A Decidable Dichotomy Theorem on Directed Graph Homomorphisms with Non-negative Weights*

Jin-Yi Cai† Xi Chen‡

April 7, 2010

Abstract

The complexity of graph homomorphism problems has been the subject of intense study. It is a long standing open problem to give a (decidable) complexity dichotomy theorem for the partition function of directed graph homomorphisms. In this paper, we prove a decidable complexity dichotomy theorem for this problem and our theorem applies to all non-negative weighted form of the problem: Given any fixed matrix $\mathbf{A}$ with non-negative entries, the partition function $Z_{\mathbf{A}}(G)$ of directed graph homomorphisms from any directed graph G is *either* tractable in polynomial time *or* #P-hard, depending on the matrix $\mathbf{A}$. The proof of the dichotomy theorem is combinatorial, but involves the definition of an infinite family of graph homomorphism problems. The proof of its decidability is algebraic using properties of polynomials.

*We thank the following colleagues for their interest and helpful comments: Martin Dyer, Alan Frieze, Leslie Goldberg, Richard Lipton, Pinyan Lu, and Leslie Valiant.

†Computer Science Department, University of Wisconsin-Madison

‡Computer Science Department, University of Southern California

1 Introduction

The complexity of counting graph homomorphisms has received much attention recently [8, 5, 3, 1, 7, 13, 6]. The problem can be defined for both *directed* and *undirected* graphs. Most results have been obtained for *undirected* graphs, while the study of complexity of the problem is significantly more challenging for *directed* graphs. In particular, Feder and Vardi showed that the decision problems defined by directed graph homomorphisms are as general as the Constraint Satisfaction Problems (CSPs), and a complexity dichotomy for the former would resolve their long standing dichotomy conjecture for all CSPs [10].

Let G and H be two graphs. We follow the standard definition of graph homomorphisms, where G is allowed to have multiple edges but no self loops; and H can have both multiple edges and self loops.[1] We say $\xi : V(G) \rightarrow V(H)$ is a graph homomorphism from G to H if $\xi(u)\xi(v)$ is an edge in $E(H)$ for all $uv \in E(G)$. Here if H is an *undirected* graph, then G is also an undirected graph; if H is *directed*, then G is also directed. The undirected problem is a special case of the directed one.

For a fixed H, we are interested in the complexity of the following integer function $Z_H(G)$: The input is a graph G, and the output is the number of graph homomorphisms from G to H. More generally, we can define $Z_{\mathbf{A}}(\cdot)$ for any fixed $m \times m$ matrix $\mathbf{A} = (A_{i,j})$:

$$Z_{\mathbf{A}}(G) = \sum_{\xi:V\rightarrow[m]} \prod_{uv\in E} A_{\xi(u),\xi(v)}, \qquad \text{for any directed graph } G = (V, E).$$

Note that the input G is a directed graph in general. However, if $\mathbf{A}$ is a symmetric matrix, then one can always view G as an undirected graph. Moreover, if $\mathbf{A}$ is a $\{0, 1\}$-matrix, then $Z_{\mathbf{A}}(\cdot)$ is exactly $Z_H(\cdot)$, where H is the graph whose adjacency matrix is $\mathbf{A}$.

Graph homomorphisms can express many interesting counting problems over graphs. For example, if we take H to be an undirected graph over two vertices $\{0, 1\}$ with an edge $(0, 1)$ and a loop $(1, 1)$ at 1, then a graph homomorphism from G to H corresponds to a Vertex Cover of G, and $Z_H(G)$ is simply the number of vertex covers of G. As another example, if H is the complete graph on k vertices without self loops, then $Z_H(G)$ is the number of k-Colorings of G. In [11], Freedman, Lovász, and Schrijver characterized what graph functions can be expressed as $Z_{\mathbf{A}}(\cdot)$.

For increasingly more general families $\mathcal{C}$ of matrices $\mathbf{A}$, the complexity of $Z_{\mathbf{A}}(\cdot)$ has been studied and *dichotomy* theorems have been proved. A *dichotomy* theorem for a given family $\mathcal{C}$ of matrices $\mathbf{A}$ states that for any $\mathbf{A} \in \mathcal{C}$, the problem of computing $Z_{\mathbf{A}}(\cdot)$ is either *in polynomial time* or *#P-hard*. A *decidable* dichotomy theorem requires that the dichotomy criterion is computably decidable: There is a finite-time classification algorithm that, given any $\mathbf{A} \in \mathcal{C}$, decides whether $Z_{\mathbf{A}}(\cdot)$ is in polynomial time or #P-hard. Most results have been obtained for undirected graphs.

Symmetric matrices A, and $Z_{\mathbf{A}}(G)$ over undirected graphs G:
In [14, 15], Hell and Nešetřil showed that given any symmetric {0,1} matrix $\mathbf{A}$, *deciding* whether $Z_{\mathbf{A}}(G) > 0$ is either in P or NP-complete. Then Dyer and Greenhill [8] showed that given any symmetric {0,1} matrix $\mathbf{A}$, the problem of computing $Z_{\mathbf{A}}(\cdot)$ is either in P or #P-complete. Bulatov and Grohe generalized their result to all non-negative symmetric matrices $\mathbf{A}$ [5].[2] They obtained an elegant dichotomy theorem which basically says that $Z_{\mathbf{A}}(\cdot)$ is in P if every *block* of $\mathbf{A}$ has rank at most one, and is #P-hard otherwise. In [13] Goldberg, Grohe, Jerrum and Thurley proved a beautiful dichotomy for all symmetric real matrices. Finally, a dichotomy theorem for all symmetric complex matrices was recently proved by Cai, Chen, and

[1] However, our results are actually stronger in that our tractability result allows for loops in G, while our hardness result holds for G without loops.

[2] More exactly, they proved a dichotomy theorem for all matrices $\mathbf{A}$ in which every entry $A_{i,j}$ is a non-negative algebraic number. Our result in this paper applies similarly to all non-negative algebraic numbers. It can be generalized to computable numbers, which will be given in an appropriate form in a field of a finite transcendence degree. But we will not discuss the details of this, and use $\mathbb{R}$ and $\mathbb{C}$ to denote these real and complex numbers.

Lu [6]. We remark that all the dichotomy theorems for symmetric matrices $\mathbf{A}$ above are *polynomial-time* decidable, meaning that given any matrix $\mathbf{A}$, one can decide in polynomial time (in the input size of $\mathbf{A}$) whether $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard.

General matrices $\mathbf{A}$, and $Z_{\mathbf{A}}(G)$ over directed graphs G:
In a paper that won the best paper award at ICALP in 2006, Dyer, Goldberg, and Paterson [7] proved a dichotomy theorem for directed graph homomorphism problems $Z_H(\cdot)$, but restricted to directed *acyclic* graphs H. They introduced the concept of *Lovász-goodness* and proved that $Z_H(\cdot)$ is in P if the graph H is *layered* [3] and *Lovász-good*, and is #P-hard otherwise. The property of Lovász-goodness turns out to be polynomial-time decidable.

In [1], Bulatov presented a sweeping dichotomy theorem for all counting Constraint Satisfaction Problems. Recently, Dyer and Richerby [9] obtained an alternative proof. The dichotomy theorem of Bulatov then implies a dichotomy for $Z_H(\cdot)$ over all directed graphs H. However, it is rather unclear whether this dichotomy theorem is decidable or not. The criterion [4] requires one to check a condition on an infinitary object (see Appendix H for details). This situation remains the same for the Dyer-Richerby proof in [9]. The decidability of the dichotomy was then left as an open problem in [2].

In this paper, we prove a dichotomy theorem for the family of all non-negative real matrices $\mathbf{A}$. We show that for every fixed $m \times m$ non-negative matrix $\mathbf{A}$, the problem of computing $Z_{\mathbf{A}}(\cdot)$ is either in P or #P-hard. Moreover, our dichotomy criterion is *decidable*: we give a finite-time algorithm which, given any non-negative matrix $\mathbf{A}$, decides whether $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard. In particular, for the family of $\{0,1\}$ matrices, our result gives an alternative dichotomy criterion to that of Bulatov [2] and Dyer-Richerby [9], which is decidable. [5]

The main difficulty we encountered in obtaining the dichotomy is due to the *abundance* of new intricate but tractable cases, when moving from acyclic graphs to general directed graphs. For example, the graph H does not have to be layered for the problem $Z_H(\cdot)$ to be tractable (see Figure 1 in Appendix A for an interesting example). Because of the generality of directed graphs, it seems impossible to have a simply stated criterion (e.g., Lovasz-goodness, as was used in the acyclic case [7]) which is both powerful enough to completely characterize all the tractable cases and also easy to check. However, we manage to find a dichotomy criterion as well as a finite-time algorithm to decide whether $\mathbf{A}$ satisfies it or not.

In particular, the dichotomy theorem of Dyer, Goldberg and Paterson [7] for the acyclic case fits into our framework as follows. In our dichotomy, we start from $\mathbf{A}$ and then define, in each round, a (possibly infinite) set of new matrices. The size of the matrices defined in round $i+1$ is strictly smaller than that of round i (so there could be at most m rounds). The dichotomy then is that $Z_{\mathbf{A}}(\cdot)$ is in P if and only if every *block* of any matrix defined in the process above is of rank 1 (see Section 1.1 and 1.2 for details). For the special acyclic case treated by Dyer, Goldberg, and Paterson [7], let $\mathbf{A}$ be the adjacency matrix of H which is acyclic and has k layers, then at most k rounds are necessary to reach a conclusion about whether $Z_{\mathbf{A}}(\cdot) = Z_H(\cdot)$ is in P or #P-hard. However, if H has k layers but is not acyclic (i.e., there are edges from layer k back to layer 1), deciding whether $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard becomes much harder in the sense that we might need $\gg k$ rounds to reach a conclusion.

After we circulated a draft of this paper, Goldberg informed us that she and coauthors [4] found a reduction from *weighted* counting CSPs with non-negative *rational* weights to the 0-1 dichotomy theorem

[3] A directed acyclic graph is *layered* if one can partition its vertices into k sets $V_1, \ldots, V_k$, for some $k \geq 1$, such that every edge goes from V_i to V_{i+1} for some $i : 1 \leq i < k$.

[4] A dichotomy criterion is a well-defined mathematical property over the family of matrices $\mathbf{A}$ being considered such that $Z_{\mathbf{A}}(\cdot)$ is in P if $\mathbf{A}$ has this property; and is #P-hard otherwise.

[5] Both our dichotomy criterion (when specialized to the $\{0,1\}$ case) and the one of Bulatov characterize $\{0,1\}$ matrices $\mathbf{A}$ with $Z_{\mathbf{A}}(\cdot)$ in P and thus, they must be equivalent, i.e., $\mathbf{A}$ satisfies our criterion if and only if it satisfies the one of Bulatov. As a corollary, our result also implies a finite-time algorithm for checking the dichotomy criterion of Bulatov [2] (as well as the version of Dyer and Richerby [9]) for the case of $\{0,1\}$ matrices $\mathbf{A}$.

of Bulatov [2]. However, the combined result still only works for non-negative rational weights and more importantly, the dichotomy is not known to be decidable.

1.1 Intuition of the Dichotomy: Domain Reduction

Let $\mathbf{A}$ be the $m \times m$ non-negative matrix being considered, and $G = (V, E)$ be the input directed graph. Before giving a more formal sketch of the proofs, we use a simple example to illustrate one of the most important ideas of this work: *domain reduction*.

For this purpose we also need to introduce the concept of *labeled directed graphs*. A labeled directed graph $\mathcal{G}$ over domain $[m] = \{1, 2, \ldots, m\}$ is a directed graph, in which every directed edge e is labeled with an $m \times m$ matrix $\mathbf{A}^{[e]}$; and every vertex v is labeled with an m-dimensional vector $\mathbf{w}^{[v]}$. Then the partition function of $\mathcal{G}$ is defined as

$$Z(\mathcal{G}) = \sum_{\xi: V \to [m]} \prod_{v \in V} w^{[v]}_{\xi(v)} \prod_{uv \in E} A^{[uv]}_{\xi(u), \xi(v)}.$$

In particular, we have $Z_{\mathbf{A}}(G) = Z(\mathcal{G}_0)$ where $\mathcal{G}_0$ has the same graph structure as G; every edge of $\mathcal{G}_0$ is labeled with the same $\mathbf{A}$; and every vertex of $\mathcal{G}_0$ is labeled with $\mathbf{1}$, the m-dimensional all-1 vector.

Roughly speaking, starting from the input G, we build (in polynomial time) a *finite* sequence of new *labeled* directed graphs $\mathcal{G}_0, \mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_h$ one by one. $\mathcal{G}_{k+1}$ is constructed from $\mathcal{G}_k$ by using the domain reduction method which we are going to describe next. On the one hand, the *domains* of these labeled graphs *shrink* along with k. This means, the size of the edge weight matrices associated with the edges of $\mathcal{G}_k$ (or equivalently, the dimension of the vectors associated with the vertices of $\mathcal{G}_k$) strictly decreases along with k. On the other hand, we have $Z(\mathcal{G}_{k+1}) = Z(\mathcal{G}_k)$ for all $k \geq 0$ and thus,

$$Z_{\mathbf{A}}(G) = Z(\mathcal{G}_0) = \ldots = Z(\mathcal{G}_h).$$

Since the domain size decreases monotonically, the number of graphs $\mathcal{G}_k$ in this sequence is at most m. To prove our dichotomy theorem, we show that, either something bad happens which forces us to stop the domain reduction process, in which case we show that $Z_{\mathbf{A}}(\cdot)$ is #P-hard; or we can keep reducing the domain size until the computation becomes trivial, in which case we show that $Z_{\mathbf{A}}(\cdot)$ is in P.

We say a matrix $\mathbf{A}$ is *block-rank*-1 if one can (separately) permute the rows and columns of $\mathbf{A}$ to get a block diagonal matrix in which every block is of rank at most 1. If $\mathbf{A}$ is not block-rank-1 we can easily show that $Z_{\mathbf{A}}(\cdot)$ is #P-hard, using the dichotomy of Bulatov and Grohe [5] for symmetric non-negative matrices (see Lemma 1). So without loss of generality, we assume $\mathbf{A}$ is block-rank-1. For example, let $\mathbf{A}$ be the 8×8 block-rank-1 non-negative matrix in Figure 2 in Appendix A with 16 positive entries, then we use $\mathcal{T} = \{(A_1, B_1), (A_2, B_2), (A_3, B_3), (A_4, B_4)\}$ to denote the *block structure* of $\mathbf{A}$, where

$$\forall s \in [4],\ A_s = \{2s-1, 2s\},\ B_1 = \{1, 3\},\ B_2 = \{5, 7\},\ B_3 = \{2, 4\}\ \text{ and }\ B_4 = \{6, 8\},$$

so that $A_{i,j} > 0$ if and only if $i \in A_s$ and $j \in B_s$, for some $s \in [4]$. Because $\mathbf{A}$ is block-rank-1, there also exist two 8-dimensional *positive* vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that

$$A_{i,j} = \alpha_i \cdot \beta_j, \quad \text{for all } (i, j) \text{ such that } i \in A_s \text{ and } j \in B_s \text{ for some } s \in [4].$$

Now let $G = (V, E)$ be the directed graph in Figure 3 (Appendix A) with $|V| = |E| = 6$. We illustrate the domain reduction process by constructing the first labeled directed graph $\mathcal{G}_1$ in the sequence as follows. To simplify the presentation, we let $\mathbf{y} \in [8]^6$ (instead of $\xi : V \to [8]$) denote an assignment, where $y_i \in [8]$ denotes the value of vertex i in Figure 3 for every $i \in [6]$.

First, let $\mathbf{y} \in [8]^6$ be any assignment with a nonzero weight: $A_{y_i,y_j} > 0$ for every edge $ij \in E$. Because $\mathbf{A}$ has the block structure $\mathcal{T}$, for every $ij \in E$, there exists a unique index $s \in [4]$ such that $y_i \in A_s$ and $y_j \in B_s$. This inspires us to introduce a new variable $x_\ell \in [4]$ for each edge $e_\ell \in E$, $\ell \in [6]$ (as shown in Figure 3). For every possible assignment of $\mathbf{x} = (x_1, x_2, \ldots, x_6) \in [4]^6$, we use $Y[\mathbf{x}]$ to denote the set of all possible assignments $\mathbf{y} \in [8]^6$ such that for every $e_\ell = ij$, $y_i \in A_{x_\ell}$ and $y_j \in B_{x_\ell}$. Now we have

$$Z_{\mathbf{A}}(G) = \sum_{\mathbf{x}\in[4]^6} \sum_{\mathbf{y}\in Y[\mathbf{x}]} \mathrm{wt}(\mathbf{y}), \quad \text{where } \mathrm{wt}(\mathbf{y}) = \prod_{ij\in E} A_{y_i,y_j}.$$

Second, we further simplify the sum above by noticing that if $x_2 \neq x_3$ in $\mathbf{x}$, then $Y[\mathbf{x}]$ must be empty because the two edges e_2 and e_3 share the same tail in G. In general, we only need to sum over the case when $x_1 = x_2 = x_3$ and $x_4 = x_5$, since otherwise the set $Y[\mathbf{x}]$ is empty. As a result,

$$Z_{\mathbf{A}}(G) = \sum_{\substack{x_1=x_2=x_3\\ x_4=x_5\\ x_6}} \sum_{\mathbf{y}\in Y[\mathbf{x}]} \mathrm{wt}(\mathbf{y}).$$

The advantage of introducing x_ℓ, $\ell \in [6]$, is that, once $\mathbf{x}$ is fixed, one can always decompose A_{y_i,y_j} as a product $\alpha_{y_i} \cdot \beta_{y_j}$, for all $\mathbf{y} \in Y[\mathbf{x}]$ and all $ij \in E$, since $\mathbf{y}$ belonging to $Y[\mathbf{x}]$ guarantees that (y_i, y_j) falls inside one of the four blocks of $\mathbf{A}$. This allows us to greatly simplify $\mathrm{wt}(\mathbf{y})$: If $\mathbf{y} \in Y[\mathbf{x}]$, then

$$\mathrm{wt}(\mathbf{y}) = A_{y_1,y_3} \cdot A_{y_1,y_2} \cdot A_{y_2,y_3} \cdot A_{y_3,y_4} \cdot A_{y_3,y_5} \cdot A_{y_5,y_6} = \alpha_{y_1}\beta_{y_3}\alpha_{y_1}\beta_{y_2}\alpha_{y_2}\beta_{y_3}\alpha_{y_3}\beta_{y_4}\alpha_{y_3}\beta_{y_5}\alpha_{y_5}\beta_{y_6}.$$

Also notice that $Y[\mathbf{x}]$, for any $\mathbf{x}$, is a direct product of subsets of $[8]$: $\mathbf{y} \in Y[\mathbf{x}]$ if and only if

$$y_1 \in L_1 = A_{x_1},\ y_2 \in L_2 = A_{x_3} \cap B_{x_1} = A_{x_1} \cap B_{x_1},\ y_3 \in L_3 = A_{x_4} \cap A_{x_5} \cap B_{x_2} \cap B_{x_3} = A_{x_4} \cap B_{x_1},$$
$$y_4 \in L_4 = B_{x_4},\ y_5 \in L_5 = A_{x_6} \cap B_{x_4},\ y_6 \in L_6 = B_{x_6}.$$

As a result, $Z_{\mathbf{A}}(G)$ now becomes

$$Z_{\mathbf{A}}(G) = \sum_{x_1,x_4,x_6} \sum_{y_i\in L_i,\, i\in[6]} \big((\alpha_{y_1})^2\alpha_{y_2}\beta_{y_2}\big) \cdot \big((\alpha_{y_3})^2(\beta_{y_3})^2\big) \cdot \beta_{y_4} \cdot (\alpha_{y_5}\beta_{y_5}) \cdot \beta_{y_6}. \tag{1}$$

Finally we construct the following *labeled* directed graph $\mathcal{G}_1$ over domain $[4]$. There are three vertices a, b and c, which correspond to x_1, x_4 and x_6, respectively; and there are only two directed edges ab and bc. We construct the vertex/edge weights as follows. The vertex weight vector of a is

$$w^{[a]}_\ell = \sum_{y_1\in A_\ell,\ y_2\in A_\ell\cap B_\ell} (\alpha_{y_1})^2\alpha_{y_2}\beta_{y_2}, \quad \text{for every } \ell \in [4];$$

the vertex weights of b and c are the same:

$$w^{[b]}_\ell = w^{[c]}_\ell = \textstyle\sum_{y\in B_\ell} \beta_y, \quad \text{for every } \ell \in [4].$$

The edge weight matrix $\mathbf{C}^{[ab]}$ of ab is

$$C^{[ab]}_{k,\ell} = \sum_{y_3\in B_k\cap A_\ell} (\alpha_{y_3})^2(\beta_{y_3})^2, \quad \text{for all } k, \ell \in [4];$$

and the edge weight matrix $\mathbf{C}^{[bc]}$ of bc is

$$C^{[bc]}_{k,\ell} = \sum_{y_5\in B_k\cap A_\ell} \alpha_{y_5}\beta_{y_5}, \quad \text{for all } k, \ell \in [4].$$

Using (1) and the definition of $Z(\mathcal{G}_1)$, it is easy to verify that $Z_{\mathbf{A}}(G) = Z(\mathcal{G}_1)$ and thus, we reduced the domain size of the problem from 8 (which is the number of rows and columns in $\mathbf{A}$), to 4 (which is the number of blocks in $\mathbf{A}$). However, we also paid a high price. Two issues are worth pointing out here:

1. Unlike in $Z_{\mathbf{A}}(G)$, different edges in $\mathcal{G}_1$ have *different* edge weight matrices in general. For example, the matrices associated with ab and bc are clearly different, for general $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. Actually, the set of matrices that may appear as an edge weight of $\mathcal{G}_1$, constructed from *all possible* directed graphs G after one round of domain reduction, is *infinite* in general.

2. Unlike in $Z_{\mathbf{A}}(G)$, we have to introduce vertex weights in $\mathcal{G}_1$. Similarly, vertices may have different vertex weight vectors, and the set of vectors that may appear as a vertex weight of $\mathcal{G}_1$, constructed from *all possible* G after one round of domain reduction, is *infinite* in general.

It is also worth noticing that, even if the matrix $\mathbf{A}$ we start with is $\{0,1\}$, the edge and vertex weights of $\mathcal{G}_1$ immediately become *rational* right after the first round of domain reduction and we have to deal with rational weights afterwards. So $\{0,1\}$-matrices are not that special under this framework.

These two issues cause us a lot of trouble because we need to carry out the domain reduction process for several times, until the computation becomes trivial. However, the reduction process above crucially used the assumption that $\mathbf{A}$ is block-rank-1 (otherwise, one cannot replace $A_{i,j}$ with $\alpha_i \cdot \beta_j$). Therefore, there is no way to continue this process if some edge weight matrix in $\mathcal{G}_1$ is not block-rank-1. To deal with this case, we show that if this happens for some G, then $Z_{\mathbf{A}}(\cdot)$ is #P-hard. Informally, we have

Theorem 1 (Informal)**.** *For any G, if one of the edge matrices in $\mathcal{G}_k$* (constructed from G after k rounds of domain reductions)*, for some $k \geq 1$, is* not *block-rank-*1*, then $Z_{\mathbf{A}}(\cdot)$ is* #P-hard*.*

The proof of Theorem 1 for $k = 1$ is relatively straight forward, because every edge weight matrix in G is $\mathbf{A}$. However, due to the two issues mentioned earlier, the edge weights and vertex weights of $\mathcal{G}_1$ are drawn from infinite sets in general, and even proving Theorem 1 for $k = 2$ is highly non-trivial.

Even with Theorem 1 which essentially gives us a dichotomy theorem for all non-negative matrices, it is still unclear whether the dichotomy is *decidable* or not. The difficulty is that, to decide whether $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard, we need to check infinitely many matrices (all the edge weight matrices that appear in the domain reduction process, from *all possible* directed graphs G) and to see whether all of them are block-rank-1. To overcome this, we give an algebraic proof using properties of polynomials. We manage to show that it is not necessary to check these matrices one by one, but only need to check whether or not the entries of $\mathbf{A}$ satisfy finitely many polynomial constraints.

1.2 Proof Sketch

Without loss of generality, we assume $\mathbf{A}$ is an $m \times m$ block-rank-1 matrix. To show that $Z_{\mathbf{A}}(\cdot)$ is either in P or #P-hard, we *define* from $\mathbf{A}$ a finite sequence of pairs:

$$(\mathfrak{X}_0, \mathfrak{Y}_0), (\mathfrak{X}_1, \mathfrak{Y}_1), \ldots, (\mathfrak{X}_h, \mathfrak{Y}_h), \qquad \text{for some } h : 0 \leq h < m,$$

where $\mathfrak{X}_0 = \{\mathbf{1}\}$, $\mathfrak{Y}_0 = \{\mathbf{A}\}$ and $\mathbf{1}$ denotes the m-dimensional all-1 vector. Each pair $(\mathfrak{X}_k, \mathfrak{Y}_k)$, $k \in [h]$, is defined from $(\mathfrak{X}_{k-1}, \mathfrak{Y}_{k-1})$. Roughly speaking, $\mathfrak{Y}_k$ (resp. $\mathfrak{X}_k$) is the set of all edge matrices (resp. vertex vectors) that could appear in $\mathcal{G}_k$, after k rounds of domain reductions. There also exist positive integers $m = m_0 > m_1 > \ldots > m_h \geq 1$ such that every $\mathfrak{Y}_k$, $k \in [h]$, is a set of $m_k \times m_k$ non-negative matrices; and every $\mathfrak{X}_k$, $k \in [h]$, is a set of m_k-dimensional non-negative vectors. Although both sets $\mathfrak{X}_k$ and $\mathfrak{Y}_k$ are *infinite* in general (which is the reason why we used the word "*define*" instead of "*construct*"), the definition of $(\mathfrak{X}_k, \mathfrak{Y}_k)$ guarantees the following two properties:

1. For every $k \in [h]$, matrices in $\mathfrak{Y}_k$ share the same *structure*: $\forall\, \mathbf{B}, \mathbf{B}' \in \mathfrak{Y}_k$, $B_{i,j} > 0 \Leftrightarrow B'_{i,j} > 0$;
2. Every matrix $\mathbf{B}$ in $\mathfrak{Y}_h$ is a *permutation* matrix.

The definition of $(\mathfrak{X}_k, \mathfrak{Y}_k)$ from $(\mathfrak{X}_{k-1}, \mathfrak{Y}_{k-1})$ can be found in Appendix C. In Appendix F, we prove that for every $k \in [h]$, if $\mathbf{B} \in \mathfrak{Y}_k$, then the problem of computing $Z_{\mathbf{B}}(\cdot)$ is polynomial-time reducible to the computation of $Z_{\mathbf{A}}(\cdot)$. From this, we obtain the hardness part of our dichotomy theorem: If for some $k \in [h]$, there exists a matrix $\mathbf{B} \in \mathfrak{Y}_k$ such that $\mathbf{B}$ is not block-rank-1, then $Z_{\mathbf{A}}(\cdot)$ is #P-hard.

Now we assume that all matrices in $\mathfrak{Y}_k$, $k \in [h]$, are block-rank-1. To finish the proof, we only need to show that if this is true, then $Z_{\mathbf{A}}(\cdot)$ is indeed in P. To this end, we use the *domain reduction* process to construct a sequence of *labeled* directed graphs $\mathcal{G}_0, \mathcal{G}_1, \ldots, \mathcal{G}_h$ such that

1. $Z(\mathcal{G}_0) = Z_{\mathbf{A}}(G)$ and $Z(\mathcal{G}_{k+1}) = Z(\mathcal{G}_k)$ for all $k : 0 \le k < h$; and
2. For every k, we have $\mathbf{A}^{[e]} \in \mathfrak{Y}_k$ for all edges e in $\mathcal{G}_k$ and $\mathbf{w}^{[v]} \in \mathfrak{X}_k$ for all vertices v in $\mathcal{G}_k$.

This sequence can be constructed in polynomial time, because the construction of $\mathcal{G}_{k+1}$ from $\mathcal{G}_k$ can be done very efficiently as described in Section 1.1, and also because the number of graphs in the sequence is at most m. By the two properties above, we have $Z_{\mathbf{A}}(G) = Z(\mathcal{G}_h)$; and every edge weight matrix $\mathbf{A}^{[e]}$ in $\mathcal{G}_h$ is a *permutation* matrix. As a result, we can compute $Z_{\mathbf{A}}(G)$ in polynomial time, since $Z(\mathcal{G}_h)$ can be computed very efficiently.

This finishes the proof of our dichotomy theorem: Given any non-negative matrix $\mathbf{A}$, the problem of computing $Z_{\mathbf{A}}(\cdot)$ is either in polynomial time or #P-hard. Moreover, to decide which case it is, one only needs to check whether the matrices in $\mathfrak{Y}_k$, $k : 0 \le k \le h$, satisfy the following condition:

The Block-Rank-1 Condition: Every matrix $\mathbf{B} \in \mathfrak{Y}_k$, $k : 0 \le k \le h$, is block-rank-1.

However, all the sets $\mathfrak{Y}_k$, $k \in [h]$, are *infinite* in general, so we cannot afford to check the matrices one by one. Instead, we express the block-rank-1 condition as a finite collection of polynomial constraints over $\mathfrak{Y}_k$. The way $(\mathfrak{X}_k, \mathfrak{Y}_k)$ is defined from $(\mathfrak{X}_{k-1}, \mathfrak{Y}_{k-1})$ allows us to show that, to check whether every matrix in $\mathfrak{Y}_k$ (or every vector in $\mathfrak{X}_k$) satisfies a certain polynomial constraint, one only needs to check a finitely many polynomial constraints for $(\mathfrak{X}_{k-1}, \mathfrak{Y}_{k-1})$. As a consequence, to check whether $\mathfrak{Y}_k$, $k \in [h]$, satisfies the block-rank-1 condition, one only needs to check a finitely many polynomial constraints for $(\mathfrak{X}_0, \mathfrak{Y}_0)$. Since $\mathfrak{X}_0 = \{\mathbf{1}\}$ and $\mathfrak{Y}_0 = \{\mathbf{A}\}$ are both finite, this can be done in a finite number of steps.

2 Preliminaries

We say $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ is a *labeled directed graph* over $[m] = \{1, \ldots, m\}$ for some positive integer m, if

1. $G = (V, E)$ is a directed graph (which may have parallel edges but no self-loops);
2. Every $v \in V$ is labeled with an m-dimensional non-negative vector $\mathcal{V}(v) \in \mathbb{R}_+^m$ as its vertex weight;
3. Every $uv \in E$ is labeled with an $m \times m$ non-negative matrix $\mathcal{E}(uv) \in \mathbb{R}_+^{m \times m}$ as its edge weight.

Let $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ be a labeled directed graph, where $G = (V, E)$. For each $v \in V$, we use $\mathbf{w}^{[v]} = \mathcal{V}(v)$ to denote its vertex weight vector; and for each $uv \in E$, we use $\mathbf{C}^{[uv]} = \mathcal{E}(uv)$ to denote its edge weight matrix. Then we define $Z(\mathcal{G})$ as follows:

$$Z(\mathcal{G}) = \sum_{\xi : V \to [m]} \text{wt}(\mathcal{G}, \xi), \quad \text{where } \text{wt}(\mathcal{G}, \xi) = \prod_{v \in V} w^{[v]}_{\xi(v)} \prod_{uv \in E} C^{[uv]}_{\xi(u), \xi(v)} \text{ denotes the } \textit{weight} \text{ of } \xi.$$

Let $\mathbf{C}$ be an $m \times m$ non-negative matrix. We are interested in the complexity of $Z_{\mathbf{C}}(\cdot)$:

$$Z_{\mathbf{C}}(G) = Z(\mathcal{G}), \qquad \text{for any directed graph } G = (V, E),$$

where $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ is the labeled directed graph with $\mathcal{V}(v) = \mathbf{1} \in \mathbb{R}_+^m$ for all $v \in V$; and $\mathcal{E}(uv) = \mathbf{C}$ for all edges $uv \in E$.

Definition 1 (Pattern and block pattern)**.** *We say $\mathcal{P}$ is an $m \times m$* pattern *if $\mathcal{P} \subseteq [m] \times [m]$. $\mathcal{P}$ is said to be* trivial *if $\mathcal{P} = \emptyset$. A non-negative $m \times m$ matrix $\mathbf{C}$ is of* pattern *$\mathcal{P}$, if for all $i, j \in [m]$, we have $C_{i,j} > 0$ if and only if $(i, j) \in \mathcal{P}$. $\mathbf{C}$ is also called a $\mathcal{P}$-*matrix*. We say $\mathcal{T}$ is an $m \times m$* block pattern *if*

1. *$\mathcal{T} = \big\{(A_1, B_1), \ldots, (A_r, B_r)\big\}$ for some $r \geq 0$;*
2. *$A_i \subseteq [m]$, $A_i \neq \emptyset$, $B_i \subseteq [m]$ and $B_i \neq \emptyset$ for all $i \in [r]$; and*
3. *$A_i \cap A_j = B_i \cap B_j = \emptyset$, for all $i \neq j \in [r]$.*

A block pattern $\mathcal{T}$ is said to be trivial *if $\mathcal{T} = \emptyset$. A block pattern $\mathcal{T}$ naturally* defines *a pattern $\mathcal{P}$, where*

$$\mathcal{P} = \big\{ (i, j) \,\big|\, \exists\, k \in [r] \text{ such that } i \in A_k \text{ and } j \in B_k \big\}.$$

We also say $\mathcal{P}$ is consistent with *$\mathcal{T}$. Finally, we say a non-negative $m \times m$ matrix $\mathbf{C}$ is of* block pattern *$\mathcal{T}$, if $\mathbf{C}$ is of pattern $\mathcal{P}$ defined by $\mathcal{T}$. $\mathbf{C}$ is also called a $\mathcal{T}$-matrix.*

Definition 2. *We say an $m \times m$ non-negative matrix $\mathbf{C}$ is* block-rank-1 *if*

1. *Either $\mathbf{C} = \mathbf{0}$ is the zero matrix (and is of block pattern $\mathcal{T} = \emptyset$); or*
2. *$\mathbf{C}$ is of block pattern $\mathcal{T}$, for some $m \times m$ block pattern $\mathcal{T} = \{(A_1, B_1), \ldots, (A_r, B_r)\}$ with $r \geq 1$; and for every $k \in [r]$, the sub-matrix of $\mathbf{C}$ induced by A_k and B_k is (exactly) rank 1.*

Let $\mathbf{C}$ be a non-negative block-rank-1 *matrix of block pattern $\mathcal{T}$, then there exists a* unique *pair $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ of non-negative m-dimensional vectors such that*

1. *For every $i \in [m]$, $\alpha_i > 0 \iff i \in \bigcup_{k \in [r]} A_k$; and $\beta_i > 0 \iff i \in \bigcup_{k \in [r]} B_k$;*
2. *$C_{i,j} = \alpha_i \cdot \beta_j$ for all $i, j \in [m]$ such that $C_{i,j} > 0$; and*
3. *$\sum_{j \in A_k} \alpha_j = 1$, for all $k \in [r]$.*

The pair $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is called the (vector) representation *of $\mathbf{C}$. Note that we have $\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0}$ when $\mathbf{C} = \mathbf{0}$.*

It is clear that $\mathcal{T}$ and $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ together uniquely determine a non-negative block-rank-1 matrix $\mathbf{C}$. The following lemma concerns the complexity of $Z_{\mathbf{C}}(\cdot)$. The proof can be found in Appendix B.

Lemma 1. *If $\mathbf{C}$ is not block-rank-1, then $Z_{\mathbf{C}}(\cdot)$ is #P-hard.*

Let $\mathcal{T}$ be an $m \times m$ *non-trivial* block pattern, where $\mathcal{T} = \{(A_1, B_1), \ldots, (A_r, B_r)\}$ for some $r \geq 1$. It defines the following $r \times r$ pattern $\mathcal{P} = \texttt{gen}(\mathcal{T})$: For all $i, j \in [r]$, $(i, j) \in \mathcal{P}$ if and only if $B_i \cap A_j \neq \emptyset$. We also define $\texttt{gen-block}(\mathcal{T})$ as follows:

1. If $\mathcal{P} = \texttt{gen}(\mathcal{T})$ is consistent with a block pattern, denoted by $\mathcal{T}'$, then $\texttt{gen-block}(\mathcal{T}) = \mathcal{T}'$;
2. Otherwise, we set $\texttt{gen-block}(\mathcal{T}) = \textit{false}$.

We note that $\mathcal{P} = \texttt{gen}(\mathcal{T})$ could be *trivial* even if $\mathcal{T}$ is *non-trivial*.

Next, we introduce a generalized version of $Z_{\mathbf{C}}(\cdot)$. Let $m \geq 1$ and $(\mathfrak{P}, \mathfrak{Q})$ be a pair in which

1. $\mathfrak{P}$ is a *finite* and nonempty set of non-negative m-dimensional vectors with $\mathbf{1} \in \mathfrak{P}$; and
2. $\mathfrak{Q}$ is a *finite* and nonempty set of $m \times m$ non-negative matrices.

We then use $Z(\cdot)$ to define the function $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ as follows:

$$Z_{\mathfrak{P},\mathfrak{Q}}(\mathcal{G}) = Z(\mathcal{G}),$$

where $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ is a labeled directed graph with $\mathcal{V}(v) \in \mathfrak{P}$ for any vertex $v \in V(G)$; and $\mathcal{E}(uv) \in \mathfrak{Q}$ for any edge $uv \in E(G)$. As an example, $Z_{\mathbf{C}}(\cdot)$ is exactly $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ with $\mathfrak{P} = \{\mathbf{1}\}$ and $\mathfrak{Q} = \{\mathbf{C}\}$.

Finally, let $m \geq 1$ and $(\mathfrak{X}, \mathfrak{Y})$ and $(\mathfrak{X}', \mathfrak{Y}')$ be two pairs such that:

1. $\mathfrak{X}$ and $\mathfrak{X}'$ are two nonempty (and possibly infinite) sets of non-negative m-dimensional vectors with $\mathbf{1} \in \mathfrak{X}$ and $\mathbf{1} \in \mathfrak{X}'$; and
2. $\mathfrak{Y}$ and $\mathfrak{Y}'$ are two nonempty (and possibly infinite) sets of non-negative $m \times m$ matrices.

Definition 3 (Reduction)**.** *We say* $(\mathfrak{X}', \mathfrak{Y}')$ *is* polynomial-time reducible *to* $(\mathfrak{X}, \mathfrak{Y})$ *if for every finite and nonempty subset* $\mathfrak{P}' \subseteq \mathfrak{X}'$ *with* $\mathbf{1} \in \mathfrak{P}'$ *and every finite and nonempty subset* $\mathfrak{Q}' \subseteq \mathfrak{Y}'$*, there exist a finite and nonempty subset* $\mathfrak{P} \subseteq \mathfrak{X}$ *with* $\mathbf{1} \in \mathfrak{P}$ *and a finite and nonempty subset* $\mathfrak{Q} \subseteq \mathfrak{Y}$*, such that* $Z_{\mathfrak{P}',\mathfrak{Q}'}(\cdot)$ *is* polynomial-time reducible *to* $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$*.*

3 Main Theorems

We prove a complexity dichotomy theorem for all counting problems $Z_{\mathbf{C}}(\cdot)$, where $\mathbf{C}$ is any non-negative matrix. Actually, our main theorem is more general.

Definition 4. *Let* $\mathcal{P}$ *be an* $m \times m$ *pattern. An* m*-dimensional non-negative vector* $\mathbf{w}$ *is said to be*

- positive: $w_i > 0$ *for all* $i \in [m]$*; and*
- $\mathcal{P}$-weakly positive: *for all* $i \in [m]$*,* $w_i > 0$ *if and only if* $(i, i) \in \mathcal{P}$*.*

We call $(\mathfrak{X}, \mathfrak{Y})$ *a* $\mathcal{P}$-pair *if*

1. $\mathfrak{X}$ *is a nonempty (and possibly infinite) set of* positive *and* $\mathcal{P}$-weakly positive *vectors with* $\mathbf{1} \in \mathfrak{X}$*;*
2. $\mathfrak{Y}$ *is a nonempty (and possibly infinite) set of* $m \times m$ *(non-negative)* $\mathcal{P}$*-matrices.*

We say it is a finite $\mathcal{P}$*-pair if both sets are finite. We normally use* $(\mathfrak{P}, \mathfrak{Q})$ *to denote a finite* $\mathcal{P}$*-pair.*

Similarly, for any $m \times m$ *block pattern* $\mathcal{T}$*, we can define* $\mathcal{T}$-weakly positive vectors *as well as* $\mathcal{T}$-pairs *by replacing the* $\mathcal{P}$ *above with the pattern defined by* $\mathcal{T}$*.*

We prove the following complexity dichotomy theorem:

Theorem 2 (Complexity Dichotomy)**.** *Let* $\mathcal{P}$ *be an* $m \times m$ *pattern, for some* $m \geq 1$*, then for any* finite $\mathcal{P}$-pair $(\mathfrak{P}, \mathfrak{Q})$*, the problem of computing* $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ *is either in polynomial time or #P-hard.*

Clearly, it gives us a dichotomy for the special case of $Z_{\mathbf{C}}(\cdot)$ when $\mathfrak{P} = \{\mathbf{1}\}$ and $\mathfrak{Q} = \{\mathbf{C}\}$. Moreover, we show that for the special case when $\mathfrak{P} = \{\mathbf{1}\}$, we can decide in a finite number of steps whether $Z_{\mathfrak{P},\mathfrak{Q}}$ is in polynomial time or #P-hard. In particular, it implies that the dichotomy for $Z_{\mathbf{C}}(\cdot)$ is decidable.

Theorem 3 (Decidability)**.** *Given any positive integer* $m \geq 1$*, an* $m \times m$ *pattern* $\mathcal{P}$*, and a finite* $\mathcal{P}$*-pair* $(\mathfrak{P}, \mathfrak{Q})$ *with* $\mathfrak{P} = \{\mathbf{1}\}$*, the problem of whether* $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ *is in polynomial time or #P-hard is decidable.*

We prove Theorem 2 and 3 in the rest of the section. The lemmas (Lemma 2, 3, and 4) used in the proof will be proved in the appendix.

3.1 Defining New Pairs: gen-pair$(\mathfrak{X},\mathfrak{Y})$

Before proving Theorem 2, we state a key lemma which will be proved in Appendix C and Appendix F.

Let $(\mathfrak{X},\mathfrak{Y})$ be a (possibly infinite) $\mathcal{T}$-pair, for some non-trivial $m\times m$ block pattern $\mathcal{T}$. Also assume that every matrix in $\mathfrak{Y}$ is block-rank-1. Then in Appendix C, we introduce an operation `gen-pair` over $(\mathfrak{X},\mathfrak{Y})$, which *defines* a new (and possibly infinite) pair $(\mathfrak{X}',\mathfrak{Y}')=\texttt{gen-pair}(\mathfrak{X},\mathfrak{Y})$.

Definition 5. *A set S of non-negative m-dimensional vectors, for some $m\ge 1$, is* closed *if $\mathbf{w}_1\circ\mathbf{w}_2\in S$ for all vectors $\mathbf{w}_1,\mathbf{w}_2\in S$, where we let $\circ$ denote the* Hadamard product *of two vectors: $\mathbf{w}_1\circ\mathbf{w}_2$ is the m-dimensional vector whose ith entry is $w_{1,i}\cdot w_{2,i}$ for all $i\in[m]$.*

In Appendix F, we prove the following lemma:

Lemma 2. *Let $(\mathfrak{X},\mathfrak{Y})$ be a $\mathcal{T}$-pair, for some non-trivial block pattern $\mathcal{T}$. Suppose every matrix in $\mathfrak{Y}$ is block-rank-1. Then $(\mathfrak{X}',\mathfrak{Y}')=\texttt{gen-pair}(\mathfrak{X},\mathfrak{Y})$ is a $\mathcal{P}'$-pair, where $\mathcal{P}'=\texttt{gen}(\mathcal{T})$; the new vector set $\mathfrak{X}'$ is* closed*; and $(\mathfrak{X}',\mathfrak{Y}')$ is* polynomial-time reducible *to $(\mathfrak{X},\mathfrak{Y})$.*

3.2 Proof of Theorem 2

Let $(\mathfrak{P},\mathfrak{Q})$ be a *finite* $\mathcal{P}$-pair, where $\mathcal{P}$ is an $m\times m$ pattern. We assume $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is not #P-hard, and we only need to show that $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is in polynomial time.

By Lemma 1, there must be a block pattern $\mathcal{T}$ consistent with $\mathcal{P}$ and all the matrices in $\mathfrak{Q}$ are block-rank-1 since otherwise $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is #P-hard, which contradicts the assumption. Therefore, we have

$\mathbf{R}_0$: $(\mathfrak{P},\mathfrak{Q})$ is a finite $\mathcal{T}$-pair, for some $m\times m$ block pattern $\mathcal{T}$; and
Every matrix in $\mathfrak{Q}$ is block-rank-1.

For convenience, we rename $(\mathfrak{P},\mathfrak{Q})$ to be $(\mathfrak{X}_0,\mathfrak{Y}_0)$ and rename m and $\mathcal{T}$ to be m_0 and $\mathcal{T}_0$, respectively.

Now we define a finite sequence of pairs using the `gen-pair` operation starting with $(\mathfrak{X}_0,\mathfrak{Y}_0)$. First, if $|A_k|=|B_k|=1$ for all k, i.e., every set A_k and B_k in $\mathcal{T}_0$ is a singleton, then the sequence has only one pair $(\mathfrak{X}_0,\mathfrak{Y}_0)$ and the definition of the sequence is complete. Note that this also includes the special case when $\mathcal{T}_0=\emptyset$ and $\mathfrak{Y}_0=\{\mathbf{0}\}$. Otherwise, in Step 1, we define a new $\mathcal{P}_1$-pair $(\mathfrak{X}_1,\mathfrak{Y}_1)$ using `gen-pair`:

$$(\mathfrak{X}_1,\mathfrak{Y}_1)=\texttt{gen-pair}(\mathfrak{X}_0,\mathfrak{Y}_0),\qquad \text{where }\ \mathcal{P}_1=\texttt{gen}(\mathcal{T}_0).$$

By Lemma 2 $(\mathfrak{X}_1,\mathfrak{Y}_1)$ is polynomial-time reducible to $(\mathfrak{X}_0,\mathfrak{Y}_0)$. This implies that $\mathcal{P}_1$ must be consistent with a block pattern, denoted by $\mathcal{T}_1$, and every matrix in $\mathfrak{Y}_1$ is block-rank-1. (Otherwise assume $\mathbf{D}\in\mathfrak{Y}_1$ is not block-rank-1, then by Lemma 1, $Z_{\mathfrak{P}_1,\mathfrak{Q}_1}(\cdot)$ is #P-hard, where $\mathfrak{P}_1=\{\mathbf{1}\}$ and $\mathfrak{Q}_1=\{\mathbf{D}\}$. It follows from Lemma 2 that there exists a finite pair $(\mathfrak{P}_0,\mathfrak{Q}_0)$ with $\mathfrak{P}_0\subseteq\mathfrak{X}_0$ and $\mathfrak{Q}_0\subseteq\mathfrak{Y}_0$ such that $Z_{\mathfrak{P}_1,\mathfrak{Q}_1}(\cdot)$ is polynomial-time reducible to $Z_{\mathfrak{P}_0,\mathfrak{Q}_0}(\cdot)$. It is also clear that $Z_{\mathfrak{P}_0,\mathfrak{Q}_0}(\cdot)$ is reducible to $Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\cdot)$ and thus, the latter is also #P-hard, which contradicts our assumption.) As a result, we have

$\mathbf{R}_1$: $\mathcal{T}_1=\texttt{gen-block}(\mathcal{T}_0)$ is an $m_1\times m_1$ block pattern, where m_1 is the number of pairs in $\mathcal{T}_0$;
$(\mathfrak{X}_1,\mathfrak{Y}_1)=\texttt{gen-pair}(\mathfrak{X}_0,\mathfrak{Y}_0)$ is a $\mathcal{T}_1$-pair, and every matrix in $\mathfrak{Y}_1$ is block-rank-1.

We also have $m_0>m_1$ since at least one of the sets in $\mathcal{T}_0$ is not a singleton.

We remark that both sets $\mathfrak{X}_1$ and $\mathfrak{Y}_1$ are generally infinite, so one can not check the matrices in $\mathfrak{Y}_1$ for the block-rank-1 property one by one. It does not matter right now because we are only proving the dichotomy theorem. However, it will become a serious problem later when we prove that the dichotomy is decidable. We have to show that the block-rank-1 property can be verified in a finite number of steps.

We then repeat the process above. After $\ell \geq 1$ steps, we get a sequence of $\ell+1$ pairs:

$$(\mathfrak{X}_0, \mathfrak{Y}_0), (\mathfrak{X}_1, \mathfrak{Y}_1), \ldots, (\mathfrak{X}_\ell, \mathfrak{Y}_\ell),$$

and $\ell+1$ block patterns $\mathcal{T}_0, \mathcal{T}_1, \ldots, \mathcal{T}_\ell$ such that

$\mathbf{R}_\ell$: For every $i \in [\ell]$, $\mathcal{T}_i = \texttt{gen-block}(\mathcal{T}_{i-1})$;
For every $i \in [\ell]$, $(\mathfrak{X}_i, \mathfrak{Y}_i) = \texttt{gen-pair}(\mathfrak{X}_{i-1}, \mathfrak{Y}_{i-1})$ is a $\mathcal{T}_i$-pair; and
For every $i \in [0:\ell]$, all the matrices in $\mathfrak{Y}_i$ are block-rank-1.

We have two cases. If every set in $\mathcal{T}_\ell$ is a singleton (including the case when $\mathcal{T}_\ell = \emptyset$ and $\mathfrak{Y}_\ell = \{\mathbf{0}\}$), then the sequence has only $\ell+1$ pairs and the definition of the sequence is complete. Otherwise in Step $\ell+1$ we apply the `gen-pair` operation again to define a new pair $(\mathfrak{X}_{\ell+1}, \mathfrak{Y}_{\ell+1})$ from $(\mathfrak{X}_\ell, \mathfrak{Y}_\ell)$.

Finally, assuming $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is not #P-hard, we get a sequence of $h+1$ pairs

$$(\mathfrak{X}_0, \mathfrak{Y}_0), (\mathfrak{X}_1, \mathfrak{Y}_1), \ldots, (\mathfrak{X}_h, \mathfrak{Y}_h), \qquad \text{for some } h \geq 0,$$

together with $h+1$ positive integers $m_0 > \ldots > m_h \geq 1$ and $h+1$ block patterns $\mathcal{T}_0, \ldots, \mathcal{T}_h$ such that

R: For every $i \in [0:h]$, $\mathcal{T}_i$ is an $m_i \times m_i$ block pattern;
For every $i \in [h]$, $\mathcal{T}_i = \texttt{gen-block}(\mathcal{T}_{i-1})$;
Either $\mathcal{T}_h = \emptyset$ is trivial or every set in $\mathcal{T}_h$ is a singleton;
For every $i \in [h]$, $(\mathfrak{X}_i, \mathfrak{Y}_i) = \texttt{gen-pair}(\mathfrak{X}_{i-1}, \mathfrak{Y}_{i-1})$ is a $\mathcal{T}_i$-pair; and
For every $i \in [0:h]$, all the matrices in $\mathfrak{Y}_i$ are block-rank-1.

Because $m_0 > \ldots > m_h \geq 1$, we also have $h < m_0 = m$.

3.2.1 Dichotomy

Now we know that if $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is not #P-hard, then there is a sequence of $h+1$ pairs for some $h : 0 \leq h < m$, which satisfies condition **(R)**. To complete the proof of Theorem 2, we show in Appendix D that

Lemma 3 (Tractability)**.** *Given a block pattern $\mathcal{T}$ and a finite $\mathcal{T}$-pair $(\mathfrak{P}, \mathfrak{Q})$, let $(\mathfrak{X}_0, \mathfrak{Y}_0), \ldots, (\mathfrak{X}_h, \mathfrak{Y}_h)$ be a sequence of pairs defined as above, with $(\mathfrak{X}_0, \mathfrak{Y}_0) = (\mathfrak{P}, \mathfrak{Q})$. Suppose it satisfies* condition **(R)***, then $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ is computable in polynomial time.*

3.3 Proof of Theorem 3

Next, we show that for the special case when $\mathfrak{X}_0 = \mathfrak{P} = \{\mathbf{1}\}$, the dichotomy criterion is decidable. First, the condition $(\mathbf{R}_0)$ can be checked easily since there are only finitely many matrices in $\mathfrak{Y}_0$.

Assume after $\ell : 0 \leq \ell < m$ steps, we get a sequence of $\ell+1$ pairs: $(\mathfrak{X}_0, \mathfrak{Y}_0), (\mathfrak{X}_1, \mathfrak{Y}_1), \ldots, (\mathfrak{X}_\ell, \mathfrak{Y}_\ell)$, together with $\ell+1$ block patterns $\mathcal{T}_0, \ldots, \mathcal{T}_\ell$. Moreover, we know that they satisfy $(\mathbf{R}_\ell)$. If every set in $\mathcal{T}_\ell$ is a singleton (including the case when $\mathcal{T}_\ell = \emptyset$), then we are done because by Lemma 3, the problem is in polynomial time. Otherwise, to prove Theorem 3, we need a finite-time algorithm to check whether every matrix in the new $\mathcal{P}$-pair $(\mathfrak{X}_{\ell+1}, \mathfrak{Y}_{\ell+1}) = \texttt{gen-pair}(\mathfrak{X}_\ell, \mathfrak{Y}_\ell)$, where $\mathcal{P} = \texttt{gen}(\mathcal{T}_\ell)$, is block-rank-1 or not. We refer to this property as the *rank property* for $\mathfrak{Y}_{\ell+1}$ and prove the following lemma in Appendix G. Theorem 3 then follows.

Lemma 4. *Given a block pattern $\mathcal{T}$ and a finite $\mathcal{T}$-pair $(\mathfrak{X}_0, \mathfrak{Y}_0)$ with $\mathfrak{X}_0 = \{\mathbf{1}\}$, let $(\mathfrak{X}_0, \mathfrak{Y}_0), \ldots, (\mathfrak{X}_\ell, \mathfrak{Y}_\ell)$ be a sequence of $\ell+1$ pairs defined as above. Suppose it satisfies* condition $(\mathbf{R}_\ell)$*. Then the* rank property *for $\mathfrak{Y}_{\ell+1}$ can be checked in a finite number of steps.*

References

[1] A. Bulatov. The complexity of the counting constraint satisfaction problem. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 646–661, 2008.

[2] A. Bulatov. The complexity of the counting constraint satisfaction problem. *ECCC Report*, TR07-093, 2009.

[3] A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 562–571, 2003.

[4] A. Bulatov, M.E. Dyer, L.A. Goldberg, M. Jalsenius, M. Jerrum, and D. Richerby. Private communication. 2009.

[5] A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2):148–186, 2005.

[6] J.-Y. Cai, X. Chen, and P. Lu. Graph homomorphisms with complex values: A dichotomy theorem. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming*, 2010.

[7] M.E. Dyer, L.A. Goldberg, and M. Paterson. On counting homomorphisms to directed acyclic graphs. *Journal of the ACM*, 54(6): Article 27, 2007.

[8] M.E. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3–4):260–289, 2000.

[9] M.E. Dyer and D. Richerby. On the complexity of #CSP. In *Proceedings of the 42th ACM Symposium on Theory of Computing*, to appear.

[10] T. Feder and M.Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1999.

[11] M. Freedman, L. Lovász, and A. Schrijver. Reflection positivity, rank connectivity, and homomorphism of graphs. *Journal of the American Mathematical Society*, 20:37–51, 2007.

[12] D. Geiger. Closed systems of function and predicates. *Pacific Journal of Mathematics*, pages 95–100, 1968.

[13] L.A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. A complexity dichotomy for partition functions with mixed signs. In *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science*, 2008.

[14] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.

[15] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.

A Figures

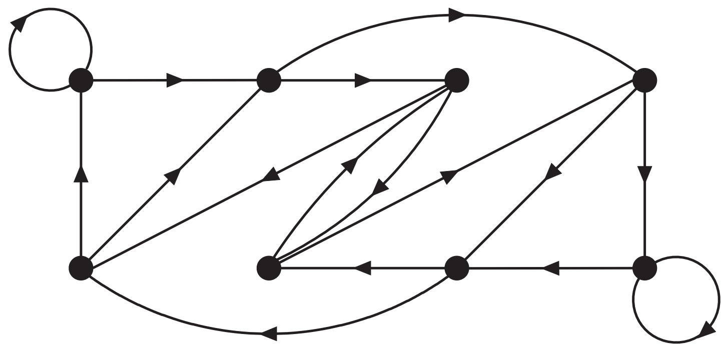

Figure 1: A directed graph H such that $Z_H(\cdot)$ is tractable

$$\mathbf{A} = \begin{pmatrix} A_{1,1} & & A_{1,3} & & & & & \\ A_{2,1} & & A_{2,3} & & & & & \\ & & & & A_{3,5} & & A_{3,7} & \\ & & & & A_{4,5} & & A_{4,7} & \\ & A_{5,2} & & A_{5,4} & & & & \\ & A_{6,2} & & A_{6,4} & & & & \\ & & & & & A_{7,6} & & A_{7,8} \\ & & & & & A_{8,6} & & A_{8,8} \end{pmatrix}$$

Figure 2: The 8×8 block-rank-1 matrix $\mathbf{A}$

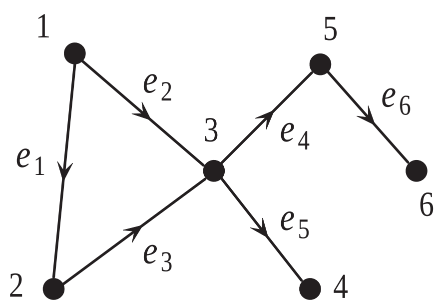

Figure 3: The input directed graph $G = (V, E)$ with $|V| = |E| = 6$

B Proof of Lemma 1

Bulatov and Grohe showed that for any $m \times m$ non-negative *symmetric* matrix $\mathbf{D}$, $Z_{\mathbf{D}}(\cdot)$ is #P-hard if $\mathbf{D}$ is not block-rank-1. Note that when $\mathbf{D}$ is symmetric, the directions of the edges in G do not affect the value of $Z_{\mathbf{D}}(G)$, so we can always assume that G is an undirected graph.

We prove Lemma 1 by giving a reduction from the symmetric case.

Let $\mathbf{C}$ be an $m \times m$ non-negative matrix, which is not block-rank-1. Without loss of generality, we may assume that $\mathbf{C}_1, \mathbf{C}_2$, the first and the second row vectors of $\mathbf{C}$, satisfy $\mathbf{C}_1 \cdot \mathbf{C}_2 > 0$; but $\mathbf{C}_1$ and $\mathbf{C}_2$ are not linearly dependent. Let $\mathbf{D}$ denote the following symmetric matrix:

$$D_{i,j} = \mathbf{C}_i \cdot \mathbf{C}_j, \qquad \text{for all } i, j \in [m].$$

By the assumption, we have $D_{1,1}, D_{1,2}, D_{2,1}, D_{2,2} > 0$ but $D_{1,1}D_{2,2} > D_{1,2}D_{2,1}$. It then follows from the result of Bulatov and Grohe that $Z_{\mathbf{D}}(\cdot)$ is #P-hard to compute.

Now we prove the #P-hardness of $Z_{\mathbf{C}}(\cdot)$ by showing a reduction from $Z_{\mathbf{D}}(\cdot)$. Let $G = (V, E)$ be an input undirected graph of $Z_{\mathbf{D}}(\cdot)$. We construct a directed graph $G' = (V', E')$ in which

$$V' = V \cup \big\{ w_e : e \in E \big\} \quad \text{ and } \quad E' = \big\{ uw_e, vw_e : e = uv \in E \big\}.$$

By the definition of $Z_{\mathbf{C}}(\cdot)$ and $Z_{\mathbf{D}}(\cdot)$, it is easy to verify that

$$Z_{\mathbf{C}}(G') = Z_{\mathbf{D}}(G), \qquad \text{for any undirected graph } G.$$

As a result, $Z_{\mathbf{D}}(\cdot)$ is polynomial-time reducible to $Z_{\mathbf{C}}(\cdot)$, and the latter is also #P-hard.

C Definition of the `gen-pair` Operation

In this section, we define the operation `gen-pair`.

Let $\mathcal{T} = \{(A_1, B_1), \ldots, (A_r, B_r)\}$ be a non-trivial $m \times m$ block pattern with $r \geq 1$. We use $\text{diag}(\mathcal{T})$ to denote the set of all $i \in [m]$ such that $i \in A_k$ and $i \in B_k$ for some $k \in [r]$. In this section, we always assume that $(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{T}$-pair such that every matrix in $\mathfrak{Y}$ is block-rank-1. This means that

1. All matrices in $\mathfrak{Y}$ are block-rank-1 and are of the same block pattern $\mathcal{T}$;
2. $\mathbf{1} \in \mathfrak{X}$ and every vector $\mathbf{w} \in \mathfrak{X}$ is either

   positive: $w_i > 0$ for all $i \in [m]$; or

   $\mathcal{T}$-weakly positive: $w_i > 0$ if and only if $i \in \text{diag}(\mathcal{T})$.

Given such a pair $(\mathfrak{X}, \mathfrak{Y})$, `gen-pair` defines a new $\mathcal{P}$-pair

$$(\mathfrak{X}', \mathfrak{Y}') = \texttt{gen-pair}(\mathfrak{X}, \mathfrak{Y}), \qquad \text{where } \mathcal{P} = \texttt{gen}(\mathcal{T}).$$

To this end we first define a pair $(\mathfrak{X}^*, \mathfrak{Y}^*)$ from $(\mathfrak{X}, \mathfrak{Y})$, which is a *generalized $\mathcal{P}$-pair* defined as follows.

Definition 6. *Let $\mathcal{P}$ be an $r \times r$ pattern with $r \geq 1$. An $r \times r$ nonnegative matrix is called a* $\mathcal{P}$-diagonal *matrix if it is a diagonal matrix and for all $i \in [r]$, its (i,i)th entry is positive if and only if $(i,i) \in \mathcal{P}$.*

We call $(\mathfrak{X}^*, \mathfrak{Y}^*)$ a generalized $\mathcal{P}$-pair *if*

1. *$\mathfrak{X}^*$ is a nonempty (and possibly infinite) set of* positive *and* $\mathcal{P}$-weakly positive *vectors with $\mathbf{1} \in \mathfrak{X}^*$;*
2. *$\mathfrak{Y}^*$ is a nonempty (and possibly infinite) set of $\mathcal{P}$-matrices and* $\mathcal{P}$-diagonal *matrices.*

*For any block pattern $\mathcal{T}$, one can define $\mathcal{T}$-*diagonal *matrices and* generalized $\mathcal{T}$-pairs *similarly, by replacing the pattern $\mathcal{P}$ above with the one defined by $\mathcal{T}$.*

We then use $(\mathfrak{X}^*, \mathfrak{Y}^*)$ to define $(\mathfrak{X}', \mathfrak{Y}')$. In this section we only show that $(\mathfrak{X}', \mathfrak{Y}')$ is a $\mathcal{P}$-pair and $\mathfrak{X}'$ is *closed*. We will give the polynomial-time reduction from $(\mathfrak{X}', \mathfrak{Y}')$ to $(\mathfrak{X}, \mathfrak{Y})$ in Appendix F.

C.1 Definition of $\mathfrak{Y}^*$

We define $\mathfrak{Y}^*$ which contains both $\mathcal{P}$-matrices and $\mathcal{P}$-diagonal matrices, where $\mathcal{P} = \texttt{gen}(\mathcal{T})$.

There are two types of matrices in $\mathfrak{Y}^*$. First, $\mathbf{D}$ is an $r \times r$ $\mathcal{P}$-matrix in $\mathfrak{Y}^*$ if there exist

1. a finite subset of matrices $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[g]}\} \subseteq \mathfrak{Y}$ with $g \geq 1$, and positive integers $s_1, \ldots, s_g$;
2. a finite subset of matrices $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[h]}\} \subseteq \mathfrak{Y}$ with $h \geq 1$, and positive integers $t_1, \ldots, t_h$;
3. a *positive* vector $\mathbf{w} \in \mathfrak{X}$,

such that: Let $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ be the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively, then

$$D_{i,j} = \sum_{x \in B_i \cap A_j} \left(\beta_x^{[1]}\right)^{s_1} \cdots \left(\beta_x^{[g]}\right)^{s_g} \cdot \left(\gamma_x^{[1]}\right)^{t_1} \cdots \left(\gamma_x^{[h]}\right)^{t_h} \cdot w_x, \qquad \text{for all } i, j \in [r].$$

The following lemma is easy to prove.

Lemma 5. *If* $\mathbf{w} \in \mathfrak{X}$ *is positive, then the matrix* $\mathbf{D}$ *defined above is a* $\mathcal{P}$-matrix*, where* $\mathcal{P} = \texttt{gen}(\mathcal{T})$*.*

Proof. Because $(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{T}$-pair, all the matrices $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[j]}$, $i \in [g]$ and $j \in [h]$, are $\mathcal{T}$-matrices and thus, $\boldsymbol{\beta}^{[i]}$ is positive over $B_1 \cup \cdots \cup B_r$ and $\boldsymbol{\gamma}^{[j]}$ is positive over $A_1 \cup \cdots \cup A_r$. Since $\mathbf{w}$ is positive, it is easy to check that $D_{i,j} > 0$ if and only if $B_i \cap A_j \neq \emptyset$. □

Second, $\mathbf{D}$ is an $r \times r$ $\mathcal{P}$-diagonal matrix in $\mathfrak{Y}^*$ if there exist

1. a finite subset of matrices $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[g]}\} \subseteq \mathfrak{Y}$ with $g \geq 1$, and positive integers $s_1, \ldots, s_g$;
2. a finite subset of matrices $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[h]}\} \subseteq \mathfrak{Y}$ with $h \geq 1$, and positive integers $t_1, \ldots, t_h$;
3. a $\mathcal{T}$*-weakly positive* vector $\mathbf{w} \in \mathfrak{X}$,

such that: Let $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ be the representation of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively, then

$$D_{i,j} = \sum_{x \in B_i \cap A_j} \left(\beta_x^{[1]}\right)^{s_1} \cdots \left(\beta_x^{[g]}\right)^{s_g} \cdot \left(\gamma_x^{[1]}\right)^{t_1} \cdots \left(\gamma_x^{[h]}\right)^{t_h} \cdot w_x, \qquad \text{for all } i, j \in [r].$$

Similarly one can show that

Lemma 6. *If* $\mathbf{w}$ *is* $\mathcal{T}$*-weakly positive, then the matrix* $\mathbf{D}$ *defined above is* $\mathcal{P}$-diagonal *where* $\mathcal{P} = \texttt{gen}(\mathcal{T})$*.*

Proof. First, we show that $\mathbf{D}$ is diagonal. Let $i \neq j$ be two distinct indices in $[r]$. If $B_i \cap A_j = \emptyset$, then $D_{i,j}$ is trivially 0. Otherwise, for every $k \in B_i \cap A_j$, we know that (k, k) is not in the pattern defined by $\mathcal{T}$ because $k \in B_i$, $k \in A_j$ but $i \neq j$. As a result, we have $w_k = 0$ which implies $D_{i,j} = 0$ for all $i \neq j \in [r]$.

Second, if $A_i \cap B_i \neq \emptyset$ then (k, k) is in the pattern defined by $\mathcal{T}$ for every $k \in A_i \cap B_i$. This implies that $w_k > 0$. As a result, we have $D_{i,i} > 0$ if and only if $A_i \cap B_i \neq \emptyset$. □

C.2 Definition of $\mathfrak{X}^*$

Now we define $\mathfrak{X}^*$. To this end, we first define $\mathfrak{X}^\#$ which is a set of r-dimensional positive and $\mathcal{P}$-weakly positive vectors. We have $\mathbf{w}^\# \in \mathfrak{X}^\#$ if and only if one of the following four cases is true:

1. $\mathbf{w}^\# = \mathbf{1}$;

2. There exist a finite subset $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[g]}\} \subseteq \mathfrak{Y}$ with $g \geq 1$, positive integers $s_1, \ldots, s_g$ and a vector $\mathbf{w} \in \mathfrak{X}$ (positive or $\mathcal{T}$-weakly positive) such that: Let $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ be the representation of $\mathbf{C}^{[i]}$, then

$$w_i^\# = \sum_{x \in A_i} \left(\alpha_x^{[1]}\right)^{s_1} \cdots \left(\alpha_x^{[g]}\right)^{s_g} \cdot w_x, \qquad \text{for all } i \in [r].$$

   It can be checked that $\mathbf{w}^\#$ is positive if $\mathbf{w}$ is positive and $\mathbf{w}^\#$ is $\mathcal{P}$-weakly positive if $\mathbf{w}$ is $\mathcal{T}$-weakly positive.

3. There exist a finite subset $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[h]}\} \subseteq \mathfrak{Y}$ with $h \geq 1$, positive integers $t_1, \ldots, t_g$ and a vector $\mathbf{w} \in \mathfrak{X}$ (positive or $\mathcal{T}$-weakly positive) such that: Let $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ be the representation of $\mathbf{D}^{[i]}$, then

$$w_i^\# = \sum_{x \in B_i} \left(\delta_x^{[1]}\right)^{t_1} \cdots \left(\delta_x^{[h]}\right)^{t_h} \cdot w_x, \qquad \text{for all } i \in [r].$$

   Similarly, it can be checked that $\mathbf{w}^\#$ is positive if $\mathbf{w}$ is positive and $\mathbf{w}^\#$ is $\mathcal{P}$-weakly positive if $\mathbf{w}$ is $\mathcal{T}$-weakly positive.

4. There exist two finite subsets $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[g]}\} \subseteq \mathfrak{Y}$ and $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[h]}\} \subseteq \mathfrak{Y}$ with $g \geq 1$ and $h \geq 1$, positive integers $s_1, \ldots, s_g, t_1, \ldots, t_h$ and a vector $\mathbf{w} \in \mathfrak{X}$ (positive or $\mathcal{T}$-weakly positive) such that: Let $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ be the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively, then

$$w_i^\# = \sum_{x \in B_i \cap A_i} \left(\beta_x^{[1]}\right)^{s_1} \cdots \left(\beta_x^{[g]}\right)^{s_g} \cdot \left(\gamma_x^{[1]}\right)^{t_1} \cdots \left(\gamma_x^{[h]}\right)^{t_h} \cdot w_x, \qquad \text{for all } i \in [r].$$

   It can be checked that $\mathbf{w}^\#$ is always a $\mathcal{P}$-weakly positive vector.

This finishes the definition of $\mathfrak{X}^\#$.

Set $\mathfrak{X}^*$ is the *closure* of $\mathfrak{X}^\#$: $\mathbf{w} \in \mathfrak{X}^*$ if and only if there exist a finite subset $\{\mathbf{w}_1, \ldots, \mathbf{w}_g\} \subseteq \mathfrak{X}^\#$ and positive integers $s_1, \ldots, s_g$ such that

$$\mathbf{w} = \left(\mathbf{w}_1\right)^{s_1} \circ \cdots \circ \left(\mathbf{w}_g\right)^{s_g},$$

where $\circ$ denotes the Hadamard product. It immediately implies that $\mathfrak{X}^*$ is closed, and any vector in it is either positive or $\mathcal{P}$-weakly positive. It is also easy to check that $(\mathfrak{X}^*, \mathfrak{Y}^*)$ is a *generalized* $\mathcal{P}$-pair.

C.3 Definition of $(\mathfrak{X}', \mathfrak{Y}')$

We use $(\mathfrak{X}^*, \mathfrak{Y}^*)$ to define $(\mathfrak{X}', \mathfrak{Y}')$ as follows.

First, $\mathfrak{Y}'$ contains exactly all the $\mathcal{P}$-matrices in $\mathfrak{Y}^*$.

The definition of $\mathfrak{X}'$ is more complicated. We have $\mathbf{w}' \in \mathfrak{X}'$ if and only if

1. $\mathbf{w}' \in \mathfrak{X}^*$; or

2. There exist

(a) a finite subset of $\mathcal{P}$-matrices $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[g]}\} \subseteq \mathfrak{Y}^*$ with $g \geq 0$ (so this set could be empty) and g positive integers $s_1, \ldots, s_g$;

(b) a finite subset of $\mathcal{P}$-diagonal matrices $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[h]}\} \subseteq \mathfrak{Y}^*$ with $h \geq 1$, and h positive integers $t_1, \ldots, t_h$;

(c) and a vector $\mathbf{w} \in \mathfrak{X}^*$ (which is either positive or $\mathcal{P}$-weakly positive),

such that $\mathbf{w}'$ satisfies

$$w'_i = w_i \cdot \left(C^{[1]}_{i,i}\right)^{s_1} \cdots \left(C^{[g]}_{i,i}\right)^{s_g} \cdot \left(D^{[1]}_{i,i}\right)^{t_1} \cdots \left(D^{[h]}_{i,i}\right)^{t_h}, \qquad \text{for any } i \in [r].$$

It can be checked that every $\mathbf{w}' \in \mathfrak{X}'$ is either positive or $\mathcal{P}$-weakly positive.

This finishes the definition of $(\mathfrak{X}', \mathfrak{Y}')$ and the `gen-pair` operation. It is easy to verify that the new pair $(\mathfrak{X}', \mathfrak{Y}')$ is a $\mathcal{P}$-pair. Moreover, since $\mathfrak{X}^*$ is closed, one can show that $\mathfrak{X}'$ is also closed. This proved the first part of Lemma 2:

Lemma 7. *Let $(\mathfrak{X}, \mathfrak{Y})$ be a $\mathcal{T}$-pair for some non-trivial block pattern $\mathcal{T}$. Suppose every matrix in $\mathfrak{Y}$ is block-rank-*1*, then $(\mathfrak{X}', \mathfrak{Y}') =$ `gen-pair`$(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{P}$-pair, where $\mathcal{P} =$ `gen`$(\mathcal{T})$, and $\mathfrak{X}'$ is* closed. *Moreover, the pair $(\mathfrak{X}^*, \mathfrak{Y}^*)$ defined from $(\mathfrak{X}, \mathfrak{Y})$ is a* generalized *$\mathcal{P}$-pair and $\mathfrak{X}^*$ is also* closed.

D Dichotomy: Tractability

In this section, we prove Lemma 3, the tractability part of the dichotomy theorem.

Let $(\mathfrak{X}_0, \mathfrak{Y}_0) = (\mathfrak{P}, \mathfrak{Q})$ be a finite $\mathcal{T}_0$-pair, for some block pattern $\mathcal{T}_0$. Let $(\mathfrak{X}_0, \mathfrak{Y}_0), \ldots, (\mathfrak{X}_h, \mathfrak{Y}_h)$ be a sequence of $h+1$ pairs for some $h \geq 0$, $m_0 > m_1 > \ldots > m_h \geq 1$ be $h+1$ positive integers, and $\mathcal{T}_0$, $\mathcal{T}_1, \ldots, \mathcal{T}_h$ be $h+1$ block patterns such that

R: For every $i \in [0:h]$, $\mathcal{T}_i$ is an $m_i \times m_i$ block pattern;
For every $i \in [h]$, $\mathcal{T}_i =$ `gen-block`$(\mathcal{T}_{i-1})$;
Either $\mathcal{T}_h = \emptyset$ is trivial or every set in $\mathcal{T}_h$ is a singleton;
For every $i \in [h]$, $(\mathfrak{X}_i, \mathfrak{Y}_i) =$ `gen-pair`$(\mathfrak{X}_{i-1}, \mathfrak{Y}_{i-1})$ is a $\mathcal{T}_i$-pair; and
For every $i \in [0:h]$, all the matrices in $\mathfrak{Y}_i$ are block-rank-1.

We need to show that $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot) = Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\cdot)$ can be computed in polynomial time.

Let $\mathcal{G}_0 = (G_0, \mathcal{V}_0, \mathcal{E}_0)$ be an input labeled directed graph of $Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\cdot)$. By definition we have $\mathcal{V}_0(v) \in \mathfrak{X}_0$ for all vertices $v \in V(G_0)$, and $\mathcal{E}_0(uv) \in \mathfrak{Y}_0$ for all edges $uv \in E(G_0)$. We further assume that the underlying undirected graph of G_0 is connected. (If G_0 is not connected, then we only need to compute $Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\cdot)$ for each undirected connected component of G_0 and multiply them to obtain $Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\mathcal{G}_0)$.)

To compute $Z_{\mathfrak{X}_0,\mathfrak{Y}_0}(\mathcal{G}_0)$, we will construct in polynomial-time a sequence of $h+1$ labeled directed graphs $\mathcal{G}_0, \ldots, \mathcal{G}_h$. We will show that these graphs have the following two properties:

$\mathbf{P}_1$: For every $\ell \in [0:h]$, $\mathcal{G}_\ell = (G_\ell, \mathcal{V}_\ell, \mathcal{E}_\ell)$ is a labeled directed graph such that $\mathcal{V}_\ell(v) \in \mathfrak{X}_\ell$ for all $v \in V(G_\ell)$; $\mathcal{E}_\ell(uv) \in \mathfrak{Y}_\ell$ for all $uv \in E(G_\ell)$; and the underlying undirected graph of G_ℓ is connected.

$\mathbf{P}_2$: $Z(\mathcal{G}_0) = Z(\mathcal{G}_1) = \cdots = Z(\mathcal{G}_h)$.

As a result, to compute $Z(\mathcal{G}_0)$, one only needs to compute $Z(\mathcal{G}_h)$. On the other hand, we do know how to compute $Z(\mathcal{G}_h)$ in polynomial time. If $\mathcal{T}_h$ is trivial, then computing $Z(\mathcal{G}_h)$ is also trivial. Otherwise, if every set in $\mathcal{T}_h$ is a singleton, then one can efficiently enumerate all possible assignments of $\mathcal{G}_h$ with non-zero weight (since the underlying undirected graph of G_h is *connected*). This allows us to compute $Z(\mathcal{G}_0) = Z(\mathcal{G}_h)$ in polynomial time.

D.1 Construction of $\mathcal{G}'$ from $\mathcal{G}$

Let $(\mathfrak{X}, \mathfrak{Y})$ be a $\mathcal{T}$-pair for some $m \times m$ non-trivial block pattern $\mathcal{T}$ such that all the matrices in $\mathfrak{Y}$ are block-rank-1. Then by Lemma 7, $(\mathfrak{X}', \mathfrak{Y}') = \texttt{gen-pair}(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{P}$-pair where $\mathcal{P} = \texttt{gen}(\mathcal{T})$.

Let $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ be a labeled directed graph such that $\mathcal{V}(v) \in \mathfrak{X}$ for all $v \in V(G)$; $\mathcal{E}(uv) \in \mathfrak{Y}$ for all $uv \in E(G)$; and the underlying undirected graph of G is connected. We further assume that G is not trivial: V is not a singleton (since for this special case, $Z(\mathcal{G})$ can be computed trivially). In this section, we show how to construct a new graph $\mathcal{G}' = (G', \mathcal{V}', \mathcal{E}')$ in polynomial time such that $\mathcal{V}'(v) \in \mathfrak{X}'$ for all $v \in V(G')$; $\mathcal{E}'(uv) \in \mathfrak{Y}'$ for all $uv \in E(G')$; the underlying undirected graph of G' is connected; and

$$Z(\mathcal{G}) = Z(\mathcal{G}'). \tag{2}$$

Then we can repeatedly apply this construction, starting from $\mathcal{G}_0$, to obtain a sequence of $h + 1$ labeled directed graphs $\mathcal{G}_0, \ldots, \mathcal{G}_h$ that satisfy both $\mathbf{P}_1$ and $\mathbf{P}_2$. Lemma 3 then follows.

Now we describe the construction of $\mathcal{G}'$. Let $G = (V, E)$ and $\mathcal{T} = \{(A_1, B_1), \ldots, (A_n, B_n)\}$ for some $n \geq 1$, then $\mathcal{P} = \texttt{gen}(\mathcal{T})$ is an $n \times n$ pattern. The construction of $\mathcal{G}'$ is divided into two steps, just like the definition of $(\mathfrak{X}', \mathfrak{Y}') = \texttt{gen-pair}(\mathfrak{X}, \mathfrak{Y})$ in Appendix C. In the first step, we construct a labeled graph $\mathcal{G}^* = (G^*, \mathcal{V}^*, \mathcal{E}^*)$ from $\mathcal{G}$ such that

1. $\mathcal{V}^*(v) \in \mathfrak{X}^*$ for all $v \in V(G^*)$; $\mathcal{E}^*(uv) \in \mathfrak{Y}^*$ for all $uv \in E(G^*)$; and the underlying undirected graph of G^* is connected, where $(\mathfrak{X}^*, \mathfrak{Y}^*)$ denotes the *generalized* $\mathcal{P}$-pair defined in Appendix C.

2. $Z(\mathcal{G}^*) = Z(\mathcal{G})$.

In the second step, we construct $\mathcal{G}'$ from $\mathcal{G}^*$ and show that $Z(\mathcal{G}') = Z(\mathcal{G}^*)$.

D.1.1 Construction of $\mathcal{G}^*$ from $\mathcal{G}$

Let $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ and $G = (V, E)$. We decompose the edge set using the following equivalence relation:

Definition 7. *Let e, e' be two directed edges in E. We say $e \sim e'$ if there exist a sequence of edges*

$$e = e_0, e_1, \ldots, e_k = e'$$

in E such that for all $i \in [0 : k-1]$, e_i and e_{i+1} share either the same head or the same tail.

We divide E into equivalence classes $R_1, \ldots, R_f$ using $\sim$:

$$E = R_1 \cup \ldots \cup R_f, \quad \text{for some } f \geq 1.$$

Because the underlying undirected graph of G is connected, there is no isolated vertex v in G and thus every vertex $v \in V$ appears as an incident vertex of some edge in at least one of the equivalence classes. This equivalence relation is useful because of the following observation.

Observation 1. For any $i \in [f]$, the subgraph spanned by R_i is connected if we view it as an undirected graph. There are three types of vertices in it:

1. Type-L: vertices which only have outgoing edges in R_i;

2. Type-R: vertices which only have incoming edges in R_i; and

3. Type-M: vertices which have both incoming and outgoing edges in R_i.

Let $\xi : V \to [m]$ be any assignment with $\text{wt}(\mathcal{G}, \xi) \neq 0$, then for any $i \in [f]$ there exists a unique $k_i \in [n]$ such that the value of every edge $uv \in R_i$ is derived from the k_i-th block of $\mathcal{T}$:

$$\xi(u) \in A_{k_i} \quad \text{and} \quad \xi(v) \in B_{k_i}.$$

Therefore, for every $i \in [f]$, there exists a unique $k_i \in [n]$ such that

1. For every Type-L vertex v in the graph spanned by R_i, $\xi(v) \in A_{k_i}$;
2. For every Type-R vertex v in the graph spanned by R_i, $\xi(v) \in B_{k_i}$; and
3. For every Type-M vertex v in the graph spanned by R_i, $\xi(v) \in A_{k_i} \cap B_{k_i}$.

Now we build $\mathcal{G}^* = (G^*, \mathcal{V}^*, \mathcal{E}^*)$, where $G^* = (V^*, E^*)$. We start with the construction of G^*. V^* is exactly $[f]$ in which the vertex $i \in [f]$ corresponds to R_i of G. For every vertex $v \in V$, if it appears in both the subgraph spanned by R_i and the one spanned by R_j for some $i \neq j \in [f]$ (note that it cannot appear in more than two such subgraphs) and if the incoming edges of v are from R_i and the outgoing edges of v are from R_j, then we add a directed edge ij in E^*. Note that E^* may have parallel edges. This finishes the construction of G^*. It is easy to verify that the underlying undirected graph of G^* is also *connected*.

The only thing left is to label the graph G^* with vertex and edge weights. For every edge in E^* we assign it the following $n \times n$ matrix $\mathbf{D}$. Assume the edge ij is created because of $v \in V$, which appears in both R_i and R_j. Let the incoming edges of v be $u_1v, \ldots, u_sv$ in R_i and the outgoing edges of v be $vw_1, \ldots, vw_t$ in R_j, where $s, t \geq 1$. We use $\mathbf{C}^{[i]} \in \mathfrak{Y}$ to denote the edge weight of u_iv, $\mathbf{D}^{[i]} \in \mathfrak{Y}$ to denote the edge weight of vw_i, and $\mathbf{w} \in \mathfrak{X}$ to denote the vertex weight of v in $\mathcal{G}$. We also use $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ to denote the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively. Then the (i,j)th entry of $\mathbf{D}$ is

$$D_{i,j} = \sum_{x \in B_i \cap A_j} \beta_x^{[1]} \cdots \beta_x^{[s]} \cdot \gamma_x^{[1]} \cdots \gamma_x^{[t]} \cdot w_x, \qquad \text{for all } i, j \in [n].$$

By the definition of `gen-pair`, it is easy to check that $\mathbf{D} \in \mathfrak{Y}^*$.

Finally, we define the vertex weight of $i \in [f]$. To this end, we first define an n-dimensional vector $\mathbf{w}^{[v]}$ for each vertex $v \in V$ that *only appears* in R_i. We then multiply (using Hadamard product) all such vectors to get the vertex weight vector of $i \in [f]$.

Let $v \in V$ be a vertex which only appears in R_i, then we have the following three cases:

1. If v is Type-L, then we use $vw_1, \ldots, vw_s$ to denote its outgoing edges. We let $\mathbf{w}$ denote the vertex weight of v in $\mathcal{G}$ and $\mathbf{C}^{[j]}$ denote the edge weight of vw_j with representation $(\boldsymbol{\alpha}^{[j]}, \boldsymbol{\beta}^{[j]})$. Then

$$w_k^{[v]} = \sum_{x \in A_k} \alpha_x^{[1]} \cdots \alpha_x^{[s]} \cdot w_x, \qquad \text{for all } k \in [n].$$

2. If v is Type-R, then we use $u_1v, \ldots, u_sv$ to denote its incoming edges. We let $\mathbf{w}$ denote the vertex weight of v in $\mathcal{G}$ and $\mathbf{C}^{[j]}$ denote the edge weight of u_jv with representation $(\boldsymbol{\alpha}^{[j]}, \boldsymbol{\beta}^{[j]})$. Then

$$w_k^{[v]} = \sum_{x \in B_k} \beta_x^{[1]} \cdots \beta_x^{[s]} \cdot w_x, \qquad \text{for all } k \in [n].$$

3. If v is Type-M, then we use $u_1v, \ldots, u_sv, vw_1, \ldots, vw_t$ to denote its edges where $s, t \geq 1$. We let $\mathbf{w}$ be the vertex weight of v in $\mathcal{G}$, $\mathbf{C}^{[j]}$ be the edge weight of u_jv with representation $(\boldsymbol{\alpha}^{[j]}, \boldsymbol{\beta}^{[j]})$, and $\mathbf{D}^{[j]}$ be the edge weight of vw_j with representation $(\boldsymbol{\gamma}^{[j]}, \boldsymbol{\delta}^{[j]})$. Then

$$w_k^{[v]} = \sum_{x \in B_k \cap A_k} \beta_x^{[1]} \cdots \beta_x^{[s]} \cdot \gamma_x^{[1]} \cdots \gamma_x^{[t]} \cdot w_x, \qquad \text{for all } k \in [n].$$

We then multiply (using Hadamard product) all the vectors $\mathbf{w}^{[v]}$ over all vertices v that only appear in R_i to get the vertex weight vector $\mathbf{w}$ of $i \in [f]$ in $\mathcal{G}^*$. By definition, it can be checked that $\mathbf{w} \in \mathfrak{X}^*$. This finishes the construction of $\mathcal{G}^*$. Next, we show that $Z(\mathcal{G}^*) = Z(\mathcal{G})$.

Let $\phi : V^* = [f] \to [n]$ be any assignment. We use Ξ_ϕ to denote

$$\left\{ \xi : V \to [m] \;\middle|\; \forall i \in [f],\ \forall\, uv \in R_i,\ \xi(u) \in A_{\phi(i)} \text{ and } \xi(v) \in B_{\phi(i)} \right\}.$$

Equivalently, ϕ defines for each vertex $v \in V$ a set $U_v \subseteq [m]$, where

1. If v appears in both the subgraph spanned by R_i and the subgraph spanned by R_j, for some $i \neq j \in [f]$; and v is Type-R in R_i and Type-L in R_j, then $U_v = B_{\phi(i)} \cap A_{\phi(j)}$;
2. Otherwise, assume v only appears in the subgraph spanned by R_i. Then
   (a) If v is Type-L, then $U_v = A_{\phi(i)}$;
   (b) If v is Type-R, then $U_v = B_{\phi(i)}$; and
   (c) If v is Type-M, then $U_v = B_{\phi(i)} \cap A_{\phi(i)}$,

such that $\xi \in \Xi_\phi \iff \xi(v) \in U_v$ for all $v \in V$. In particular, $\Xi_\phi = \emptyset$ if $U_v = \emptyset$ for some $v \in V$.

By Observation 1, if $\text{wt}(\mathcal{G}, \xi) \neq 0$ then $\xi \in \Xi_\phi$ for some unique ϕ. For any $v \in V$, we let $\mathbf{w}^{[v]}$ denote its vertex weight in $\mathcal{G}$; and for any $uv \in E$, we let $\mathbf{D}^{[uv]}$ denote its edge weight in $\mathcal{G}$, with representation $(\boldsymbol{\alpha}^{[uv]}, \boldsymbol{\beta}^{[uv]})$. Then by the definition of Ξ_ϕ, we have for all $\xi \in \Xi_\phi$,

$$D^{[uv]}_{\xi(u),\xi(v)} = \alpha^{[uv]}_{\xi(u)} \cdot \beta^{[uv]}_{\xi(v)}, \qquad \text{for all } uv \in E.$$

Therefore, we have the following equation:

$$\sum_{\xi \in \Xi_\phi} \text{wt}(\mathcal{G}, \xi) = \sum_{\xi \in \Xi_\phi} \left(\prod_{v \in V} w^{[v]}_{\xi(v)} \prod_{uv \in E} \alpha^{[uv]}_{\xi(u)} \cdot \beta^{[uv]}_{\xi(v)} \right).$$

This sum can be written as a product:

$$\sum_{\xi \in \Xi_\phi} \text{wt}(\mathcal{G}, \xi) = \prod_{v \in V} H_v,$$

in which for every $v \in V$, the factor H_v is a sum over $\xi(v) \in U_v$.

By the construction of $\mathcal{G}^*$, we can show that

$$\text{wt}(\mathcal{G}^*, \phi) = \sum_{\xi \in \Xi_\phi} \text{wt}(\mathcal{G}, \xi) = \prod_{v \in V} H_v. \tag{3}$$

This follows from the following observations:

1. If v appears in both the subgraph spanned by R_i and the subgraph spanned by R_j, for some $i \neq j \in [n]$, and this v defines an edge $ij \in E^*$, then the edge weight of this edge ij in $\mathcal{G}^*$ with respect to ϕ is exactly H_v;
2. For every $i \in [n]$, we let $V_i \subseteq V$ denote the set of vertices that only appear in the subgraph spanned by R_i. We also let $\mathbf{w}$ denote the vertex weight of $i \in [n]$ in $\mathcal{G}^*$. Then we have

$$w_{\xi(i)} = \prod_{v \in V_i} H_v.$$

As a result, it follows from (3) that

$$Z(\mathcal{G}^*) = \sum_{\phi} \text{wt}(\mathcal{G}^*, \phi) = \sum_{\phi} \sum_{\xi \in \Xi_\phi} \text{wt}(\mathcal{G}, \xi) = Z(\mathcal{G}).$$

D.1.2 Construction of $\mathcal{G}'$ from $\mathcal{G}^*$

Let $\mathcal{G}^* = (G^*, \mathcal{V}^*, \mathcal{E}^*)$ be the labeled directed graph constructed above, where $G^* = (V^*, E^*)$. We know that $\mathcal{V}^*(v) \in \mathfrak{X}^*$ for all $v \in V^*$; $\mathcal{E}^*(uv) \in \mathfrak{Y}^*$ for all $uv \in E^*$; and the underlying undirected graph of G^* is connected. Since $(\mathfrak{X}^*, \mathfrak{Y}^*)$ is a generalized $\mathcal{P}$-pair, every $\mathbf{D} \in \mathfrak{Y}^*$ is either a $\mathcal{P}$-matrix or a $\mathcal{P}$-diagonal matrix.

We will build a new labeled directed graph $\mathcal{G}' = (G', \mathcal{V}', \mathcal{E}')$ with $G' = (V', E')$ such that $\mathcal{V}'(v) \in \mathfrak{X}'$ for all $v \in V'$; $\mathcal{E}'(uv) \in \mathfrak{Y}'$ for all $uv \in E'$; the underlying undirected graph of G' is connected; and

$$Z(\mathcal{G}') = Z(\mathcal{G}^*).$$

Let $E^* = E_0 \cup E_1$, where E_0 consists of the edges in E^* whose weight is a $\mathcal{P}$-matrix and E_1 consists of the edges in E^* whose weight is a $\mathcal{P}$-diagonal matrix. We decompose the vertex set V^* of G^* using the following equivalence relation $\sim$.

Definition 8. *Let v, v' be two distinct vertices in V^*. $v \sim v'$ if v and v' are connected by E_1 (which is viewed as a set of undirected edges here).*

By using $\sim$, we divide V^* into equivalence classes $V_1, \ldots, V_g$ for some $g \geq 1$. This relation is useful because of the following observation:

Observation 2. Let $\phi : V^* \to [n]$ be an assignment with non-zero weight: $\text{wt}(\mathcal{G}^*, \phi) \neq 0$. Then for any $i \in [g]$, there exists a unique $k_i \in [n]$ such that $\phi(v) = k_i$ for all $v \in V_i$.

Now we construct $\mathcal{G}' = (G', \mathcal{V}', \mathcal{E}')$. First we construct $G' = (V', E')$. V' is exactly $[g]$ in which vertex $i \in [g]$ corresponds to V_i. For every edge $uv \in E_0$ such that $u \in V_i$, $v \in V_j$, and $i \neq j \in [g]$, we add an edge from i to j in G'. This finishes the construction of G'. It is easy to verify that the underlying undirected graph of G' is also *connected*.

Finally, we assign vertex and edge weights. For each edge ij in G', suppose it is created because of $uv \in E_0$. Then the edge weight of ij is the same as that of uv. As a result, all the edge weight matrices of $\mathcal{G}'$ come from $\mathfrak{Y}'$ (since by definition of `gen-pair`, $\mathfrak{Y}'$ contains all the $\mathcal{P}$-matrices in $\mathfrak{Y}^*$).

We define the vertex weights of $\mathcal{G}'$ as follows. If $V_i = \{v\}$ is a singleton, then the vertex weight of i in $\mathcal{G}'$ is the same as the weight of v in $\mathcal{G}^*$. Otherwise, we let $v_1, \ldots, v_r$ be the vertices in V_i with $r > 1$, let $e_1, \ldots, e_s$ be the edges in E_1 with both vertices in V_i for some $s \geq 1$, and let $e'_1, \ldots, e'_t$ be the edges in E_0 with both vertices in V_i for some $t \geq 0$. We use $\mathbf{w}^{[j]} \in \mathfrak{X}^*$ to denote the vertex weight of v_j in $\mathcal{G}'$ $\mathbf{C}^{[j]} \in \mathfrak{Y}^*$ to denote the $\mathcal{P}$-diagonal matrix of e_j and $\mathbf{D}^{[j]} \in \mathfrak{Y}^*$ to denote the $\mathcal{P}$-matrix of e'_j. Then we assign the following vertex weight vector $\mathbf{w}$ to $i \in V'$:

$$w_k = w_k^{[1]} \cdots w_k^{[r]} \cdot C_{k,k}^{[1]} \cdots C_{k,k}^{[s]} \cdot D_{k,k}^{[1]} \cdots D_{k,k}^{[t]}, \qquad \text{for every } k \in [n].$$

By definition, we have $\mathbf{w} \in \mathfrak{Y}'$. Using Observation 2, it is also easy to verify that $Z(\mathcal{G}') = Z(\mathcal{G}^*)$.

This completes the proof of Lemma 3.

E Reduction: Normalized Matrices are Free to Use

To give a polynomial-time reduction from $(\mathfrak{X}',\mathfrak{Y}')=\texttt{gen-pair}(\mathfrak{X},\mathfrak{Y})$ to $(\mathfrak{X},\mathfrak{Y})$, we need to first prove a technical lemma on *normalized* block-rank-1 matrices.

Let $\mathbf{C}$ be an $m\times m$ block-rank-1 matrix of block pattern $\mathcal{T}$ and representation $(\boldsymbol{\alpha},\boldsymbol{\beta})$, where $\mathcal{T}=\{(A_1,B_1),\ldots,(A_r,B_r)\}$ for some $r\ge 1$. By definition, $\boldsymbol{\alpha}$ satisfies

$$\sum_{j\in A_i}\alpha_j=1,\qquad \text{for all } i\in[r].$$

We say $\mathbf{C}'$ is the *normalized* version of $\mathbf{C}$ if it is an $m\times m$ block-rank-1 matrix of block pattern $\mathcal{T}$ and representation $(\boldsymbol{\alpha},\boldsymbol{\delta})$, where

$$\delta_j=\frac{\beta_j}{\sum_{k\in B_i}\beta_k},\qquad \text{for all } j\in B_i \text{ and } i\in[r],$$

so that $\boldsymbol{\delta}$ also satisfies

$$\sum_{j\in B_i}\delta_j=1,\qquad \text{for all } i\in[r].$$

Let $(\mathfrak{P},\mathfrak{Q})$ be a finite $\mathcal{T}$-pair for some non-trivial $m\times m$ block pattern $\mathcal{T}$, and

$$\mathfrak{Q}=\big\{\mathbf{C}^{[1]},\ldots,\mathbf{C}^{[s]}\big\},$$

in which every $\mathbf{C}^{[i]}$ is block-rank-1 and has representation $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\beta}^{[i]})$. For each $i\in[s]$, we let $\mathbf{D}^{[i]}$ denote the normalized version of $\mathbf{C}^{[i]}$ with representation $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\delta}^{[i]})$, and

$$\mathfrak{Q}'=\big\{\mathbf{C}^{[1]},\ldots,\mathbf{C}^{[s]},\mathbf{D}^{[1]},\ldots,\mathbf{D}^{[s]}\big\}.$$

In this section, we prove the following technical lemma:

Lemma 8. *$Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ and $Z_{\mathfrak{P},\mathfrak{Q}'}(\cdot)$ are computationally equivalent.*

Proof. In the proof, we use two levels of interpolations and Vandermonde systems.

We start with some notation. Let $\mathcal{G}=(G,\mathcal{V},\mathcal{E})$ be the input labeled directed graph of $Z_{\mathfrak{P},\mathfrak{Q}'}(\cdot)$ with $G=(V,E)$. For $v\in V$, we use $\mathbf{w}^{[v]}\in\mathfrak{P}$ to denote its vertex weight. We use $E_i\subseteq E$, $i\in[s]$, to denote the set of edges labeled with $\mathbf{C}^{[i]}$, and $F_i\subseteq E$, $i\in[s]$, to denote the set of edges labeled with $\mathbf{D}^{[i]}$. For every assignment $\xi:V\to[m]$, we define

$$\text{vw}(\xi)=\prod_{v\in V}w^{[v]}_{\xi(v)},\qquad \text{cw}(\xi)=\prod_{i\in[s]}\prod_{uv\in E_i}C^{[i]}_{\xi(u),\xi(v)},\qquad \text{dw}(\xi)=\prod_{i\in[s]}\prod_{uv\in F_i}D^{[i]}_{\xi(u),\xi(v)}.$$

Note that a product over an empty set is equal to 1.
Then we need to compute the following sum

$$Z_{\mathfrak{P},\mathfrak{Q}'}(\mathcal{G})=\sum_{\xi}\text{vw}(\xi)\cdot\text{cw}(\xi)\cdot\text{dw}(\xi).$$

For all $a\in[s]$ and $b\in[r]$, we use $K^{[a]}_b>0$ to denote the number such that

$$C^{[a]}_{i,j}=K^{[a]}_b\cdot D^{[a]}_{i,j},\qquad \text{for all } i\in A_b \text{ and } j\in B_b.$$

Actually, this gives us the following equation

$$C^{[a]}_{i,j} = K^{[a]}_b \cdot D^{[a]}_{i,j}, \qquad \text{for all } i \in A_b \text{ and } j \in [m],$$

since $\mathbf{C}^{[a]}$ and $\mathbf{D}^{[a]}$ have the same block pattern $\mathcal{T}$. Then we use $\text{kw}(\xi)$, where $\xi : V \to [m]$, to denote

$$\text{kw}(\xi) = \prod_{a\in[s]} \left(\prod_{uv\in F_a \text{ with } \xi(u)\in A_b} K^{[a]}_b \right).$$

We use X to denote the set of all possible values of $\text{kw}(\xi)$:

$$X = \big\{ \text{kw}(\xi) \,\big|\, \xi : V \to [m] \big\}.$$

It can be checked that $|X|$ is polynomial in $|E|$ since both s and r are considered as constants here. We use L to denote $|X|$.

For all $k \in [0 : L-1]$, we build a new graph $\mathcal{G}^{[k]} = (G^{[k]}, \mathcal{V}^{[k]}, \mathcal{E}^{[k]})$, where $G^{[k]} = (V^{[k]}, E^{[k]})$:

1. $V \subseteq V^{[k]}$ and every $v \in V$ is labeled with the same vertex weight as in $\mathcal{G}$;

2. For all $i \in [s]$ and $uv \in E_i$, we add one edge $uv \in E^{[k]}$ and label it with the same matrix $\mathbf{C}^{[i]}$;

3. For all $i \in [s]$ and all $e = uv \in F_i$, we add $L-k$ parallel edges from u to v with $\mathbf{C}^{[i]}$ as their edge weights; we also add $2k$ new vertices $u_{e,j}$ and $v_{e,j}$, $j \in [k]$, to $V^{[k]}$; we add one edge from u to $u_{e,j}$ and one edge from $v_{e,j}$ to v for all $j \in [k]$, all of which are labeled with $\mathbf{C}^{[i]}$. For each new vertex, we assign $\mathbf{1}$ as its vertex weight.

It is clear that $\mathcal{G}^{[k]}$ can be constructed in polynomial time and is a valid input of $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$.

Fix $k \in [0 : L-1]$. For every assignment $\phi : V \to [m]$, we let Ξ_ϕ denote the set of all $\xi : V^{[k]} \to [m]$ such that $\xi(v) = \phi(v)$ for all $v \in V$. We also define

$$\text{wt}^{[k]}(\phi) = \sum_{\xi\in\Xi_\phi} \text{wt}(\mathcal{G}^{[k]}, \xi).$$

Then we have the following equation

$$Z_{\mathfrak{P},\mathfrak{Q}}(\mathcal{G}^{[k]}) = \sum_{\xi:V^{[k]}\to[m]} \text{wt}(\mathcal{G}^{[k]}, \xi) = \sum_{\phi:V\to[m]} \text{wt}^{[k]}(\phi).$$

By the construction, we show that

$$\text{wt}^{[k]}(\phi) = \text{vw}(\phi) \cdot \text{cw}(\phi) \cdot \Big(\text{dw}(\phi)\Big)^L \cdot \Big(\text{kw}(\phi)\Big)^{L+k}, \qquad \text{for all } k \in [0 : L-1]. \tag{4}$$

First, we have

$$\text{wt}^{[k]}(\phi) = \text{vw}(\phi) \cdot \text{cw}(\phi) \cdot \sum_{\xi\in\Xi_\phi} \left(\prod_{i\in[s]} \left(\prod_{e=uv\in F_i} \Big(C^{[i]}_{\xi(u),\xi(v)} \Big)^{L-k} \left(\prod_{j\in[k]} C^{[i]}_{\xi(u),\xi(u_{e,j})} C^{[i]}_{\xi(v_{e,j}),\xi(v)} \right) \right) \right). \tag{5}$$

For each edge $e = uv \in F_i$ for some $i \in [s]$, there must exist an index $b_e \in [r]$ such that $\phi(u) \in A_{b_e}$ and $\phi(v) \in B_{b_e}$; otherwise both sides of (4) are 0 and we are done. In this case, the sum in (5) becomes

$$\prod_{i\in[s]}\left(\prod_{e=uv\in F_i}\left(K^{[i]}_{b_e}\cdot D^{[i]}_{\xi(u),\xi(v)}\right)^{L-k}\left(\sum_{x\in B_{b_e}} C^{[i]}_{\xi(u),x}\right)^k\left(\sum_{x\in A_{b_e}} C^{[i]}_{x,\xi(v)}\right)^k\right). \tag{6}$$

By the definition of $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\delta}^{[i]})$, we have

$$\sum_{x\in B_{b_e}} C^{[i]}_{\xi(u),x}=\alpha^{[i]}_{\xi(u)}\sum_{x\in B_{b_e}}\beta^{[i]}_x=\alpha^{[i]}_{\xi(u)}\cdot K^{[i]}_{b_e}\qquad\text{and}\qquad\sum_{x\in A_{b_e}} C^{[i]}_{x,\xi(v)}=\beta^{[i]}_{\xi(v)}.$$

As a result, (6) becomes

$$\prod_{i\in[s]}\left(\prod_{e=uv\in F_i}\left(K^{[i]}_{b_e}\cdot D^{[i]}_{\xi(u),\xi(v)}\right)^{L-k}\left(\alpha^{[i]}_{\xi(u)}\cdot K^{[i]}_{b_e}\right)^k\left(\beta^{[i]}_{\xi(v)}\right)^k\right)=\prod_{i\in[s]}\left(\prod_{e=uv\in F_i}\left(K^{[i]}_{b_e}\right)^{L+k}\left(D^{[i]}_{\xi(u),\xi(v)}\right)^L\right).$$

This finishes the proof of equation (4).

Since L is polynomial in the input size, we can use $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ as an oracle to compute

$$\sum_{\phi:V\to[m]}\text{vw}(\phi)\cdot\text{cw}(\phi)\cdot\Big(\text{dw}(\phi)\Big)^L\cdot\Big(\text{kw}(\phi)\Big)^{L+k},\qquad\text{for all }k\in[0:L-1].$$

in a polynomial number of steps.

For every $x\in X$, we use Φ_x to denote the set of $\phi:V\to[m]$ with $\text{kw}(\phi)=x$, then we computed

$$\sum_{x\in X}\left(\sum_{\phi\in\Phi_x}\text{vw}(\phi)\cdot\text{cw}(\phi)\cdot\Big(\text{dw}(\phi)\Big)^L\right)\cdot x^{L+k},\qquad\text{for all }k\in[0:L-1].$$

Because $x>0$ for all $x\in X$, we can solve this Vandermonde system and obtain

$$\sum_{\phi\in\Phi_x}\text{vw}(\phi)\cdot\text{cw}(\phi)\cdot\Big(\text{dw}(\phi)\Big)^L,\qquad\text{for each }x\in X,$$

in a polynomial number of steps.

It is also clear that the whole process can be repeated for any $L'\ge L$ with

$$L'\le L+\text{poly(input size)},$$

and we can use $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ as an oracle to compute

$$\sum_{\phi\in\Phi_x}\text{vw}(\phi)\cdot\text{cw}(\phi)\cdot\Big(\text{dw}(\phi)\Big)^{L'},\qquad\text{for all }x\in X\text{ and }L\le L'\le L+\text{poly(input size)},$$

in a polynomial number of steps.

Next we use Y to denote the set of all possible values of $\text{dw}(\phi)$, $\phi:V\to[m]$ (note it is possible that $0\in Y$). Again, $|Y|$ is polynomial and we use M to denote $|Y|$. For every $x\in X$, we can compute

$$\sum_{\phi\in\Phi_x}\text{vw}(\phi)\cdot\text{cw}(\phi)\cdot\Big(\text{dw}(\phi)\Big)^{L+k},\qquad\text{for all }k\in[0:M-1].$$

Let $\Phi_{x,y}$ denote the set of ϕ with $\text{kw}(\phi)=x$ and $\text{dw}(\phi)=y$. Solving this Vandermonde system, we get

$$\sum_{\phi\in\Phi_{x,y}} \text{vw}(\phi)\cdot\text{cw}(\phi), \qquad \text{for all } x\in X \text{ and } 0<y\in Y.$$

Finally, using all these items, we can compute $Z_{\mathfrak{P},\mathfrak{Q}'}(\mathcal{G})$ in a polynomial number of steps:

$$Z_{\mathfrak{P},\mathfrak{Q}'}(\mathcal{G}) = \sum_{x\in X,\, 0<y\in Y}\left(\sum_{\phi\in\Phi_{x,y}} \text{vw}(\phi)\cdot\text{cw}(\phi)\right)\cdot y.$$

This proves the lemma since the other direction from $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ to $Z_{\mathfrak{P},\mathfrak{Q}'}(\cdot)$ is trivial. □

F Polynomial-Time Reduction from $(\mathfrak{X}',\mathfrak{Y}')$ to $(\mathfrak{X},\mathfrak{Y})$

Let $(\mathfrak{X},\mathfrak{Y})$ be a $\mathcal{T}$-pair, where $\mathcal{T}$ is a non-trivial $m\times m$ block pattern $\mathcal{T}=\{(A_1,B_1),\ldots,(A_r,B_r)\}$ with $r\ge 1$ and every matrix in $\mathfrak{Y}$ is block-rank-1. Let $\mathcal{P}$ be the $r\times r$ pattern where $\mathcal{P}=\texttt{gen}(\mathcal{T})$ and $(\mathfrak{X}',\mathfrak{Y}')$ be the $\mathcal{P}$-pair generated from $(\mathfrak{X},\mathfrak{Y})$ using the `gen-pair` operation: $(\mathfrak{X}',\mathfrak{Y}')=\texttt{gen-pair}(\mathfrak{X},\mathfrak{Y})$. We also use $(\mathfrak{X}^*,\mathfrak{Y}^*)$ to denote the generalized $\mathcal{P}$-pair defined in Appendix C.

In this section, we prove that $(\mathfrak{X}',\mathfrak{Y}')$ is polynomial-time reducible to $(\mathfrak{X},\mathfrak{Y})$. To this end, we first reduce $(\mathfrak{X}',\mathfrak{Y}')$ to $(\mathfrak{X}^*,\mathfrak{Y}^*)$, and then reduce $(\mathfrak{X}^*,\mathfrak{Y}^*)$ to $(\mathfrak{X},\mathfrak{Y})$. The first step is trivial, so we will only give a polynomial-time reduction from $(\mathfrak{X}^*,\mathfrak{Y}^*)$ to $(\mathfrak{X},\mathfrak{Y})$ below.

Let $\mathfrak{P}^*=\{\mathbf{p}^{[i]}: i\in[s]\}$ be a finite subset of vectors in $\mathfrak{X}^*$ with $\mathbf{1}\in\mathfrak{P}^*$ and $\mathfrak{Q}^*=\{\mathbf{F}^{[i]}: i\in[t]\}$ be a finite subset of matrices in $\mathfrak{Y}^*$. By the definition of `gen-pair`, they can be generated by a finite subset $\mathfrak{P}=\{\mathbf{w}^{[i]}: i\in[h]\}\subseteq\mathfrak{X}$ with $\mathbf{1}\in\mathfrak{P}$ and a finite subset $\mathfrak{Q}=\{\mathbf{C}^{[i]}: i\in[g]\}\subseteq\mathfrak{Y}$ in the following sense. (We let $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\beta}^{[i]})$ denote the representation of $\mathbf{C}^{[i]}$ for every $i\in[g]$.)

For every matrix $\mathbf{F}\in\mathfrak{Q}^*$, there exists a $(2g+1)$-tuple

$$\Big(k\in[h];\mathbf{k}=(k_1,\ldots,k_g);\boldsymbol{\ell}=(\ell_1,\ldots,\ell_g)\Big),$$

where $k_i,\ell_i\ge 0$, $\mathbf{k}\ne\mathbf{0}$ and $\boldsymbol{\ell}\ne\mathbf{0}$, such that

$$F_{i,j}=\sum_{x\in B_i\cap A_j}\left(\beta_x^{[1]}\right)^{k_1}\cdots\left(\beta_x^{[g]}\right)^{k_g}\cdot\left(\alpha_x^{[1]}\right)^{\ell_1}\cdots\left(\alpha_x^{[g]}\right)^{\ell_g}\cdot w_x^{[k]}. \tag{7}$$

This $(2g+1)$-tuple is also call the (not necessarily unique) representation of $\mathbf{F}$ with respect to $(\mathfrak{P},\mathfrak{Q})$.

For every $\mathbf{p}\in\mathfrak{P}^*$, there exist three finite (and possibly empty) sets $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_3$ of tuples, where every tuple in $\mathcal{S}_1$ and $\mathcal{S}_2$ is of the form

$$\Big(k\in[h];\mathbf{k}=(k_1,\ldots,k_g)\Big)$$

with $k_i\ge 0$ and $\mathbf{k}\ne\mathbf{0}$, and every tuple in $\mathcal{S}_3$ is of the form

$$\Big(k\in[h];\mathbf{k}=(k_1,\ldots,k_g);\boldsymbol{\ell}=(\ell_1,\ldots,\ell_g)\Big)$$

with $k_i,\ell_i\ge 0$, $\mathbf{k}\ne\mathbf{0}$ and $\boldsymbol{\ell}\ne\mathbf{0}$. Every tuple in $\mathcal{S}_1$ gives us a vector whose ith entry, $i\in[r]$, is equal to

$$\sum_{x\in A_i}\left(\alpha_x^{[1]}\right)^{k_1}\cdots\left(\alpha_x^{[g]}\right)^{k_g}\cdot w_x^{[k]};$$

every tuple in $\mathcal{S}_2$ gives us a vector whose ith entry, $i \in [r]$, is equal to

$$\sum_{x \in B_i} \left(\beta_x^{[1]}\right)^{k_1} \cdots \left(\beta_x^{[g]}\right)^{k_g} \cdot w_x^{[k]};$$

and every $(2g+1)$-tuple in $\mathcal{S}_3$ gives us a vector whose ith entry, $i \in [r]$, is equal to

$$\sum_{x \in B_i \cap A_i} \left(\beta_x^{[1]}\right)^{k_1} \cdots \left(\beta_x^{[g]}\right)^{k_g} \cdot \left(\alpha_x^{[1]}\right)^{\ell_1} \cdots \left(\alpha_x^{[g]}\right)^{\ell_g} \cdot w_x^{[k]}.$$

Vector $\mathbf{p}$ is then the Hadamard product of all these vectors.

We remark that all the exponents k_i, ℓ_i in the equations above are considered as constants, because both $(\mathfrak{P}, \mathfrak{Q})$ and $(\mathfrak{P}^*, \mathfrak{Q}^*)$ are fixed. We now prove the following lemma.

Lemma 9. *$Z_{\mathfrak{P}^*, \mathfrak{Q}^*}(\cdot)$ is polynomial-time reducible to $Z_{\mathfrak{P}, \mathfrak{Q}}(\cdot)$.*

F.1 Proof Sketch

We first give a proof sketch. Again, we will use interpolations and Vandermonde systems.

First, by Lemma 8, we only need to give a reduction from $Z_{\mathfrak{P}^*, \mathfrak{Q}^*}(\cdot)$ to $Z_{\mathfrak{P}, \mathfrak{R}}(\cdot)$, where

$$\mathfrak{R} = \left\{ \mathbf{C}^{[i]}, \mathbf{D}^{[i]} : i \in [g] \right\}$$

contains both $\mathbf{C}^{[i]}$ and its *normalized* version $\mathbf{D}^{[i]}$, $i \in [g]$.

Let $\mathcal{G} = (G, \mathcal{V}, \mathcal{E})$ be an input labeled graph of $Z_{\mathfrak{P}^*, \mathfrak{Q}^*}(\cdot)$, where $G = (V, E)$. For every assignment $\xi : V \to [r]$, we will define $\text{nvw}(\xi) > 0$. Moreover, let X be the set of all possible values of $\text{nvw}(\xi)$, and $L = |X|$, then L is polynomially bounded. For every $k \in [L]$, we will build a new labeled directed graph $\mathcal{G}^{[k]}$ from $\mathcal{G}$. $\mathcal{G}^{[k]}$ is a valid input graph of $Z_{\mathfrak{P}, \mathfrak{R}}(\cdot)$ (with domain $[m]$) and satisfies

$$Z_{\mathfrak{P}, \mathfrak{R}}(\mathcal{G}^{[k]}) = \sum_{\xi : V \to [r]} \text{wt}(\mathcal{G}, \xi) \cdot \Big(\text{nvw}(\xi)\Big)^k. \tag{8}$$

For each $x \in X$, we use Ξ_x to denote the set of all $\xi : V \to [r]$ with $\text{nvw}(\xi) = x$. Then by solving the Vandermonde system which consists of equations (8) for $k = 1, 2, \ldots, L$, we can compute

$$\sum_{\xi \in \Xi_x} \text{wt}(\mathcal{G}, \xi), \qquad \text{for every } x \in X,$$

which allow us to compute in polynomial time

$$Z_{\mathfrak{P}^*, \mathfrak{Q}^*}(\mathcal{G}) = \sum_{\xi : V \to [r]} \text{wt}(\mathcal{G}, \xi) = \sum_{x \in X} \left(\sum_{\xi \in \Xi_x} \text{wt}(\mathcal{G}, \xi) \right).$$

F.2 Construction of $\mathcal{G}^{[k]}$

We start with the construction of $\mathcal{G}^{[1]} = (G^{[1]}, \mathcal{V}^{[1]}, \mathcal{E}^{[1]})$. It will become clear that the construction can be generalized to get $\mathcal{G}^{[k]}$ for every $k \in [L]$.

Let $V = [n]$, then the vertex set $V^{[1]}$ of $G^{[1]} = (V^{[1]}, E^{[1]})$ will be defined as a union:

$$V^{[1]} = R_1 \cup R_2 \cup \cdots \cup R_n,$$

where R_k corresponds to vertex $k \in V$ and any edge $uv \in E^{[1]}$ will be between two vertices $u, v \in V^{[1]}$ such that $u, v \in R_k$ for some unique $k \in [n]$. R_i and R_j, $i \neq j \in [n]$, are not necessarily disjoint and there could be vertices shared by (at most) two different sets R_i and R_j. We further divide the vertices of R_i, $i \in [n]$, into three types: In the subgraph of $G^{[1]}$ spanned by R_i,

1. The Type-L vertices only have outgoing edges;
2. The Type-R vertices only have incoming edges; and
3. The Type-M vertices have both incoming and outgoing edges.

When adding a new vertex, we will also specify which type it is. The construction also guarantees that the underlying undirected graph spanned by every R_i is connected.

F.2.1 Construction of $G^{[1]} = (V^{[1]}, E^{[1]})$

We start with the vertex set $V^{[1]}$.

1. First, for every $i \in [n]$ and $a \in [g]$, we add a new Type-L vertex $u_{i,a}$ in R_i and add a new Type-R vertex $w_{i,a}$ in R_i. All these vertices appear in R_i only.
2. Second, for every $e = ij \in E$, where $i, j \in [n]$, we add a vertex $v_e \in R_i \cap R_j$, which is a Type-R vertex in R_i and a Type-L vertex in R_j.
3. Finally, for every $i \in V$ let $\mathbf{p} \in \mathfrak{P}^*$ be its vertex weight in $\mathcal{G}$. Then by the discussion earlier, it can be generated from $(\mathfrak{P}, \mathfrak{Q})$ using three finite sets of tuples $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$. For each tuple $\mathbf{s}$ in $\mathcal{S}_1$ we add a new Type-L vertex $v_{i,\mathbf{s}}$ in R_i; for each tuple $\mathbf{s}$ in $\mathcal{S}_2$, we add a new Type-R vertex in R_i; and for each tuple $\mathbf{s}$ in $\mathcal{S}_3$ we add a new Type-M vertex in R_i. All these vertices appear in R_i only.

We will add some more vertices later. Now we start to create edges, and assign edge/vertex weights.

First, for every $i \in [n]$, we add $2g$ edges to connect $u_{i,a}$ and $w_{i,a}$, $a \in [g]$:

1. For every $a \in [g]$, add one edge from $u_{i,a}$ to $w_{i,a}$, and label the edge with $\mathbf{C}^{[1]}$;
2. For every $a \in [g]$, add one edge from $u_{i,a}$ to $w_{i,a+1}$ (with $w_{i,g+1} = w_{i,1}$), and label it with $\mathbf{C}^{[1]}$;
3. For every $a \in [g]$, the vertex weight vector of both $u_{i,a}$ and $w_{i,a}$ is the all-one vector $\mathbf{1}$.

Second, for each edge $e = ij \in E$, we add the incident edges of $v_e \in R_i \cap R_j$ as follows. Assume the edge weight matrix of ij in $\mathcal{G}$ is generated by $(\mathfrak{P}, \mathfrak{Q})$ using the following $(2g+1)$-tuple:

$$\big(k \in [h]; \mathbf{k} = (k_1, \ldots, k_g); \boldsymbol{\ell} = (\ell_1, \ldots, \ell_g)\big),$$

where $k_i, \ell_i \geq 0$, $\mathbf{k} \neq \mathbf{0}$ and $\boldsymbol{\ell} \neq \mathbf{0}$. Then we add the following incident edges of v_e:

1. For each $b \in [g]$, we add k_b parallel edges from $u_{i,b}$ to v_e in R_i, all of which are labeled with $\mathbf{C}^{[b]}$;
2. For each $b \in [g]$, we add ℓ_b parallel edges from v_e to $w_{j,b}$ in R_j, all of which are labeled with $\mathbf{C}^{[b]}$;
3. Assign the vertex weight vector $\mathbf{w}^{[k]} \in \mathfrak{P}$ to v_e.

Finally, for every vertex $i \in V$ we use $\mathbf{p}$ to denote its vertex weight in $\mathcal{G}$. Assume $\mathbf{p}$ is generated by $(\mathfrak{P}, \mathfrak{Q})$ using three finite sets $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$ of tuples. For each $\mathbf{s} = (k \in [h]; \mathbf{k} = (k_1, \ldots, k_g))$ in $\mathcal{S}_1$ with $k_i \geq 0$ and $\mathbf{k} \neq \mathbf{0}$, we already added a Type-L vertex $v_{i,\mathbf{s}}$ in R_i (which appears in R_i only). We add the following incident edges of $v_{i,\mathbf{s}}$:

1. For each $b \in [g]$, add k_b parallel edges from $v_{i,\mathbf{s}}$ to $w_{i,b}$ in R_i, all of which are labeled with $\mathbf{C}^{[b]}$;
2. Assign the vertex weight vector $\mathbf{w}^{[k]} \in \mathfrak{P}$ to $v_{i,\mathbf{s}}$.

For every $\mathbf{s} = (k \in [h]; \mathbf{k} = (k_1, \ldots, k_g))$ in $\mathcal{S}_2$, we already added a Type-R vertex $v_{i,\mathbf{s}} \in R_i$. We add the following incident edges of $v_{i,\mathbf{s}}$ in R_i:

1. For each $b \in [g]$, add k_b parallel edges from $u_{i,b}$ to $v_{i,\mathbf{s}}$ in R_i, all of which are labeled with $\mathbf{C}^{[b]}$;
2. Assign the vertex weight vector $\mathbf{w}^{[k]} \in \mathfrak{P}$ to $v_{i,\mathbf{s}}$.

For every tuple $\mathbf{s} = (k \in [h]; \mathbf{k} = (k_1, \ldots, k_g); \boldsymbol{\ell} = (\ell_1, \ldots, \ell_g))$ in $\mathcal{S}_3$, we already added a Type-M vertex $v_{i,\mathbf{s}}$ in R_i. We add the following incident edges of $v_{i,\mathbf{s}}$ in R_i:

1. For every $b \in [g]$, add k_b parallel edges from $u_{i,b}$ to $v_{i,\mathbf{s}}$, all of which are labeled with $\mathbf{C}^{[b]}$;
2. For every $b \in [g]$, add ℓ_b parallel edges from $v_{i,\mathbf{s}}$ to $w_{i,b}$, all of which are labeled with $\mathbf{C}^{[b]}$; and
3. Assign the vertex weight vector $\mathbf{w}^{[k]} \in \mathfrak{P}$ to $v_{i,\mathbf{s}}$.

It can be checked that the (undirected) subgraph spanned by R_i, for all $i \in [n]$, is connected.

This almost finishes the construction. The only thing left is to add some more vertices and edges so that the out-degree of $u_{i,a}$ and the in-degree of $w_{i,a}$ are the same for all $i \in [n]$ and $a \in [g]$.

To this end, we notice that for all $i \in [n]$ and $a \in [g]$, both the out-degree of $u_{i,a}$ and the in-degree of $w_{i,a}$ constructed so far are linear in the maximum degree of G, because all the parameters k_i, ℓ_i and the sets $\mathcal{S}_i$ are considered as constants. As a result, we can pick a large enough positive integer $M \geq 2$ which is linear in the maximum degree of G, such that

$$M \geq \text{the out-degree of } u_{i,a} \text{ and the in-degree of } w_{i,a} \text{ constructed so far, for all } i \text{ and } a.$$

We now add vertices and edges so that the out-degree of $u_{i,a}$ and the in-degree of $w_{i,a}$ all become M.

Let $i \in [n]$ and $a \in [g]$. Assume the current out-degree of $u_{i,a}$ is $k \leq M$. Then we add $M - k$ new Type-R vertices in R_i and add one edge from $u_{i,a}$ to each of these vertices. The vertex weights of all the new vertices are $\mathbf{1}$, and the edge weights of all the new edges are $\mathbf{D}^{[a]}$ (recall that we are allowed to use the normalized version $\mathbf{D}^{[a]}$ of $\mathbf{C}^{[a]}$, and this is actually the only place we use it).

Similarly, assume the current in-degree of $w_{i,a}$ is $k \leq M$. Then we add $M - k$ new Type-L vertices in R_i and add one edge from each of these vertices to $w_{i,a}$. The vertex weights of all the new vertices are $\mathbf{1}$ while the edge weights of all the new edges are $\mathbf{C}^{[a]}$.

This finishes the construction of the new labeled directed graph $\mathcal{G}^{[1]} = (G^{[1]}, \mathcal{V}^{[1]}, \mathcal{E}^{[1]})$.

F.3 Proof of Equation (8)

We start with the definition of $\text{nvw}(\xi)$, for any assignment $\xi : V = [n] \to [r]$.

First, for each $a \in [g]$, we let $\boldsymbol{\mu}^{[a]}$ denote the following positive r-dimensional vector:

$$\mu_i^{[a]} = \sum_{x \in A_i} \left(\alpha_x^{[1]}\right)^2 \cdot \left(\alpha_x^{[a]}\right)^{M-2}, \qquad \text{for every } i \in [r].$$

For every $a \in [g]$, we let $\boldsymbol{\nu}^{[a]}$ denote the following positive r-dimensional vector:

$$\nu_i^{[a]} = \sum_{x \in B_i} \left(\beta_x^{[1]}\right)^2 \cdot \left(\beta_x^{[a]}\right)^{M-2}, \qquad \text{for every } i \in [r].$$

Finally, we define $\text{nvw}(\xi)$ as follows:

$$\text{nvw}(\xi) = \prod_{i\in[n]} \prod_{a\in[g]} \mu^{[a]}_{\xi(i)} \cdot \nu^{[a]}_{\xi(i)}, \qquad \text{for any } \xi : V = [n] \to [r].$$

It is easy to check that $\text{nvw}(\xi) > 0$ and the number of possible values of $\text{nvw}(\xi)$ is polynomial in n.

Now we prove equation (8) for $k = 1$:

$$Z_{\mathfrak{P},\mathfrak{R}}(\mathcal{G}^{[1]}) = \sum_{\xi:V\to[r]} \text{wt}(\mathcal{G}, \xi) \cdot \text{nvw}(\xi). \tag{9}$$

Let ξ be an assignment from V to $[r]$. We use Φ_ξ to denote the set of all assignments $\phi : V^{[1]} \to [m]$ such that for every edge uv in the subgraph spanned by R_i, $i \in [n]$, we have

$$\phi(u) \in A_{\xi(i)} \quad \text{and} \quad \phi(v) \in B_{\xi(i)}.$$

In other words, for all $i \in [n]$ and $v \in R_i$, if v a Type-L vertex then $\phi(v) \in A_{\xi(i)}$; if v is a Type-R vertex then $\phi(v) \in B_{\xi(i)}$; and if v is a Type-M of R_i, then $\phi(v) \in A_{\xi(i)} \cap B_{\xi(i)}$. Equivalently, we can associate every vertex $v \in V^{[1]}$ with a subset $U_v \subseteq [m]$, where

1. If v appears in both R_i and R_j for some $i \neq j \in V = [n]$, and v is Type-R in R_i and Type-L in R_j, then $U_v = B_{\xi(i)} \cap A_{\xi(j)}$;
2. Otherwise, assume v only appears in R_i for some $i \in V = [n]$. Then
   (a) If v is Type-L, then $U_v = A_{\xi(i)}$;
   (b) If v is Type-R, then $U_v = B_{\xi(i)}$; and
   (c) If v is Type-M, then $U_v = B_{\xi(i)} \cap A_{\xi(i)}$,

such that $\phi \in \Phi_\xi$ if and only if $\phi(v) \in U_v$ for all $v \in V^{[1]}$. In particular, $\Phi_\xi = \emptyset$ iff $U_v = \emptyset$ for some v.

By the construction, we know the subgraph spanned by R_i is *connected*, for any $i \in [n]$. It implies that $\text{wt}(\mathcal{G}^{[1]}, \phi) \neq 0$ only if $\phi \in \Phi_\xi$ for a unique $\xi : V \to [r]$. As a result, we have

$$Z_{\mathfrak{P},\mathfrak{R}}(\mathcal{G}^{[1]}) = \sum_{\phi} \text{wt}(\mathcal{G}^{[1]}, \phi) = \sum_{\xi} \sum_{\phi\in\Phi_\xi} \text{wt}(\mathcal{G}^{[1]}, \phi),$$

and to prove (9) we only need to show that

$$\sum_{\phi\in\Phi_\xi} \text{wt}(\mathcal{G}^{[1]}, \phi) = \text{wt}(\mathcal{G}, \xi) \cdot \text{nvw}(\xi), \qquad \text{for any assignment } \xi : V = [n] \to [r].$$

We use $\mathbf{w}^{[v]}$ to denote the weight of $v \in V^{[1]}$, E_i to denote the set of edges in $E^{[1]}$ labeled with $\mathbf{C}^{[i]}$, and F_i to denote the set of edges in $E^{[1]}$ labeled with $\mathbf{D}^{[i]}$, then we have

$$\sum_{\phi\in\Phi_\xi} \text{wt}(\mathcal{G}^{[1]}, \phi) = \sum_{\phi\in\Phi_\xi} \left(\prod_{v\in V^{[1]}} \mathbf{w}^{[v]}_{\phi(v)} \prod_{i\in[g]} \left(\prod_{uv\in E_i} C^{[i]}_{\phi(u),\phi(v)} \right) \left(\prod_{uv\in F_i} D^{[i]}_{\phi(u),\phi(v)} \right) \right).$$

By the definition of Φ_ξ, if $\Phi_\xi \neq \emptyset$, then every $\phi \in \Phi_\xi$ satisfies

$$C^{[i]}_{\phi(u),\phi(v)} = \alpha^{[i]}_{\phi(u)} \cdot \beta^{[i]}_{\phi(v)} \quad \text{and} \quad D^{[i]}_{\phi(u),\phi(v)} = \alpha^{[i]}_{\phi(u)} \cdot \delta^{[i]}_{\phi(v)},$$

where $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\delta}^{[i]})$ is the representation of $\mathbf{D}^{[i]}$. As a result, we have

$$\sum_{\phi\in\Phi_\xi} \text{wt}(\mathcal{G}^{[1]}, \phi) = \sum_{\phi\in\Phi_\xi} \left(\prod_{v\in V^{[1]}} \mathbf{w}^{[v]}_{\phi(v)} \prod_{i\in[g]} \left(\prod_{uv\in E_i} \alpha^{[i]}_{\phi(u)} \cdot \beta^{[i]}_{\phi(v)} \right) \left(\prod_{uv\in F_i} \alpha^{[i]}_{\phi(u)} \cdot \delta^{[i]}_{\phi(v)} \right) \right),$$

Because $\phi \in \Phi_\xi$ iff $\phi(v) \in U_v$ for all v, we can express this sum of products as a product of sums:

$$\prod_{v\in V^{[1]}} H_v,$$

in which every H_v, $v \in V^{[1]}$, is a sum over $\phi(v) \in U_v$.

Finally, we show the following equation:

$$\prod_{v\in V^{[1]}} H_v = \text{wt}(\mathcal{G}, \xi) \cdot \text{nvw}(\xi). \tag{10}$$

This follows from the construction of $\mathcal{G}^{[1]}$ and the following observations:

1. For each $v_e \in R_i \cap R_j$, which is added because of edge $ij \in E$, it can be checked that the sum H_{v_e} over $U_{v_e} = B_{\xi(i)} \cap A_{\xi(j)}$ is exactly $F_{\xi(i),\xi(j)}$, where $\mathbf{F}$ is the weight of ij in $\mathcal{G}$ (as defined in (7)).
2. Let $\mathbf{p}$ denote the vertex weight of $i \in V$, which is generated using $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$. Then we have

$$p_{\xi(i)} = \prod_{\mathbf{s}\in\mathcal{S}_1} H_{v_{i,\mathbf{s}}} \prod_{\mathbf{s}\in\mathcal{S}_2} H_{v_{i,\mathbf{s}}} \prod_{\mathbf{s}\in\mathcal{S}_3} H_{v_{i,\mathbf{s}}}.$$

3. For all $i \in [n]$ and $a \in [g]$, we have

$$\mu^{[a]}_{\xi(i)} = H_{u_{i,a}} \quad \text{and} \quad \nu^{[a]}_{\xi(i)} = H_{w_{i,a}}.$$

4. Finally, it can be checked that $H_v = 1$ for all other vertices in $V^{[1]}$, which is the reason we need to use the normalized matrices $\mathbf{D}^{[a]}$ in the construction.

F.3.1 Construction of $\mathcal{G}^{[k]}$

We can similarly construct $\mathcal{G}^{[k]}$ for every $k \in [L]$.

The only difference is that, instead of $u_{i,a}$ and $w_{i,a}$, we add the following $2kg$ vertices in R_i:

$$u_{i,j,a} \text{ and } w_{i,j,a}, \qquad \text{for all } j \in [k] \text{ and } a \in [g].$$

We also connect these vertices by adding $4kg$ edges, whose underlying undirected graph is a cycle. All these edges are labeled with $\mathbf{C}^{[1]}$. We also add extra vertices and edges so that the out-degree of $u_{i,j,a}$ and the in-degree of $v_{i,j,a}$ are M for all $i \in [n]$, $j \in [k]$ and $a \in [g]$. It then can be proved similarly that

$$Z_{\mathfrak{P},\mathfrak{R}}(\mathcal{G}^{[k]}) = \sum_{\xi:V\to[r]} \text{wt}(\mathcal{G}, \xi) \cdot \Big(\text{nvw}(\xi)\Big)^k.$$

This completes the proof of Lemma 2.

G Decidability

In this section, we show that the rank condition is decidable in a finite number of steps.

G.1 A Technical Lemma

We prove a very useful technical lemma.

Lemma 10. *Let $L, n, m \geq 1$ be positive integers. For every $i \in [L]$, let $\{a_1^{[i]}, \ldots, a_n^{[i]}\}$ be a sequence of n positive numbers; and let $\{b_1^{[i]}, \ldots, b_m^{[i]}\}$ be a sequence of m positive numbers. If*

$$\sum_{i\in[n]} \prod_{j\in[L]} \left(a_i^{[j]}\right)^{k_j} = \sum_{i\in[m]} \prod_{j\in[L]} \left(b_i^{[j]}\right)^{k_j}, \qquad \text{for all } k_1, k_2, \ldots, k_L \geq 1,$$

then $m = n$ and there exists a one-to-one correspondence π from $[n]$ to itself such that

$$a_i^{[j]} = b_{\pi(i)}^{[j]}, \qquad \text{for all } i \in [n] \text{ and } j \in [L].$$

Proof. We prove it by induction on L. The base case when $L = 1$ is trivial.

Assume the lemma is true for $L - 1 \geq 1$. Without loss of generality, we assume that $\{a_1^{[L]}, \ldots, a_n^{[L]}\}$ and $\{b_1^{[L]}, \ldots, b_m^{[L]}\}$ are already sorted:

$$a_1^{[L]} \geq \ldots \geq a_n^{[L]} > 0 \qquad \text{and} \qquad b_1^{[L]} \geq \ldots \geq b_m^{[L]} > 0.$$

We let $s \geq 1$ and $t \geq 1$ be the two maximum integers such that

$$a_1^{[L]} = a_2^{[L]} = \cdots = a_s^{[L]} = a > 0 \qquad \text{and} \qquad b_1^{[L]} = b_2^{[L]} = \cdots = b_t^{[L]} = b > 0.$$

First it is easy to show that $a = b$. Otherwise assume $a > b$, then we set $k_1 = \ldots = k_{L-1} = 1$, divide $(a)^{k_L}$ from both sides, and let k_L go to infinity. It is easy to check that the left side converges to

$$\sum_{i\in[s]} \prod_{j\in[L-1]} a_i^{[j]} > 0,$$

while the right side converges to 0, which contradicts the assumption.

Second, we fix $k_1, \ldots, k_{L-1}$ to be any positive integers, divide $(a)^{k_L} = (b)^{k_L}$ from both sides and let k_L go to infinity. It is easy to check that the left side converges to

$$\sum_{i\in[s]} \prod_{j\in[L-1]} \left(a_i^{[j]}\right)^{k_j},$$

while the right hand side converges to

$$\sum_{i\in[t]} \prod_{j\in[L-1]} \left(b_i^{[j]}\right)^{k_j}.$$

So these two sums are equal for all $k_1, \ldots, k_{L-1} \geq 1$. Then we apply the inductive hypothesis to claim that $s = t$ and there exists a permutation π from $[s]$ to itself such that

$$a_i^{[j]} = b_{\pi(i)}^{[j]}, \qquad \text{for all } j \in [L-1] \text{ and } i \in [s]. \tag{11}$$

It is also easy to see that for any $i \in [s]$, (11) also holds for $j = L$.

We then repeat the whole process after removing the first s elements from the $2L$ sequences. □

Additionally, we also need the following simple lemma in the proof.

Lemma 11. *Let $m \geq 1$ be an integer and $(P_1, P_2, \ldots,)$ be a sequence of subsets of $[m]$. If for any finite subset $\{i_1, \ldots, i_k\} \subset \mathbb{N}$, $P_{i_1} \cap P_{i_2} \cap \cdots \cap P_{i_k} \neq \emptyset$, then there exists a $j \in [m]$ such that $j \in P_i$ for all i.*

Proof. If for every $j \in [m]$, there exists some $i_j \geq 1$ such that $j \notin P_{i_j}$, then the finite intersection

$$\bigcap_{j=1}^{m} P_{i_j} = \emptyset,$$

which contradicts the assumption. ☐

G.2 Matrix and Vector Polynomials

Let $(\mathfrak{X}, \mathfrak{Y})$ be a generalized $\mathcal{P}$-pair, for some $m \times m$ pattern $\mathcal{P}$. So every vector $\mathbf{w} \in \mathfrak{X}$ is either *positive* or *$\mathcal{P}$-weakly positive* and every $\mathbf{D} \in \mathfrak{Y}$ is either a *$\mathcal{P}$-matrix* or a *$\mathcal{P}$-diagonal matrix*. Note that if $\mathfrak{Y}$ only has $\mathcal{P}$-matrices, then $(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{P}$-pair. The definitions below also apply to $\mathcal{P}$-pairs.

We say f is a *$\mathcal{P}$-matrix polynomial* if f is a polynomial over variables

$$\Big\{ x_{i,j} : (i,j) \in \mathcal{P} \Big\}$$

with integer coefficients and zero constant term. We say $\mathfrak{Y}$ satisfies f if for every $\mathcal{P}$-matrix $\mathbf{D} \in \mathfrak{Y}$, we have $f(\mathbf{D}) = 0$, in which we substitute $x_{i,j}$ by $D_{i,j} > 0$ for all $(i,j) \in \mathcal{P}$. We also say $(\mathfrak{X}, \mathfrak{Y})$ satisfies f if $\mathfrak{Y}$ satisfies f.

We say f is a *$\mathcal{P}$-diagonal matrix polynomial* if f is a polynomial over variables

$$\Big\{ x_i : (i,i) \in \mathcal{P} \Big\}$$

with integer coefficients and zero constant term. We say $\mathfrak{Y}$ satisfies f if every $\mathcal{P}$-diagonal matrix $\mathbf{D} \in \mathfrak{Y}$ satisfies $f(\mathbf{D}) = 0$. We also say $(\mathfrak{X}, \mathfrak{Y})$ satisfies f if $\mathfrak{Y}$ satisfies f.

We say g is an *m-vector polynomial* if g is a polynomial over variables

$$\Big\{ y_i : i \in [m] \Big\}$$

with integer coefficients and zero constant term. Similarly, we say $\mathfrak{X}$ satisfies g if every positive vector $\mathbf{w} \in \mathfrak{X}$ satisfies $g(\mathbf{w}) = 0$. We also say $(\mathfrak{X}, \mathfrak{Y})$ satisfies g if $\mathfrak{X}$ satisfies g.

Finally, we say g is a *$\mathcal{P}$-weakly positive vector polynomial* if g is a polynomial over variables

$$\Big\{ y_i : (i,i) \in \mathcal{P} \Big\}$$

with integer coefficients and zero constant term. We say $\mathfrak{X}$ satisfies g if every $\mathcal{P}$-weakly positive vector $\mathbf{w} \in \mathfrak{X}$ satisfies $g(\mathbf{w}) = 0$. We also say $(\mathfrak{X}, \mathfrak{Y})$ satisfies g if $\mathfrak{X}$ satisfies g.

Let F be a finite set of $\mathcal{P}$-matrix, $\mathcal{P}$-diagonal matrix, m-vector, and $\mathcal{P}$-weakly positive vector polynomials. Then we say $(\mathfrak{X}, \mathfrak{Y})$ satisfies F if $(\mathfrak{X}, \mathfrak{Y})$ satisfies every polynomial $f \in F$.

Similarly, given any block pattern $\mathcal{T}$, we can define $\mathcal{T}$-matrix polynomials, $\mathcal{T}$-diagonal matrix polynomials, and $\mathcal{T}$-weakly positive vector polynomials for $\mathcal{T}$-pairs and *generalized* $\mathcal{T}$-pairs.

We remark that, for the case when $(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{T}$-pair, to check whether $\mathfrak{Y}$ satisfies the rank condition (i.e., every matrix $\mathbf{D} \in \mathfrak{Y}$ is block-rank-1), one only needs to check whether $\mathfrak{Y}$ satisfies all the $\mathcal{T}$-matrix polynomials $f_{i,i',j,j'}$ of the following form

$$f_{i,i',j,j'}(\mathbf{x}) = x_{i,j} \cdot x_{i',j'} - x_{i,j'} \cdot x_{i',j}, \qquad \text{where } i, i' \in A_k \text{ and } j, j' \in B_k \text{ for some } k \in [r].$$

G.3 Checking Matrix and Vector Polynomials

Now let $(\mathfrak{X},\mathfrak{Y})$ be a $\mathcal{T}$-pair for some non-trivial $m\times m$ block pattern $\mathcal{T}=\{(A_1,B_1),\ldots,(A_r,B_r)\}$ with $r\ge 1$. We also assume that every matrix in $\mathfrak{Y}$ is block-rank-1, and $\mathfrak{X}$ is *closed*.

We can apply the gen-pair operation to get a new $\mathcal{P}$-pair

$$(\mathfrak{X}',\mathfrak{Y}')=\texttt{gen-pair}(\mathfrak{X},\mathfrak{Y}),\qquad \text{where } \mathcal{P}=\texttt{gen}(\mathcal{T}).$$

We also let $(\mathfrak{X}^*,\mathfrak{Y}^*)$ denote the *generalized* $\mathcal{P}$-pair defined in Appendix C. By definition, $\mathfrak{X}^*$ is also closed.

In this section, we first show that to check whether $(\mathfrak{X}^*,\mathfrak{Y}^*)$ satisfies a matrix or vector polynomial, one only needs to check finitely many polynomials for $(\mathfrak{X},\mathfrak{Y})$. One can prove a similar relation between $(\mathfrak{X}',\mathfrak{Y}')$ and $(\mathfrak{X}^*,\mathfrak{Y}^*)$. As a result, to check whether $(\mathfrak{X}',\mathfrak{Y}')$ satisfies a polynomial or not, we only need to check finitely many polynomials for $(\mathfrak{X},\mathfrak{Y})$.

We start with the following lemma.

Lemma 12. *Let f be a $\mathcal{P}$-matrix or $\mathcal{P}$-diagonal matrix polynomial. Then one can construct a finite set $\{F_1,\ldots,F_L\}$ in a finite number of steps, in which every F_i, $i\in[L]$, is a finite set of $\mathcal{T}$-matrix, m-vector, and $\mathcal{T}$-weakly positive vector polynomials, such that*

$$(\mathfrak{X}^*,\mathfrak{Y}^*)\ \textit{satisfies}\ f\iff \exists\, i\in[L]\ \forall g\in F_i,\ \big[(\mathfrak{X},\mathfrak{Y})\ \textit{satisfies}\ g\big].$$

Proof. We first prove the case when f is a $\mathcal{P}$-matrix polynomial.

If f is the zero polynomial, then the lemma follows by setting $L=1$ and F_1 to be the set consists of the zero polynomial only. From now on we assume that f is not the zero polynomial.

Let $\{\mathbf{C}^{[1]},\ldots,\mathbf{C}^{[s]}\}$ and $\{\mathbf{D}^{[1]},\ldots,\mathbf{D}^{[t]}\}$ be two finite subsets of $\mathcal{T}$-matrices in $\mathfrak{Y}$ and $\{\mathbf{w}^{[1]},\ldots,\mathbf{w}^{[h]}\}$ be a finite subset of *positive* vectors in $\mathfrak{X}$, where $s,t,h\ge 1$. We also let $(\boldsymbol{\alpha}^{[i]},\boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]},\boldsymbol{\delta}^{[i]})$ denote the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively. By the definition of $\mathfrak{Y}^*$ and the assumption that $\mathfrak{Y}$ is *closed*, we can construct from every $(s+t+h)$-tuple

$$\mathbf{p}=\big(k_1,\ldots,k_s,\ell_1,\ldots,\ell_t,e_1,\ldots,e_h\big),\qquad \text{where } k_i,\ell_i,e_i\ge 1,$$

the following $\mathcal{P}$-matrix $\mathbf{C}^{[\mathbf{p}]}$ in $\mathfrak{Y}^*$: the (i,j)th entry of $\mathbf{C}^{[\mathbf{p}]}$ is

$$\sum_{x\in B_i\cap A_j}\Big(\beta^{[1]}_x\Big)^{k_1}\cdots\Big(\beta^{[s]}_x\Big)^{k_s}\cdot\Big(\gamma^{[1]}_x\Big)^{\ell_1}\cdots\Big(\gamma^{[t]}_x\Big)^{\ell_t}\cdot\Big(w^{[1]}_x\Big)^{e_1}\cdots\Big(w^{[h]}_x\Big)^{e_h},\qquad \text{for all } i,j\in[r]. \tag{12}$$

This follows from the fact that the Hadamard product of $(\mathbf{w}^{[1]})^{e_1},\ldots,(\mathbf{w}^{[h]})^{e_h}$ is actually a vector in $\mathfrak{X}$, because $\mathfrak{X}$ is known to be *closed*.

Now we assume $(\mathfrak{X}^*,\mathfrak{Y}^*)$ satisfies f, then by definition we must have

$$f(\mathbf{C}^{[\mathbf{p}]})=0,\qquad \text{for all } \mathbf{p}, \tag{13}$$

since $\mathbf{C}^{[\mathbf{p}]}$ is a $\mathcal{P}$-matrix in $\mathfrak{Y}^*$. By combining (13) and (12) and *rearranging* terms, we have

$$\sum_{i\in[n_1]}\left(\prod_{j\in[s]}\Big(f_i\big(\beta^{[j]}_1,\ldots,\beta^{[j]}_m\big)\Big)^{k_j}\right)\left(\prod_{j\in[t]}\Big(f_i\big(\gamma^{[j]}_1,\ldots,\gamma^{[j]}_m\big)\Big)^{\ell_j}\right)\left(\prod_{j\in[h]}\Big(f_i\big(w^{[j]}_1,\ldots,w^{[j]}_m\big)\Big)^{e_j}\right)$$

$$=\sum_{i\in[n_2]}\left(\prod_{j\in[s]}\Big(g_i\big(\beta^{[j]}_1,\ldots,\beta^{[j]}_m\big)\Big)^{k_j}\right)\left(\prod_{j\in[t]}\Big(g_i\big(\gamma^{[j]}_1,\ldots,\gamma^{[j]}_m\big)\Big)^{\ell_j}\right)\left(\prod_{j\in[h]}\Big(g_i\big(w^{[j]}_1,\ldots,w^{[j]}_m\big)\Big)^{e_j}\right)$$

for all $\mathbf{p}$. In the equation above, n_1 and n_2 are two non-negative integers. For all $i \in [n_1]$ and $j \in [n_2]$, both $f_i(x_1, \ldots, x_m)$ and $g_j(x_1, \ldots, x_m)$ are monomials in $x_1, \ldots, x_m$. Also note that all the monomials f_i, g_j only depend on the $\mathcal{P}$-matrix polynomial f but do not depend on the choices of $\mathbf{p}$ and the subsets $\{\mathbf{C}^{[1]}, \ldots, \mathbf{C}^{[s]}\}$, $\{\mathbf{D}^{[1]}, \ldots, \mathbf{D}^{[t]}\}$, and $\{\mathbf{w}^{[1]}, \ldots, \mathbf{w}^{[h]}\}$. Moreover, because we assumed that f is not the zero polynomial, at least one of n_1 and n_2 is nonzero.

It follows directly from Lemma 10 that if $(\mathfrak{X}^*, \mathfrak{Y}^*)$ satisfies f, then we must have $n_1 = n_2$ which we denote by n. (If $n_1 \neq n_2$, then we already know that $f(\mathbf{C}^{[\mathbf{p}]}) = 0$ cannot hold for all $\mathbf{p}$. The lemma then follows by setting $L = 1$ and F_1 to be the set consisting of the following m-vector polynomial: $g(\mathbf{x}) = x_1$ so that $(\mathfrak{X}, \mathfrak{Y})$ does not satisfy F_1.) Moreover, by Lemma 10, if $(\mathfrak{X}^*, \mathfrak{Y}^*)$ satisfies f then there also exists a permutation π from $[n]$ to itself such that

$$f_i\big(\beta_1^{[j]}, \ldots, \beta_m^{[j]}\big) = g_{\pi(i)}\big(\beta_1^{[j]}, \ldots, \beta_m^{[j]}\big), \qquad \text{for all } j \in [s] \text{ and } i \in [n];$$

$$f_i\big(\gamma_1^{[j]}, \ldots, \gamma_m^{[j]}\big) = g_{\pi(i)}\big(\gamma_1^{[j]}, \ldots, \gamma_m^{[j]}\big), \qquad \text{for all } j \in [t] \text{ and } i \in [n]; \text{ and}$$

$$f_i\big(w_1^{[j]}, \ldots, w_m^{[j]}\big) = g_{\pi(i)}\big(w_1^{[j]}, \ldots, w_m^{[j]}\big), \qquad \text{for all } j \in [h] \text{ and } i \in [n].$$

Since all the discussion above and all the monomials f_i, g_i do not depend on the choice of the three subsets, we can apply Lemma 11 to claim that if $(\mathfrak{X}^*, \mathfrak{Y}^*)$ satisfies f, then there must exist a (universal) permutation π from $[n]$ to itself such that for all $\mathbf{D} \in \mathfrak{X}$ (since $(\mathfrak{X}, \mathfrak{Y})$ is a $\mathcal{T}$-pair, $\mathbf{D}$ is a $\mathcal{T}$-matrix),

$$f_i(\alpha_1, \ldots, \alpha_m) - g_{\pi(i)}(\alpha_1, \ldots, \alpha_m) = 0, \qquad \text{for all } i \in [n] \text{ and}$$
$$f_i(\beta_1, \ldots, \beta_m) - g_{\pi(i)}(\beta_1, \ldots, \beta_m) = 0, \qquad \text{for all } i \in [n],$$

where $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is the representation of $\mathbf{D}$; and for every positive vector $\mathbf{w} \in \mathfrak{Y}$,

$$f_i(w_1, \ldots, w_m) - g_{\pi(i)}(w_1, \ldots, w_m) = 0, \qquad \text{for all } i \in [n].$$

It is also easy to check that these conditions are sufficient.

Furthermore, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ can be expressed by the positive entries of $\mathbf{D}$ as follows. For every $i \in A_k$, where $k \in [r]$, let d be the smallest index in B_k, then we have

$$\alpha_i = \frac{D_{i,d}}{\sum_{j \in A_k} D_{j,d}}.$$

For every $i \in B_k$, where $k \in [r]$, let d be the smallest index in A_k, then $\beta_i = D_{d,i}/\alpha_d$. Now it is easy to see that for every permutation π from $[n]$ to itself, we can construct a finite set F_π of $\mathcal{T}$-matrix and m-vector polynomials, such that, if $(\mathfrak{X}^*, \mathfrak{Y}^*)$ satisfies f then $(\mathfrak{X}, \mathfrak{Y})$ satisfies F_π for some π.

The case when f is a $\mathcal{P}$-diagonal matrix polynomial can be proved similarly. The only difference is that every F_π is now a finite set of $\mathcal{T}$-matrix and $\mathcal{T}$-weakly positive vector polynomials. ☐

It also follows directly by definition that $\mathfrak{Y}'$ satisfies a $\mathcal{P}$-matrix polynomial if and only if $\mathfrak{Y}^*$ satisfies the same polynomial, because $\mathfrak{Y}'$ contains precisely all the $\mathcal{P}$-matrices in $\mathfrak{Y}^*$. Next, we deal with vector polynomials.

Lemma 13. *Let g be an r-vector or a $\mathcal{P}$-weakly positive vector polynomial. One can construct a finite set $\{G_1, \ldots, G_L\}$ in a finite number of steps, in which every G_i is a finite set of $\mathcal{T}$-matrix, m-vector, and $\mathcal{T}$-weakly positive vector polynomials, such that*

$$(\mathfrak{X}^*, \mathfrak{Y}^*) \text{ satisfies } g \iff \exists\, i \in [L]\ \forall f \in G_i,\ \big[(\mathfrak{X}, \mathfrak{Y}) \text{ satisfies } f\big].$$

Proof. We only prove the case when g is $\mathcal{P}$-weakly positive. The other case can be proved similarly.

Again, we assume that g is not the zero polynomial.

Recall that when defining $\mathfrak{X}^*$ in Appendix C, we first define $\mathfrak{X}^\#$ and $\mathfrak{X}^*$ is then the closure of $\mathfrak{X}^\#$: $\mathbf{w}$ is a $\mathcal{P}$-weakly positive vector in $\mathfrak{X}^*$ if and only if there exist a finite and possibly empty subset of positive vectors $\{\mathbf{w}^{[1]},\ldots,\mathbf{w}^{[s]}\}\subseteq\mathfrak{X}^\#$ for some $s\ge 0$, a finite and nonempty subset of $\mathcal{P}$-weakly positive vectors $\{\mathbf{u}^{[1]},\ldots,\mathbf{u}^{[t]}\}\subseteq\mathfrak{X}^\#$ for some $t\ge 1$, and positive integers $k_1,\ldots,k_s,\ell_1,\ldots,\ell_t$, such that

$$\mathbf{w}=\big(\mathbf{w}^{[1]}\big)^{k_1}\circ\cdots\circ\big(\mathbf{w}^{[s]}\big)^{k_s}\circ\big(\mathbf{u}^{[1]}\big)^{\ell_1}\circ\ldots\circ\big(\mathbf{u}^{[t]}\big)^{\ell_t}.$$

To prove Lemma 13, we first construct a finite set $\{F_1,\ldots,F_M\}$, in which every F_i is a finite set of r-vector and $\mathcal{P}$-weakly positive vector polynomials, such that

$$\mathfrak{X}^*\text{ satisfies }g\iff\exists\, i\in[M]\ \forall f\in F_i,\ \big[\mathfrak{X}^\#\text{ satisfies }f\big].\tag{14}$$

To this end, we let $\{\mathbf{w}^{[1]},\ldots,\mathbf{w}^{[s]}\}$ be a finite subset of positive vectors in $\mathfrak{X}^\#$; and $\{\mathbf{u}^{[1]},\ldots,\mathbf{u}^{[t]}\}$ be a finite subset of $\mathcal{P}$-weakly positive vectors in $\mathfrak{X}^\#$, with $s\ge 0$ and $t\ge 1$. Then from any tuple

$$\mathbf{p}=\big(k_1,\ldots,k_s,\ell_1,\ldots,\ell_t\big),\qquad\text{where }k_i,\ell_i\ge 1,$$

we get a $\mathcal{P}$-weakly positive vector $\mathbf{w}^{[\mathbf{p}]}\in\mathfrak{X}^*$, where

$$\mathbf{w}^{[\mathbf{p}]}=\big(\mathbf{w}^{[1]}\big)^{k_1}\circ\cdots\circ\big(\mathbf{w}^{[s]}\big)^{k_s}\circ\big(\mathbf{u}^{[1]}\big)^{\ell_1}\circ\cdots\circ\big(\mathbf{u}^{[t]}\big)^{\ell_t}.$$

Assume $\mathfrak{X}^*$ satisfies g, then we have $g(\mathbf{w}^{[\mathbf{p}]})=0$ for all $\mathbf{p}$. Combining these two equations, we have

$$\sum_{i\in[n_1]}\left(\prod_{j\in[s]}\Big(f_i\big(\mathbf{w}^{[j]}\big)\Big)^{k_j}\right)\left(\prod_{j\in[t]}\Big(f_i\big(\mathbf{u}^{[j]}\big)\Big)^{\ell_j}\right)=\sum_{i\in[n_2]}\left(\prod_{j\in[s]}\Big(g_i\big(\mathbf{w}^{[j]}\big)\Big)^{k_j}\right)\left(\prod_{j\in[t]}\Big(g_i\big(\mathbf{u}^{[j]}\big)\Big)^{\ell_j}\right)$$

for all $\mathbf{p}$. In the equation, $f_i(\mathbf{x})$ and $g_i(\mathbf{x})$ are both monomials over x_i, $(i,i)\in\mathcal{P}$. Again, f_i and g_i only depend on the polynomial g but do not depend on the choices of $\mathbf{p}$ and the two subsets $\{\mathbf{w}^{[1]},\ldots,\mathbf{w}^{[s]}\}$ and $\{\mathbf{u}^{[1]},\ldots,\mathbf{u}^{[t]}\}$.

Because g is not the zero polynomial, one of n_1 and n_2 must be positive, and we have the following two cases. If $n_1\ne n_2$, then by Lemma 10, $\mathfrak{X}^*$ cannot satisfy g and (14) follows by setting $L=1$ and F_1 to be the set consists of the following r-vector polynomial: $f(\mathbf{x})=x_1$.

Otherwise, we have $n_1=n_2>0$, which we denote by n. It follows from Lemma 10 and Lemma 11 that if $\mathfrak{X}^*$ satisfies g, then there exists a universal permutation π from $[n]$ to itself such that for every positive and $\mathcal{P}$-weakly positive vector $\mathbf{w}\in\mathfrak{X}^\#$,

$$f_i(\mathbf{w})=g_{\pi(i)}(\mathbf{w}),\qquad\text{for all }i\in[n].$$

As a result, we can construct F_π for each π, and $\mathfrak{X}^*$ satisfies g if and only if $\mathfrak{X}^\#$ satisfies F_π for some π.

In the second step, we show that for any r-vector or $\mathcal{P}$-weakly positive vector polynomial f, one can construct $\{F_1,\ldots,F_L\}$ in a finite number of steps, in which each F_i is a finite set of $\mathcal{T}$-matrix, m-vector and $\mathcal{T}$-weakly positive vector polynomials, such that, $\mathfrak{X}^\#$ satisfies f if and only if $(\mathfrak{X},\mathfrak{Y})$ satisfies F_i for some $i\in[L]$. The idea of the proof is very similar to the proof of Lemma 12 so we omit it here.

Lemma 13, for the case when g is $\mathcal{P}$-weakly positive, then follows by combing these two steps. $\square$

We can also prove the following lemma similarly.

Lemma 14. *Let g be an r-vector or a $\mathcal{P}$-weakly positive vector polynomial. Then one can construct a finite set $\{G_1, \ldots, G_L\}$ in a finite number of steps, in which every G_i, $i \in [L]$, is a finite set of $\mathcal{P}$-matrix, $\mathcal{P}$-diagonal matrix, r-vector, and $\mathcal{P}$-weakly positive vector polynomials, such that*

$$(\mathfrak{X}', \mathfrak{Y}') \text{ satisfies } g \iff \exists\, i \in [L]\ \forall f \in G_i,\ \big[(\mathfrak{X}^*, \mathfrak{Y}^*) \text{ satisfies } f\big].$$

G.4 Decidability of the Rank Condition

Finally, we use these lemmas to prove Lemma 4, the decidability of the rank condition.

We start with the following simple observation. Let $F = \{f_1, \ldots, f_s\}$ be a finite set of matrix and vector polynomials. For each $i \in [s]$, there is a finite set $\{F_{i,1}, \ldots, F_{i,L_i}\}$ in which every $F_{i,j}$ is some finite set of polynomials, and we have the following statement:

$$(\mathfrak{X}', \mathfrak{Y}') \text{ satisfies } f_i \iff \exists\, j \in [L_i]\ \forall f \in F_{i,j},\ \big[(\mathfrak{X}, \mathfrak{Y}) \text{ satisfies } f\big].$$

Then the conjunction of these statements over $f_i \in F$, $i \in [s]$, can be expressed in the same form: One can construct from $\{F_{i,j} : i \in [s], j \in [L_i]\}$ a new finite set $\{G_1, \ldots, G_L\}$ in which every G_j is some finite set of polynomials, such that

$$\forall f \in F,\ \big[(\mathfrak{X}', \mathfrak{Y}') \text{ satisfies } f\big] \iff \exists\, j \in [L]\ \forall g \in G_j,\ \big[(\mathfrak{X}, \mathfrak{Y}) \text{ satisfies } g\big].$$

Now we prove Lemma 4. After $\ell \geq 0$ steps, we get a sequence of $\ell + 1$ pairs

$$(\mathfrak{X}_0, \mathfrak{Y}_0), (\mathfrak{X}_1, \mathfrak{Y}_1), \ldots, (\mathfrak{X}_\ell, \mathfrak{Y}_\ell),$$

which satisfies condition $(\mathbf{R}_\ell)$. Since we assumed that $\mathfrak{X}_0 = \{\mathbf{1}\}$, every $\mathfrak{X}_i$ in the sequence is *closed*.

We show how to check whether every matrix $\mathbf{D} \in \mathfrak{Y}_{\ell+1}$, where

$$(\mathfrak{X}_{\ell+1}, \mathfrak{Y}_{\ell+1}) = \texttt{gen-pair}(\mathfrak{X}_\ell, \mathfrak{Y}_\ell),$$

is block-rank-1 or not. To this end we first check whether $\mathcal{P} = \texttt{gen}(\mathcal{T}_\ell)$ is consistent with a block pattern or not. If not, then we conclude that $\mathfrak{Y}_{\ell+1}$ does not satisfy the rank condition.

Otherwise, we use $\mathcal{T}_{\ell+1}$ to denote the block pattern consistent with $\mathcal{P}$. To check the rank condition, it is equivalent to check whether $\mathfrak{Y}_{\ell+1}$ satisfies the following $\mathcal{P}$-matrix polynomials:

$$f_{i,i',j,j'}(\mathbf{x}) = x_{i,j} \cdot x_{i',j'} - x_{i,j'} \cdot x_{i',j}, \qquad \text{where } i, i' \in A_k \text{ and } j, j' \in B_k \text{ for some } k \in [r]$$

and $(A_1, B_1), \ldots, (A_r, B_r)$ are the pairs in $\mathcal{T}_{\ell+1}$.

By Lemma 12-14, we can construct a finite set $\{F_1, \ldots, F_L\}$ in which every F_i is a finite set of

$$\mathcal{T}_\ell\text{-matrix, } m_\ell\text{-vector, and } \mathcal{T}_\ell\text{-weakly positive vector polynomials}$$

such that

$$\mathfrak{Y}_{\ell+1} \text{ satisfies the rank condition if and only if } (\mathfrak{X}_\ell, \mathfrak{Y}_\ell) \text{ satisfies } F_i \text{ for some } i \in [L].$$

If $\ell = 0$, then we are done, since $(\mathfrak{X}_0, \mathfrak{Y}_0)$ is finite and we can check all the polynomials in F_i for all $i \in [L]$ in a finite number of steps. Otherwise, $\ell \geq 1$ and we can use Lemma 12–14 and the observation above to construct, for each F_i, a finite set $\{F_{i,1}, \ldots, F_{i,L_i}\}$ in which every $F_{i,j}$ is a finite set of

$\mathcal{T}_{\ell-1}$-matrix, $m_{\ell-1}$-vector, and $\mathcal{T}_{\ell-1}$-weakly positive vector polynomials

such that

$(\mathfrak{X}_\ell, \mathfrak{Y}_\ell)$ satisfies F_i if and only if $(\mathfrak{X}_{\ell-1}, \mathfrak{Y}_{\ell-1})$ satisfies $F_{i,j}$ for some $j \in [L_i]$.

We repeat this process until we reach the finite pair $(\mathfrak{X}_0, \mathfrak{Y}_0)$. So the checking procedure looks like a huge tree of depth $\ell + 1$. Every leaf v of the tree is associated with a finite set F_v of

$\mathcal{T}_0$-matrix, m_0-vector, and $\mathcal{T}_0$-weakly positive vector polynomials.

Set $\mathfrak{Y}_{\ell+1}$ satisfies the rank condition if and only if $(\mathfrak{X}_0, \mathfrak{Y}_0)$ satisfies F_v for some leaf v of the tree.

H The Dichotomy for the $\{0, 1\}$ Case

We briefly describe the dichotomy criterion of Bulatov [2].

A finite *relational structure* $\mathcal{H}$ over a finite set of relational symbols $R_1, R_2, \ldots, R_k$, each of which has a fixed arity, is a non-empty set H together with an interpretation of these relational symbols $R_1^{\mathcal{H}}$, $R_2^{\mathcal{H}}, \ldots, R_k^{\mathcal{H}}$ which are relations on H of the corresponding arities. For graph homomorphism (i.e., $\mathcal{H}$-coloring), we start with a single binary relation, namely the edge relation E on H. A relation R is said to be *pp-definable* in $\mathcal{H}$, if it can be expressed by the relations $R_i^{\mathcal{H}}$, $1 \leq i \leq k$, together with the binary Equality predicate on $\mathcal{H}$, conjunction, and existential quantifiers.

A mapping f from H^m to H, for some $m \geq 1$, is called a *polymorphism* of $\mathcal{H}$ if it satisfies the following condition: For any relation $R \in \{R_1^{\mathcal{H}}, R_2^{\mathcal{H}}, \ldots, R_k^{\mathcal{H}}\}$ of arity n, for any m tuples in H^n:

$$(a_{1,1}, \ldots, a_{1,n}), \ldots, (a_{m,1}, \ldots, a_{m,n}) \in H^n,$$

if each $(a_{i,1}, \ldots, a_{i,n}) \in R$ for all $i : 1 \leq i \leq m$, then

$$\big(f(a_{1,1}, \ldots, a_{m,1}), \ldots, f(a_{1,n}, \ldots, a_{m,n})\big) \in R.$$

A relational structure $\mathcal{H}$ defines a *universal algebra* $\mathbf{A}$, where the universe is H and the set of all polymorphisms are its operations. A theorem of Geiger [12] then states that, for finite $\mathcal{H}$, a relation on H is invariant under all polymorphisms iff it is pp-definable.

A pp-definable binary equivalence relation is called a *congruence*. A subalgebra is a unary pp-definable relation (subset) together with the restrictions of the given relations. One can easily define direct product algebras and homomorphic images (quotient algebra modulo a congruence). A class of universal algebras closed under quotient, subalgebra and direct product is called a *variety*. The class of algebras that are homomorphic images of subalgebras of direct powers of some universal algebra is called the variety generated by it (HSP theorem).

A Mal'tsev polymorphism m is a ternary polymorphism satisfying $m(x, x, y) = y$ and $m(x, y, y) = x$, for all $x, y \in H$. Having a Mal'tsev polymorphism is a necessary condition for tractability.

Now start with the relational structure $\mathcal{H}$ with a single edge relation E, then add to it all the unary relations $\{C_h \mid h \in H\}$, where $C_h = \{(h)\}$, we obtain a relational structure denoted by $\mathcal{H}_{\text{id}}$. Then the polymorphisms of $\mathcal{H}_{\text{id}}$ define the universal algebra called the *full idempotent reduct*. These are the idempotent polymorphisms of $\mathcal{H}$: $f(x, \ldots, x) = x$.

Congruences form a lattice. Given any two congruences α and β, we let $A_1, \ldots, A_s$ and $B_1, \ldots, B_t$ be the equivalence classes of α and β respectively, then the $s \times t$ matrix $M(\alpha, \beta)$ has (i, j) entry $|A_i \cap B_j|$.

The tractability criterion of Bulatov can now be stated: Start with $\mathcal{H}_{\mathrm{id}}$ and take the full idempotent reduct. The #CSP problem defined by $\mathcal{H}$ is tractable iff every finite algebra $\mathbf{A}$ in the *variety* generated by this full idempotent reduct satisfies the following condition: For any two congruences α and β in $\mathbf{A}$, the rank$(M(\alpha, \beta))$ is equal to the number of equivalence classes of $\alpha \vee \beta$, the join congruence of α and β.

The reason it is difficult to show that this dichotomy criterion is decidable is because it talks about *all* finite algebras $\mathbf{A}$ in the *variety* generated by the full idempotent reduct of $\mathcal{H}_{\mathrm{id}}$. This variety is infinite, containing arbitrarily large arities over H. Thus, even though in graph homomorphism we are given only a binary relation, the process of forming the variety produces arbitrarily large arities, and this criterion is a condition involving infinitely many relations.