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

Jin-Yi Cai*
Computer Sciences Department
University of Wisconsin-Madison
Madison, USA
jyc@cs.wisc.edu
and
Beijing University, Beijing, China

Xi Chen*
Computer Science Department
University of Southern California
Los Angeles, USA
csxichen@gmail.com

***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 A with non-negative algebraic 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 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.**

***Keywords*-graph homomorphism; dichotomy; decidability.**

I. Introduction

The complexity of counting graph homomorphisms has received great attention recently [9], [5], [3], [1], [8], [13], [7]. The problem can be defined for both *directed* and *undirected* graphs. Most results have been obtained for the *undirected* case, 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 actually 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 [11].

Let G and H be two graphs. We then 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 $\xi : V(G) \to V(H)$ 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 $\mathbf{A} = (A_{i,j})$:

$$Z_{\mathbf{A}}(G) = \sum_{\xi : V \to [m]} \prod_{uv \in E} A_{\xi(u),\xi(v)},$$

for any directed graph $G = (V, E)$. Note that the input G is a directed graph in general. However, if $\mathbf{A}$ is symmetric, 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 and important counting problems over graphs. For example, if H is the undirected graph on two vertices $\{0, 1\}$ with an edge $(0, 1)$ and a loop $(1, 1)$, then a homomorphism from G to H corresponds to a Vertex Cover of G, and $Z_H(G)$ is exactly the number of vertex covers. 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 [12] 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 computational complexity of $Z_{\mathbf{A}}(\cdot)$ has been studied and *dichotomy* theorems have been proved. For a given family $\mathcal{C}$ of matrices $\mathbf{A}$, a dichotomy theorem 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 $\mathbf{A}$, and $Z_{\mathbf{A}}(G)$ over undirected graphs G:

In [14] and [15], Hell and Nešetřil showed that given any symmetric $\{0, 1\}$ matrix $\mathbf{A}$, *deciding* whether $Z_{\mathbf{A}}(G) > 0$ is

*Research supported by NSF grants CCF-0830488 and CCF-0914969.

*Part of the research done while the author was a postdoc at Princeton University. Supported by NSF grants CCF-0832797, DMS-0635607 and a USC startup fund.

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

either in P or NP-complete. Dyer and Greenhill [9] studied the counting version of $Z_{\mathbf{A}}(\cdot)$. They 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 Lu [7]. Thurley also proved a dichotomy theorem for all Hermitian matrices [16].

We remark that all the dichotomy theorems for *symmetric* matrices 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 A, and $Z_{\mathbf{A}}(G)$ over directed graphs G:

In a paper [8] that won the best paper award at ICALP in 2006, Dyer, Goldberg, and Paterson 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 directed graph H is *layered*[3] and *Lovász-good*, and is #P-hard otherwise. The property of Lovász-goodness turns out to be polynomial-time checkable.

In [1], Bulatov presented a sweeping dichotomy theorem for all counting Constraint Satisfaction Problems. Recently Dyer and Richerby [10] obtained an alternative proof of the theorem. 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 is decidable or not. The criterion[4] requires one to check a condition over an infinitary object, and this situation remains the same for the Dyer-Richerby proof [10]. The decidability of the dichotomy was then left as an open problem in [2].

In this paper, we prove a dichotomy theorem for $Z_{\mathbf{A}}(\cdot)$ over all non-negative algebraic matrices $\mathbf{A}$. We show that, for any 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 $\mathbf{A}$, decides whether computing $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard. In particular, for the family of $\{0,1\}$ matrices $\mathbf{A}$, our result gives an alternative dichotomy criterion[5] to that of Bulatov [2] as well as Dyer and Richerby [10], which is decidable.

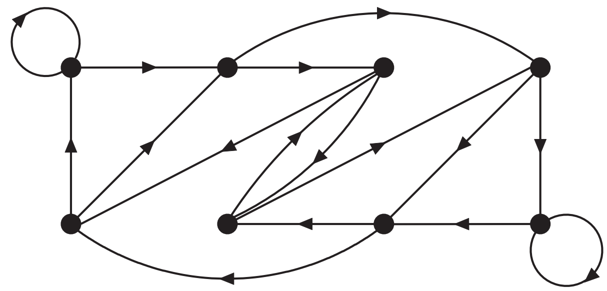

Figure 1. A directed graph H such that $Z_H(\cdot)$ is tractable

The main difficulty we encountered in obtaining the dichotomy theorem is due to the abundance of new intricate but tractable cases, when moving from acyclic graphs H to general directed graphs. For example, H does not have to be *layered* for the problem $Z_H(\cdot)$ to be tractable (see Figure 1 above for an example). Due to 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 [8]) 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 [8] for the acyclic case fits into our framework as follows. In our dichotomy theorem, 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 if $\mathbf{A}$ is m-by-m). The dichotomy is then

> $Z_{\mathbf{A}}(\cdot)$ is in P if every *block* of any matrix defined in the process above is of rank 1; otherwise it is #P-hard (see Section I-A and I-B for details).

For the special acyclic case treated by Dyer, Goldberg, and Paterson, 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 to layer 1), then deciding whether $Z_{\mathbf{A}}(\cdot)$ is in P or #P-hard becomes much harder in the sense that we might need $\gg k$ (though $\leq m$) rounds to reach a conclusion.

[2] More exactly, they proved a dichotomy theorem for symmetric 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, and throughout the paper we use $\mathbb{R}$ to denote the set of real algebraic numbers.

[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 [1] characterize $\{0,1\}$ matrices $\mathbf{A}$ with $Z_{\mathbf{A}}(\cdot)$ in P. As a result, 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 implies a finite time algorithm for checking the dichotomy criterion of Bulatov (as well as the version of Dyer and Richerby [10]) for the case of $\{0,1\}$ matrices.

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

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

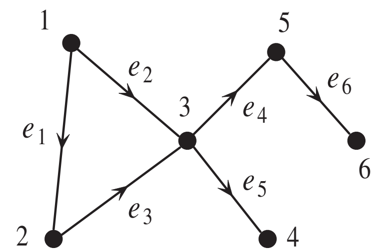

Figure 3. The input directed graph G

After circulating a draft of this paper, Goldberg informed us that she and coauthors found a reduction from *weighted* counting CSPs with *non-negative rational* weights [4] to the 0-1 dichotomy theorem of Bulatov [2]. In particular, it implies a dichotomy for all $Z_{\mathbf{A}}(\cdot)$, where $\mathbf{A}$ is any nonnegative rational matrix. However, the combined result only works for non-negative rational weights and the dichotomy is still not known to be decidable.

A. Intuition of our dichotomy theorem: Domain reduction

Let $\mathbf{A} = (A_{i,j})$ be the $m \times m$ non-negative matrix being considered, and let $G = (V, E)$ 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 $[m] = \{1, \ldots, m\}$ 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

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

Clearly we have $Z_{\mathbf{A}}(G) = Z(\mathcal{G}_0)$ if the underlying directed graph of $\mathcal{G}_0$ is exactly G; every edge of $\mathcal{G}_0$ is labeled with the same matrix $\mathbf{A}$; 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 shall describe next. On the one hand, the *domains* of these labeled graphs shrink along with k: The size of the 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 in the sequence $\mathcal{G}_0, \mathcal{G}_1, \ldots, \mathcal{G}_h$ is at most m. To prove our dichotomy theorem, we show that,

1) either something bad happens for some input graph G and forces us to stop the domain reduction process, in which case we show that $Z_{\mathbf{A}}(\cdot)$ is #P-hard;
2) or for any input G, we can keep reducing the domain size until the computation of $Z_{\mathbf{A}}(G)$ becomes trivial, in which case we show that $Z_{\mathbf{A}}(\cdot)$ is in P.

Now we start to describe the domain reduction method.

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 then we can already show the #P-hardness of $Z_{\mathbf{A}}(\cdot)$, by using the dichotomy theorem of Bulatov and Grohe [5] for symmetric non-negative matrices (see Lemma 1). Without loss of generality we assume $\mathbf{A}$ is block-rank-1.

In the rest of this section, we let $\mathbf{A}$ be the 8×8 block-rank-1 and non-negative matrix shown in Figure 2, 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 describe 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 $A_{i,j} > 0$ iff $i \in A_s$ and $j \in B_s$, for some $s \in [4]$. Since $\mathbf{A}$ is block-rank-1, there also exist two 8-dimensional *positive* vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that

$$A_{i,j} = \alpha_i \cdot \beta_j, \quad \text{for all } A_{i,j} > 0.$$

Let $G = (V, E)$ be the directed graph shown in Figure 3 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 use $\mathbf{y} \in [8]^6$ (instead of $\xi : V \to [8]$) to denote an assignment, where $y_i \in [8]$ denotes the value of vertex i in Figure 3 for each $i \in [6]$.

First of all, let $\mathbf{y} \in [8]^6$ be an 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 each $ij \in E$, 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 of the edges $e_\ell \in E$, $\ell \in [6]$ (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 $y_j \in B_{x_\ell}$. This gives us

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

Next, 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, one only needs to sum over $\mathbf{x}$ in which $x_1 = x_2 = x_3$ 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 variables $\mathbf{x}$ is that, once $\mathbf{x}$ is fixed, one can always decompose A_{y_i,y_j} as the product of α_{y_i} and β_{y_j}, for all $\mathbf{y} \in Y[\mathbf{x}]$ and for all $ij \in E$, because $\mathbf{y}$ 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 $\mathbf{y} \in Y[\mathbf{x}]$ for some $\mathbf{x}$, then

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

Also notice that $Y[\mathbf{x}]$, for any $\mathbf{x} \in [4]^6$, is a direct product of subsets of $[8]$: $\mathbf{y} \in Y[\mathbf{x}]$ if and only if

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

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

$$\sum_{x_1,x_4,x_6} \sum_{y_i\in L_i,\, i\in[6]} \left(\alpha_{y_1}^2\alpha_{y_2}\beta_{y_2}\right) \cdot \left(\alpha_{y_3}^2\beta_{y_3}^2\right) \cdot \beta_{y_4} \cdot \left(\alpha_{y_5}\beta_{y_5}\right) \cdot \beta_{y_6}.$$

Finally we build 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 two directed edges ab and bc. The vertex weight vector of a is

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

the vertex weights of b and c are the same:

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

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

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

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

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

By the definition of $Z(\cdot)$, 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 $\mathbf{A}$ that 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 in our proof.

The two issues above caused 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 demonstrated here crucially used the assumption that $\mathbf{A}$ is a block-rank-1 matrix (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 ever happens *for any* input directed graph G, then $Z_{\mathbf{A}}(\cdot)$ is #P-hard. Informally, we have

Theorem 1 (Informal)**.** *For any directed graph 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 the case $k = 1$ is relatively straight forward, because every edge weight matrix in $\mathcal{G}_1$ is derived from $\mathbf{A}$ in some way. However, due of the issues mentioned above, both the set of edge weight matrices and the set of vertex weight vectors that may appear in $\mathcal{G}_1$ (over all possible input graphs G) are *infinite* in general, and even proving Theorem 1 for $k = 2$ is highly non-trivial.

Theorem 1 essentially gives us a dichotomy theorem for all non-negative matrices. However, 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.

B. Proof Sketch

Without loss of generality, we assume that $\mathbf{A}$ is a nonnegative block-rank-1 matrix. To show that $Z_{\mathbf{A}}(\cdot)$ is either in P or #P-hard, we use the following two steps.

First we *define* from $\mathbf{A}$ a finite sequence of pairs:

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

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 may appear in $\mathcal{G}_k$, after k rounds of domain reductions. There also exist positive integers

$$m=m_0>m_1>\ldots>m_h\ge 1$$

such that $\mathfrak{Y}_k$, $k\in[h]$, is a set of $m_k\times m_k$ non-negative matrices; and $\mathfrak{X}_k$, $k\in[h]$, is a set of m_k-dimensional non-negative vectors. While 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 each $k\in[h]$, all matrices in $\mathfrak{Y}_k$ share the same *structure*:

$$\text{For all } \mathbf{B},\mathbf{B}'\in\mathfrak{Y}_k,\ B_{i,j}>0\Leftrightarrow B'_{i,j}>0;$$

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 Section IV. In Section III-A, Lemma 2, we prove that for every $k\in[h]$ and every $\mathbf{B}\in\mathfrak{Y}_k$, the problem of computing $Z_{\mathbf{B}}(\cdot)$ is polynomial-time reducible to the computation of $Z_{\mathbf{A}}(\cdot)$. The hardness part of our dichotomy theorem then follows: 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.

Next we assume that all matrices $\mathbf{B}\in\mathfrak{Y}_k$, $k\in[h]$, are block-rank-1. To finish the proof of the dichotomy, 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 $\mathcal{G}_1,\ldots,\mathcal{G}_h$ such that

1) $Z(\mathcal{G}_1)=Z_{\mathbf{A}}(G)$ and $Z(\mathcal{G}_{k+1})=Z(\mathcal{G}_k)$ for all $k: 1\le k<h$; and
2) For every $k\in[h]$, 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 I-A, and also because the number of graphs h 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, we only need to check whether the matrices in $\mathfrak{Y}_k$, $k\in[h]$, satisfy the following condition:

> **The Block-Rank-1 Condition**: Every matrix in $\mathfrak{Y}_k$, $k\in[h]$, is block-rank-1.

However, as mentioned earlier, all the sets $\mathfrak{Y}_k$, $k\in[h]$, are infinite in general and thus, one 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 $(\mathfrak{X}_{k-1},\mathfrak{Y}_{k-1})$. Therefore, to check whether $\mathfrak{Y}_k$, $k\in[h]$, satisfies the block-rank-1 condition, we only need 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.

II. 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 vertex $v\in V$ is labeled with an m-dimensional non-negative vector $\mathcal{V}(v)\in\mathbb{R}_+^m$ as its vertex weight;
3) Every edge $uv\in E$ is labeled with an $m\times m$ (not necessarily symmetric) 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, with $G=(V,E)$. For each $v\in V$, we let $\mathbf{w}^{[v]}=\mathcal{V}(v)$ denote its vertex weight vector; and for each $uv\in E$, we let $\mathbf{C}^{[uv]}=\mathcal{E}(uv)$ denote its edge weight matrix. Then we define the function $Z(\mathcal{G})$ as follows:

$$Z(\mathcal{G})=\sum_{\xi:V\to[m]}\mathrm{wt}(\mathcal{G},\xi),$$

where

$$\mathrm{wt}(\mathcal{G},\xi)=\prod_{v\in V} w^{[v]}_{\xi(v)}\prod_{uv\in E} C^{[uv]}_{\xi(u),\,\xi(v)}$$

denotes the *weight* of the assignment ξ.

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

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

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 $uv\in E$.

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) $A_i\subseteq[m]$, $A_i\neq\emptyset$, $B_i\subseteq[m]$, $B_i\neq\emptyset$ *for all $i\in[r]$;*
3) $A_i\cap A_j=B_i\cap B_j=\emptyset$, *for all $i\neq j\in[r]$.*

$\mathcal{T}$ 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 $\mathbf{C}=\mathbf{0}$ (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 and 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 $i\in[m]$, we have*
$$\alpha_i>0 \iff i\in\textstyle\bigcup_{k\in[r]}A_k \quad \text{and}$$
$$\beta_i>0 \iff i\in\textstyle\bigcup_{k\in[r]}B_k$$
2) $C_{i,j}=\alpha_i\cdot\beta_j$ *for all $i,j\in[m]$ such that $C_{i,j}>0$;*
3) $\sum_{j\in A_i}\alpha_j=1$, *for all $i\in[r]$.*

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.

Lemma 1 below concerns the complexity of $Z_{\mathbf{C}}(\cdot)$. The proof can be found in the appendix of the full version [6].

Lemma 1. *If $\mathbf{C}$ is not block-rank-1, $Z_{\mathbf{C}}(\cdot)$ 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 $r\geq 1$. Then it defines the following $r\times r$ pattern $\mathcal{P}=\mathsf{gen}(\mathcal{T})$: 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 $\mathsf{gen\text{-}block}(\mathcal{T})$ as follows:

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

We note that $\mathcal{P}=\mathsf{gen}(\mathcal{T})$ could be *trivial*, even if $\mathcal{T}$ itself 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 $v\in V(G)$; and $\mathcal{E}(uv)\in\mathfrak{Q}$ for any $uv\in E(G)$. 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 any finite and nonempty subset $\mathfrak{P}'\subseteq\mathfrak{X}'$ with $\mathbf{1}\in\mathfrak{P}'$ and any 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 $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$.*

III. 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: $\forall\, i \in [m]$, $w_i > 0$ *iff* $(i,i) \in \mathcal{P}$.

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 also say $(\mathfrak{X}, \mathfrak{Y})$ *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, this 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, this 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 Theorem 3 in the rest of this section. The proofs of Lemma 2, 3, and 4 can be found in the full version [6].

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

First we state a key lemma in the proof of Theorem 2.

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 Section IV, we introduce an operation gen-pair over $(\mathfrak{X}, \mathfrak{Y})$, which *defines* a new (and possibly infinite) pair

$$(\mathfrak{X}', \mathfrak{Y}') = \textsf{gen-pair}(\mathfrak{X}, \mathfrak{Y}).$$

Definition 5. *A set* S *of nonnegative* m*-dimensional vectors 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* i*th entry is* $w_{1,i} w_{2,i}$ *for all* $i \in [m]$.

The proof of Lemma 2 can be found in [6].

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

B. *Proof of Theorem 2*

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

By Lemma 1, there must be an $m \times m$ 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

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

Next we define a finite sequence of pairs using the gen-pair operation, starting with $(\mathfrak{X}_0, \mathfrak{Y}_0)$.

First, if $|A_i| = |B_i| = 1$ for all i, i.e., every set A_i and B_i in $\mathcal{T}_0$ is a singleton, then the sequence has only one pair $(\mathfrak{X}_0, \mathfrak{Y}_0)$, and the definition of this 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:

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

By Lemma 2, we know $(\mathfrak{X}_1, \mathfrak{Y}_1)$ 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

$$\mathfrak{P}_1 = \{\mathbf{1}\} \ \text{ and } \ \mathfrak{Q}_1 = \{\mathbf{D}\}.$$

It then follows from Lemma 2 that there exists a finite pair $(\mathfrak{P}_0, \mathfrak{Q}_0)$ where

$$\mathfrak{P}_0 \subseteq \mathfrak{X}_0 \ \text{ and } \ \mathfrak{Q}_0 \subseteq \mathfrak{Y}_0,$$

such that $Z_{\mathfrak{P}_1,\mathfrak{Q}_1}(\cdot)$ is reducible to $Z_{\mathfrak{P}_0,\mathfrak{Q}_0}(\cdot)$. On the other hand, it is clear that $Z_{\mathfrak{P}_0,\mathfrak{Q}_0}(\cdot)$ is polynomial-time 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

R$_1$: $\mathcal{T}_1 = \textsf{gen-block}(\mathcal{T}_0)$ is an $m_1 \times m_1$ block pattern, where m_1 is the number of pairs in $\mathcal{T}_0$; $(\mathfrak{X}_1, \mathfrak{Y}_1) = \textsf{gen-pair}(\mathfrak{X}_0, \mathfrak{Y}_0)$ 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 here. 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 $i \in [\ell]$, $\mathcal{T}_i = \textsf{gen-block}(\mathcal{T}_{i-1})$ and $(\mathfrak{X}_i, \mathfrak{Y}_i) = \textsf{gen-pair}(\mathfrak{X}_{i-1}, \mathfrak{Y}_{i-1})$ is a $\mathcal{T}_i$-pair; 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 $\textsf{gen-pair}$ operation again to define a new pair $(\mathfrak{X}_{\ell+1}, \mathfrak{Y}_{\ell+1})$.

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

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

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 $i \in [h]$, $\mathcal{T}_i = \textsf{gen-block}(\mathcal{T}_{i-1})$ and $(\mathfrak{X}_i, \mathfrak{Y}_i) = \textsf{gen-pair}(\mathfrak{X}_{i-1}, \mathfrak{Y}_{i-1})$ is a $\mathcal{T}_i$-pair; Either $\mathcal{T}_h = \emptyset$ is trivial or every set in $\mathcal{T}_h$ is a singleton; 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 $h < m_0 = m$.

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 dichotomy theorem we show in the full version [6] that

Lemma 3 (Tractability). *Let $(\mathfrak{P}, \mathfrak{Q})$ be a $\mathcal{T}$-pair where $\mathcal{T}$ is a block pattern. 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 $Z_{\mathfrak{P},\mathfrak{Q}}(\cdot)$ can be computed in polynomial time.*

This finishes the proof of Theorem 2.

C. Proof of Theorem 3

Next, we show that for the special case when $\mathfrak{X}_0 = \mathfrak{P} = \{\mathbf{1}\}$, the dichotomy theorem is decidable.

First, the condition ($\mathbf{R}_0$) can be checked easily since there are only finitely many matrices in $\mathfrak{Y}_0$.

Now 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

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

where $\mathcal{P} = \textsf{gen}(\mathcal{T}_\ell)$, is block-rank-1 or not. We refer to this property as the *rank property* for $\mathfrak{Y}_{\ell+1}$.

We prove the following lemma concerning decidability in the full version [6]. Theorem 3 then follows directly.

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

IV. Definition of the gen-pair Operation

In this section, we define the operation $\textsf{gen-pair}$.

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

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: $w_i > 0$ iff $i \in \text{diag}(\mathcal{T})$.

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

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

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 $r \geq 1$. Then an $r \times r$ non-negative 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 also show that $(\mathfrak{X}', \mathfrak{Y}')$ is a $\mathcal{P}$-pair with $\mathfrak{X}'$ being *closed*. The other part of Lemma 2, that there is a polynomial-time reduction from $(\mathfrak{X}', \mathfrak{Y}')$ to $(\mathfrak{X}, \mathfrak{Y})$, can be found in [6].

A. Definition of $\mathfrak{Y}^*$

We start with the definition of set $\mathfrak{Y}^*$ which contains both $\mathcal{P}$-matrices and $\mathcal{P}$-diagonal matrices, where $\mathcal{P} = \mathsf{gen}(\mathcal{T})$.

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

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

for all $i, j \in [r]$, where $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ are the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively. 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 $\mathcal{P} = \mathsf{gen}(\mathcal{T})$.*

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

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

for all $i, j \in [r]$, where $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ are the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively. Similarly,

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

Proof: First, we show that $\mathbf{D}$ is a diagonal matrix. 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 $A_i \cap B_i \neq \emptyset$. ■

B. 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
$$w_i^{\#} = \sum_{x \in A_i} \left(\alpha_x^{[1]}\right)^{s_1} \cdots \left(\alpha_x^{[g]}\right)^{s_g} \cdot w_x, \quad \text{for all } i \in [r],$$
where $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ is the representation of $\mathbf{C}^{[i]}$. 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
$$w_i^{\#} = \sum_{x \in B_i} \left(\delta_x^{[1]}\right)^{t_1} \cdots \left(\delta_x^{[h]}\right)^{t_h} \cdot w_x, \quad \text{for all } i \in [r],$$
where $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ is the representation of $\mathbf{D}^{[i]}$. 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
$$w_i^{\#} = \sum_{x \in B_i \cap A_i} \left(\beta_x^{[1]}\right)^{s_1} \cdots \left(\beta_x^{[g]}\right)^{s_g} \left(\gamma_x^{[1]}\right)^{t_1} \cdots \left(\gamma_x^{[h]}\right)^{t_h} w_x,$$
for all $i \in [r]$, where $(\boldsymbol{\alpha}^{[i]}, \boldsymbol{\beta}^{[i]})$ and $(\boldsymbol{\gamma}^{[i]}, \boldsymbol{\delta}^{[i]})$ are the representations of $\mathbf{C}^{[i]}$ and $\mathbf{D}^{[i]}$, respectively. 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 then 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

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

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

We now 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

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

   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, because $\mathfrak{X}^*$ is closed, one can show that $\mathfrak{X}'$ is also closed. This proved the first part of Lemma 2.

V. Acknowledgements

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

References

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

[2] ——, "The complexity of the counting constraint satisfaction problem," *ECCC Report*, no. 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*, 2003, pp. 562–571.

[4] A. Bulatov, M. Dyer, L. Goldberg, M. Jalsenius, M. Jerrum, and D. Richerby, "The complexity of weighted and unweighted #CSP," *arXiv*, no. 1005.2678, 2010.

[5] A. Bulatov and M. Grohe, "The complexity of partition functions," *Theoretical Computer Science*, vol. 348, no. 2, pp. 148–186, 2005.

[6] J.-Y. Cai and X. Chen, "A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights," *arXiv*, no. 1008.0915, 2010.

[7] 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, pp. 275–286.

[8] M. Dyer, L. Goldberg, and M. Paterson, "On counting homomorphisms to directed acyclic graphs," *Journal of the ACM*, vol. 54, no. 6, Article 27, 2007.

[9] M. Dyer and C. Greenhill, "The complexity of counting graph homomorphisms," *Random Structures & Algorithms*, vol. 17, no. 3–4, pp. 260–289, 2000.

[10] M. Dyer and D. Richerby, "On the complexity of #CSP," in *Proceedings of the 42th ACM Symposium on Theory of Computing*, 2010, pp. 725–734.

[11] T. Feder and M. Vardi, "The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory," *SIAM Journal on Computing*, vol. 28, no. 1, pp. 57–104, 1999.

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

[13] L. 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, pp. 493–504.

[14] P. Hell and J. Nešetřil, "On the complexity of H-coloring," *Journal of Combinatorial Theory, Series B*, vol. 48, no. 1, pp. 92–110, 1990.

[15] ——, *Graphs and Homomorphisms*. Oxford University Press, 2004.

[16] M. Thurley, "The complexity of partition functions on Hermitian matrices," *arXiv*, no. 1004.0992, 2010.