3 + O(x^4))^2 = 4x^4 + 8x^5 + O(x^6).$$

The terms $k = i, j$ are $O(x^3)O(x^4) = O(x^7)$; for every other k, at least one of ik, kj is a nonedge, giving $O(x^4)O(x^2) = O(x^6)$. Thus

$$(Q^2)_{ij} = 8x^4 + 16x^5 + O(x^6).$$

All entries of Q are $O(x^2)$, so the convergent matrix logarithm series gives $\log(I + Q) = Q - Q^2/2 + O(x^6)$. Consequently

$$\begin{aligned}(\log \Omega(E))_{ij} &= 4x^4 + (8 + 2\varepsilon)x^5 - \tfrac{1}{2}(8x^4 + 16x^5) + O(x^6) \\ &= 2\varepsilon x^5 + O(x^6).\end{aligned}$$

For every $ab \in E(R)$, $(\log \Omega(E))_{ab} = 2x^2 + 2x^3 + O(x^4) > 0$ for small $x > 0$, so all edges of R remain. Since $\Omega(E) \in \mathcal{C}$ for small positive rational x, $\varepsilon = 1$ would add the nonedge ij to its logarithmic support and violate the maximum edge count. Hence $\varepsilon = 0$.

Finally, an induced path a, b, c in the neighbourhood of a vertex w would give the nonadjacent pair a, c two adjacent common neighbours, w and b, a contradiction. Each neighbourhood therefore has no induced three-vertex path. A connected noncomplete component would contain such a path as the first three vertices of a shortest path between nonadjacent vertices. Thus every component of the neighbourhood is a clique. $\square$

4.2 Entropy, positive definiteness and Cartesian products

The distance kernel is initially PD only near zero. We use the entropy derivatives in Section 2.4 to identify it with a unique PD optimizer for every $0 < x < 1$. This analytic extension then turns graph distance into squared Euclidean distance in Lemma 4.6; it requires no oracle for the optimizer.

Lemma 4.5 (Positive definiteness throughout the parameter interval)**.** *Let R be a finite connected simple graph. If $\log E_R(x)$ vanishes at all nonedges for all sufficiently small positive x, then $E_R(x)$ is PD for every $0 < x < 1$.*

Proof. Fix $x \in (0, 1)$, put $q = |V(R)|$, and consider the compact convex set

$$\mathcal{F}_x = \{H \succeq 0 : H_{ii} = 1,\ H_{ij} = x\ (ij \in E(R))\}.$$

The PSD cone and the affine entry constraints are closed and convex, so $\mathcal{F}_x$ is closed and convex. Its two-by-two principal minors give

$$H \succeq 0, \quad H_{ii} = H_{jj} = 1 \quad \Longrightarrow \quad 1 - H_{ij}^2 \geq 0 \quad \Longrightarrow \quad |H_{ij}| \leq 1.$$

Thus $\mathcal{F}_x$ is bounded, and the Heine–Borel theorem in the finite-dimensional space of real symmetric matrices makes it compact. There is a strictly positive definite feasible point $P_x = (1 - x)I + xJ$. Writing $\lambda_1(H), \ldots, \lambda_q(H)$ for the eigenvalues, minimize the continuous function

$$\Phi(H) = \sum_{j=1}^{q} (\lambda_j(H) \log \lambda_j(H) - \lambda_j(H)), \qquad 0 \log 0 = 0.$$

For PD H, this equals $\operatorname{Tr}(H \log H - H)$; the eigenvalue formula defines it also on the singular boundary. On $\mathcal{F}_x$, $\operatorname{Tr} H = q$, and hence

$$-\Phi(H) = -\operatorname{Tr}(H \log H) + q.$$

Minimizing Φ therefore maximizes the von Neumann entropy

$$-\operatorname{Tr}(H \log H) = -\sum_j \lambda_j(H) \log \lambda_j(H),$$

with the same spectral convention on the boundary.

A minimizer cannot be singular. If H has a kernel, along $H_\epsilon = (1 - \epsilon)H + \epsilon P_x$ the eigenvalues arising from that kernel are $\epsilon c_j + O(\epsilon^2)$, with $c_j > 0$. They change Φ by $(\sum_j c_j)\epsilon \log \epsilon + O(\epsilon)$, whereas all other eigenvalues contribute only $O(\epsilon)$. This decreases Φ for small positive ϵ.

The gradient is $\log H$. For $0 \neq V = V^{\mathsf{T}}$,

$$\langle V, D \log_H [V] \rangle_F = \int_0^\infty \|(H + sI)^{-1/2} V (H + sI)^{-1/2}\|_F^2 \, ds > 0.$$

Thus Φ is strictly convex on the PD cone and has a unique feasible minimizer $H(x)$. Its first-order condition is

$$(\log H(x))_{ij} = 0 \quad \text{for all } i \neq j \text{ with } ij \notin E(R),$$

because precisely these symmetric entries may vary freely.

To obtain analyticity on all of $(0,1)$, let

$$\mathcal{S}_R = \{C = C^\mathsf{T} : C_{ij} = 0 \text{ if } i \neq j,\ ij \notin E(R)\}$$

and let π_R be its Frobenius-orthogonal projection. Writing A_R for the adjacency matrix of R, the minimizer solves

$$C(x) = \log H(x) \in \mathcal{S}_R, \qquad F(C(x)) = I + xA_R, \qquad F(C) = \pi_R(e^C).$$

For $0 \neq V \in \mathcal{S}_R$,

$$\langle V, DF_C[V]\rangle_F = \int_0^1 \operatorname{Tr}(Ve^{(1-s)C}Ve^{sC})\, ds > 0.$$

In an orthogonal eigenbasis of C, the integrand is a sum of squares of entries of V with positive coefficients. If $DF_C[V] = 0$, the displayed positive quadratic form forces $V = 0$. Thus $\ker DF_C = \{0\}$, and rank–nullity on the same finite-dimensional domain and codomain makes $DF_C : \mathcal{S}_R \to \mathcal{S}_R$ invertible. The analytic implicit-function theorem gives a local analytic solution at every $x \in (0,1)$; uniqueness patches them into $H(x)$.

Near zero, $E_R(x)$ is itself PD and satisfies the first-order condition, so it equals $H(x)$. The identity theorem for real analytic functions, entry by entry on the connected interval, gives $E_R(x) = H(x)$ throughout $(0,1)$. □

Lemma 4.6 (Cartesian product characterization)**.** *Let R be a finite connected simple graph. Suppose that $E_R(x) \succeq 0$ for every $0 < x < 1$, that each neighbourhood is a disjoint union of cliques, and that each distance-two pair has exactly two nonadjacent common neighbours. Then*

$$R \cong K_{s_1} \square \cdots \square K_{s_d}, \qquad s_r \geq 2.$$

The one-vertex graph is the empty product.

Proof. We first represent graph distance as a squared Euclidean distance; then local squares translate entire cliques and determine product coordinates. Let $\Delta = (d_R(i,j))_{i,j\in V(R)}$, $q = |V(R)|$, and $P = I - J/q$, the orthogonal projector onto $\mathbf{1}^\perp = \{z : z^\mathsf{T}\mathbf{1} = 0\}$. The expansion

$$E_R(x) = J - (1-x)\Delta + O((1-x)^2)$$

and $E_R(x) \succeq 0$ imply

$$z^\mathsf{T}\mathbf{1} = 0 \implies 0 \leq \lim_{x\uparrow 1} \frac{z^\mathsf{T}E_R(x)z}{1-x} = -z^\mathsf{T}\Delta z.$$

Hence $K = -P\Delta P/2 \succeq 0$. A Gram representation of K gives vectors $f(v) \in \mathbb{R}^q$ such that

$$\|f(u) - f(v)\|^2 = d_R(u,v). \tag{4.3}$$

Indeed, double centring and $\Delta_{ii} = 0$ give $K_{uu} + K_{vv} - 2K_{uv} = \Delta_{uv}$. Thus f is injective.

An induced four-cycle u, v, z, w, u becomes a parallelogram. Indeed $a = f(v) - f(u)$ and $b = f(w) - f(u)$ are orthogonal unit vectors. If $c = f(z) - f(u)$, the prescribed distances give $\|c\|^2 = 2$ and $\langle c, a\rangle = \langle c, b\rangle = 1$. The projection $a + b$ already has squared norm two, so

$$f(z) = f(v) + f(w) - f(u). \tag{4.4}$$

Every edge uv belongs to a unique maximal clique A, by the neighbourhood assumption. For $w \in N_R(u) \setminus A$, the vertices v, w are nonadjacent and have, besides u, exactly one common neighbour z. The latter is nonadjacent to u, so these vertices form an induced four-cycle. This construction and its reverse give a bijection

$$N_R(u) \setminus A \longrightarrow N_R(v) \setminus A, \qquad f(z) - f(v) = f(w) - f(u).$$

Extend this bijection by $u \mapsto v$. On each maximal clique through u other than A, it acts by the translation $f(w) \mapsto f(w) + f(v) - f(u)$, so its image is a clique through v. If that image had an extra vertex, the inverse bijection would enlarge the original maximal clique. Thus entire maximal cliques correspond, with none added or omitted.

Fix a base vertex o, translate so that $f(o) = 0$, and list its maximal cliques as $C_1, \ldots, C_d$. Let $S_r = f(C_r)$ and $U_r = \operatorname{span} S_r$. Each S_r is the vertex set of a regular simplex of unit side, containing zero. The spaces U_r are pairwise orthogonal: vertices in two different base cliques have distance two, whereas their squared norms are one.

We now keep track of all neighbours, rather than just selected directions. At a point $z = \sum_r z_r$, with $z_r \in S_r$, the vertex sets of the maximal cliques through it, in the Euclidean representation, are precisely

$$z + (S_r - z_r), \qquad 1 \le r \le d. \tag{4.5}$$

Here $z + (S_r - z_r) = \{z + s - z_r : s \in S_r\}$. The assertion holds at o. When moving along clique i, that clique stays unchanged, while every other clique is translated by (4.4); the bijection of entire maximal cliques just proved ensures there are no extra or missing cliques. Thus the assertion holds inductively along every path. Connectedness and the preceding induction give one inclusion in

$$f(V(R)) = \left\{ \sum_{r=1}^{d} z_r : z_r \in S_r \right\}.$$

For the reverse inclusion, (4.5) permits changing any one z_r to any other element of S_r. Orthogonality makes each coordinate representation unique, and

$$\left\| \sum_r z_r - \sum_r z_r' \right\|^2 = \sum_r \|z_r - z_r'\|^2 = |\{r : z_r \neq z_r'\}|.$$

By (4.3), adjacency is therefore exactly one coordinate change, giving $R \cong \square_{r=1}^d K_{|S_r|}$. $\square$

Theorem 4.7 (Maximum support is a Cartesian product)**.** *The logarithmic support R of maximum edge count is a Cartesian product of complete graphs, whose sizes need not be equal.*

Proof. Apply Lemma 4.5 to Lemma 4.2, and then apply Lemma 4.6 using Lemma 4.4. $\square$

4.3 Tensor decomposition in common coordinates

Proposition 4.8 (Tensor decomposition in common coordinates)**.** *Assume $R \cong Q_d = \square_{r=1}^d K_2$, and fix a graph isomorphism $\iota : R \to Q_d$. Every $N \in \mathcal{A}$ with $N \ge 0$ and $N \succ 0$ has, in these same coordinates, a factorization*

$$N = \gamma \bigotimes_{r=1}^{d} C_r, \qquad \gamma > 0,$$

where the C_r are nonnegative PD 2×2 matrices. If N has connected support, every C_r is strictly positive. If N has algebraic entries, all factors can be chosen algebraic. The assertion also holds when source constraints have specified endpoint domains. No definiteness assumption on the original source constraints is required.

Proof. Put $B = \log N$. We first prove $B_{ij} = 0$ for $i \neq j$ with $ij \notin E(R)$. For small positive rational t, η, Lemma 3.3 supplies $M^{t/2}$ and $N^{\eta t}$. Symmetric series composition then makes

$$A(t, \eta) = M^{t/2} N^{\eta t} M^{t/2}$$

available. It is PD, although entrywise positivity is not automatic. Take its entrywise square $H = A^{\circ 2}$; this is PD and nonnegative. Since $\log A = t(L + \eta B) + O(t^3)$, at every nonzero off-diagonal position of $L + \eta B$,

$$(\log H)_{ij} = t^2 (L + \eta B)_{ij}^2 + O(t^3).$$

Choose a sufficiently small fixed rational $\eta > 0$ to retain every edge of R. A nonzero off-diagonal entry B_{ij} with $ij \notin E(R)$ would then increase the edge count. Consequently $B_{ij} = 0$ for every nonedge ij, $i \neq j$. This step also works if N is disconnected, since the retained edges of M make H connected.

We next prove that $B_{z,z+e_r}$ depends only on the direction r. Here $0 \in \{0,1\}^d$ is the all-zero string, e_r has only its r-th bit equal to one, and addition of strings is modulo two. Both L and B now have off-diagonal support on the cube. At distance h, the series for $A(t, \eta)$ has no term below t^h; its coefficient at t^h is continuous in η and is positive at $\eta = 0$. Thus in an open region of small positive t, η, $A(t, \eta)$ is entrywise positive and its logarithm has positive weights on all cube edges. For rational parameters, $A(t, \eta) \in \mathcal{C}$, and the maximum edge count forces its logarithm to vanish at all nonedges. Lemma 4.3 therefore applies to both diagonals of every square.

For a cyclic square $v_1 v_2 v_3 v_4 v_1$, label its edge entries of $\log A(t, \eta)$ as in Figure 2:

$$\begin{aligned} a &= (\log A)_{v_1 v_2}, \quad & b &= (\log A)_{v_2 v_3}, \\ c &= (\log A)_{v_3 v_4}, \quad & e &= (\log A)_{v_4 v_1}. \end{aligned}$$

All four weights are positive. The two diagonal pairs give

$$v_1 \leftrightarrow v_3 : \quad ab = ec, \qquad v_2 \leftrightarrow v_4 : \quad ae = bc,$$

and hence $a = c$, $b = e$. Thus opposite edge weights of $\log A(t, \eta)$ agree on every square. By continuity these equalities hold throughout the open parameter region.

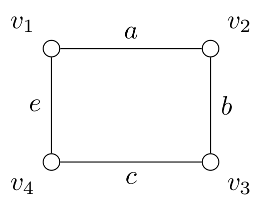

Figure 2: Cyclic edge weights of $\log A(t, \eta)$ on a cube square. The opposite pairs are (a, c) and (b, e).

Divide the opposite-edge equalities by t and let $t \downarrow 0$. Since $\log A(t, \eta) = t(L + \eta B) + O(t^3)$, and the resulting equalities hold for an interval of η, their coefficients of η give

$$B_{z,z+e_r} = B_{z+e_s, z+e_s+e_r} \qquad (r \neq s).$$

Every pair of edges in direction r is joined by such squares; symmetry also permits exchanging their endpoints. Therefore

$$b_r := B_{0,e_r}, \qquad B_{z,z+e_r} = b_r \quad (z \in \{0,1\}^d,\ 1 \leq r \leq d).$$

Putting $\phi(z) = B_{z,z}$, the entries of B are precisely

$$B_{z,w} = \begin{cases} \phi(z), & w = z, \\ b_r, & w = z + e_r \quad (1 \le r \le d), \\ 0, & \text{otherwise.} \end{cases}$$

Define the bit-flip matrices by

$$(X_r)_{z,w} = \mathbf{1}_{\{w=z+e_r\}}, \qquad X_r = I_2^{\otimes(r-1)} \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \otimes I_2^{\otimes(d-r)}.$$

The middle factor flips bit r; the identity factors preserve all other bits. Comparing entries gives

$$B = \operatorname{diag}\phi + \sum_{r=1}^{d} b_r X_r. \tag{4.6}$$

It remains to exclude interactions between distinct coordinates in ϕ. For a fixed rational $x \in (0,1)$, let

$$C = \log E_{Q_d}(x) = cI + b\sum_r X_r, \qquad c = \frac{d}{2}\log(1-x^2), \quad b = \frac{1}{2}\log\frac{1+x}{1-x} > 0.$$

Consider $e^{tC/2}e^{uB}e^{tC/2}$. As in the preceding positivity argument, there is an open sector $t > 0,\ 0 < u < \eta_0 t$, sufficiently near zero, on which this matrix is entrywise positive and its logarithm retains every edge of R. Its nonedge logarithmic entries vanish at rational parameters by the maximum edge count, hence throughout the sector and identically in their analytic expansions at zero. With $[A,B] = AB - BA$, the coefficient of t^2u is

$$-\frac{1}{24}[C,[C,B]].$$

For completeness, the products C^2B, CBC, and BC^2 have coefficients $-1/24$, $1/12$, and $-1/24$, respectively. These coefficients follow by expanding the three exponentials and the matrix logarithm about I to total degree three. Since $X_rX_s = X_sX_r$, we have $[C, \sum_r b_rX_r] = 0$; the double commutator depends only on $\operatorname{diag}\phi$. At opposite square vertices i, j, with intermediate vertices v, w, its entry is

$$2b^2(\phi_i + \phi_j - \phi_v - \phi_w).$$

Vanishing of the nonedge coefficient therefore gives

$$\phi(z) + \phi(z+e_r+e_s) - \phi(z+e_r) - \phi(z+e_s) = 0 \quad (r \ne s),$$

with bit addition modulo two. For $z_r = 0$, this says that the increment $\phi(z+e_r) - \phi(z)$ is unchanged by flipping any other bit. Changing those bits successively gives

$$\phi(z+e_r) - \phi(z) = \phi(e_r) - \phi(0) \qquad (z_r = 0).$$

Starting at 0 and changing each required bit from zero to one yields

$$\phi(z) = \phi(0) + \sum_{r=1}^{d}(\phi(e_r) - \phi(0))z_r.$$

Putting $D_r = \left(\begin{smallmatrix} 0 & b_r \\ b_r & \phi(e_r)-\phi(0) \end{smallmatrix}\right)$, equation (4.6) gives

$$B = \phi(0)I + \sum_{r=1}^{d} I_2^{\otimes(r-1)} \otimes D_r \otimes I_2^{\otimes(d-r)}.$$

The summands act on different tensor factors and commute, so

$$N = e^B = e^{\phi(0)} \bigotimes_{r=1}^{d} e^{D_r}.$$

Each factor e^{D_r} is PD with positive diagonal. An entry of N changing only bit r shows that the factor's off-diagonal entry is nonnegative; a zero entry disconnects the support along that coordinate. Thus all factors are strictly positive when Γ_N is connected.

Normalizing each factor's 00 entry to one gives, in the fixed coordinates,

$$\gamma = N_{0,0}, \qquad C_r = \begin{pmatrix} 1 & N_{0,e_r}/N_{0,0} \\ N_{0,e_r}/N_{0,0} & N_{e_r,e_r}/N_{0,0} \end{pmatrix}, \qquad N = \gamma \bigotimes_{r=1}^{d} C_r.$$

These factors are algebraic when N is. Every construction used joint availability with the prescribed endpoint domains; no definiteness assumption on the original source was used. ☐

5 Hardness from unequal diagonals in a tensor factor

The tensor identity evaluates a product of Boolean partition functions, but hardness requires recovering a factor from that product. The algebraic-function tools of Section 2.5 distinguish unequal-diagonal parameter classes and bound exceptional sample values; two finite interpolations implement the resulting separation.

Theorem 5.1 (Tensor products of Boolean-domain matrices)**.** *Let $N = C_1 \otimes \cdots \otimes C_d$, where each C_r is a strictly positive algebraic PD 2×2 matrix. If every factor has equal diagonal entries,* Pl-GH(N) *is polynomial-time computable. If at least one factor has unequal diagonal entries,* Pl-GH(N) *is* #P*-hard.*

Proof. Use unit background weights. The tensor identity

$$Z_N(G) = \prod_{r=1}^{d} Z_{C_r}(G)$$

and the planar zero-field Ising algorithm give the easy case. For hardness, fix G with $m = |E(G)| \geq 1$; edgeless inputs are directly computable. The main task is to recover a Boolean factor with unequal diagonal entries from the full tensor. We first retain all copies in its parameter class, then recover one copy by a fixed-degree positive root.

We normalize the factors and construct a parameterized family. After permuting each pair of colors, write $C_r = \left(\begin{smallmatrix} u_r & v_r \\ v_r & z_r \end{smallmatrix}\right)$ with $u_r \geq z_r > 0$. Set

$$\theta_r = \sqrt{u_r/z_r} \geq 1, \qquad w_r = v_r/\sqrt{u_r z_r} \in (0,1), \qquad \gamma = \prod_{r=1}^{d} \sqrt{u_r z_r} > 0.$$

Then

$$N = \gamma \bigotimes_{r=1}^{d} \begin{pmatrix} \theta_r & w_r \\ w_r & \theta_r^{-1} \end{pmatrix}, \qquad \gamma > 0, \quad \theta_r \geq 1, \quad 0 < w_r < 1. \tag{5.1}$$

All parameters are algebraic. A factor has unequal diagonal entries precisely when $\theta_r > 1$. The logarithm of N is a Kronecker sum with positive edge weights on Q_d. Indeed, if a positive PD 2×2 matrix has off-diagonal entry $v > 0$ and eigenvalues $\alpha_+ > \alpha_- > 0$, the off-diagonal entry of its logarithm is

$$v \, \frac{\log \alpha_+ - \log \alpha_-}{\alpha_+ - \alpha_-} > 0.$$

The scalar part does not affect the off-diagonal support. The general simulation assertion of Lemma 4.2 therefore makes $E_{Q_d}(x)$ jointly available.

The diagonal matrix of source loop weights is

$$D = \operatorname{diag}(N_{ii}) = \gamma \bigotimes_{r=1}^{d} \operatorname{diag}(\theta_r, \theta_r^{-1}).$$

For $k \in \mathbb{Z}_{\geq 0}$, adding k source self-loops at each endpoint realizes $D^k N D^k$. With $t = 2k+1$,

$$\gamma^{-t} D^k N D^k = \bigotimes_{r=1}^{d} \begin{pmatrix} \theta_r^t & w_r \\ w_r & \theta_r^{-t} \end{pmatrix},$$

because the off-diagonal entry of each factor is $\theta_r^k w_r \theta_r^{-k} = w_r$. The known scalar is removed from evaluations: m' occurrences contribute $\gamma^{tm'}$. In the same cube coordinates,

$$E_{Q_d}(x)_{z,z'} = x^{|\{r : z_r \neq z'_r\}|} = \prod_{r=1}^{d} x^{\mathbf{1}_{\{z_r \neq z'_r\}}}, \qquad E_{Q_d}(x) = \begin{pmatrix} 1 & x \\ x & 1 \end{pmatrix}^{\otimes d}.$$

Parallel composition identifies the two terminals and keeps the internal vertices disjoint, so it multiplies the signatures entrywise. Consequently the jointly simulated matrix is

$$F_t(x) = (\gamma^{-t} D^k N D^k) \circ E_{Q_d}(x).$$

The identity

$$\begin{pmatrix} \theta_r^t & w_r \\ w_r & \theta_r^{-t} \end{pmatrix} \circ \begin{pmatrix} 1 & x \\ x & 1 \end{pmatrix} = \begin{pmatrix} \theta_r^t & w_r x \\ w_r x & \theta_r^{-t} \end{pmatrix}$$

gives

$$F_t(x) = \bigotimes_{r=1}^{d} B_r(x), \qquad B_r(x) = \begin{pmatrix} \theta_r^t & w_r x \\ w_r x & \theta_r^{-t} \end{pmatrix}. \tag{5.2}$$

For rational $0 < x < 1$ all these matrices are strictly positive and PD, since $\det B_r(x) = 1 - w_r^2 x^2 > 0$. Thus $F_t(x)$ is jointly available from N, uniformly for rational $0 < x < 1$, for each fixed positive odd t.

We next distinguish the unequal-diagonal parameter classes through their eigenvalue functions. Partition the coordinate indices into classes of identical pairs (θ_r, w_r), and choose one representative g from each class. Denote the set of representatives by $\mathcal{G}$. For $g \in \mathcal{G}$, let

$$I_g = \{r : (\theta_r, w_r) = (\theta_g, w_g)\}, \qquad n_g = |I_g|;$$

write $g(r)$ for the representative of the class containing r. For every class put

$$c_g = \frac{\theta_g^t + \theta_g^{-t}}{2}, \quad a_g = \frac{\theta_g^t - \theta_g^{-t}}{2}.$$

For classes with $\theta_g > 1$, also put $\beta_g = a_g / w_g > 0$. Choose a fixed positive odd integer t for which all these β_g's are distinct. Such a choice exists: equal θ's with different w's already give different β's. For $h_g = \log\theta_g \neq h_h = \log\theta_h$, both positive, write $\sinh z = (e^z - e^{-z})/2$, $\cosh z = (e^z + e^{-z})/2$, and $\coth z = \cosh z/\sinh z$. Then

$$\frac{d}{dt}\log\frac{\sinh(th_g)}{\sinh(th_h)} = h_g\coth(th_g) - h_h\coth(th_h).$$

This has constant nonzero sign because

$$\frac{d}{dz}(z\coth z) = \frac{\sinh z\cosh z - z}{\sinh^2 z} > 0 \quad (z > 0).$$

Indeed the numerator vanishes at zero and has derivative $2\sinh^2 z > 0$. Each pair therefore excludes at most one $t > 0$; finitely many pairs exclude only finitely many positive odd integers.

For the representative factor $B_g(x)$, define its eigenvalues and their ratio by

$$\lambda_g^{\pm}(x) = c_g \pm \sqrt{a_g^2 + w_g^2x^2}, \qquad \delta_g(x) = \frac{\lambda_g^-(x)}{\lambda_g^+(x)}. \tag{5.3}$$

On $0 < x < 1$ choose the positive real branch of $\sqrt{a_g^2 + w_g^2x^2}$. Since

$$c_g^2 - (a_g^2 + w_g^2x^2) = 1 - w_g^2x^2 > 0, \qquad c_g > 0,$$

we have $0 < \sqrt{a_g^2 + w_g^2x^2} < c_g$, hence $\lambda_g^+ > \lambda_g^- > 0$. If $\theta_g = 1$, then $a_g = 0, c_g = 1$, $\sqrt{w_g^2x^2} = w_gx$ on this interval, and $\lambda_g^{\pm}(x) = 1 \pm w_gx$. If $\theta_g > 1$, the factorization

$$a_g^2 + w_g^2x^2 = w_g^2(x - \mathrm{i}\beta_g)(x + \mathrm{i}\beta_g), \qquad \mathrm{i}^2 = -1,$$

shows that $\sqrt{a_g^2 + w_g^2x^2}$ has branch points $x = \pm\mathrm{i}\beta_g$ at two distinct simple zeros. By the choice of t, neither point is a branch point of any other class. Suppose there is an identity of algebraic functions

$$\prod_{g\in\mathcal{G}} \delta_g(x)^{z_g} = 1, \qquad z_g \in \mathbb{Z}. \tag{5.4}$$

A loop around a branch point unique to class g, with $\theta_g > 1$, exchanges the two eigenvalue branches of that class. More precisely, analytic continuation around $x = \mathrm{i}\beta_g$ gives

$$\sqrt{a_g^2 + w_g^2x^2} \longmapsto -\sqrt{a_g^2 + w_g^2x^2}, \qquad \lambda_g^+ \longleftrightarrow \lambda_g^-,$$

$$\delta_g \longmapsto \delta_g^{-1}, \qquad \delta_h \longmapsto \delta_h \quad (h \neq g).$$

All copies in class g undergo the same exchange. Dividing the continued identity by the original gives $\delta_g^{2z_g} = 1$. Since $0 < \delta_g(x) < 1$ on $0 < x < 1$, this forces $z_g = 0$. Thus

$$\prod_{g\in\mathcal{G}} \delta_g^{z_g} \equiv 1 \quad \Longrightarrow \quad z_g = 0 \quad \text{for every } g \text{ with } \theta_g > 1.$$

Continuation is used only to prove identities; all oracle query parameters remain positive real algebraic numbers.

The eigenvalues of $F_t(x)$, indexed by $\epsilon \in \{0,1\}^d$, are

$$\Lambda_\epsilon(x) = \left(\prod_{r=1}^{d} \lambda_{g(r)}^+(x)\right)\prod_{r=1}^{d} \delta_{g(r)}(x)^{\epsilon_r}.$$

Choose $\epsilon^{(1)}, \dots, \epsilon^{(m)} \in \{0,1\}^d$, allowing repetitions. In $\epsilon_r^{(a)}$, the superscript (a) labels the factor in a product of m eigenvalues, and r labels its tensor coordinate. The values $0, 1$ select $\lambda_{g(r)}^+$ and $\lambda_{g(r)}^-$, respectively; both eigenvalues are positive. Thus

$$\prod_{a=1}^m \Lambda_{\epsilon^{(a)}}(x) = \left(\prod_{r=1}^d \lambda_{g(r)}^+(x)\right)^m \prod_{h\in\mathcal{G}} \delta_h(x)^{k_h}, \qquad k_h = \sum_{a=1}^m \sum_{r\in I_h} \epsilon_r^{(a)}.$$

In two identically equal products of length m, the common first factor cancels. The preceding implication gives

$$\prod_{a=1}^m \Lambda_{\epsilon^{(a)}} \equiv \prod_{a=1}^m \Lambda_{\widetilde{\epsilon}^{(a)}} \quad \Longrightarrow \quad k_h = \widetilde{k}_h \quad (\theta_h > 1).$$

Hence a product identity preserves the total number of smaller-branch choices in every unequal-diagonal class.

We can therefore retain one such class g and replace all other coordinates by I_2. Define $T_g(x)$ in the original coordinate order by

$$\big(T_g(x)\big)_{z,z'} = \prod_{r\in I_g} \big(B_r(x)\big)_{z_r,z_r'} \prod_{r\notin I_g} \mathbf{1}_{\{z_r=z_r'\}}, \qquad z, z' \in \{0,1\}^d. \tag{5.5}$$

Equivalently, after listing the coordinates in I_g first, it is $B_g(x)^{\otimes n_g} \otimes I_2^{\otimes(d-n_g)}$. It has the same tensor eigenbasis as $F_t(x)$. The corresponding product of m eigenvalues of $T_g(x)$ is

$$(\lambda_g^+(x))^{n_g m} \delta_g(x)^{k_g}.$$

Consequently, identically equal source products have equal target products: the combined contribution of an indistinguishable source class admits a well-defined replacement weight. Equality at one isolated parameter value requires a separate check.

Two finite interpolations turn these function identities into a polynomial-time reduction. The inner interpolation varies path length at a rational sample with no accidental product equalities; the outer interpolation recovers the target value at $x_0 = 1/2$. Put $q = 2^d$ and

$$P_m = \binom{m+q-1}{q-1}.$$

The formal products to consider are

$$p_{\boldsymbol{\nu}}(x) = \prod_{\epsilon\in\{0,1\}^d} \Lambda_\epsilon(x)^{\nu_\epsilon}, \qquad \nu_\epsilon \in \mathbb{Z}_{\geq 0}, \quad \sum_\epsilon \nu_\epsilon = m.$$

There are at most P_m such products; some may be identical as functions of x. Let K be a fixed real number field containing all the constant parameters, put $\mathcal{G}_+ = \{g \in \mathcal{G} : \theta_g > 1\}$, and set $b = |\mathcal{G}_+|$. All the products lie in

$$K(x, \sqrt{a_g^2 + w_g^2 x^2} : g \in \mathcal{G}_+),$$

an extension of degree at most 2^b over $K(x)$. A nonzero difference $p_{\boldsymbol{\nu}} - p_{\boldsymbol{\nu}'}$ is integral over $K[x]$. Its field norm down to $K(x)$ is therefore a nonzero polynomial. Its degree is at most $D_m = 2^b dm$: each algebraic conjugate over $K(x)$ is a sum of products of at most dm factors growing at most linearly at infinity. Thus two nonidentical product functions can take the same value at at most D_m real parameters. Call a parameter exceptional if any such pair takes equal values there. There are at most

$$B_m = \binom{P_m}{2} D_m$$

exceptional parameters altogether.

Take a rational sample set $\mathcal{X}_m \subset (0,1)$ with

$$|\mathcal{X}_m| = B_m + n_g m + 1, \qquad \mathcal{X}_m = \left\{ \frac{j}{B_m + n_g m + 2} : 1 \le j \le B_m + n_g m + 1 \right\}.$$

Each sample has bit length $O(\log m)$; at least $n_g m + 1$ are nonexceptional. Compute

$$c(x) = |\{p_{\boldsymbol{\nu}}(x) : \textstyle\sum_\epsilon \nu_\epsilon = m\}|, \qquad \mathcal{X}_m^* = \{x \in \mathcal{X}_m : c(x) = \max_{y \in \mathcal{X}_m} c(y)\}.$$

Identical functions agree at every sample. At a nonexceptional sample $c(x)$ equals the number of distinct product functions; at an exceptional sample it is strictly smaller. Thus precisely the nonexceptional samples are retained, without computing an infinite multiplicative-relation lattice.

For a retained x, replace every input edge by an ℓ-edge path of $F_t(x)$. The spectral expansion is

$$Z_{F_t(x)^\ell}(G) = \sum_\rho A_\rho(x) \rho^\ell,$$

where ρ ranges over the $c(x) \le P_m$ distinct positive values of $p_{\boldsymbol{\nu}}(x)$. Each ρ is a product of m source eigenvalues; $A_\rho(x)$ is the sum of the spectral contraction coefficients with that product. These coefficients need not be nonnegative. Vandermonde interpolation at $\ell = 1, \ldots, c(x)$ recovers all $A_\rho(x)$.

For any representative $\boldsymbol{\nu}$ with $p_{\boldsymbol{\nu}}(x) = \rho$, define

$$k_g(\rho) = \sum_{\epsilon \in \{0,1\}^d} \nu_\epsilon \sum_{r \in I_g} \epsilon_r, \qquad R_\rho(x) = (\lambda_g^+(x))^{n_g m} \delta_g(x)^{k_g(\rho)}.$$

Since x is nonexceptional, equality of two sampled products implies identity of their product functions. The previously proved invariance of k_g for the chosen class g therefore makes both $k_g(\rho)$ and $R_\rho(x)$ independent of the representative.

The matrices $F_t(x)$ and $T_g(x)$ have the same tensor eigenbasis. Expand each edge matrix in the associated rank-one spectral projectors, and then sum over vertex colors. For each assignment of eigenvalue indices to edges, the contraction of projector entries is the same in both expansions; only the eigenvalue product changes. By (5.5), the target product in the group indexed by ρ is $R_\rho(x)$. Thus

$$v_x = Z_{T_g(x)}(G) = \sum_\rho A_\rho(x) R_\rho(x).$$

Keeping the recovered $A_\rho(x)$, computing $R_\rho(x)$, and summing the displayed expression gives v_x.

Since each of the m edges contributes degree at most n_g,

$$\deg_x Z_{T_g(x)}(G) \le n_g m, \qquad |\mathcal{X}_m^*| \ge n_g m + 1.$$

Polynomial interpolation from the retained v_x's therefore recovers $Z_{T_g(x_0)}(G)$ at the fixed $x_0 = 1/2$. At each sample, take

$$K_x = K\big(\sqrt{a_g^2 + w_g^2 x^2} : g \in \mathcal{G}_+\big).$$

Although these extensions K_x/K have uniformly bounded degree, their compositum need not have bounded degree. No such compositum is used: $T_g(x)$ has entries in K, hence $v_x \in K$. After each inner interpolation, express its value in K using

$$v_x = \frac{\mathrm{Tr}_{K_x/K}(v_x)}{[K_x : K]}.$$

The trace is fixed-dimensional linear algebra in the sample field. Perform the outer interpolation entirely in K.

All counts of samples and powers are polynomial in m, with exponents depending only on N. The sample eigenvalues have bounded degree and logarithmic height $O(\log m)$. Products of m eigenvalues of $F_t(x)$ have logarithmic height $O(m \log m)$. The Vandermonde systems have polynomial dimension and use polynomial powers, so their determinants and solutions have polynomial logarithmic height. Exact algebraic comparison and arithmetic therefore take polynomial bit complexity. Joint simulation of the finitely many operation layers preserves this bound.

Finally we remove the known contribution of the identity factors and recover one copy of the retained Boolean factor. Let $c(G)$ count the connected components of G, including isolated vertices. Since $Z_{I_2}(G) = 2^{c(G)}$,

$$Z_{T_g(x_0)}(G) = 2^{(d-n_g)c(G)} Z_{B_g(x_0)}(G)^{n_g}.$$

Consequently,

$$Z_{B_g(x_0)}(G) = \left(2^{-(d-n_g)c(G)} Z_{T_g(x_0)}(G)\right)^{1/n_g} > 0,$$

using the unique positive root. Its degree n_g is fixed; exact algebraic root recovery returns the value in the fixed parameter field. Finally,

$$B_g(x_0) > 0, \qquad \det B_g(x_0) > 0, \qquad \theta_g^t \neq \theta_g^{-t},$$

so Proposition 2.2(i) makes Pl-GH($B_g(x_0)$) #P-hard. The construction proves

$$\text{Pl-GH}(B_g(1/2)) \leq_T \text{Pl-GH}(N),$$

which completes the hardness direction. □

6 The positive definite dichotomy

The maximum-support argument leaves Cartesian factors of any size. We first isolate a Potts factor when a clique has size at least three; the remaining factors are Boolean and fall under Theorem 5.1.

Theorem 6.1 (Connected nonnegative PD matrices)**.** *Let M be a fixed nonnegative symmetric algebraic PD matrix with connected support. The problem* Pl-GH(M) *is polynomial-time computable if, up to one simultaneous row and column permutation,*

$$M = \gamma \bigotimes_{r=1}^{d} \begin{pmatrix} a_r & b_r \\ b_r & a_r \end{pmatrix}, \qquad \gamma > 0, \quad a_r > b_r > 0. \tag{6.1}$$

The empty tensor product covers size one. For every other such matrix, Pl-GH(M) *is* #P*-hard.*

Proof. Take the fixed family $\mathcal{A}$ generated from M by the operations of Section 4. It contains M, is closed under those operations, and every finite subfamily is jointly available by composing its finitely many constructions. Choose a logarithmic support of maximum edge count among the nonnegative PD matrices in $\mathcal{A}$ with connected support. Theorem 4.7 gives $R = K_{s_1} \square \cdots \square K_{s_d}$. Its available distance kernels factor as

$$E_R(x) = \bigotimes_r W_{s_r}(x), \qquad W_s(x) = (1-x)I_s + xJ_s, \qquad x \in \mathbb{Q}_{>0}.$$

Availability is uniform in the bit length of x, by Lemma 4.2. Discard trivial Cartesian factors, so each $s_r \geq 2$. For each occurring size k, put $n_k = |\{r : s_r = k\}|$. Suppose a size $s \geq 3$ occurs. We will apply Lemma 3.10 to simulate, after reordering coordinates,

$$H_s = (I_s + J_s)^{\otimes n_s} \otimes J_\ell, \qquad \ell = \prod_{r : s_r \neq s} s_r,$$

where an empty product is 1. The size-s coordinates retain a Potts factor, while J_ℓ contributes only a known counting factor. In the original coordinates,

$$(H_s)_{uv} = 2^{\sum_{r:s_r=s} \mathbf{1}_{\{u_r=v_r\}}}.$$

Each selected coordinate contributes 2 on its diagonal and 1 off its diagonal; every other coordinate contributes 1. In particular, $(H_s)_{uu} = 2^{n_s}$. For m positions (u_a, v_a), the target product is 2^{d_s}, where

$$d_s = \sum_{a=1}^{m} \sum_{r:s_r=s} \mathbf{1}_{\{(u_a)_r=(v_a)_r\}}.$$

We therefore need a source family whose product identities preserve d_s.

The entries of $E_R(t)$ record only distance and do not distinguish which coordinate sizes contribute to it. To make those sizes appear, join two $E_R(t)$-edges in series, with unit weight on the middle vertex. The resulting matrix family is the matrix product

$$F(t) = E_R(t)E_R(t) = E_R(t)^2 = \bigotimes_{r=1}^{d} W_{s_r}(t)^2,$$
$$F_{uv}(t) = \sum_{z \in V(R)} E_R(t)_{uz} E_R(t)_{zv}.$$

For a factor of size k, summing over its middle color gives

$$(W_k(t)^2)_{ii} = 1 + (k-1)t^2 =: p_k(t),$$
$$(W_k(t)^2)_{ij} = t + t + (k-2)t^2 =: q_k(t) \qquad (i \neq j).$$

In the first line the middle color i contributes 1, and the other $k-1$ colors contribute t^2 each. In the second line the middle colors i, j contribute t each, and the other $k-2$ colors contribute t^2 each. These sums introduce the dependence on k.

Apply Lemma 3.10 with the fixed source M, the family $F(t)$ above, target $B = H_s$, and a fixed number field K containing the entries of M. Its four conditions hold as follows.

(i) The family satisfies $F(t) \in \mathrm{Sym}_q(\mathbb{Q}[t])$ and $\deg_t F_{uv} \leq 2d$, so it has type (P).

(ii) Lemma 4.2 uniformly simulates $E_R(n)$ from the fixed source for every integer $n \geq 1$. Replacing each marked edge by the two-edge path above then simulates $F(n)$. The cost is polynomial in $|\Omega| + \mathrm{bit}(n)$, and hence in $|\Omega| + n$, as required.

(iii) For $n \geq 1$, every entry $n^{d_R(u,v)}$ is positive. Thus each $F_{uv}(n)$ is a sum of positive products, and $n_0 = 1$ suffices.

(iv) For each occurring size k, define

$$d_k = \sum_{a=1}^{m} \sum_{r:s_r=k} \mathbf{1}_{\{(u_a)_r=(v_a)_r\}}.$$

Tensor factorization gives

$$\prod_{a=1}^{m} F_{u_a v_a}(t) = \prod_k p_k(t)^{d_k} q_k(t)^{mn_k - d_k}.$$

Since $k \geq 2$, the quadratic $p_k(t) = 1 + (k-1)t^2$ has no real root and is irreducible in $\mathbb{R}[t]$. Distinct k give nonassociate polynomials: their constant terms are all 1, but their quadratic coefficients differ. Furthermore,

$$q_2(t) = 2t, \qquad q_k(t) = t(2 + (k-2)t) \quad (k > 2)$$

has only real linear factors. No q_k therefore contributes a factor p_s. Unique factorization shows that an identity of source products preserves every exponent d_k, and in particular d_s. The target products 2^{d_s} are consequently equal.

The lemma now gives joint simulation of H_s from M, and hence Pl-GH$(H_s) \leq_T$ Pl-GH(M). To make the final replacement explicit, let $m = |E(G)| \geq 1$ and choose the separating integer n_m from the lemma's proof. Grouping color assignments by the source entry product λ gives

$$Z_{F(n_m)^{\circ h}}(G) = \sum_\lambda A_\lambda \lambda^h, \qquad Z_{H_s}(G) = \sum_\lambda A_\lambda 2^{d_s(\lambda)}.$$

The choice of n_m ensures that equal sampled products come from identical product functions, so the preceding factorization makes $d_s(\lambda)$ well-defined. The lemma recovers A_λ by Vandermonde interpolation using parallel copies, retains those coefficients, and substitutes $2^{d_s(\lambda)}$ in the final sum. If $m = 0$, the value is directly $q^{|V(G)|}$.

Finally, tensor factorization of partition functions and $Z_{J_\ell}(G) = \ell^{|V(G)|}$ give

$$Z_{H_s}(G) = \ell^{|V(G)|} Z_{I_s+J_s}(G)^{n_s},$$
$$Z_{I_s+J_s}(G) = \left(\frac{Z_{H_s}(G)}{\ell^{|V(G)|}}\right)^{1/n_s} > 0.$$

The Potts matrix has strictly positive entries, so the unique positive root is the required value. Since n_s is a fixed positive integer, exact algebraic root recovery and the output-field conversion of Remark 3.4 have polynomial bit cost. Thus

$$\text{Pl-GH}(I_s + J_s) \leq_T \text{Pl-GH}(H_s) \leq_T \text{Pl-GH}(M),$$

and Proposition 2.2(ii) gives hardness for $s \geq 3$.

If this reduction does not give hardness, the maximum support is Q_d. Proposition 4.8 now factors the original M, in the same coordinates, into strictly positive PD algebraic 2×2 matrices. If any factor has unequal diagonal entries, Theorem 5.1 applies. Otherwise the factorization is (6.1). Conversely, a tensor product satisfies $Z_{\otimes_r C_r}(G) = \prod_r Z_{C_r}(G)$; each 2×2 factor with equal diagonal entries is computable by the planar zero-field Ising algorithm. The scalar contributes the known factor $\gamma^{|E(G)|}$. □

Corollary 6.2 (Common coordinates when the source is not hard)**.** *Fix a family $\mathcal{A}$ of symmetric matrices on q colours, with the same-domain, joint availability and closure conditions stated at the start of Section 4. Suppose the source counting problem is not #P-hard and $\mathcal{A}$ contains a connected nonnegative PD matrix. Then $q = 2^d$, and there is one identification of the colour set with $\{0,1\}^d$ in which every nonnegative PD matrix in $\mathcal{A}$ is a positive scalar multiple of a tensor product of 2×2 matrices, with the empty product allowed. Every connected such matrix has factors with equal diagonal entries. Moreover, every strictly positive symmetric $N \in \mathcal{A}$, including singular or indefinite N, has in these same coordinates*

$$N = c\,T(\boldsymbol{\rho}), \qquad c > 0, \quad \rho_r > 0.$$

Proof. Choose the maximum among the nonnegative PD matrices in $\mathcal{A}$ with connected support. The reduction to a Potts problem in the preceding proof rules out every Cartesian factor of size at least three. Apply Proposition 4.8 to the resulting fixed cube. For a connected nonnegative PD matrix in $\mathcal{A}$, a factor with unequal diagonal entries would contradict Theorem 5.1 and the assumption that the source problem is not #P-hard. For the last assertion, the available distance kernel gives

$$H(x) = N \circ E_{Q_d}(x) \longrightarrow \mathrm{diag}(N_{ii}) \succ 0 \quad (x \downarrow 0).$$

For a sufficiently small fixed rational $x > 0$, $H(x) > 0$ and $H(x) \succ 0$. The preceding assertions give $H(x) = cT(\boldsymbol{\eta})$, with $0 < \eta_r < 1$, in the fixed coordinates. Comparing entries yields

$$N = cT(\eta_1/x, \dots, \eta_d/x).$$

This establishes a matrix identity from the simulated $H(x)$. □

7 The zero–one dichotomy

The tensor Gram construction of Lemma 3.6 turns a zero–one matrix with distinct nonzero rows into a PD matrix. Theorem 6.1 then either gives hardness or forces integer entries into a tensor form incompatible with their row sums.

Theorem 7.1 (Zero–one matrices)**.** *Let $A \in \mathrm{Sym}_q(\{0,1\})$ be fixed. If every support component is a fully looped complete graph, a loopless complete bipartite graph, or a loopless isolated vertex, then* Pl-GH$(A) \in$ FP*. Otherwise it is* #P*-complete. With any fixed positive algebraic vertex weights, the same structural criterion gives membership in* FP *or* #P*-hardness.*

Proof. Let C be the vertex set of a connected component of Γ_A. By Lemma 3.5, hardness of the principal submatrix problem Pl-GH$(A[C])$ implies hardness of Pl-GH(A). We therefore restrict to $A[C]$ and again call this connected matrix A. Zero rows are isolated vertices of Γ_A and are elementary to handle. By Corollary 3.8, it is enough to prove hardness after identical rows have been merged. Call the resulting matrix again A, and write q for its size. Its support remains connected, although the support of A^2 need not be connected. Replicating identical rows and columns in any of the three listed structures preserves that class.

Its nonzero rows are pairwise nonproportional. For $p = \max(1, q-1)$, Lemma 3.6 shows that

$$Q = (A^2)^{\circ p} \succ 0.$$

This is a planar gadget: first replace an edge by a two-edge path, and then take p parallel copies. Let $\mathcal{D}$ be the collection of vertex sets of the connected components of Γ_{A^2}. Since Q and A^2 have the same zero entries, these are also precisely the components of Γ_Q. Thus, after the same simultaneous row and column permutation,

$$A^2 \cong \bigoplus_{D\in\mathcal{D}} (A^2)[D], \qquad Q \cong \bigoplus_{D\in\mathcal{D}} Q[D], \qquad (A^2)_{D,D'} = Q_{D,D'} = 0 \quad (D \neq D').$$

Here a connected block of A^2 or Q means the corresponding principal submatrix on D. If some $Q[D]$ is outside the tractable class of Theorem 6.1, that block is hard, and the component and gadget reductions finish the proof.

Suppose every connected block $Q[D]$, $D \in \mathcal{D}$, has the form in (6.1). Normalize its Boolean factors by their diagonals and take positive entrywise p-th roots. Since $Q[D] = ((A^2)[D])^{\circ p}$, each connected block of A^2 has, in its own tensor coordinates, the form

$$(A^2)[D] = c \bigotimes_{r=1}^{\ell} \begin{pmatrix} 1 & \rho_r \\ \rho_r & 1 \end{pmatrix}, \qquad c > 0, \quad 0 < \rho_r < 1. \tag{7.1}$$

The 2×2 matrices on the right are tensor factors of the single block $(A^2)[D]$; the direct sum above separates different support components. Each tensor factor has eigenvalues $1 + \rho_r > 0$ and $1 - \rho_r > 0$. Hence $(A^2)[D] \succ 0$, and its diagonal is constant within D; the constant may initially depend on D. The entrywise root is used only in this identity, not as an oracle.

An edge of Γ_{A^2} records a walk of length two in Γ_A. Thus the sets $D \in \mathcal{D}$ are the equivalence classes of vertices joined by even-length walks in Γ_A. In a connected bipartite graph these are exactly the two bipartition sides. In a connected nonbipartite graph, an odd cycle (including a loop) allows the parity of a connecting walk to be changed, so there is only one such class. If Γ_A is nonbipartite, $\mathcal{D} = \{[q]\}$, so the constant diagonal of A^2 covers all colors, and

$$(A^2)_{ii} = \sum_j A_{ij}^2 = \sum_j A_{ij} = k \quad \Longrightarrow \quad A\mathbf{1} = k\mathbf{1}, \qquad k \in \mathbb{Z}_{>0}.$$

Here a loop in the support graph contributes one to the row sum; this is distinct from the input-degree convention used for vertex-weight algorithms. If Γ_A is bipartite, write

$$A = \begin{pmatrix} 0 & V \\ V^{\mathsf{T}} & 0 \end{pmatrix}, \qquad A^2 = VV^{\mathsf{T}} \oplus V^{\mathsf{T}}V.$$

The two connected blocks of A^2 are VV^{T} and $V^{\mathsf{T}}V$, on the two sides of Γ_A. If V has order $n_L \times n_R$, their positive definiteness in (7.1) gives

$$n_L = \operatorname{rank} V = n_R.$$

Their constant diagonals give degrees $k_L, k_R > 0$ on the two sides. Counting cross edges yields $n_L k_L = n_R k_R$, hence $k_L = k_R = k$ and again $A\mathbf{1} = k\mathbf{1}$.

For comparison, merging identical rows in the three listed basic structures gives, respectively,

support before merging	quotient	quotient squared
fully looped complete graph	$[1]$	$[1]$
loopless complete bipartite graph	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	$[1] \oplus [1]$
loopless isolated vertex	$[0]$	$[0]$

Every connected block of the squared quotient is therefore 1×1; the zero-row case was already handled at the start of the proof. Conversely, if every connected block of A^2 is a singleton, then $\langle A_i, A_j \rangle = 0$ for $i \neq j$, so distinct rows of A have disjoint supports. A row with two nonzero entries in columns i, j would, by symmetry, give rows i, j a common nonzero position. Thus every row sum is at most one. Connectedness and the absence of zero rows leave only $[1]$ or $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$. Consequently, if A is outside the three basic classes, there is $D \in \mathcal{D}$ with $|D| \geq 2$.

Choose such a connected block $B = (A^2)[D]$ of A^2. Its tensor form (7.1) has $|D| = 2^\ell \geq 2$, hence $\ell \geq 1$. Since

$$A^2\mathbf{1} = k^2\mathbf{1}, \qquad (A^2)_{D,[q]\setminus D} = 0,$$

the integer matrix B has diagonal k and row sum k^2. Thus (7.1) becomes

$$B = k \bigotimes_{r=1}^{\ell} \begin{pmatrix} 1 & \rho_r \\ \rho_r & 1 \end{pmatrix}, \qquad \ell \geq 1, \quad 0 < \rho_r < 1.$$

An entry differing in only coordinate r shows that $\rho_r = s_r/t_r$ is rational, with $\gcd(s_r, t_r) = 1$ and $1 \leq s_r < t_r$. For every subset $S \subseteq [\ell]$, integrality gives

$$k \prod_{r \in S} \frac{s_r}{t_r} \in \mathbb{Z}.$$

For a prime π and $n \in \mathbb{Z}_{>0}$, define

$$v_\pi(n) := \max\{a \in \mathbb{Z}_{\geq 0} : \pi^a \mid n\}, \qquad n = \pi^{v_\pi(n)} u, \quad u \in \mathbb{Z}_{>0}, \quad \pi \nmid u.$$

Set $S_\pi = \{r \in [\ell] : \pi \mid t_r\}$. Coprimality gives $\pi \nmid s_r$ for $r \in S_\pi$, while the integer-entry condition gives

$$\prod_{r \in S_\pi} t_r \;\mid\; k \prod_{r \in S_\pi} s_r.$$

Since prime exponents add under multiplication,

$$\sum_{r \in S_\pi} v_\pi(t_r) \leq v_\pi(k) + \sum_{r \in S_\pi} v_\pi(s_r) = v_\pi(k),$$
$$v_\pi\left(\prod_{r=1}^{\ell} t_r\right) = \sum_{r=1}^{\ell} v_\pi(t_r) = \sum_{r \in S_\pi} v_\pi(t_r) \leq v_\pi(k).$$

The second equality uses $v_\pi(t_r) = 0$ for $r \notin S_\pi$. Selecting S_π separately for each prime prevents cancellation with numerators from other coordinates. The inequality for every prime therefore implies

$$\prod_{r=1}^{\ell} t_r \mid k. \tag{7.2}$$

Set $\tau = \prod_r t_r$, so $k \geq \tau$. The row sum gives a conflicting inequality:

$$k = \prod_r \left(1 + \frac{s_r}{t_r}\right), \qquad k\tau = \prod_r (t_r + s_r) < \prod_r t_r^2 = \tau^2.$$

For each r, $t_r \geq 2$ and $s_r \leq t_r - 1$ give $t_r + s_r \leq 2t_r - 1 < t_r^2$. Since $\ell \geq 1$, multiplying these inequalities is strict, and yields $k < \tau$, contradicting (7.2). If $\ell = 0$, both products are 1; the nonzero singleton basic cases instead have $k = \tau = 1$ and produce no contradiction. Thus, for a zero–one A outside the three basic classes, the assumption that every connected block $Q[D]$ has the form (6.1) is impossible. A hard block of Q reduces to the corresponding block of A^2 by replacing each edge with p parallel edges. Thus at least one connected block of A^2 is hard, and its problem reduces to that for A by the component and two-edge-path reductions.

The algorithms for the three listed structures count the allowed colors independently, with the two orientations of each connected bipartite input considered separately. They are also explicit in Proposition 2.6. Positive vertex weights are summed within twin classes. Corollary 3.8 then transfers the unweighted hardness of a quotient outside those classes to every such weighting. For a component in a listed class, the same direct algorithms replace cardinalities by weight sums.

Finally, unweighted zero–one homomorphism counting belongs to #P: a nondeterministic machine guesses a color for every input vertex and checks all edges. This proves completeness in the remaining cases. General algebraic weighted answers need not be integers, so the weighted conclusion is stated as hardness rather than membership in #P. $\square$

8 Strictly positive matrices and proportional rows

This section removes the positive-definiteness hypothesis. Throughout, matrices and auxiliary weights are fixed, and all reductions are polynomial-time Turing reductions. We initially work over the real algebraic numbers. We use $W(\rho)$ and $T(\boldsymbol{\rho})$ from (1.1) and the weighted counting notation of Section 2.1. Positive unary powers from Lemma 3.1 normalize diagonal entries; finite moments then recover the multiplicities lost when equal normalized rows are merged.

Theorem 8.1 (Strictly positive matrices)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{>0})$ have fixed algebraic entries. If, after a simultaneous permutation of rows and columns,*

$$M = (aa^{\mathsf{T}}) \otimes T(\boldsymbol{\rho}), \qquad a \in \mathbb{R}^k_{>0}, \quad \rho_j > 0, \tag{8.1}$$

then Pl-GH(M) *is in* FP*. Otherwise it is* #P*-hard. The dimensions in the displayed form satisfy $q = k2^d$. Factors with $\rho_j = 1$ may be absorbed into the rank-one factor.*

We use two established operations, retaining all other fixed constraint types in each mixed instance. For an available positive conditional unary u and every fixed $z \in \mathbb{Q}$, Lemma 3.1 supplies u^z. Attaching it at both ends of each marked C-edge gives

$$C_{ij} \longmapsto u_i^z C_{ij} u_j^z.$$

This uses one unary type at the incident half-edges, not a new type for each vertex degree. An input loop has two such occurrences.

For a symmetric matrix with nonzero pairwise nonproportional rows, Lemma 3.7 supplies, under positive background weights $D = \mathrm{diag}(w)$,

$$D^{-1} \xRightarrow{\text{Lemma 3.1}} D^r, \qquad u_r(i) = (D^r)_{ii} = w_i^r \quad (r \in \mathbb{Q} \text{ fixed}).$$

Each u_r is implemented by a D^r-loop at its vertex. Adding one D^{-1}-loop at each vertex gives $w_i(D^{-1})_{ii} = 1$, hence unit background weights while retaining the available unaries. Zero rows must first be removed, with isolated input vertices evaluated by their prescribed weight sum. The strictly positive matrices used below have no zero rows.

Lemma 8.2 (Unit diagonal and distinct rows)**.** *Let $B \in \mathrm{Sym}_q(\mathbb{R}_{>0})$ have algebraic entries and $B_{ii} = 1$, and suppose that its rows are distinct. Either* Pl-GH(B) *is* #P*-hard, or, after a simultaneous color permutation,*

$$B = T(\boldsymbol{\rho}), \qquad \rho_j > 0, \quad \rho_j \neq 1.$$

In the latter case B is nonsingular.

Proof. If $B_i = cB_j$, positivity gives $c > 0$, while symmetry and unit diagonals give $1 = cB_{ij}$ and $B_{ij} = c$. Thus $c^2 = 1$ and $c = 1$, contrary to distinctness. With $p = \max(1, q-1)$, Lemma 3.6 therefore gives the actual planar two-step/parallel signature

$$Q = (B^2)^{\circ p} \succ 0.$$

Fix the family $\mathcal{A}$ generated from B by the operations of Section 4. Every finite subfamily is jointly available: its finitely many constructions can be composed while retaining the other constraint labels. The family is closed under those operations and contains both B and Q. Take a matrix whose logarithmic support has the maximum number of edges among the connected nonnegative positive-definite members of $\mathcal{A}$. This collection is nonempty because it contains Q. The argument for a support of maximum size in Theorem 6.1 and Corollary 6.2 apply to this fixed family: their source need not itself be positive definite. Unless they already give a hardness reduction, this support is a Boolean cube Q_d, and its distance kernel

$$E(x)_{s,t} = x^{d_H(s,t)} = T(x, \ldots, x)_{s,t}$$

belongs to $\mathcal{A}$ in those same color coordinates. Since $H(x) = B \circ E(x)$ satisfies $H(0) = I$, continuity gives

$$H(x) > 0, \qquad H(x) \succ 0 \quad \text{for all sufficiently small } x \in \mathbb{Q}_{>0}.$$

Each such $H(x)$ belongs to $\mathcal{A}$. Proposition 4.8 expresses $H(x)$ as a tensor product of positive-definite Boolean-domain matrices in these same coordinates. The hardness alternative in Theorem 6.1 excludes any nonconstant diagonal. Hence each Boolean factor has equal diagonal entries. Normalizing the factors and using $H(x)_{ii} = 1$ gives $H(x) = T(\boldsymbol{\eta})$, with $0 < \eta_j < 1$. Consequently

$$B_{s,t} = \prod_{j:s_j \neq t_j} \frac{\eta_j}{x}.$$

Entrywise comparison gives $\rho_j = \eta_j/x > 0$, possibly greater than 1. If any $\rho_j = 1$, the two rows differing only in coordinate j would coincide. All the factors are therefore nonsingular. $\square$

Lemma 8.3 (Vertex weights on Boolean tensor products)**.** *Let $B = T(\boldsymbol{\rho})$ be as in Lemma 8.2, and let w be a fixed positive algebraic vector. If w is constant,* Pl-GH(B, w) *is in* FP*. If w is nonconstant, it is* #P*-hard.*

Proof. Put $D = \mathrm{diag}(w)$. The rows of B are nonzero and pairwise nonproportional. Lemma 3.7 supplies unit background weights, jointly with the unary $f(i) = \sqrt{w_i}$. Attaching f at a vertex means identifying its unique marked vertex with that vertex; color i then contributes the factor $f(i)$. This is a jointly simulated constraint. Equivalently, a $D^{1/2}$-loop contributes $(D^{1/2})_{ii} = f(i)$.

Take a B-edge with endpoints u, v of colors i, j, and attach one copy of f at each endpoint. Its external vertices remain u, v, and its signature is

$$\widetilde{B}_{ij} = f(i) B_{ij} f(j) = \sqrt{w_i} B_{ij} \sqrt{w_j}, \qquad \widetilde{B} = D^{1/2} B D^{1/2}.$$

Identify the right endpoint of one such decorated edge with the left endpoint of another, forming a middle vertex z. Both copies of f at z are retained: if its color is k, their product is w_k. The external endpoints retain one copy each. Summing over all middle colors, with background weight 1, gives

$$\begin{aligned} S_{ij} &= \sum_k (\sqrt{w_i} B_{ik} \sqrt{w_k})(\sqrt{w_k} B_{kj} \sqrt{w_j}) \\ &= \sqrt{w_i} \left(\sum_k B_{ik} w_k B_{kj} \right) \sqrt{w_j}, \\ S &= \widetilde{B}^2 = D^{1/2} B D B D^{1/2}. \end{aligned}$$

The middle factor D comes from the two explicit unary constraints. Since B is nonsingular and symmetric, $\widetilde{B}$ is also nonsingular and symmetric; hence $S \succ 0$. Its entries are positive. All operations use a fixed number of constraint types.

The matrix

$$Q = B^{\circ 2} = T(\rho_1^2, \ldots, \rho_d^2) > 0, \qquad Q\mathbf{1} = R\mathbf{1}, \quad R = \prod_j (1 + \rho_j^2) > 0$$

satisfies $S_{ii} = w_i (Qw)_i$. If $W = \max_i w_i > \min_i w_i = m$, choose a, b attaining these extrema. Positivity gives

$$S_{aa} = W(Qw)_a > WmR, \qquad S_{bb} = m(Qw)_b < mWR.$$

Thus S has a nonconstant diagonal and is hard by Theorem 6.1; its hardness transfers to the weighted source. If $w = c\mathbf{1}$ with $c > 0$, separation of the tensor color coordinates gives

$$Z_{B,w}(G) = c^{|V(G)|} Z_B(G) = c^{|V(G)|} \prod_{j=1}^{d} Z_{W(\rho_j)}(G).$$

Every $W(\rho_j)$ is strictly positive, has equal diagonal entries, and is nonsingular since $\rho_j \neq 1$. Proposition 2.2(i) and the zero-field algorithm in Section 2.8 evaluate each factor in polynomial time, including when $\rho_j > 1$. The number d is fixed, and $d = 0$ uses the empty product 1. The fixed-field arithmetic of Section 2.5 computes the products and $c^{|V(G)|}$ in polynomial bit complexity, proving membership in FP. $\square$

Proof of Theorem 8.1. Write $d_i = M_{ii}$ and define

$$C_{ij} = \frac{M_{ij}}{\sqrt{d_i d_j}}.$$

A source self-loop realizes $d_i > 0$. The endpoint replacement $M_{ij} \mapsto d_i^{-1/2} M_{ij} d_j^{-1/2}$ therefore supplies C, jointly with every fixed unary power d_i^t. The matrix C has diagonal 1. Merge its equal-row classes $I_1, \dots, I_r$ to obtain a symmetric positive matrix B with distinct rows. For each fixed integer $t \geq 0$, the source M simulates

$$\text{Pl-GH}(B, w(t)), \qquad w_s(t) = \sum_{i \in I_s} d_i^t. \tag{8.2}$$

This is an exact summation over the colors in each class, so the quotient weights cannot be discarded. If $r = 1$, then $C = J$ and M has rank one, which is already of the required form.

For $r > 1$, vertex-weight removal transfers hardness of the unweighted quotient B to (8.2). Lemma 8.2 therefore gives an invertible matrix $B = T(\boldsymbol{\rho})$ unless M is hard. Lemma 8.3 then shows that every vector $w(t)$ must be constant unless M is hard.

Only finitely many moments are needed. Let $\delta_1, \dots, \delta_\ell$ be the distinct values among the d_i, and let n_{sj} be the multiplicity of δ_j in class I_s. Equality of $w_s(t)$ and $w_{s'}(t)$ for $0 \leq t < \ell$ gives

$$\sum_{j=1}^{\ell} (n_{sj} - n_{s'j}) \delta_j^t = 0 \quad (0 \leq t < \ell).$$

The coefficient matrix $(\delta_j^t)_{0 \leq t < \ell,\, 1 \leq j \leq \ell}$ has determinant $\prod_{i<j}(\delta_j - \delta_i) \neq 0$. Hence $n_{sj} = n_{s'j}$ for every s, s', j: all classes have the same multiset of diagonal values. If not, one of these ℓ fixed moments yields a hardness reduction. Otherwise arrange the common multiset as $d_1, \dots, d_k$ and put $a_j = \sqrt{d_j}$. Ordering colors within every class in this way gives

$$M_{(s,j),(t,l)} = B_{s,t} a_j a_l,$$

which is (8.1).

Conversely, for the displayed form and every planar input G,

$$Z_M(G) = \left(\prod_{v \in V(G)} \sum_{i=1}^{k} a_i^{\deg_G(v)} \right) \prod_{j=1}^{d} Z_{W(\rho_j)}(G).$$

Each loop contributes twice to $\deg_G(v)$; an isolated vertex therefore contributes k to the first product. The remaining factors use the zero-field planar Ising algorithm, as detailed in Proposition 2.6. $\square$

9 Positive rectangular blocks

For a fixed matrix $V \in \mathrm{Mat}_{p,q}(\mathbb{R}_{>0})$ with algebraic entries, set

$$A_V = \begin{pmatrix} 0 & V \\ V^\mathsf{T} & 0 \end{pmatrix}.$$

Its row and column colour sets are $X = [p]$ and $Y = [q]$, regarded as disjoint labelled copies. The two Gram matrices first determine cube coordinates on the two sides. The Fourier tools of Section 2.6 and a planar gadget with parallel edges then force a compatible pairing of those coordinates.

Theorem 9.1 (Positive rectangular blocks). *The problem* Pl-GH(A_V) *is in* FP *if, after independent row and column permutations,*

$$V = (ab^{\mathsf{T}}) \otimes T(\boldsymbol{\rho}), \qquad a \in \mathbb{R}^k_{>0}, \quad b \in \mathbb{R}^l_{>0}, \quad \rho_j > 0. \tag{9.1}$$

It is #P*-hard otherwise. The dimensions in* (9.1) *satisfy* $p = k2^d$ *and* $q = l2^d$.

9.1 Prescribed bipartitions and planar reductions

For a bipartite input with a specified bipartition $L \cup R$, consider the version that counts assignments with $L \to X$ and $R \to Y$. Both the bipartition and this correspondence with X, Y are prescribed. With fixed positive algebraic weights μ on X and ν on Y, denote this value by

$$Z_{V,\mu,\nu}(G; L, R) = \sum_{\substack{\sigma: L \to X \\ \tau: R \to Y}} \prod_{u \in L} \mu_{\sigma(u)} \prod_{v \in R} \nu_{\tau(v)} \prod_{uv \in E(G),\, u \in L,\, v \in R} V_{\sigma(u),\tau(v)}.$$

For fixed V, μ, ν, evaluating this function on planar inputs with specified sides is polynomial-time Turing equivalent to Pl-GH$(A_V, (\mu, \nu))$, where

$$A_V = \begin{pmatrix} 0 & V \\ V^{\mathsf{T}} & 0 \end{pmatrix} \quad \text{on the color set } X \sqcup Y.$$

The standard problem sums over colors in $X \sqcup Y$ without a prescribed side for an input vertex. Both reductions retain the fixed weights μ, ν; the unweighted case is $\mu = \nu = \mathbf{1}$, for which we abbreviate the prescribed value by $Z_V(G; L, R)$.

For a connected bipartite input with at least one edge,

$$Z_{A_V,(\mu,\nu)}(G) = Z_{V,\mu,\nu}(G; L, R) + Z_{V,\mu,\nu}(G; R, L).$$

Every nonzero assignment has exactly one of these two orientations; the two summands need not agree. This computes the standard problem from the prescribed one. A nonbipartite connected input has value zero. An isolated vertex contributes $\sum_{i \in X} \mu_i + \sum_{j \in Y} \nu_j$ to the standard problem, or the weight sum on its specified side to the prescribed problem. Disconnected inputs are evaluated by multiplying their component values.

For the reverse reduction, choose a root in L of a connected input containing an edge. Lemma 3.5, applied to the color subset X, recovers from the standard oracle the sum with this root's color restricted to X. All nonzero edges of A_V cross X, Y; connectedness therefore forces $L \to X$ and $R \to Y$. The recovered value is exactly $Z_{V,\mu,\nu}(G; L, R)$. The root restriction is supplied by that oracle reduction.

Every auxiliary gadget below has a prescribed part, X or Y, at each terminal. It expands to a bipartite planar graph whose assignment to these parts is fixed by one root per connected component. Thus the one-root reduction suffices without restrictions at arbitrarily many roots. Under unit background weights, the two-edge paths give

$$VV^{\mathsf{T}} \quad (X\text{–}X), \qquad V^{\mathsf{T}}V \quad (Y\text{–}Y).$$

With background vectors μ, ν, these signatures are instead $V \operatorname{diag}(\nu) V^{\mathsf{T}}$ and $V^{\mathsf{T}} \operatorname{diag}(\mu) V$. Replacing every edge of any planar input by either such two-terminal gadget gives a legal instance with prescribed bipartition. The two terminals can be kept on the outer face. Series and parallel composition, the gadget of Lemma 4.4, and rooted attachments preserve this property.

On each part we use the family generated by all the prescribed-domain simulations above, with the closure conditions of Section 4, and maximize logarithmic support within that family. This permits the Gram matrices and later mixed constructions to share the same coordinates. Every finite subfamily is jointly available by composition of its finite constructions. Each oracle query finally expands to a planar bipartite instance with edge matrix A_V and the fixed background weights. Its prescribed orientation is recovered from the corresponding standard oracle by the one-root reduction above.

9.2 Nonproportional rows and columns

Proposition 9.2 (Positive matrices with nonproportional rows and columns)**.** *Let $B \in \mathrm{Mat}_{p,q}(\mathbb{R}_{>0})$ have fixed algebraic entries, with no two distinct rows proportional and no two distinct columns proportional. Either* Pl-GH(A_B) *is* #P*-hard, or B is square of order 2^d and, after independent row and column permutations, is a positive scalar multiple of $T(\boldsymbol{\rho})$ with $0 < \rho_j < 1$.*

Proof. Use unit background weights. Lemma 3.6 supplies positive exponents for the two Gram matrices; a common multiple, using the Schur product theorem for further entrywise powers, gives a fixed integer h such that

$$Q_X = (BB^\mathsf{T})^{\circ h} \succ 0, \qquad Q_Y = (B^\mathsf{T} B)^{\circ h} \succ 0.$$

These gadgets have both terminals in X or both terminals in Y. On each part take a logarithmic support with the maximum number of edges among the entire jointly available collection just described. If a matrix on one part yields the hardness alternative of the preceding sections, its hardness transfers to A_B by the reduction for prescribed bipartitions. Otherwise both supports of maximum size are Boolean cubes, and Proposition 4.8 places every available positive-definite matrix in Boolean tensor form in the corresponding fixed coordinates. Theorem 6.1 forces equal diagonal entries in each Boolean factor. In particular Q_X and Q_Y are positive-definite zero-field tensors. Their entrywise positive hth roots remain positive-definite zero-field tensors: every ratio $0 < \rho < 1$ becomes $\rho^{1/h}$. Thus both BB^T and $B^\mathsf{T} B$ are positive definite. The matrix B has full row and full column rank, so it is square, and the two cube dimensions agree. Write its order as $n = 2^d$.

Set $G_X = BB^\mathsf{T}$ and $G_Y = B^\mathsf{T} B$. Each is a positive scalar times a zero-field tensor. Since

$$W(\rho)\mathbf{1} = (1+\rho)\mathbf{1}, \qquad T(\boldsymbol{\rho})\mathbf{1} = \left(\prod_j (1+\rho_j)\right)\mathbf{1},$$

each Gram matrix has $\mathbf{1} > 0$ as an eigenvector. By the Perron–Frobenius theorem in Section 2.3, its eigenvalue is the largest and its eigenspace is one-dimensional. Moreover, B is now invertible, so

$$B^{-1} G_X B = G_Y.$$

Thus the two largest eigenvalues agree; call their common value $\lambda > 0$. We have

$$G_X\mathbf{1} = G_Y\mathbf{1} = \lambda\mathbf{1},$$
$$G_X(B\mathbf{1}) = B(G_Y\mathbf{1}) = \lambda B\mathbf{1}.$$

Since $\ker(G_X - \lambda I) = \mathrm{span}\{\mathbf{1}\}$ and $B\mathbf{1} > 0$, it follows that $B\mathbf{1} = a\mathbf{1}$ for some $a > 0$. The transpose argument gives $B^\mathsf{T}\mathbf{1} = b\mathbf{1}$, $b > 0$. Summing all entries of the $n \times n$ matrix B yields

$$na = \mathbf{1}^\mathsf{T} B\mathbf{1} = nb, \qquad a = b =: s_0 > 0.$$

Finally, $G_X\mathbf{1} = BB^\mathsf{T}\mathbf{1} = s_0^2\mathbf{1}$, so

$$B\mathbf{1} = B^\mathsf{T}\mathbf{1} = s_0\mathbf{1}, \qquad \lambda = s_0^2.$$

This fixes the constant Fourier mode and the normalization below. The case $d = 0$ is immediate, so assume $d \geq 1$.

For each part use its fixed Boolean coordinates, with $x, y \in \{0,1\}^d$ and Fourier indices $S, T \subseteq [d]$. Its orthogonal Walsh matrix is

$$(H_X)_{x,S} = n^{-1/2}\chi_S(x), \quad (H_Y)_{y,T} = n^{-1/2}\chi_T(y), \qquad \chi_S(x) = (-1)^{\sum_{i\in S} x_i}.$$

Define $C = s_0^{-1}H_X^\mathsf{T} B H_Y$. The Gram tensor on the X side gives

$$CC^\mathsf{T} = \operatorname{diag}(\theta_S^2), \qquad \theta_S = \prod_{i\in S}\theta_i, \quad 0 < \theta_i < 1, \quad \theta_\varnothing = 1. \tag{9.2}$$

The empty-set row and column of C have their sole nonzero entry, equal to 1, at the empty set, because, with $e_\varnothing$ the standard basis vector at Fourier index $\varnothing$, $H_X e_\varnothing = H_Y e_\varnothing = n^{-1/2}\mathbf{1}$ and $B\mathbf{1} = B^\mathsf{T}\mathbf{1} = s_0\mathbf{1}$. Writing $\Theta = \operatorname{diag}(\theta_S)$ gives

$$C = \Theta O, \qquad O = \Theta^{-1}C, \qquad OO^\mathsf{T} = I.$$

The Walsh transform expresses these fixed matrices in an eigenbasis; all simulated constraints remain in the original colour coordinates.

Fourier degree is preserved. For every fixed rational $z \in (0,1)$, the Y-side distance kernel with parameter $(1-z)/(1+z)$, scaled by $((1+z)/2)^d$, supplies

$$K_Y(z) = \bigotimes_{i=1}^{d} \frac{1}{2}\begin{pmatrix} 1+z & 1-z \\ 1-z & 1+z \end{pmatrix}, \qquad 0 < z < 1.$$

In Walsh coordinates,

$$H_Y^\mathsf{T} K_Y(z) H_Y = \operatorname{diag}(z^{|T|}) =: \Delta(z).$$

The matrix $BK_Y(z)B^\mathsf{T}$ is available, positive and positive definite, so in the fixed X coordinates it is a zero-field tensor unless the source is already hard. Consequently

$$s_0^{-2}H_X^\mathsf{T} BK_Y(z)B^\mathsf{T} H_X = C\Delta(z)C^\mathsf{T} = \Theta O\Delta(z)O^\mathsf{T}\Theta$$

is diagonal, and so is $O\Delta(z)O^\mathsf{T}$, since Θ is an invertible diagonal matrix. For a fixed rational $z \in (0,1)$ the eigenvalues belonging to different degrees are distinct. Each row S of O is therefore supported on the columns of one degree, denoted by $h(S)$.

The row supports of the fixed matrix O determine $h(S)$ independently of z. Since $h(\varnothing) = 0$, the eigenvalue at $\varnothing$ of $BK_Y(z)B^\mathsf{T}/s_0^2$ is 1, and the zero-field tensor identity gives

$$\theta_S^2 z^{h(S)} = \prod_{i\in S}(\theta_i^2 z^{h(\{i\})}).$$

Cancel $\theta_S^2 = \prod_{i\in S}\theta_i^2 > 0$ and use $0 < z < 1$ to obtain $h(S) = \sum_{i\in S} h(\{i\})$. The degree-zero eigenspace has dimension one and is already occupied by the empty-set row. Thus $h(\{i\}) \geq 1$, whereas $h([d]) \leq d$. All these singleton degrees are 1, and $h(S) = |S|$. We have proved

$$C_{S,T} = \begin{cases} \theta_S (U_{|S|})_{S,T}, & |S| = |T|, \\ 0, & |S| \neq |T|, \end{cases} \tag{9.3}$$

Here U_j is an orthogonal matrix of order $\binom{d}{j}$, indexed by the j-element subsets of $[d]$. We write $(U_j)_{S,*}$ for its entire row at S, and use the Euclidean norm on these row vectors. No distinctness of the coordinate parameters is used. In particular, this argument does not yet make the matrices U_j diagonal.

The next gadget may give a singular Gram matrix. The final assertion of Corollary 6.2 applies to every available strictly positive symmetric matrix on a part, so this Gram matrix still has a zero-field tensor form in the fixed coordinates.

Parallel squaring imposes a norm equality. Let $K_X(z)$ be the noise matrix given by the formula for $K_Y(z)$ in the fixed X coordinates. For each rational $0 < \varepsilon < 1$, form the available matrix $B_\varepsilon = K_X(\varepsilon)B$ and put

$$F_\varepsilon(x, y) = \frac{n}{s_0}(B_\varepsilon)_{x,y} = \sum_{|S|=|T|} \varepsilon^{|S|}\theta_S(U_{|S|})_{S,T}\chi_S(x)\chi_T(y). \tag{9.4}$$

We also write F_ε for the n-by-n matrix with these entries. In particular $F_0 \equiv 1$. Two parallel copies give $F_\varepsilon^{\circ 2}$. Let its double Fourier coefficients be

$$D_{S,T}(\varepsilon) = \frac{1}{n^2}\sum_{x,y} F_\varepsilon(x, y)^2\chi_S(x)\chi_T(y).$$

Define the Gram matrix

$$L_\varepsilon = \frac{1}{n^2}F_\varepsilon^{\circ 2}(F_\varepsilon^{\circ 2})^\mathsf{T}.$$

It is strictly positive and symmetric, so Corollary 6.2 makes it a zero-field tensor in the fixed X coordinates. Figure 3 shows the planar construction. Before the known scalar normalization, its entry is

$$n^2(L_\varepsilon)_{x,x'} = \sum_{y\in\{0,1\}^d} F_\varepsilon(x, y)^2F_\varepsilon(x', y)^2.$$

Thus parallel edges turn the entrywise product into a Fourier convolution, and series composition turns those coefficients into squared row norms. The tensor form will impose equality in a norm bound, determining the unknown coordinate pairing.

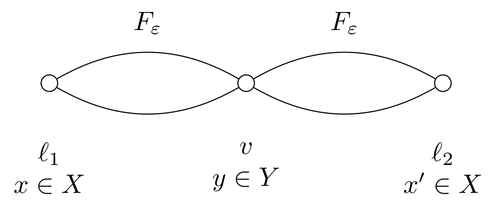

Figure 3: Two parallel copies on each side of one internal vertex. Every edge has the simulated X–Y signature F_ε; reading its endpoints from Y to X uses the transpose. Summing the internal colour gives n^2L_ε.

The normalization above gives exactly

$$H_X^\mathsf{T}F_\varepsilon^{\circ 2}H_Y = nD(\varepsilon), \qquad H_X^\mathsf{T}L_\varepsilon H_X = D(\varepsilon)D(\varepsilon)^\mathsf{T}.$$

Thus this product is diagonal. Its diagonal entries $\mu_S(\varepsilon) = \sum_T D_{S,T}(\varepsilon)^2$ satisfy, for every nonempty $S \subseteq [d]$,

$$\mu_S\mu_\varnothing^{|S|-1} = \prod_{i\in S}\mu_{\{i\}}. \tag{9.5}$$

This identity holds for infinitely many positive rational ε, hence as a polynomial identity. Comparing its coefficients at 0 requires no simulation of the degenerate $\varepsilon = 0$ constraint. Since $\mu_\varnothing(0) = 1$ and

$$D_{\{i\},\{j\}}(\varepsilon) = 2\varepsilon\theta_i(U_1)_{ij} + O(\varepsilon^2),$$

while $D_{\{i\},T}(\varepsilon) = O(\varepsilon^2)$ for $|T| \neq 1$, orthogonality gives

$$\mu_{\{i\}}(\varepsilon) = 4\theta_i^2\varepsilon^2 \sum_j (U_1)_{ij}^2 + O(\varepsilon^3) = 4\theta_i^2\varepsilon^2 + O(\varepsilon^3).$$

Substitution into (9.5) gives, for $|S| = r$,

$$\mu_S(\varepsilon) = 4^r\theta_S^2\varepsilon^{2r} + O(\varepsilon^{2r+1}). \tag{9.6}$$

The degree-two equality. Assume $d \geq 2$ and write $U = U_1$. For two-element sets define

$$P_2(U)_{\{i,k\},\{j,l\}} = U_{ij}U_{kl} + U_{il}U_{kj}.$$

The degree-two coefficient of the row $\{i,k\}$ in D is

$$2\varepsilon^2\theta_i\theta_k\big((U_2)_{\{i,k\},*} + P_2(U)_{\{i,k\},*}\big).$$

It has no support outside degree two. The only other possible column degree at this order is zero, whose coefficient vanishes because $\sum_j U_{ij}U_{kj} = 0$. Equation (9.6) therefore implies

$$\big\|(U_2)_{\{i,k\},*} + P_2(U)_{\{i,k\},*}\big\|^2 = 4. \tag{9.7}$$

For orthogonal unit rows u, v of U, direct expansion gives

$$\begin{aligned}\sum_{j<l}(u_jv_l + u_lv_j)^2 &= \left(\sum_j u_j^2\right)\left(\sum_j v_j^2\right) + \left(\sum_j u_jv_j\right)^2 - 2\sum_j u_j^2v_j^2 \\ &= 1 - 2\sum_j u_j^2v_j^2 \leq 1.\end{aligned}$$

Each row of U_2 has norm 1. For the rows at $\{i,k\}$,

$$2 = \|(U_2)_{\{i,k\},*} + P_2(U)_{\{i,k\},*}\| \leq 1 + \|P_2(U)_{\{i,k\},*}\| \leq 2.$$

Both inequalities are equalities. The preceding expansion forces $U_{ij}U_{kj} = 0$ for every $i \neq k$ and every j. The supports of the d rows of the square orthogonal matrix U are pairwise disjoint and nonempty. Each therefore contains exactly one coordinate, with value 1 or -1. Thus U is a signed permutation. Its action comes from a genuine permutation of the Y coordinates and bit flips of those coordinates. Relabel these colors so that $U_1 = I$. Then $P_2(I) = I$, and equality again forces $U_2 = I$. For $d = 1$, $U_1 = [\pm 1]$ is normalized by a bit flip directly.

Induction on the degree. Suppose $r \geq 3$ and $U_j = I$ for $1 \leq j < r$. Fix $|S| = r$, and let e_S be the standard row vector at S in the space indexed by the r-element subsets of $[d]$. In the square of (9.4), a term contributing to row S has two X-indices I, J with $I \triangle J = S$, where $\triangle$ denotes symmetric difference. Its order is $\varepsilon^{|I|+|J|}$, at least ε^r. Equality holds precisely when I, J are disjoint and their union is S. The two ordered decompositions containing the empty set contribute $2(U_r)_{S,*}$. The other $2^r - 2$ decompositions use only lower-degree identity matrices and contribute $(2^r - 2)e_S$. The coefficient is therefore

$$\theta_S\big(2(U_r)_{S,*} + (2^r - 2)e_S\big),$$

supported in column degree r. Equation (9.6) gives

$$\big\|2(U_r)_{S,*} + (2^r - 2)e_S\big\|^2 = 4^r.$$

Put $a = 2^r - 2 > 0$ and $u = (U_r)_{S,*}$. Since $\|u\| = \|e_S\| = 1$,

$$(2 + a)^2 = \|2u + ae_S\|^2 = 4 + a^2 + 4a\langle u, e_S\rangle \quad \Longrightarrow \quad \langle u, e_S\rangle = 1.$$

Thus $u = e_S$ for every S, proving $U_r = I$ and completing the induction.

We conclude that $C = \operatorname{diag}(\theta_S)$ and hence

$$B = s_0 \bigotimes_{i=1}^{d} \frac{1}{2}\begin{pmatrix} 1 + \theta_i & 1 - \theta_i \\ 1 - \theta_i & 1 + \theta_i \end{pmatrix}.$$

Taking $\rho_i = (1 - \theta_i)/(1 + \theta_i) \in (0, 1)$ and absorbing the positive scalar proves the proposition. □

9.3 Vertex weights and repeated rows and columns

Lemma 9.3 (Vertex weights on the two parts)**.** *Let B have nonproportional distinct rows and columns and the tensor form in Proposition 9.2. Equip its left and right colors with fixed positive algebraic vertex weights μ and ν, respectively. The planar problem with prescribed bipartition is in* FP *if each vector is constant. Otherwise it is* #P*-hard. The two constants need not be equal.*

Proof. Put $D_\mu = \operatorname{diag}(\mu)$ and $D_\nu = \operatorname{diag}(\nu)$. All rows of A_B are nonzero and pairwise nonproportional: within a part this follows from the hypotheses on B; between parts their nonempty supports are disjoint. Lemma 3.7 supplies unit background weights, and Lemma 3.1 supplies their fixed square-root unaries. At the two endpoints of each B-edge these give $\widetilde{B} = D_\mu^{1/2} B D_\nu^{1/2}$. The two ordinary same-side path matrices are

$$S_X = \widetilde{B}\widetilde{B}^{\mathsf{T}} = D_\mu^{1/2} B D_\nu B^{\mathsf{T}} D_\mu^{1/2},$$
$$S_Y = \widetilde{B}^{\mathsf{T}}\widetilde{B} = D_\nu^{1/2} B^{\mathsf{T}} D_\mu B D_\nu^{1/2}.$$

The tensor form makes B, and hence $\widetilde{B}$, square and nonsingular. Thus both Grams are positive and PD. If either has nonconstant diagonal, Theorem 6.1 gives hardness.

It remains to consider constant diagonals. Put $Q = B^{\circ 2}$; it has the same constant row and column sum $R_0 > 0$. The diagonal conditions are

$$\mu_i(Q\nu)_i = c_X, \qquad \nu_j(Q^{\mathsf{T}}\mu)_j = c_Y.$$

Writing their common order as n, the Gram identities give

$$nc_X = \operatorname{Tr} S_X = \operatorname{Tr} S_Y = nc_Y, \qquad c_X = c_Y = c > 0.$$

If μ is nonconstant and $U = \max_i \mu_i$, positivity gives $(Q^{\mathsf{T}}\mu)_j < R_0 U$ for every j, hence $\nu_j > c/(R_0 U)$. At a row with $\mu_i = U$ this implies $\mu_i(Q\nu)_i > U R_0 c/(R_0 U) = c$, a contradiction. Thus $\mu = a\mathbf{1}$ for some $a > 0$, and the column equations give $\nu_j = c/(aR_0)$ for every j. Conversely, if $\mu_i = a$ and $\nu_j = b$ for all colors, they contribute the known factor $a^{|L|}b^{|R|}$ on an input with prescribed bipartition $L \cup R$. □

Proof of Theorem 9.1. For necessity, suppose the source problem is not #P-hard. For the original positive rectangular matrix V, define

$$a_i = \left(\sum_j V_{ij}^2\right)^{1/2}, \qquad b_j = \left(\sum_i V_{ij}^2\right)^{1/2}, \qquad C_{ij} = \frac{V_{ij}}{a_i b_j}.$$

Equivalently, $C = D_a^{-1} V D_b^{-1}$, where $D_a = \operatorname{diag}(a_i)$ and $D_b = \operatorname{diag}(b_j)$. The conditional weight of two parallel edges from a root to one internal vertex is a_i^2 on the left, and b_j^2 on the right. Rational powers of these unary weights simulate a_i^{-1} and b_j^{-1} at edge occurrences, so C is jointly available with fixed powers of these unaries. If $V_i = tV_{i'}$, positivity gives $t > 0$ and $a_i = ta_{i'}$, so $C_i = C_{i'}$. Conversely, $C_i = tC_{i'}$ implies $V_i = (ta_i/a_{i'})V_{i'}$; taking row norms gives $t = 1$. The same argument with column norms gives

$$C_i \propto C_{i'} \iff C_i = C_{i'}, \qquad C_{*,j} \propto C_{*,j'} \iff C_{*,j} = C_{*,j'}.$$

Merge all repeated row classes I_s and repeated column classes J_t of C, producing a positive rectangular matrix B. Deleting repeated coordinates cannot create a new proportionality relation, so B has no proportional distinct rows or columns. The source simulates, for every fixed pair of nonnegative integers m, n, the quotient with prescribed bipartition and vertex weights

$$\mu_s(m) = \sum_{i \in I_s} a_i^{2m}, \qquad \nu_t(n) = \sum_{j \in J_t} b_j^{2n}. \tag{9.8}$$

In particular, exponent zero retains the class cardinalities. Summing the original colours within each class gives these positive background weights on the quotient.

If the unweighted quotient is hard, vertex-weight removal transfers its hardness to each of these weighted quotients and thence to the source. Otherwise Proposition 9.2 expresses B as a scalar multiple of $T(\boldsymbol{\rho})$. Lemma 9.3 shows that every $\mu(m)$ must be constant across the left classes and every $\nu(n)$ constant across the right classes, unless the source is hard. Let $\alpha_1, \ldots, \alpha_u$ be the distinct values of a_i^2, and let n_{sv} count the occurrences of α_v in I_s. For any two left classes,

$$\sum_{v=1}^{u} (n_{sv} - n_{s'v})\alpha_v^m = 0 \quad (0 \le m < u).$$

The Vandermonde determinant is nonzero, so $n_{sv} = n_{s'v}$ for every s, s', v. Positivity of the a_i gives a common multiset of a_i in all left classes. Applying the same argument independently to the distinct b_j^2, with exponents from zero to one less than their number, gives a common multiset of b_j in all right classes.

Write these common positive vectors as $a \in \mathbb{R}_{>0}^k$ and $b \in \mathbb{R}_{>0}^l$, and order each class by its common multiset. The original matrix becomes

$$V_{(s,i),(t,j)} = B_{s,t} a_i b_j.$$

Absorbing the positive scalar of B into a yields (9.1). The Vandermonde argument applies to every fixed domain size.

Conversely, suppose $V = (ab^{\mathsf{T}}) \otimes T(\boldsymbol{\rho})$. Listing each color by its side, amplitude index and Boolean vector gives the simultaneous permutation identity

$$A_V \cong \begin{pmatrix} 0 & ab^{\mathsf{T}} \\ ba^{\mathsf{T}} & 0 \end{pmatrix} \otimes T(\boldsymbol{\rho}).$$

Proposition 2.6 therefore gives Pl-GH(A_V) $\in \mathsf{FP}$. Explicitly, for a specified input bipartition L, R,

$$Z_V(G; L, R) = \left(\prod_{u\in L}\sum_{i=1}^{k} a_i^{\deg_G(u)}\right)\left(\prod_{v\in R}\sum_{j=1}^{l} b_j^{\deg_G(v)}\right)\prod_{h=1}^{d} Z_{W(\rho_h)}(G).$$

The amplitude and Boolean color sums separate. Each factor is computable by the algorithms in Section 2.8; summing the two orientations and handling isolated vertices as in Section 9.1 evaluates the standard problem. This proves the sufficient direction as well. ☐

10 Completion for nonnegative matrices

Support interpolation reduces the necessity direction to the zero–one classification and the two positive-block theorems. For sufficiency, Proposition 2.6 already evaluates the listed forms by tensor products and component sums.

Proof of Theorem 1.1 for algebraic weights. For the nonnegative matrix M, Corollary 3.2 gives

$$P(M) = \mathbf{1}_{M>0}, \qquad \text{Pl-GH}(P(M)) \leq_T \text{Pl-GH}(M).$$

Theorem 7.1 gives hardness unless every connected support component is a fully looped clique, a loopless complete bipartite graph, or a loopless isolated vertex.

A fully looped clique is a strictly positive symmetric block, classified by Theorem 8.1. A complete bipartite component is A_V for a positive rectangular V, and its necessity is Theorem 9.1. A loopless isolated color is the block $[0]$. A support component with vertex set C and hard Pl-GH($M[C]$) gives Pl-GH($M[C]$) $\leq_T$ Pl-GH(M) by Lemma 3.5. On a connected input the one-root restriction keeps every nonzero assignment within C. This proves all hardness assertions.

For the algorithmic direction, every listed block and their direct sum satisfy Proposition 2.6. Formula (2.7) combines their values on connected input components, including isolated vertices and the empty input. Hence Pl-GH(M) $\in \mathsf{FP}$. ☐

11 One fixed positive vertex-weight vector

11.1 Weighted quotients and algebraic reductions

We prove Theorem 1.3, the positive-vertex-weight version of Theorem 1.1. Throughout this section, $M \geq 0$ and $w > 0$ have fixed real algebraic entries. The pair $(R, \bar{w})$ is the quotient of the original matrix in Definition 1.2; equation (1.5) sums weights over its actual identical-row classes K_a. By contrast, the classes I_s (and J_t on a second part) below belong to a normalized matrix C, and their quotient B is an intermediate matrix used to determine the Ising coordinates.

Sections 11.2 and 11.3 prove necessity; Section 11.4 gives the algorithms and completes the proof. We use the following algebraic reductions from Section 3.

(i) *Root restrictions.* For a fixed matrix and positive background vector, Lemma 3.5 extracts a support component. On a connected bipartite input, one root restricted to one color part prescribes the whole orientation, as in Section 9.1. All background weights are retained.

(ii) *Joint replacements.* Lemma 3.1 retains all other fixed edge and unary types while replacing one type. In particular, for a positive unary f and fixed $r \in \mathbb{Q}$,

$$\prod_{v=1}^{m} f(i_v) = \prod_{v=1}^{m} f(j_v) \implies \prod_{v=1}^{m} f(i_v)^r = \prod_{v=1}^{m} f(j_v)^r.$$

Thus f^r is jointly available. Attaching these unaries at edge endpoints multiplies conditional entries and does not change the background weight at a vertex.

(iii) *Weight removal.* If a symmetric matrix B has no zero row and no proportional row pair, Lemma 3.7 supplies unit background weights and retains every previously available constraint. With $D = \mathrm{diag}(u)$, it also supplies the unaries u_i^r for a fixed finite set of rational r. We apply this only after verifying the row conditions for the relevant quotient; these conditions need not hold for M.

(iv) *Support hardness.* Support replacement from Corollary 3.2 retains w, and Theorem 7.1 already includes positive algebraic vertex weights. Consequently

$$\begin{gathered}\text{Pl-GH}(P(M), w) \leq_T \text{Pl-GH}(M, w),\\ P(M) \text{ has a nonbasic component}\\ \implies \quad \text{Pl-GH}(M, w) \text{ is \#P-hard.}\end{gathered}$$

Here $P(M)_{ij} = \mathbf{1}_{\{M_{ij} \neq 0\}}$, and the basic components are fully looped cliques, loopless complete bipartite graphs and loopless isolated vertices.

Every use involves a fixed finite number of algebraic constraint types. Section 2.5 and Remark 3.4 supply the bit bounds and conversion between fixed number fields. The following three consequences identify the unweighted matrices needed after operation (iii).

Lemma 11.1 (Consequences of the unweighted classification)**.** *All matrix entries in this lemma are fixed real algebraic numbers.*

(a) *If $B > 0$ is symmetric, has diagonal 1, and has distinct rows, then either* Pl-GH(B) *is hard or $B \cong T(\boldsymbol{\rho})$ with every $\rho_j \neq 1$.*

(b) *If a positive rectangular B has no proportional distinct rows or columns, then either* Pl-GH($\begin{pmatrix} 0 & B \\ B^\mathsf{T} & 0 \end{pmatrix}$) *is hard or, under independent row and column permutations, $B = cT(\boldsymbol{\rho})$ for $c > 0$ and every $\rho_j \neq 1$.*

(c) *Every strictly positive positive-definite symmetric matrix with a nonconstant diagonal has a hard unweighted planar problem.*

Proof. For (a), the positive form in the unweighted theorem is $aa^\mathsf{T} \otimes T$. Unit diagonal forces every entry of a to be 1, and distinct rows force its length to be 1 and exclude $W(1)$ factors. For (b), the basic bipartite factor must have only one color on each side, since two would give proportional rows or columns; again $W(1)$ is excluded. For (c), positivity gives a single positive support block and full rank forces the rank-one factor to have size one. Such a tensor has constant diagonal. These are also the alternatives in Lemma 8.2, Proposition 9.2, and Theorem 6.1, respectively. Theorem 1.1 is used here only in its already proved algebraic form. □

11.2 Necessity for positive components

This is the weighted counterpart of the normalization and moment argument in Theorem 8.1: the multiplicity of an amplitude is replaced by the sum of its vertex weights. The reductions in Section 11.1 first discard nonbasic supports by hardness and restrict to one remaining component. A zero component is already of the required form. Each other component is strictly positive or has positive complete bipartite cross block.

Let $M > 0$, and define

$$d_i = M_{ii} > 0, \qquad a_i = \sqrt{d_i}, \qquad C_{ij} = \frac{M_{ij}}{a_i a_j}.$$

A loop supplies the conditional unary d_i. Operation (ii) supplies $d_i^{-1/2}$ at both endpoints of an M-edge, hence C, jointly with each fixed unary d_i^t. The background at every original vertex is still w_i.

Let I_s be the identical-row classes of C, and let B be their quotient. It is symmetric, positive, and has diagonal 1. Its distinct rows cannot be proportional: $B_s = cB_t$ gives $1 = cB_{st}$ and $B_{st} = c$, so $c = 1$. Our goal is to show that each amplitude has the same total vertex weight in every class I_s. List the distinct values of d_i as $\delta_1, \ldots, \delta_m$, and put

$$p_{sj} := \sum_{\substack{i \in I_s \\ d_i = \delta_j}} w_i, \qquad \alpha_j := \sqrt{\delta_j}.$$

Empty sums are zero. Attaching d_i^t at each vertex and then summing over its class, as in Corollary 3.8, gives

$$\text{Pl-GH}(B, u(t)) \le_T \text{Pl-GH}(M, w), \qquad u_s(t) = \sum_{i \in I_s} w_i d_i^t = \sum_{j=1}^{m} p_{sj} \delta_j^t. \tag{11.1}$$

This is equation (8.2) with the original weights retained. Only $0 \le t < m$ will be used.

Operation (iii) applies to B because its rows are nonzero and nonproportional. Hence $\text{Pl-GH}(B) \le_T \text{Pl-GH}(B, u(0))$. Lemma 11.1(a) gives hardness or $B \cong T(\boldsymbol{\rho})$ with $\rho_r > 0$, $\rho_r \ne 1$. In the latter case Lemma 8.3, the algebraic positive-core weight criterion, gives

$$u(t) \notin \mathbb{R}_{>0}\mathbf{1} \implies \text{Pl-GH}(B, u(t)) \text{ is \#P-hard.}$$

Its proof uses unit backgrounds and two explicit square-root unaries at the identified middle vertex; the background $u(t)$ is not counted a second time. Thus nonhardness of the source forces $u_s(t) = u_{s'}(t)$ for every s, s' and $0 \le t < m$.

Set $V = (\delta_j^t)_{0 \le t < m,\, 1 \le j \le m}$ and $p_s = (p_{s1}, \ldots, p_{sm})^{\mathsf{T}}$. Then

$$V(p_s - p_{s'}) = 0, \qquad \det V = \prod_{i<j} (\delta_j - \delta_i) \ne 0,$$

so

$$p_{sj} = p_{s'j} =: \lambda_j \qquad \text{for every } s, s', j. \tag{11.2}$$

If these equalities fail, at least one of the displayed m moments is nonconstant and gives the preceding hardness reduction. Since $w_i > 0$ and every δ_j occurs, $\lambda_j > 0$; hence every I_s contains every amplitude α_j.

Write $s(i)$ for the class containing i. The identity

$$M_{ij} = a_i a_j B_{s(i),s(j)}$$

implies

$$M_i = M_j \iff s(i) = s(j) \text{ and } a_i = a_j.$$

Indeed, the right-hand side implies equality of all entries. Conversely, equal rows of a symmetric matrix have equal diagonal entries, hence equal a_i, a_j; dividing by these and by each column amplitude gives equal rows of C. Thus the classes K_a of the original matrix are precisely the nonempty sets $I_s \cap \{i : a_i = \alpha_j\}$. Reorder these classes by (j, s) to obtain

$$R = (\alpha\alpha^{\mathsf{T}}) \otimes T(\boldsymbol{\rho}), \qquad \bar{w}_{(j,s)} = \lambda_j, \qquad \bar{w} = \lambda \otimes \mathbf{1}.$$

This is (1.6). The argument includes $d = 0$: then $B = [1]$ and there is only one normalized row class.

11.3 Necessity for bipartite components

We follow the two-sided normalization in Theorem 9.1, now counting each internal vertex with its prescribed weight. Let

$$A_V = \begin{pmatrix} 0 & V \\ V^{\mathsf{T}} & 0 \end{pmatrix}, \quad V > 0, \qquad w^X = w|_X, \quad w^Y = w|_Y.$$

The superscripts specify the two color parts. Section 9.1 gives the Turing equivalence between Pl-GH$(A_V, (w^X, w^Y))$ and the problem with a prescribed input bipartition $L \to X$, $R \to Y$. The first is the sum of the two orientations on a connected bipartite input with an edge. The second is recovered by one root restriction; it is not obtained by dividing that sum by two. Isolated vertices contribute the sum of weights on their allowed color set.

Take two parallel V-edges from a root of color $i \in X$ to one internal vertex of color $j \in Y$. Omitting the root background weight, its conditional contribution is $w_j^Y V_{ij}^2$. Summing this vertex, and doing the same with the sides interchanged, gives the positive unaries

$$a_i^2 = \sum_{j \in Y} w_j^Y V_{ij}^2, \qquad b_j^2 = \sum_{i \in X} w_i^X V_{ij}^2, \qquad C_{ij} = \frac{V_{ij}}{a_i b_j}. \tag{11.3}$$

Operation (ii) in Section 11.1 supplies their fixed rational powers; a_i^{-1} and b_j^{-1} attached at the two endpoints supply C. Every internal vertex of the parallel-edge gadget retains its background weight. All other fixed constraints remain jointly available.

If two rows of V are proportional, their weighted norms are proportional by the same positive factor, so their rows in C coincide. Conversely,

$$C_i = tC_{i'} \implies V_i = t(a_i/a_{i'})V_{i'} \implies a_i = t(a_i/a_{i'})a_{i'} \implies t = 1.$$

The same reasoning applies to columns. Merge the repeated rows I_s and columns J_t of C to obtain B. It has nonproportional distinct rows and columns: deleting repeated coordinates does not change a proportionality relation. The symmetric matrix $A_B = \left(\begin{smallmatrix} 0 & B \\ B^{\mathsf{T}} & 0 \end{smallmatrix}\right)$ has no zero row; across parts its row supports are disjoint.

On a prescribed input bipartition, the left and right unary exponents can be chosen independently. Summing the normalized classes gives

$$\mu_s(m) = \sum_{i \in I_s} w_i^X a_i^{2m}, \qquad \nu_t(n) = \sum_{j \in J_t} w_j^Y b_j^{2n}. \tag{11.4}$$

For each fixed pair $m, n \geq 0$, this supplies the prescribed problem with matrix B and weights $\mu(m), \nu(n)$. By the two-orientation formula it also supplies Pl-GH$(A_B, (\mu(m), \nu(n)))$. Operation (iii) removes these positive quotient weights while retaining other available constraints. Lemma 11.1(b) therefore gives hardness or $B = cT(\boldsymbol{\rho})$, with $c > 0$ and all $\rho_r \neq 1$. In the latter case Lemma 9.3 forces $\mu(m)$ and $\nu(n)$ to be constant on their own parts unless the source is hard.

List the distinct amplitudes as $\alpha_1, \dots, \alpha_u$ on X and $\beta_1, \dots, \beta_v$ on Y, and set

$$p_{sh} = \sum_{\substack{i \in I_s \\ a_i = \alpha_h}} w_i^X, \qquad q_{th} = \sum_{\substack{j \in J_t \\ b_j = \beta_h}} w_j^Y.$$

Then $\mu_s(m) = \sum_h p_{sh} (\alpha_h^2)^m$ and $\nu_t(n) = \sum_h q_{th} (\beta_h^2)^n$. Use only the finite pairs

$$\{(m, 0) : 0 \le m < u\} \ \cup \ \{(0, n) : 0 \le n < v\}.$$

For $V_X = ((\alpha_h^2)^m)_{0 \le m < u,\, 1 \le h \le u}$, constancy gives $V_X(p_s - p_{s'}) = 0$. Its determinant is $\prod_{h < h'} (\alpha_{h'}^2 - \alpha_h^2) \neq 0$. The corresponding matrix V_Y at the distinct β_h^2 is also invertible. Consequently

$$p_{sh} = \lambda_h > 0 \qquad \text{independent of } s, \tag{11.5}$$

$$q_{th} = \eta_h > 0 \qquad \text{independent of } t. \tag{11.6}$$

Positivity follows because each amplitude occurs somewhere and every original vertex weight is positive. No equality between λ and η is required. Failure of either family of equalities is detected by one of the finite pairs above and gives hardness.

On X, equality of original rows is equivalent to equality of the normalized class and amplitude: equal V rows have equal weighted norm and equal C rows, while the converse follows from $V_{ij} = a_i b_j B_{s(i),t(j)}$. The same holds on Y. An X row and a Y row cannot coincide because their nonempty supports are disjoint. Hence the original identical-row quotient has cross matrix and weights

$$(\alpha \beta^{\mathsf{T}}) \otimes cT(\boldsymbol{\rho}), \qquad \bar{w}^X = \lambda \otimes \mathbf{1}, \quad \bar{w}^Y = \eta \otimes \mathbf{1}.$$

Absorbing c into α gives (1.7) and completes necessity for all support components. The two parts allow independent color permutations. In particular, with $X = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$,

$$W(\rho) X = \rho W(\rho^{-1}).$$

Thus flipping one Boolean coordinate on one bipartite side replaces $\rho > 1$ by $\rho^{-1} \in (0, 1)$, absorbing ρ into an amplitude. The side weights are unchanged because they are independent of that coordinate.

11.4 The algorithm and exact complexity

We now prove sufficiency and finish Theorem 1.3. For a positive component in (1.6), put

$$F(t) = \sum_{i=1}^{k} \lambda_i a_i^t \quad (t \in \mathbb{Z}_{\ge 0}).$$

Write a color as $(i, x) \in [k] \times \{0, 1\}^d$. For a fixed Boolean assignment $x : V(G) \to \{0, 1\}^d$, summation over the amplitude indices $i : V(G) \to [k]$ gives

$$\sum_i \prod_v \lambda_{i(v)} a_{i(v)}^{\deg_G(v)} = \prod_v \sum_{i=1}^{k} \lambda_i a_i^{\deg_G(v)}.$$

The remaining edge product depends only on x. Summing each Boolean coordinate separately, as in (2.6), gives the weighted counterpart of (2.4):

$$Z_{R_C, \bar{w}_C}(G) = \prod_{v \in V(G)} F(\deg_G(v)) \prod_{j=1}^{d} Z_{W(\rho_j)}(G). \tag{11.7}$$

This includes loops and isolated vertices: the latter contribute $2^d F(0)$, exactly the sum of the weights in that component.

For a bipartite component define

$$F_X(t) = \sum_i \lambda_i a_i^t, \qquad F_Y(t) = \sum_j \eta_j b_j^t.$$

On a connected input H containing an edge, the value is zero unless H is bipartite. For parts L, R, the two prescribed orientations contribute separately, as in (2.5). The amplitude sums acquire the coefficients λ_i, η_j:

$$Z_{R_C,\bar{w}_C}(H) = \left[\prod_{v\in L} F_X(\deg v) \prod_{v\in R} F_Y(\deg v) + \prod_{v\in L} F_Y(\deg v) \prod_{v\in R} F_X(\deg v)\right] \cdot \prod_{j=1}^{d} Z_{W(\rho_j)}(H). \tag{11.8}$$

For K_1, the value is $2^d(F_X(0) + F_Y(0))$. A zero singleton of weight z contributes z on K_1 and zero on a connected input containing an edge.

Each connected input component chooses one support component C of R; different input components choose independently. Thus the weighted version of (2.7), for input components $H_1, \ldots, H_t$, is

$$Z_{M,w}(G) = \prod_{h=1}^{t} \left(\sum_C Z_{R_C,\bar{w}_C}(H_h) \right). \tag{11.9}$$

Each Ising factor is a planar zero-field computation, as in Proposition 2.6. There are a fixed number of support components and tensor factors. Degrees, including loop multiplicity, are bounded by twice the input edge count. The displayed powers, products and sums have polynomial bit complexity in the fixed field representation.

All constants in this proof are algebraic: normalization uses arithmetic and positive square roots; the fixed unary exponents are rational; quotient weights are finite sums. The parameters in a tractable form can be chosen from entries by the formulas in Proposition 2.6. A fixed number field contains these constants and the original weights. Section 2.5 and Remark 3.4 give polynomial bit bounds and conversion to $\mathbb{Q}(M, w)$. The finite moment sets above require only a fixed number of joint simulation stages. Together with necessity, this completes Theorem 1.3 in the algebraic bit model.

11.5 Finite criteria and nonnegative vertex weights

The proof gives the following finite criterion for a fixed algebraic pair (M, w).

(i) Check the components of $P(M)$. A nonbasic component is hard by Section 11.1(iv); a zero component is evaluated directly.

(ii) On each positive component form $C_{ij} = M_{ij}/\sqrt{M_{ii}M_{jj}}$. On each bipartite component use the weighted norms (11.3). These operations keep the actual background weights until the class sums are taken.

(iii) Merge the repeated rows of C, or its repeated rows and columns, to form B. Apply the alternatives of Lemma 11.1 to test the required tensor form.

(iv) For a positive component compare p_{sj} in (11.2); for a bipartite component compare (11.5) and (11.6) separately. The finite moment sets used above suffice to detect a failure.

(v) If all tests pass, the identified original row classes give $(R, \bar{w})$ in Theorem 1.3; evaluate it by (11.7)–(11.9).

These are finite tests using exact algebraic arithmetic. If necessary, one tests all possible core color orderings and dimension factorizations. Their cost concerns recognition of the fixed pair (M, w) and may be exponential in its order; the polynomial-time claim in the theorem concerns evaluation as the input graph varies.

Corollary 11.2 (One fixed nonnegative vertex-weight vector)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ and $w \in \mathbb{R}^q_{\geq 0}$ have fixed algebraic entries. Put $S = \{i : w_i > 0\}$. If $S \neq \varnothing$, form the identical-row quotient of $(M[S], w|_S)$. The conditions of Theorem 1.3 on this quotient give a polynomial-time algorithm for* Pl-GH(M, w)*; if they fail, the problem is* #P*-hard. If $S = \varnothing$, the answer is* 1 *on the empty input and* 0 *on every nonempty input.*

Proof. For every graph G, any assignment using a color outside S has zero vertex-weight product. Therefore, when $S \neq \varnothing$,

$$Z_{M,w}(G) = Z_{M[S],w|_S}(G).$$

All remaining vertex weights are strictly positive, so Theorem 1.3 applies. If S is empty, every assignment on a nonempty vertex set has zero weight, while the empty assignment has weight one. The identity uses the original output field; conversion to or from the fixed field generated by the retained weights is supplied by the fixed-number-field conversion in Remark 3.4. □

Extending the weighted classification to signed vertex weights would require replacements for the positive root inner product, strict Gram inequalities and positive-mass argument used above.

12 Consequences and applications

We collect consequences of the two structural dichotomies, separating their statements in the Introduction from the short arguments needed here. Algebraic weights use the main theorems directly. Arbitrary fixed real Boltzmann weights use Theorems A.6 and A.12 in the representation specified in Appendix A. Parameters and state counts are fixed; the input graph alone varies.

12.1 Unweighted target graphs

Corollary 12.1 (The planar and general-graph tractable targets coincide)**.** *Let H be a fixed finite undirected graph, with loops permitted and without parallel edges. Counting homomorphisms from planar graphs to H is in* FP *if every component of H is a fully looped clique, a loopless complete bipartite graph, or a loopless isolated vertex. Otherwise it is* #P*-complete under polynomial-time Turing reductions. The list of tractable targets is the same for unrestricted input graphs.*

Proof. Apply Theorem 7.1 to the adjacency matrix of H. For completeness, a positive block of a zero–one matrix is all ones, and a bipartite block with a positive rectangular cross block has every cross entry equal to one. These are exactly the stated targets. A homomorphism is a polynomially checkable witness, giving membership in #P. The unrestricted-input classification is the theorem of Dyer and Greenhill [DG00, DG04]. □

Thus planarity preserves hardness for every unweighted target outside the classical list, even though it makes additional nonnegative weighted targets tractable. For example, loopless cliques on at least three vertices remain hard, whereas their strictly positive two-state Ising counterparts can be easy.

12.2 Planar bipartite inputs and two-sublattice models

Corollary 12.2 (Domain doubling and prescribed sublattices)**.** *For fixed symmetric $M \geq 0$ and $w > 0$, define*

$$D_M = \begin{pmatrix} 0 & M \\ M & 0 \end{pmatrix}, \qquad \widetilde{w} = \begin{pmatrix} w \\ w \end{pmatrix}.$$

Evaluation of $Z_{M,w}$ on planar bipartite inputs is polynomial-time Turing equivalent to Pl-GH$(D_M, \widetilde{w})$*. Consequently it has the complete dichotomy obtained by applying Theorem 1.3 to $(D_M, \widetilde{w})$.*

More generally, let $V \geq 0$ be a fixed rectangular matrix between color sets X, Y, and let $\mu > 0, \nu > 0$ be fixed weights on the two sets. On a planar input with prescribed bipartition $L \cup R$, evaluation of $Z_{V,\mu,\nu}(G; L, R)$ has the classification of Theorem 1.3 applied to

$$A_V = \begin{pmatrix} 0 & V \\ V^{\mathsf{T}} & 0 \end{pmatrix}, \qquad w_V = \begin{pmatrix} \mu \\ \nu \end{pmatrix}.$$

In both statements, unit weights give the corresponding unweighted classification, using Theorem 1.1.

Proof. For every bipartite graph G,

$$Z_{D_M, \widetilde{w}}(G) = 2^{c(G)} Z_{M,w}(G). \tag{12.1}$$

Each connected component has two choices of which input side maps to the first copy of the color set; after that choice, the color sum is exactly $Z_{M,w}$ on that component. An isolated vertex also has the factor two. A nonbipartite component has zero partition function for D_M. Testing bipartiteness, applying (12.1), and multiplying component values gives both reductions, preserving planarity.

For rectangular V, Section 9.1 proves the equivalence with Pl-GH(A_V, w_V) by a one-root restriction on each connected component. In the forward direction one sums the two possible orientations; unlike (12.1), those two values need not agree. Isolated vertices contribute the weight sum on their prescribed side. Applying the main dichotomies completes the classification. The fixed-real reductions follow from the same identities and the root-restriction results of Appendix A. □

Remark 12.3 (A strict enlargement of the tractable class). Let $0 < r < 1$ be algebraic, let $T = W(r) \otimes W(r)$, and let P interchange the two Boolean coordinates. Then

$$M = PT = TP = \begin{pmatrix} 1 & r & r & r^2 \\ r & r^2 & 1 & r \\ r & 1 & r^2 & r \\ r^2 & r & r & 1 \end{pmatrix}$$

is symmetric and strictly positive. Relabeling the colors on one side of a bipartite input changes M to $MP = T$, giving a polynomial time algorithm there. On general planar inputs, however, M is hard: it has full rank and $\det M = -(1 - r^2)^4 < 0$. A strictly positive full-rank tractable four-state matrix must be a positive scalar multiple of two Ising factors, whose determinant is positive. Hence Theorem 1.1 excludes M.

12.3 Clock models and coupled Ising systems

The next lemma gives a useful obstruction for interactions with constant diagonal. Its logarithm is taken *entry by entry*; it is different from the spectral matrix logarithm used in Section 4.

Lemma 12.4 (A centered entrywise-logarithm obstruction)**.** *Let $M > 0$ have constant diagonal and distinct rows, and let $w > 0$. If (M, w) belongs to the tractable class of Theorem 1.3, then, after relabeling, for some $d \geq 0$,*

$$q = 2^d, \qquad M = \gamma \bigotimes_{j=1}^{d} W(\rho_j), \qquad w = \lambda \mathbf{1}, \quad \gamma, \lambda > 0, \quad \rho_j > 0,\ \rho_j \neq 1.$$

Writing $C_{uv} = \log M_{uv}$ and $H_q = I - J/q$, we have

$$\operatorname{rank}(H_q C H_q) = d. \tag{12.2}$$

Conversely, every displayed tensor form with constant w is tractable.

Proof. The matrix is one positive support component, and distinct rows make its identical-row quotient unchanged. Its tractable form is $(aa^{\mathsf{T}}) \otimes T(\boldsymbol{\rho})$. Since every diagonal entry of T is one, constant diagonal forces all entries of a to agree. More than one such amplitude would produce identical rows, so a has length one. The weight condition then requires constant w.

Index colors by $x \in \{0,1\}^d$ and write $\chi_j(x) = (-1)^{x_j}$. For some scalar c,

$$C = cJ - \frac{1}{2} \sum_{j=1}^{d} (\log \rho_j) \chi_j \chi_j^{\mathsf{T}}.$$

The χ_j are nonzero mutually orthogonal vectors perpendicular to $\mathbf{1}$. All coefficients are nonzero, proving (12.2). The converse is the Ising product algorithm. □

Corollary 12.5 (Every fixed-state standard clock model)**.** *Let $q \geq 2$ and $K \in \mathbb{R}$ be fixed, with colors $0, \ldots, q-1$ and*

$$C(q, K)_{ij} = \exp\!\left(K \cos \frac{2\pi(i-j)}{q}\right).$$

If $K \neq 0$, its planar partition function with unit vertex weights is in FP *for $q = 2, 4$ and is #P-hard for every other q. With a fixed positive vertex-weight vector w, the tractable cases additionally require w to be constant; every other case is #P-hard. If $K = 0$, every $w > 0$ gives a polynomial-time problem.*

Proof. For $K \neq 0$ the diagonal is constant, and each row has its unique maximum at its diagonal if $K > 0$, or unique minimum there if $K < 0$. Thus the rows are distinct. For $q \geq 3$, put $c_i = \cos(2\pi i/q)$ and $s_i = \sin(2\pi i/q)$. The entrywise logarithm is

$$L = K(cc^{\mathsf{T}} + ss^{\mathsf{T}}).$$

The vectors c, s are nonzero and orthogonal, and both sum to zero. Consequently $H_q L H_q = L$ has rank two. If the weighted problem belongs to the tractable class, Lemma 12.4 forces $q = 2^2 = 4$ and constant w. For the two remaining state counts,

$$C(2, K) = e^K W(e^{-2K}), \qquad C(4, K) \cong e^K W(e^{-K}) \otimes W(e^{-K}).$$

For $q = 4$ use the cyclic Gray ordering $00, 01, 11, 10$. These give the algorithms when w is constant. The lemma excludes nonconstant w also when $q = 2$. The main weighted dichotomy supplies hardness in every excluded case. For $K = 0$ the interaction is all ones and the answer is $(\sum_i w_i)^{|V(G)|}$. □

The result covers each fixed state count and both signs of the coupling. The familiar two-Ising representation explains the four-state exception, while the obstruction excludes every other state count at once. The clock model has been studied through stochastic estimation [AM14] and through partition-function zeros obtained numerically [Kim17]. These computational approaches concern different tasks from the exact evaluation classified here. In particular, planar hardness alone does not establish hardness on rectangular square-lattice inputs.

Corollary 12.6 (Character criterion for coupled Ising interactions)**.** *Let n be fixed, let $\mathcal{A} \subseteq \mathbb{F}_2^n \setminus \{0\}$, and let $J_a \in \mathbb{R} \setminus \{0\}$ for $a \in \mathcal{A}$. For colors $x, y \in \mathbb{F}_2^n$, put*

$$M_{xy} = \exp\left(\sum_{a\in\mathcal{A}} J_a(-1)^{a\cdot x}(-1)^{a\cdot y}\right).$$

With unit vertex weights the planar evaluation problem is tractable if $\mathcal{A}$ is linearly independent over $\mathbb{F}_2$, and is #P-hard otherwise. For a fixed vector $w > 0$, define

$$F(x) = (a\cdot x)_{a\in\mathcal{A}}, \qquad \bar{w}_z = \sum_{x:F(x)=z} w_x \quad (z \in \operatorname{im} F).$$

The weighted problem is tractable exactly in the listed structural cases: $\mathcal{A}$ is linearly independent and $\bar{w}_z$ is constant on $\operatorname{im} F$. Outside these cases it is #P-hard.

Proof. Let $r = \operatorname{rank}_{\mathbb{F}_2} \mathcal{A}$ and $m = |\mathcal{A}|$. Two rows of M agree exactly when $F(x) = F(x')$: after taking logarithms, independence of the distinct characters $y \mapsto (-1)^{a\cdot y}$ proves the reverse implication. The quotient therefore has 2^r distinct rows and constant diagonal. Its entrywise logarithm is

$$L = \sum_{a\in\mathcal{A}} J_a \chi_a \chi_a^{\mathsf{T}},$$

where the χ_a are distinct nonconstant characters on $\mathbb{F}_2^n/\ker F$. They are mutually orthogonal and have zero sum, so the centered logarithm has rank m. Lemma 12.4 requires $m = r$, precisely linear independence, and constant aggregated weights. If $m = r$, the values of F range over all $\mathbb{F}_2^m$, and the quotient is

$$\exp\left(\sum_{a\in\mathcal{A}} J_a\right) \bigotimes_{a\in\mathcal{A}} W(e^{-2J_a}).$$

Constant $\bar{w}$ gives the algorithm. Unit original weights have constant fiber sum 2^{n-r}. When $\mathcal{A}$ is empty the quotient has one state and the condition is vacuous, as required. □

For an n-layer Ashkin–Teller interaction [GW81], ordinary layer couplings have labels e_i, while four-spin couplings between layers i, j have labels $e_i + e_j$. If every ordinary layer coupling is nonzero, adding any nonzero four-spin coupling creates a dependence and gives planar hardness. If only four-spin couplings are present, their labels are independent exactly when the graph of coupled layers is a forest: a cycle sums to zero, whereas a leaf eliminates its incident coefficient in any putative dependence. Thus forests give tractable unit-weight systems and cycles give hard ones. This refers to the fixed graph of layers, distinct from the variable planar input graph.

12.4 Fixed single-site potentials

The aggregated-weight condition has a direct interpretation for states with identical interactions. If $w_i = e^{-\beta U_i}$ with $\beta > 0$, a class K of identical rows has effective potential

$$U_K^{\mathrm{eff}} = -\beta^{-1} \log \sum_{i \in K} e^{-\beta U_i}.$$

Theorem 1.3 requires this effective potential to be independent of the nondegenerate Ising coordinates, within each amplitude and each bipartite side. Unequal microscopic potentials can therefore preserve tractability when their class sums agree. This is the appropriate field condition when colors have interaction degeneracies.

Corollary 12.7 (Spin-one Blume–Capel model)**.** *For $s, t \in \{-1, 0, 1\}$, let*

$$M_{st} = e^{Kst}, \qquad w_s = e^{-Ds^2 + Fs},$$

where K, D, F are fixed finite real numbers. Exact planar evaluation is #P-hard for every $K \neq 0$, and is polynomial-time computable when $K = 0$.

Proof. For $K \neq 0$, M is strictly positive, has distinct rows, and, putting $x = e^K$, has determinant

$$\det M = (x-1)^3(x+1)/x^2 \neq 0.$$

A positive tractable matrix of full rank has no rank-one factor of size greater than one, so its order must be a power of two. Order three is impossible, regardless of the positive weights. Theorem 1.3, or its fixed-real extension, proves hardness. For $K = 0$ the value is $(\sum_s w_s)^{|V(G)|}$. □

The parameters D and F are the crystal and magnetic fields, with inverse temperature absorbed into them. Studies of this model include series expansions for its phase diagram [BP18]. Its exact classification is compatible with approximation and sampling algorithms in restricted regimes: for example, algorithmic Pirogov–Sinai theory treats appropriate low-temperature lattice systems [HPR18]. For Potts interactions with external fields, the corresponding partition functions also have a graph-polynomial formulation [EM11].

12.5 Graphs of fixed bounded genus

Corollary 12.8 (The same dichotomies at bounded orientable genus)**.** *Fix $g \geq 0$. Replace planar inputs in either main theorem by graphs of orientable genus at most g, supplied with an embedding. The tractable structural classes and hardness alternatives are unchanged.*

Proof. The decompositions in Section 2.8 and their positive-vertex-weight counterparts in Section 11.4 hold on any input graph. In particular, (11.7), (11.8), and (11.9) separate the amplitude sums, the two bipartite orientations, and the input components without using planarity. For each input component, they leave a fixed number of zero-field Ising factors. A bipartite target component contributes zero on an input component with a loop. For a positive target component, a loop has already contributed twice to the degree in each amplitude factor. Its remaining Ising contribution is the constant diagonal entry of $W(\rho)$, namely 1, so these Ising loops may be deleted. Isolated vertices and parallel edges are retained.

It therefore suffices to evaluate one such Ising factor on a loopless graph H embedded in an orientable surface of genus $h \leq g$. For $n = |V(H)|$, $m = |E(H)|$, and $x = (1-\rho)/(1+\rho)$, the expansion from Section 2.8 reads

$$Z_{W(\rho)}(H) = 2^n \left(\frac{1+\rho}{2}\right)^m P_H(x), \qquad P_H(t) = \sum_{\substack{A \subseteq E(H) \\ \deg_A(v) \text{ even for all } v}} t^{|A|}.$$

We must allow $x < 0$, which occurs when $\rho > 1$. To do so, use the signed Pfaffian formula of [Pha20, Theorem 2.3]. Its terminal graph has one vertex for each half-edge of H: short edges join half-edges at the same vertex of H, and long edges join the two halves of each original edge. Give every short edge the formal weight t and every long edge weight 1. The fixed orientations prescribed by that theorem give skew-adjacency matrices $A_j(t)$ and signs $c_j \in \{-1, 1\}$ such that, for $t > 0$,

$$P_H(t) = \frac{\epsilon_0}{2^h} \sum_{j=1}^{4^h} c_j \operatorname{Pf}(A_j(t)).$$

The signs depend only on the graph, its embedding, the orientations, and the vertex order. The overall sign $\epsilon_0 \in \{-1, 1\}$ is determined by the reference perfect matching consisting of all long edges, equivalently by the normalization $P_H(0) = 1$. Every entry of $A_j(t)$ is a polynomial with integer coefficients, so both sides of the displayed identity are polynomials in t. Their equality for $t > 0$ therefore proves equality for all real t. Substituting $t = x$ preserves the sign of the high-temperature weight, just as in the planar polynomial identity of Section 2.8.

Thus each Ising factor is evaluated using at most 4^g Pfaffians. The matrices have polynomial size, and their entries at $t = x$ belong to the fixed field generated by ρ; no additional square roots are needed. For fixed g, the amplitude operations and Pfaffian evaluations consequently have polynomial bit complexity under the same exact arithmetic conventions. The hard cases already hold on the planar subclass, whose embeddings can be computed in polynomial time. $\square$

References

[AM14] A. Al-Bashabsheh and Y. Mao. On stochastic estimation of partition function, 2014. https://arxiv.org/abs/1401.7273.

[BTR20] C. Boutillier, B. de Tilière, and K. Raschel. The Z-invariant Ising model via dimers, 2020. URL https://arxiv.org/abs/1612.09082v3. Version 3, 27 October 2020; Section 2.1.3, pp. 7–8.

[Bul08] A. A. Bulatov. The complexity of the counting constraint satisfaction problem. In *Automata, Languages and Programming (ICALP 2008), Part I*, volume 5125 of *Lecture Notes in Computer Science*, pages 646–661. Springer, 2008. doi: 10.1007/978-3-540-70575-8_53.

[BG05] A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2–3):148–186, 2005. doi: 10.1016/j.tcs.2005.09.011.

[BP18] P. Butera and M. Pernici. The Blume–Capel model for spins $S = 1$ and 3/2 in dimensions $d = 2$ and 3. *Physica A*, 507:22–66, 2018. doi: 10.1016/j.physa.2018.05.010.

[CC17] J.-Y. Cai and X. Chen. Complexity of counting CSP with complex weights. *Journal of the ACM*, 64(3), Article 19, 1–39, 2017. doi: 10.1145/2822891.

[CCL13] J.-Y. Cai, X. Chen, and P. Lu. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM Journal on Computing*, 42(3):924–1029, 2013. doi: 10.1137/110840194.

[CCL16] J.-Y. Cai, X. Chen, and P. Lu. Nonnegative weighted #CSP: An effective complexity dichotomy. *SIAM Journal on Computing*, 45(6):2177–2198, 2016. doi: 10.1137/15M1032314.

[CF16] J.-Y. Cai and Z. Fu. Holographic algorithm with matchgates is universal for planar #CSP over Boolean domain, 2016. https://arxiv.org/abs/1603.07046.

[CLX09] J.-Y. Cai, P. Lu, and M. Xia. Holant problems and counting CSP. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009)*, pages 715–724. ACM, 2009. doi: 10.1145/1536414.1536511.

[CLX17] J.-Y. Cai, P. Lu, and M. Xia. Holographic algorithms with matchgates capture precisely tractable planar #CSP. *SIAM Journal on Computing*, 46(3):853–889, 2017. doi: 10.1137/16M1073984. Preliminary version in FOCS 2010, pages 427–436.

[CM23] J.-Y. Cai and A. 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. ACM, 2023. doi: 10.1145/3564246.3585173. Full version: https://arxiv.org/abs/2302.08570v1. Theorem numbers refer to this arXiv version.

[CM24] J.-Y. Cai and A. Maran. Polynomial and analytic methods for classifying complexity of planar graph homomorphisms, 2024. https://arxiv.org/abs/2412.17122v1. Theorem numbers refer to this version.

[CFGW22] J.-Y. Cai, Z. Fu, H. Guo, and T. Williams. FKT is not universal—a planar Holant dichotomy for symmetric constraints. *Theory of Computing Systems*, 66:143–308, 2022. doi: 10.1007/s00224-021-10032-1.

[CMY26] J.-Y. Cai, A. Maran, and B. Young. Planar graph homomorphisms: A dichotomy and a barrier from quantum groups, 2026. URL https://arxiv.org/abs/2601.23198v1. Version 1, 30 January 2026.

[CH96] N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996. doi: 10.1006/inco.1996.0016.

[DG00] M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3–4):260–289, 2000. https://pages.cs.wisc.edu/~jyc/880/Dyer-Greenhill.pdf.

[DG04] M. Dyer and C. Greenhill. Corrigendum: The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 25:346–352, 2004. https://pages.cs.wisc.edu/~jyc/880/Dyer-Greenhill-corrigendum.pdf.

[DR13] M. Dyer and D. Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM Journal on Computing*, 42(3):1245–1274, 2013. doi: 10.1137/100811258.

[EM11] J. A. Ellis-Monaghan and I. Moffatt. The Tutte–Potts connection in the presence of an external magnetic field. *Advances in Applied Mathematics*, 47(4):772–782, 2011. doi: 10.1016/j.aam.2011.02.004. Extended version: https://arxiv.org/abs/1005.5470.

[Fu17] Z. Fu. On blockwise symmetric matchgate signatures and higher domain #CSP, 2017. URL https://arxiv.org/abs/1707.00373v1. Version 1.

[GGJT10] L. A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM Journal on Computing*, 39(7):3336–3402, 2010. doi: 10.1137/090757496.

[GW81] G. S. Grest and M. Widom. N-color Ashkin–Teller model. *Physical Review B*, 24:6508–6515, 1981. doi: 10.1103/PhysRevB.24.6508.

[GW20] H. Guo and T. Williams. The complexity of planar Boolean #CSP with complex weights. *Journal of Computer and System Sciences*, 107:1–27, 2020. doi: 10.1016/j.jcss.2019.07.005.

[HPR18] T. Helmuth, W. Perkins, and G. Regts. Algorithmic Pirogov–Sinai theory, 2018. https://arxiv.org/abs/1806.11548.

[Hig08] N. J. Higham. *Functions of Matrices: Theory and Computation.* Society for Industrial and Applied Mathematics, 2008. Print ISBN 9780898716467; electronic ISBN 9780898717778. doi: 10.1137/1.9780898717778.

[HJ13] R. A. Horn and C. R. Johnson. *Matrix Analysis.* Cambridge University Press, 2nd edition, 2013 (electronic edition; print edition published in 2012). doi: 10.1017/CBO9781139020411.

[Kas61] P. W. Kasteleyn. The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209–1225, 1961. doi: 10.1016/0031-8914(61)90063-5.

[Kas63] P. W. Kasteleyn. Dimer statistics and phase transitions. *Journal of Mathematical Physics*, 4(2):287–293, 1963. doi: 10.1063/1.1703953.

[Kim17] D.-H. Kim. Partition function zeros of the p-state clock model in the complex temperature plane. *Physical Review E*, 96:052130, 2017. doi: 10.1103/PhysRevE.96.052130.

[PST24] D. Pavlov, B. Sturmfels, and S. Telen. Gibbs manifolds. *Information Geometry*, 7(Suppl 2):691–717, 2024. doi: 10.1007/s41884-023-00111-2. First published online in 2023.

[Pav24] D. Pavlov. Logarithmically sparse symmetric matrices. *Beiträge zur Algebra und Geometrie*, 65:907–922, 2024. doi: 10.1007/s13366-024-00753-y.

[Pha20] A. M. Pham. A Pfaffian formula for the Ising partition function of surface graphs, 2020. https://arxiv.org/abs/2003.02967. doi: 10.1088/1742-5468/aba1e4.

[TF61] H. N. V. Temperley and M. E. Fisher. Dimer problem in statistical mechanics—an exact result. *Philosophical Magazine*, 6(68):1061–1063, 1961. doi: 10.1080/14786436108243366.

[Val08] L. G. Valiant. Holographic algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008. doi: 10.1137/070682575.

[Ver05] D. Vertigan. The computational complexity of Tutte invariants for planar graphs. *SIAM Journal on Computing*, 35(3):690–712, 2005. doi: 10.1137/S0097539704446797.

A Arbitrary fixed real weights

This appendix extends the two algebraic classifications, Theorems 1.1 and 1.3, to arbitrary fixed real weights. Exact interpolation requires equality tests for entry products, solutions of linear systems, and conversion to the output field specified by the source problem. We first establish those operations, then the unweighted classification and the reductions needed for the positive-vertex-weight classification.

Following [CMY26, Appendix A], a fixed finitely generated real field is presented as

$$F = \mathbb{Q}(X_1, \dots, X_d)(\beta), \qquad z = \sum_{j=0}^{e-1} z_j(X)\beta^j, \quad z_j(X) \in \mathbb{Q}(X).$$

Algebraic independence means that no nonzero polynomial over $\mathbb{Q}$ vanishes at $(X_1, \dots, X_d)$. These elements form a transcendence basis: the remaining generators are algebraic over $\mathbb{Q}(X_1, \dots, X_d)$. Here e is the degree of the fixed minimal polynomial of β over $\mathbb{Q}(X)$ (the monic polynomial of least positive degree vanishing at β), and every source weight has a fixed expression in this basis. Weighted and mixed sources include all vertex and constraint weights. Compatible presentations for auxiliary constants, extensions and subfields may be hardwired into a fixed reduction; none varies with the input graph. All algorithms below use exact bit encodings.

Signs of fixed nonzero entries $c_1, \dots, c_t$ are also fixed data. For $n \in \mathbb{Z}^t$, write $c^n = \prod_\alpha c_\alpha^{n_\alpha}$ and $\epsilon_\alpha = \operatorname{sgn}(c_\alpha)$. Then

$$\operatorname{sgn}(c^n) = \epsilon^n, \qquad c^n = c^{n'} \implies \epsilon^n = \epsilon^{n'}.$$

Thus product classes have well-defined sign tags without ordering transcendental field elements. Field arithmetic and output descent are proved in Lemma A.1. We first establish field arithmetic and interpolation, then approximate nonnegative weights by rational weights while preserving product relations. Relative closedness of the structural class keeps a sufficiently close approximation outside that class. We then extend support transfer and positive-weight removal to signed edges. These are reductions and sufficient hardness criteria; the complete signed-edge classification remains open. Section A.5 proves the full nonnegative-edge classification with one positive vertex-weight vector.

A.1 Exact arithmetic and interpolation

Lemma A.1 (Exact arithmetic and output descent in fixed real fields)**.** *Let F be a fixed finitely generated real field with the presentation specified above. Products of polynomially many fixed elements and solutions of polynomial-size nonsingular linear systems whose entries have polynomial degree and coefficient bit length can be computed with polynomial bit complexity. Partition functions of polynomial-size instances over a fixed constraint language have polynomial-size exact output encodings; this is an output-size bound, not a polynomial-time evaluation assertion. The same assertions hold in any fixed finite algebraic extension.*

If a computed element is known to lie in a specified fixed finitely generated subfield $F_0 \subseteq F$, it can be converted in polynomial time to the prescribed representation of F_0, using fixed compatible presentations. This includes conversion of a recovered support count to its ordinary integer output.

Proof. Arithmetic and interpolation systems. Write $F = \mathbb{Q}(X)(\beta)$ with fixed $d = |X|$ and extension degree e. Multiplication in $1, \beta, \dots, \beta^{e-1}$ has fixed rational-function structure constants. Choose a common fixed denominator for those constants and the finitely many source entries. For a product of L source entries, reduction in this basis gives coordinates $p(X)/q(X)$ satisfying

$$\deg p, \deg q = O(L), \qquad \text{bits(coefficients of } p, q) = \text{poly}(L).$$

Indeed, multiplication by a source entry is a fixed matrix operation: each step increases degrees by a constant and coefficient bit lengths by a constant plus the logarithm of a polynomial in the current degree. Since there are only $\binom{D+d}{d} = \text{poly}(D)$ monomials of total degree at most D, these are polynomial-size dense encodings.

Field operations and exact equality are polynomial-time on these encodings; inversion uses the e-dimensional multiplication matrix over $\mathbb{Q}(X)$. An N-dimensional linear system over F becomes an eN-dimensional system over $\mathbb{Q}(X)$. Clear its denominators and apply fraction-free elimination. Determinant degree and coefficient bounds make the solution size polynomial in N and in the entry degree and bit bounds. This supplies an algorithm without enumerating determinant permutations. A fixed finite algebraic extension merely replaces e by another fixed degree, so the same argument applies there.

Length of an exact answer. List all fixed nonzero source entries and background weights as $c_1, \dots, c_t$. For an instance Ω with n vertices and L binary or unary occurrences, grouping colorings by occurrence counts gives

$$Z(\Omega) = \sum_{\substack{\alpha \in \mathbb{Z}_{\geq 0}^t \\ \sum_{j=1}^t \alpha_j = L+n}} b_\alpha c^\alpha, \qquad 0 \leq b_\alpha \leq q^n.$$

Terms using zero entries have vanished. The number of exponent vectors is $\binom{L+n+t-1}{t-1}$, and each integer b_α has $O(n)$ bits because q, t are fixed. The preceding product bounds thus give polynomial-size encodings even when signed terms cancel. This groups colorings to bound output size; it does not give an algorithm to obtain the coefficients b_α.

Conversion to the prescribed output field. For a finite extension E/F, $\text{Tr}_{E/F}(z)$ is the trace of the F-linear multiplication map $x \mapsto zx$ on E; see Remark 3.4. For $z \in F$ this map is zI, so an element already known to lie in F satisfies

$$z = \frac{1}{[E:F]} \text{Tr}_{E/F}(z).$$

The trace is obtained from a fixed-dimensional multiplication matrix. Now consider the general inclusion $F_0 \subseteq F$. Fix a transcendence basis $Y_1, \dots, Y_b$ for F/F_0 and a compatible presentation making F finite algebraic over $F_0(Y_1, \dots, Y_b)$. These choices and the changes between the fixed presentations are constants of the reduction. Such changes substitute fixed rational expressions for finitely many generators and reduce in a fixed algebraic basis, so they preserve polynomial degree and bit bounds. For the given $z \in F_0$, compute

$$z = \frac{\text{Tr}_{F/F_0(Y)}(z)}{[F : F_0(Y)]} = \frac{p(Y)}{q(Y)}, \qquad p, q \in F_0[Y], \quad q \neq 0.$$

Algebraic independence implies $p = zq$. If $[Y^\alpha]q \neq 0$, where $[Y^\alpha]$ denotes coefficient extraction, then

$$z = \frac{[Y^\alpha]p}{[Y^\alpha]q} \in F_0.$$

Exact equality locates a nonzero denominator coefficient; no order test is needed. The fixed presentation of F_0 encodes this ratio. This also covers the case that F/F_0 is not algebraic, as can happen when an auxiliary positive real vertex weight is transcendental over the output field of the unweighted counting problem.

For an integer k in the original presentation of F, all β^j coefficients for $j > 0$ vanish, and the remaining coordinate satisfies $p(X) = kq(X)$. The matching-coefficient ratio therefore returns the ordinary integer encoding, including zero when $p = 0$. Rational outputs are handled identically. Every descent uses the specified fixed output subfield; no numerical evaluation of the transcendence variables or discovery of an unspecified subfield is required. $\square$

Corollary A.2 (Root restrictions over a fixed real field)**.** *Lemma 3.5 holds for fixed $A \in \mathrm{Sym}_q(\mathbb{R})$, $w \in \mathbb{R}^q_{>0}$ and $X \subseteq [q]$ in the representation above. In particular, with fixed rooted planar graphs F_j and fixed $c_j \in \mathbb{Q}(A, w)$,*

$$\sum_{\substack{\sigma:V(G)\to[q]\\ \sigma(r)\in X}} \prod_v w_{\sigma(v)} \prod_{uv\in E(G)} A_{\sigma(u),\sigma(v)} = \sum_j c_j Z_{A,w}(G \vee F_j).$$

Support-component restrictions and prescribed bipartite orientations therefore have the same Turing reductions, also with a fixed finite mixed language.

Proof. In Lemma 3.5, the rooted signatures span a subspace of $\mathbb{R}^q$ and $D = \mathrm{diag}(w) \succ 0$ defines its inner product. A finite rooted-signature basis is fixed problem data. Its Gram matrix H and the projection vector g lie in $F = \mathbb{Q}(A, w)$, and $H \succ 0$, so $c = H^{-1}g \in F^s$. The gluing identity and its single root weight are unchanged. Lemma A.1 evaluates this fixed linear combination and supplies any required subfield conversion in polynomial time. The same reasoning includes retained constraint weights in F. $\square$

Theorem A.3 (Signed joint interpolation over a fixed real field)**.** *Work with a fixed real constraint language and a specified fixed positive background vector in the field model above. Let C be a jointly available real symmetric binary constraint, and let C' be another fixed real symmetric matrix. Assume that*

(i) $C_{ij} = 0$ implies $C'_{ij} = 0$;

(ii) whenever two products of equally many nonzero entries of C are equal, the corresponding products of entries of C' are equal.

Then C' is jointly available. Several such replacement types are jointly available together, and the analogous statement holds for unary constraints. All the constants involved are represented in a common fixed finitely generated real field; additional finite algebraic extensions are allowed as fixed data. The replacement may introduce further zero entries.

Proof. This is the fixed-real counterpart of Lemma 3.1. Use its marked-occurrence construction, with s the common number of C' occurrences to be replaced. For $s = 0$ no replacement is needed; for $s > 0$ and $C = 0$, (i) makes the requested value zero. Otherwise, if $c_1, \ldots, c_t$ are the nonzero entry values of C, parallel multiplicity $h \geq 1$ gives

$$F(h) = \sum_{\lambda\in\Lambda_s} A_\lambda \lambda^h, \qquad N = |\Lambda_s| \leq \binom{s+t-1}{t-1},$$

$$\Lambda_s = \{c^n : n \in \mathbb{Z}^t_{\geq 0},\ \sum_{j=1}^t n_j = s\}.$$

For a coloring using no zero source entry on the marked occurrences, let λ be the product of its s source entries. The coefficient A_λ is the sum, over this class of colorings, of all unmarked edge factors, explicit unaries, and background weights. These factors are retained; they may have signs and may cancel. Original zero-entry terms vanish both here and after replacement by (i). The proof of Lemma 3.1 is algebraic: after exact field equality merges the bases, its Vandermonde determinant $\prod_j \lambda_j \prod_{i<j}(\lambda_j - \lambda_i)$ is nonzero over this field as well. The values $h = 1, \dots, N$ therefore recover all A_λ, regardless of signs or cancellation. By (ii), the replacement product $\theta(\lambda)$ is well-defined, possibly zero. More explicitly, for any representative list of positions $(i_1, j_1), \dots, (i_s, j_s)$ with nonzero C entries,

$$\lambda = \prod_{a=1}^{s} C_{i_a j_a}, \qquad \theta(\lambda) = \prod_{a=1}^{s} C'_{i_a j_a}, \qquad Z_{C' \text{ on marked occurrences}} = \sum_{\lambda} A_\lambda \theta(\lambda).$$

Condition (ii) makes the middle expression independent of the representative. The finite number t of source entry types makes N polynomial in the input length.

There are N queries, each with at most sN copies on the marked edges. Lemma A.1 gives polynomial bit cost for product grouping, oracle encodings and the interpolation system. If C' introduces constants beyond the source field, use the fixed field generated by both sets of constants and their fixed embeddings, then descend the result to its prescribed output field by that lemma. The planar parallel substitutions, unary attachments and successive elimination of a fixed number of types are exactly those of Lemma 3.1, including its treatment of loops. If C is jointly simulated, append its simulation as a further fixed stage. These operations preserve jointness and polynomial complexity. □

A.2 Approximation and the nonnegative classification

Let $\mathcal{T}_q$ be the structural family in Theorem 1.1. To transfer hardness to fixed real $M \geq 0$, we seek a rational symmetric M' satisfying

$$\|M' - M\|_{\max} < \delta, \qquad P(M') = P(M),$$
$$\prod_{a=1}^{m} M_{i_a j_a} = \prod_{a=1}^{m} M_{i'_a j'_a} \neq 0 \quad \Longrightarrow \quad \prod_{a=1}^{m} M'_{i_a j_a} = \prod_{a=1}^{m} M'_{i'_a j'_a}.$$

Here $\|N\|_{\max} = \max_{i,j} |N_{ij}|$, and the implication is required for every length m. The approximation preserves signs as well when entries are signed. For $M \notin \mathcal{T}_q$, relative closedness on its fixed support gives a $\delta > 0$ such that

$$M' \geq 0, \quad P(M') = P(M), \quad \|M' - M\|_{\max} < \delta \quad \Longrightarrow \quad M' \notin \mathcal{T}_q.$$

Choosing a rational M' with the stated product property then gives

$$\text{Pl-GH}(M') \text{ is \#P-hard}, \qquad \text{Pl-GH}(M') \leq_T \text{Pl-GH}(M).$$

Hardness of M' follows from the proved algebraic classification; the reduction requires precisely the product implication above. The approximation is chosen once for fixed M and yields an exact reduction.

Lemma A.4 (Relation-preserving approximation)**.** *For every fixed $m \in (\mathbb{R} \setminus \{0\})^s$ and every $\delta > 0$, there is $m' \in (\mathbb{Q} \setminus \{0\})^s$ such that*

$$\max_{1 \leq \alpha \leq s} |m'_\alpha - m_\alpha| < \delta, \qquad \operatorname{sgn}(m'_\alpha) = \operatorname{sgn}(m_\alpha) \quad (1 \leq \alpha \leq s),$$
$$m^z = 1 \implies (m')^z = 1 \quad (z \in \mathbb{Z}^s).$$

In particular, positive m has positive rational approximations. Applying this to the nonzero upper-triangular entries of a fixed real symmetric M, preserving its zero positions and symmetry, gives a rational symmetric M' jointly available from M. Thus Pl-GH$(M') \leq_T$ Pl-GH(M) *in the fixed-real-weight model. The joint replacement retains any specified fixed positive background weights and may coexist with other fixed source or replacement types.*

Proof. Write $m_\alpha = \epsilon_\alpha a_\alpha$, where $a_\alpha > 0$ and $\epsilon_\alpha \in \{-1, 1\}$ are fixed. Let

$$\Lambda = \{z \in \mathbb{Z}^s : \prod_{\alpha=1}^{s} a_\alpha^{z_\alpha} = 1\}.$$

Taking logarithms turns each multiplicative relation into $z \cdot \log a = 0$. A subgroup of $\mathbb{Z}^s$ is finitely generated; choose integer generators $z^{(1)}, \ldots, z^{(b)}$ of Λ. Its orthogonal complement is the solution space of the integer system

$$\Lambda^\perp = \{x \in \mathbb{R}^s : z^{(k)} \cdot x = 0 \ (1 \le k \le b)\}.$$

Gaussian elimination over $\mathbb{Q}$ gives a rational basis. Multiplying each basis vector by a nonzero integer to clear its denominators gives integer vectors $u_1, \ldots, u_h$ that remain a basis over $\mathbb{R}$. Since $\log a \in \Lambda^\perp$, write $\log a = \sum_j t_j u_j$. Choose positive rationals r_j tending to e^{t_j} and put

$$a'_\alpha = \prod_j r_j^{(u_j)_\alpha}, \qquad m'_\alpha = \epsilon_\alpha a'_\alpha.$$

Integer exponents make a' positive rational, and continuity gives $m' \to m$ with the prescribed signs. For $z \in \Lambda$,

$$(a')^z = \prod_j r_j^{u_j \cdot z} = 1, \qquad m^z = 1 \iff z \in \Lambda \text{ and } \epsilon^z = 1 \implies (m')^z = 1.$$

When $h = 0$, $\log a = 0$ and the empty product gives $a' = a = \mathbf{1}$. Choosing the r_j sufficiently close proves the assertion for each $\delta > 0$; when $m > 0$, all signs are $+1$.

For the matrix replacement, if two nonzero products are $m^n = m^{n'}$, their exponent difference $z = n - n'$ satisfies $m^z = 1$, hence $(m')^n = (m')^{n'}$. Zero products remain zero, and mirrored upper-triangular entries receive the same value. For Theorem A.3, take $C = M$ and $C' = M'$, with any specified fixed positive background vector. Source zero positions stay zero, and the displayed implication verifies condition (ii) for each common length. Thus M' is jointly available and Pl-GH$(M') \le_T$ Pl-GH(M). Other constraints and background weights remain in its coefficients; Lemma A.1 returns the prescribed output encoding, in particular a rational encoding for unweighted M'. If $M = 0$, no replacement is necessary. The implication only goes from source relations to target relations; new relations in m' are allowed. The chosen approximation, the relation generators and the basis vectors are fixed data depending on M. The construction does not require computing a relation basis from a variable real-number input. It approximates fixed parameters to obtain an exact oracle reduction, rather than approximating a partition function. □

Lemma A.5 (Limits with fixed support)**.** *For a fixed support pattern and fixed domain size, the tractable families in Theorem 1.1 are relatively closed in the space of matrices positive on that support. That is, a convergent sequence in these families whose limit is positive on the same support has its limit in these families.*

Proof. For a symmetric support pattern P, the relative space is

$$\mathcal{S}_P := \{M \in \mathrm{Sym}_q(\mathbb{R}) : M_{ij} > 0 \text{ if } P_{ij} = 1, \quad M_{ij} = 0 \text{ if } P_{ij} = 0\}.$$

We show that $\mathcal{T}_q \cap \mathcal{S}_P$ is closed in $\mathcal{S}_P$. For nonbasic support this intersection is empty. Otherwise it suffices to examine the positive symmetric and positive rectangular blocks. There are finitely many dimension factorizations and color permutations. For a fixed factorization $q = k2^d$, order the colors of a positive symmetric block by $(i, x) \in [k] \times \{0, 1\}^d$. Let 0 denote the zero bit string and e_r the bit string with its only nonzero coordinate at r. Fix $i_0 \in [k]$ and set

$$R_{ij} = M_{(i,0),(j,0)}, \qquad \rho_r = \frac{M_{(i_0,0),(i_0,e_r)}}{M_{(i_0,0),(i_0,0)}},$$

For this fixed ordering, the proposed tensor form is equivalent to

$$\operatorname{rank} R = 1, \qquad M_{(i,x),(j,y)} = R_{ij} \prod_{r:x_r \neq y_r} \rho_r \quad \text{for all } i, j, x, y.$$

Since $R > 0$, its rank is one exactly when all its 2-by-2 minors vanish; symmetry then gives $R = aa^{\mathsf{T}}$ with $a_i = R_{ii_0}/\sqrt{R_{i_0i_0}} > 0$. All denominators in the displayed identities are positive. Clearing them therefore gives equivalent polynomial equalities.

For a rectangular block, fix $p = k2^d$, $q = l2^d$ and independent row and column orderings by $[k] \times \{0,1\}^d$ and $[l] \times \{0,1\}^d$. Fix $i_0 \in [k]$ and $j_0 \in [l]$, and define

$$R_{ij} = V_{(i,0),(j,0)}, \qquad \rho_r = \frac{V_{(i_0,0),(j_0,e_r)}}{V_{(i_0,0),(j_0,0)}}.$$

For these orderings the corresponding equivalence is

$$\operatorname{rank} R = 1, \qquad V_{(i,x),(j,y)} = R_{ij} \prod_{r:x_r \neq y_r} \rho_r \quad \text{for all } i, j, x, y.$$

Here vanishing 2-by-2 minors give the positive factorization $R = ab^{\mathsf{T}}$, for example $a_i = R_{ij_0}/R_{i_0j_0}$ and $b_j = R_{i_0j}$. Again the conditions become polynomial equalities after clearing positive denominators. They include $\rho_r = 1$, hence retain all degeneracies whose entries stay positive. Each fixed-ordering family is relatively closed. There are finitely many orderings and dimension factorizations, so their union is relatively closed. A fixed support fixes its connected blocks; requiring each block to belong to its corresponding family preserves relative closedness. □

Theorem A.6 (Nonnegative matrices with fixed real entries)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ be fixed in the representation of this appendix. If every support component has a form in Theorem 1.1, then* Pl-GH(M) $\in$ FP*; otherwise* Pl-GH(M) *is* #P*-hard. Thus Theorem 1.1 extends to arbitrary fixed nonnegative real entries.*

Proof. Let $M \geq 0$ be fixed. If $P(M)$ has a nonbasic component, Theorem 7.1 and the fixed-field support interpolation of Theorem A.3 give hardness through Pl-GH($P(M)$) $\leq_T$ Pl-GH(M), with the bit bounds of Lemma A.1. Otherwise let $\mathcal{T}$ be the structural family in Theorem 1.1, within the space of matrices positive on this fixed support, and suppose $M \notin \mathcal{T}$. Lemma A.5 supplies a relative neighborhood U of M with $U \cap \mathcal{T} = \varnothing$. Lemma A.4 gives rational $M' \in U$ with Pl-GH(M') $\leq_T$ Pl-GH(M). Since $M' \notin \mathcal{T}$, the algebraic dichotomy makes M' hard and proves hardness of M. Thus the reduction uses a fixed rational matrix in the same relative open complement of $\mathcal{T}$.

On the tractable side, the tensor factors are obtained from fixed matrix entries by arithmetic and fixed square roots. The formulas in Proposition 2.6 and the planar Ising algorithm run in a fixed algebraic extension F' of the original finitely generated output field $F = \mathbb{Q}(M_{ij} : i, j)$, under the adopted representation. The actual partition function belongs to F. If intermediate factor computations use F', descend the answer by

$$Z_M(G) = \frac{\mathrm{Tr}_{F'/F}(Z_M(G))}{[F' : F]}.$$

The extension degree is fixed. Lemma A.1 makes these operations and the descent polynomial in the input size, completing the nonnegative dichotomy in its prescribed output field. □

A.3 Signed transfers and positive vertex weights

We now keep a fixed positive background vector while replacing signed edge weights. The resulting support criterion is sufficient for hardness and will also remove the nonbasic supports in the real positive-vertex-weight theorem.

Corollary A.7 (Support, magnitude and sign in the fixed-real model)**.** *For every fixed real symmetric M, define*

$$P(M)_{ij} = \mathbf{1}_{\{M_{ij}\neq 0\}}, \qquad |M|_{ij} = |M_{ij}|, \qquad S(M)_{ij} = \operatorname{sgn}(M_{ij}), \quad \operatorname{sgn}(0) = 0.$$

The three matrices $P(M), |M|, S(M)$ are jointly available from M with any specified fixed positive real background vector w. In particular,

$$\begin{aligned} \text{Pl-GH}(P(M), w) \leq_T \text{Pl-GH}(|M|, w) \leq_T \text{Pl-GH}(M, w), \\ \text{Pl-GH}(M, w) \equiv_T \text{Pl-GH}(\{|M|, S(M)\}, w). \end{aligned}$$

If $M \neq 0$, the indicators of the smallest nonzero magnitude and of the largest magnitude are jointly available as well.

Proof. For $s > 0$ marked occurrences, a nonzero source product class $\lambda = c^n$ has the replacement products

$$\theta_P(\lambda) = 1, \qquad \theta_{|M|}(\lambda) = \epsilon^n \lambda = |\lambda|, \qquad \theta_S(\lambda) = \epsilon^n = \operatorname{sgn}(\lambda).$$

The sign-tag identity at the start of the section makes all three functions well-defined on merged classes; each preserves original zero terms. For each target, take $C = M$, $C' = P(M), |M|$, or $S(M)$ and the same common occurrence count s in Theorem A.3. The displayed θ verifies its second condition; source zero positions stay zero. Recovering the same A_λ and summing $\sum_\lambda A_\lambda \theta(\lambda)$ supplies each target with background w, without ordering transcendental elements. The reverse mixed reduction uses the parallel-edge identity $M = |M| \circ S(M)$. Neither operation changes the background weights.

For $M \neq 0$, fix its smallest and largest nonzero magnitudes $a_-, a_+ > 0$. The additional replacements are

$$\theta_-(\lambda) = \mathbf{1}_{\{|\lambda| = a_-^s\}}, \qquad \theta_+(\lambda) = \mathbf{1}_{\{|\lambda| = a_+^s\}}.$$

Since every factor magnitude lies in $[a_-, a_+]$, equality at either endpoint occurs exactly when every factor has that endpoint magnitude. These are therefore the desired indicator products. The numbers a_-, a_+ and their positions are fixed data; testing the displayed equalities uses field arithmetic. Original zero entries stay zero. Support alone also follows by even thickening through $M^{\circ 2}$.

These are replacements of edge products, not absolute values of the partition sum. The equivalence concerns the mixed language, for which separate homogeneous algorithms need not suffice. Its extremal indicators select magnitudes, not numerical signed extrema. ☐

Theorem A.8 (Positive vertex-weight removal for fixed real matrices)**.** *Let $R \in \operatorname{Sym}_q(\mathbb{R})$ be fixed, with no zero row and with no two rows proportional over $\mathbb{R}$. For every fixed positive real vector w, in the prescribed fixed-field representation,*

$$\text{Pl-GH}(R) \leq_T \text{Pl-GH}(R, w).$$

Set $D = \operatorname{diag}(w)$. The weighted source jointly supplies the binary matrix D^r for every fixed $r \in \mathbb{Q}$, and hence the unary $u_r : [q] \to \mathbb{R}_{>0}$, $u_r(i) = w_i^r$, by a D^r-loop. Adding one D^{-1}-loop at every vertex simulates the same mixed constraints with unit background weights.

Proof. Set $D = \operatorname{diag}(w)$ and $p = \max(1, q-1)$. Lemma 3.6 and the weighted gadgets in the proof of Lemma 3.7 give

$$\begin{aligned} K &= (RDR)^{\circ p} \succ 0, && T = D^{1/2} K D^{1/2} \succ 0, \\ N_h &= K(DK)^{h-1} = D^{-1/2} T^h D^{-1/2} && (h \geq 1). \end{aligned}$$

The construction is unchanged: a two-edge path includes one internal weight D, parallel copies give K, and chaining gives N_h with terminal weights omitted. Tensor Gram positivity permits odd p and negative entries of K, but excludes every proportional row pair, including negative proportionality.

To justify the same interpolation over the current field, let F contain R, w and all retained constraint entries, and write $T = \sum_{a=1}^{t} \lambda_a P_a$, with distinct $\lambda_a > 0$. The fixed real extension

$$E = F(\sqrt{w_1}, \ldots, \sqrt{w_q}, \lambda_1, \ldots, \lambda_t), \qquad P_a = \prod_{b \neq a} \frac{T - \lambda_b I}{\lambda_a - \lambda_b}$$

is finite algebraic: each eigenvalue satisfies the characteristic polynomial of the fixed matrix T. Roots and embeddings are fixed data; the displayed formula also places every projector in E. On s marked occurrences, the bases are $\prod_a \lambda_a^{n_a}$, $\sum_a n_a = s$, exactly as in Lemma 3.7. Its Vandermonde calculation, with the bit bounds of Lemma A.1, evaluates at $h = 0$ and jointly supplies

$$N_0 = D^{-1/2} \Big(\sum_a P_a \Big) D^{-1/2} = D^{-1}.$$

The matrices $D^{-1/2} P_a D^{-1/2}$ remain in the coefficients. Thus it is the positive spectrum of T, rather than any spectral assumption on R, that justifies this step. No variable field-order test is required. The case $s = 0$ needs no interpolation.

Let $\mathcal{F}$ be the current finite jointly available mixed constraint language containing R. For a $\mathcal{F}$-instance Ω, form Ω° by adding a D^{-1}-loop at every vertex and retaining all original constraints. For every assignment,

$$\prod_{v \in V(G)} w_{\sigma(v)} (D^{-1})_{\sigma(v), \sigma(v)} = 1, \qquad Z_{\mathcal{F} \cup \{D^{-1}\}, w}(\Omega^\circ) = Z_{\mathcal{F}, \mathbf{1}}(\Omega).$$

Thus, including isolated vertices,

$$\begin{aligned} \text{Pl-GH}(\mathcal{F}, \mathbf{1}) &\leq_T \text{Pl-GH}(\mathcal{F} \cup \{D^{-1}\}, w) \\ &\leq_T \text{Pl-GH}(\mathcal{F}, w) \leq_T \text{Pl-GH}(R, w). \end{aligned}$$

For fixed $r \in \mathbb{Q}$, define $D^r = \operatorname{diag}(w_1^r, \ldots, w_q^r)$. Its nonzero products satisfy

$$\prod_{\ell=1}^{s} (D^r)_{i_\ell i_\ell} = \left(\prod_{\ell=1}^{s} (D^{-1})_{i_\ell i_\ell} \right)^{-r}.$$

The off-diagonal zero positions are preserved. Theorem A.3 therefore jointly supplies D^r, using only a further fixed finite algebraic extension. A D^r-loop at a vertex contributes the unary $u_r(i) = (D^r)_{ii} = w_i^r$.

Recovered values in F descend from E by relative trace. The final unweighted value belongs to $F_0 = \mathbb{Q}(R)$; for a mixed constraint problem use its prescribed output field F_0 instead. Apply the general subfield descent in Lemma A.1: if F/F_0 has a transcendental part, descend first to $F_0(Y)$ and then use the matching-coefficient ratio to obtain the F_0 encoding. All these operations have polynomial bit cost with fixed type and exponent sets. $\square$

Corollary A.9 (Support hardness with fixed real edge and positive vertex weights)**.** *Let M be fixed real symmetric and let w be any fixed positive real vector, both represented in the prescribed fixed-field model. If a support component of $P(M)$ is neither a fully looped complete graph, nor a loopless complete bipartite graph, nor a loopless isolated vertex, then* Pl-GH(M, w) *is* #P*-hard. For a zero–one matrix M, these three basic forms still give a polynomial-time algorithm with any fixed positive real vertex weights. Hence the positive-vertex-weight statement of Theorem 7.1 extends from algebraic to fixed real weights. In particular the unweighted support-hardness conclusion holds for arbitrary fixed real signed edge weights.*

Proof. Put $A = P(M)$. Corollary A.7 gives Pl-GH$(A, w) \leq_T$ Pl-GH(M, w) with the same actual color weights. For the hardness assertion, A has a nonbasic component and is therefore nonzero. Delete zero rows and matching columns, then merge identical rows to obtain B. For each retained class I_a put $\bar{w}_a = \sum_{i \in I_a} w_i > 0$. The twin identity proved in Corollary 3.8 is a finite sum identity, so it also holds for these fixed real weights. For an input G, let k be its number of isolated vertices and let G_+ be obtained by deleting them. It gives

$$Z_{B,\bar{w}}(G) = \left(\sum_a \bar{w}_a \right)^k Z_{A,w}(G_+).$$

Every vertex of G_+ is incident to an edge, including a possible loop, so its zero-row colors contribute nothing. The formula uses the quotient's own isolated-vertex factor and divides by no partition value. These are actual twins of A; they are not asserted to be numerical twins of M.

The nonzero distinct rows of a zero–one matrix cannot be proportional over $\mathbb{R}$. Theorem A.8 therefore gives

$$\text{Pl-GH}(B) \leq_T \text{Pl-GH}(B, \bar{w}) \leq_T \text{Pl-GH}(A, w) \leq_T \text{Pl-GH}(M, w).$$

Merging or expanding twins preserves each of the three listed basic support forms: this also follows from their description by pairs of row supports being identical or disjoint. Thus a nonbasic component of A leaves a nonbasic component of B. The unchanged unweighted Theorem 7.1 therefore makes Pl-GH(B) hard. This invokes its established integer-valued endpoint, rather than the positive-real-weight extension being proved.

For a basic zero–one component, the direct algorithms of Theorem 7.1 replace class cardinalities by the corresponding fixed weight sums. A loopless complete bipartite component contributes the sum of its two orientations on a connected bipartite input and zero on a nonbipartite input; fully looped complete components allow independent colors, and zero colors contribute only at isolated input vertices. Support-component contributions add on a connected input, and input-component contributions multiply. The arithmetic and the required output representation have the bounds of Lemma A.1. For the zero matrix $M = 0_{q \times q}$,

$$Z_{M,w}(G) = \begin{cases} 0, & E(G) \neq \varnothing, \\ \left(\sum_{i=1}^q w_i\right)^{|V(G)|}, & E(G) = \varnothing. \end{cases}$$

The empty input has value one.

The same field lemma converts the hard endpoint's answers to integers. The result includes unit and positive algebraic background weights. The transferred obstruction is the nonbasic zero–one support. Signed weightings on the three remaining support forms require further analysis. □

A.4 Weights on the nondegenerate Ising core

The next two lemmas extend the already weighted algebraic Lemmas 8.3 and 9.3 to fixed real constants. Their unweighted structural input is Theorem A.6; arithmetic, root restrictions and weight removal have already been proved in this appendix. All parameters in these lemmas are fixed in this model.

Lemma A.10 (Positive symmetric core)**.** *Let $B = T(\boldsymbol{\rho})$ with $\rho_j > 0$, $\rho_j \neq 1$, and let $u > 0$. Constant u gives a polynomial-time problem. Nonconstant u gives* Pl-GH(B, u) *hard.*

Proof. The matrix B is nonsingular and has no proportional rows. Set $D = \operatorname{diag}(u)$. Theorem A.8 provides unit background weights while retaining the unary $u_i^{1/2}$. Attach $f(i) = \sqrt{u_i}$ at each external endpoint. Identifying one endpoint from each of two decorated edges retains two copies of f at their middle vertex. With unit background there, the entry is

$$S_{ij} = \sum_k (\sqrt{u_i} B_{ik} \sqrt{u_k})(\sqrt{u_k} B_{kj} \sqrt{u_j}) = \sqrt{u_i} \left(\sum_k B_{ik} u_k B_{kj} \right) \sqrt{u_j}.$$

Thus $\widetilde{B} = D^{1/2} B D^{1/2}$ and

$$S = \widetilde{B}^2 = D^{1/2} B D B D^{1/2} > 0, \qquad S \succ 0.$$

These simulations use only a fixed number of jointly available constraint types. Put $Q = B^{\circ 2}$. Every row sum of Q equals $r_0 = \prod_j (1 + \rho_j^2) > 0$, and $S_{ii} = u_i (Qu)_i$. If $U = \max_i u_i > m = \min_i u_i$, with extrema at a, b, strict positivity gives

$$S_{aa} > U m r_0, \qquad S_{bb} < m U r_0.$$

Hence S has nonconstant diagonal. Theorem A.6 makes it hard: positive definiteness forces the rank-one factor of a tractable positive block to have size one, leaving constant diagonal. If $u = c\mathbf{1}$, its only effect is $c^{|V(G)|}$, and

$$Z_{B,u}(G) = c^{|V(G)|} \prod_j Z_{W(\rho_j)}(G).$$

Proposition 2.6 supplies the factor algorithms; Lemma A.1 supplies their real-field bit bounds and output conversion. ☐

Lemma A.11 (Positive bipartite core)**.** *Let $B = cT(\boldsymbol{\rho})$, $c > 0$, with all $\rho_j > 0$ and $\rho_j \neq 1$. Give its row and column colour sets positive weights μ, ν. The bipartite problem is tractable when each vector is constant on its own part, and is hard otherwise. The two constants need not agree.*

Proof. The corresponding symmetric block matrix $A_B = \left(\begin{smallmatrix} 0 & B \\ B^{\mathsf{T}} & 0 \end{smallmatrix}\right)$ has no zero or proportional rows. Put $D_\mu = \operatorname{diag}(\mu)$ and $D_\nu = \operatorname{diag}(\nu)$. Theorem A.8 and the square-root unaries yield, with unit backgrounds, $\widetilde{B} = D_\mu^{1/2} B D_\nu^{1/2}$. At a middle Y vertex, the two retained square-root unaries multiply to ν_j, while an X middle vertex contributes μ_i. Thus, with unit backgrounds,

$$(S_X)_{ii'} = \sqrt{\mu_i \mu_{i'}} \sum_j B_{ij} \nu_j B_{i'j},$$

$$(S_Y)_{jj'} = \sqrt{\nu_j \nu_{j'}} \sum_i B_{ij} \mu_i B_{ij'}.$$

The corresponding same-side matrices are

$$S_X = \widetilde{B} \widetilde{B}^{\mathsf{T}}, \qquad S_Y = \widetilde{B}^{\mathsf{T}} \widetilde{B}.$$

Both are strictly positive and positive definite. Prescribing the colour side at one input root suffices to orient every connected bipartite instance, so these same-side simulations are valid planar reductions from arbitrary planar inputs by Corollary A.2. If either diagonal is nonconstant, Theorem A.6 gives hardness, since a tractable strictly positive positive-definite matrix has constant diagonal.

Suppose both diagonals are constant. With $Q = B^{\circ 2}$, every row and column sum is the same positive value r_0. The equal traces of the two Grams imply that the two diagonal constants coincide:

$$\mu_i(Q\nu)_i = h, \qquad \nu_j(Q^{\mathsf{T}}\mu)_j = h, \qquad h > 0.$$

If μ is nonconstant and $U = \max_i \mu_i$, then $(Q^{\mathsf{T}}\mu)_j < r_0 U$ for every j. Consequently $\nu_j > h/(r_0 U)$ for every j. At a row attaining U this forces $\mu_i(Q\nu)_i > h$, a contradiction. Thus μ is constant; the second displayed equations then force ν to be constant. The converse is the known factor $\mu_0^{|L|}\nu_0^{|R|}$ on an input with prescribed bipartition $L \cup R$, times $Z_B(G)$. For Pl-GH$(A_B, (\mu, \nu))$, sum the two orientations as in Section 9.1; nonbipartite inputs have value zero, and isolated vertices contribute $\sum_i \mu_i + \sum_j \nu_j$. The factor $c^{|E(G)|}$ in $Z_B(G)$ and the Ising algorithms of Proposition 2.6 have the bit bounds of Lemma A.1. $\square$

A.5 The classification with fixed real vertex weights

Theorem A.12 (One fixed positive real vertex-weight vector)**.** *Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ and $w \in \mathbb{R}^q_{>0}$ be fixed in the field representation of this appendix. Form the original identical-row quotient $(R, \bar{w})$ as in Definition 1.2. If every component of $(R, \bar{w})$ has one of the forms (i)–(iii) in Theorem 1.3, then* Pl-GH$(M, w) \in \mathsf{FP}$*; otherwise it is* #P*-hard. Thus the matrix and aggregated-weight conditions of Theorem 1.3 remain a complete dichotomy for arbitrary fixed nonnegative real edge weights and positive real vertex weights.*

Proof. We verify the operations of the algebraic proof in Section 11 in their order of use. Let $F = \mathbb{Q}(M, w)$; include the finitely many retained constraint values in a common fixed field when a mixed instance is used.

Support, components, and normalization. Corollary A.9 excludes nonbasic supports. Its proof uses the established unweighted zero–one theorem and Theorem A.8, independently of the present classification. Corollary A.2 restricts to a component and prescribes a bipartite orientation with the original background weights. The identical-row sum identity (1.5) is a finite sum identity over $\mathbb{R}$.

In a positive component the loop unary is $d_i = M_{ii} > 0$. For $r \in \mathbb{Q}$ fixed and m marked unary occurrences, take $C(i) = d_i$, $C'(i) = d_i^r$ in Theorem A.3. There are no source zeros, and the target product is $\theta(\lambda) = \lambda^r$ for $\lambda = \prod_{a=1}^m d_{i_a} > 0$; hence both conditions hold. Endpoint attachment gives $C_{ij} = M_{ij}/\sqrt{d_i d_j}$ while retaining the moment unaries d_i^t and the original background w. In a bipartite component the parallel-edge gadget gives exactly (11.3): its internal vertex contributes w_j^Y or w_i^X. These positive sums lie in F. Their fixed rational powers satisfy the same product condition, so the normalized cross matrix and the separate left and right moment unaries are jointly available. The required positive roots lie in a fixed finite algebraic extension of F.

Unweighted cores and the restrictions on core weights. The row arguments in Sections 11.2 and 11.3 use only positivity and exact equalities. The resulting symmetric B, or the block matrix A_B, has no zero or proportional rows. Thus Theorem A.8 removes the positive quotient background while retaining the fixed moment constraints. Theorem A.6 gives the real versions of the three alternatives in Lemma 11.1. For the positive unit-diagonal core, $aa^{\mathsf{T}} \otimes T$ and distinct rows force $a = [1]$; for the rectangular core, nonproportional rows and columns force a single amplitude on each side, hence $B = cT$. Distinctness excludes

$W(1)$, so both cores are nonsingular. Lemmas A.10 and A.11 now apply. Their square-root unary constructions and the Gram matrices S, S_X, S_Y have been verified above with unit backgrounds; nonconstant core weights therefore yield hardness.

Finite moments and the original quotient. The positive block uses only $0 \le t < m$, where m is the number of distinct diagonal values. The bipartite block uses only $(m, 0)$ and $(0, n)$ in the finite ranges specified in Section 11.3. All entries of these moment vectors are finite sums of original weights times integer powers of fixed amplitudes. The Vandermonde determinants at distinct positive real amplitudes are nonzero. The same equations (11.2), (11.5) and (11.6) follow by exact field elimination. If an equality fails, one of these finitely many core-weight problems is hard and reduces to the source. Otherwise positivity makes every common amplitude mass positive, and the identification of original row classes gives exactly the forms in Theorem 1.3. This is the finite weighted-moment argument itself, with arithmetic in the fixed real field.

Algorithms and output encodings. The identities (11.7)–(11.9) hold for real entries. The known planar Ising algorithms and their amplitude sums take polynomially many operations; determinant or Pfaffian formulas give polynomial degree and coefficient bounds in the fixed generators. The factor parameters are ratios of entries and fixed positive square roots, as in Proposition 2.6; the class masses are finite sums of the given vertex weights. Thus factor computations take place in a fixed finite algebraic extension E/F, and their final value lies in F. Lemma A.1 converts it to the prescribed F encoding, for example by

$$Z_{M,w}(G) = \frac{\mathrm{Tr}_{E/F}(Z_{M,w}(G))}{[E : F]}.$$

For the hardness reductions, an intermediate unweighted problem may have a smaller output field than F even by a transcendental extension. The general subfield conversion in that lemma applies to these oracle values. All exponents, auxiliary types and fields are fixed; the query number, query size and bit cost are polynomial in the input graph size. This proves both directions in the claimed output model. $\square$

Corollary A.13 (Nonnegative real vertex weights)**.** *Corollary 11.2 holds for fixed $M \in \mathrm{Sym}_q(\mathbb{R}_{\ge 0})$ and $w \in \mathbb{R}^q_{\ge 0}$ in the representation of this appendix. For $S = \{i : w_i > 0\} \ne \varnothing$, the conditions of Theorem A.12 on the identical-row quotient of $(M[S], w|_S)$ give the tractable cases; all other cases are #P-hard. If $S = \varnothing$, the value is 1 on the empty graph and 0 otherwise.*

Proof. Assignments using a color outside S have zero vertex product, so $Z_{M,w}(G) = Z_{M[S],w|_S}(G)$ when $S \ne \varnothing$. Apply Theorem A.12 after this deletion and then form the actual identical-row quotient. The empty-set case follows from the empty-assignment convention. Conversion between $\mathbb{Q}(M[S], w|_S)$ and $\mathbb{Q}(M, w)$ uses Lemma A.1, preserving polynomial bit cost. $\square$

A.6 Combined consequences for signed edge weights

Let $\mathcal{C}_q$ be the nonnegative structural class of Theorem 1.1, now allowing fixed real parameters. The following consequence collects the signed reductions and their support and magnitude hardness criteria.

Corollary A.14 (Real edge weights and positive vertex weights)**.** *Fix $M \in \mathrm{Sym}_q(\mathbb{R})$ and $w \in \mathbb{R}^q_{>0}$ in the prescribed fixed-real representation. Then*

$$\text{Pl-GH}(P(M), w) \le_T \text{Pl-GH}(|M|, w) \le_T \text{Pl-GH}(M, w). \tag{A.1}$$

If a component of $P(M)$ is neither a fully looped complete graph, nor a loopless complete bipartite graph, nor a loopless isolated vertex, then Pl-GH(M, w) *is* #P*-hard. For unit vertex weights there is the additional implication*

$$|M| \notin \mathcal{C}_q \quad \Longrightarrow \quad \text{Pl-GH}(M) \textit{ is } \#\text{P-}\textit{hard.}$$

For $M \in \mathrm{Sym}_q(\{0,1\})$, the three listed component forms give Pl-GH$(M, w) \in \mathsf{FP}$ *for every fixed positive real w; for every other such matrix,* Pl-GH(M, w) *is* #P*-hard.*

Proof. The reductions and the weighted zero–one and support conclusions are Corollaries A.7 and A.9. For the magnitude implication, apply Theorem A.6 to $|M|$ and then (A.1) with $w = \mathbf{1}$. □