Counting 4 × 4 Matrix Partitions of Graphs[*]

Martin Dyer[†]   Leslie Ann Goldberg[‡]   David Richerby[‡]

20 June 2016

Abstract

Given a symmetric matrix $M \in \{0, 1, *\}^{D \times D}$, an M-partition of a graph G is a function from $V(G)$ to D such that no edge of G is mapped to a 0 of M and no non-edge to a 1. We give a computer-assisted proof that, when $|D| = 4$, the problem of counting the M-partitions of an input graph is either in FP or is #P-complete. Tractability is proved by reduction to the related problem of counting list M-partitions; intractability is shown using a gadget construction and interpolation. We use a computer program to determine which of the two cases holds for all but a small number of matrices, which we resolve manually to establish the dichotomy. We conjecture that the dichotomy also holds for $|D| > 4$. More specifically, we conjecture that, for any symmetric matrix $M \in \{0, 1, *\}^{D \times D}$, the complexity of counting M-partitions is the same as the related problem of counting list M-partitions.

1   Introduction

Let M be a symmetric matrix in $\{0, 1, *\}^{D \times D}$. An M-*partition* of an undirected graph $G = (V, E)$ is a partition of V into parts labeled by the elements of D (some of which may be empty). The partition is represented as a function $\sigma \colon V \to D$ where $\sigma(v)$ is the part of vertex v. It satisfies the following property: For all pairs of distinct vertices u and v,

- $M_{\sigma(u),\sigma(v)} \in \{1, *\}$ if $(u, v) \in E$ and

- $M_{\sigma(u),\sigma(v)} \in \{0, *\}$ if $(u, v) \notin E$.

Thus, if $M_{i,j} = 0$, no edges are permitted between vertices in parts i and j and, if $M_{i,j} = 1$, then all edges must be present between the two parts. If $M_{i,j} = *$, there is no restriction on edges between parts i and j. Note that self-loops play no role — the property applies only to pairs of distinct vertices u and v.

M-partitions were introduced by Feder, Hell, Klein and Motwani [5, 6] to study graph partition problems arising in the proof of the strong perfect graph conjecture, such as recognising skew cutsets, clique-cross partitions, two-clique cutsets and Winkler partitions. A skew

cutset of a connected graph $G = (V, E)$ is a pair of disjoint, non-empty sets $A, B \subset V$ such that $A \cup B$ is a cutset (deleting the vertices in A and B disconnects the graph) and G contains every possible edge between A and B. Skew cutsets correspond to M-partitions for

$$
M = \begin{array}{c c} & \begin{array}{c c c c} A & B & C & D \end{array} \\ \begin{array}{c} A \\ B \\ C \\ D \end{array} & \begin{pmatrix} * & 1 & * & * \\ 1 & * & * & * \\ * & * & * & 0 \\ * & * & 0 & * \end{pmatrix} \end{array}.
$$

The rows (and columns) correspond to parts A, B, C and D, respectively. Consider an M-partition in which every part is non-empty. $M_{A,B} = 1$ so G must contain every edge between those two parts. The rest of the graph must be assigned to parts C and D but, with no edges allowed between those parts, each of them must be a non-empty union of components of $G - (A \cup B)$. Therefore, the partition corresponds to a skew cutset. Clique-cross partitions, two-clique cutsets and Winkler partitions also correspond to M-partition problems for 4×4 matrices M; see [6] for both the definition of these problems and the corresponding matrices.

We study the problem of counting M-partitions, which was introduced by Hell, Hermann and Nevisi [8].

Name. $\#M$-PARTITIONS.

Instance. A graph G.

Output. $Z_M(G)$, the number of M-partitions of G.

Note that the matrix M is considered as a parameter and is not part of the input. For the decision problem of determining whether an M-partition of some graph exists, it is conventional to require every part to be non-empty since, otherwise, the problem is trivial whenever there is a $*$ on the diagonal (as is the case above). Counting, however, includes all M-partitions of the graph, including those where some parts may be empty. Hell, Hermann and Nevisi [8] show that, for any 2×2 or 3×3 matrix M, the problem $\#M$-PARTITIONS is either in FP or is $\#$P-complete. Our main result is an extension of this dichotomy to 4×4 matrices.

Theorem 1. *Let M be a symmetric matrix in $\{0, 1, *\}^{4 \times 4}$. Then $\#M$-PARTITIONS is either in FP or is $\#$P-complete.*

Thus, we completely resolve the complexity of counting M-partitions for 4×4 matrices, including all the examples above.

We explain the criterion that determines whether $\#M$-PARTITIONS is in FP or $\#$P-complete for a given symmetric 4×4 matrix M in the next section. Doing this requires the related concept of *list M-partitions*, also due to Feder et al. [6]. Here, each vertex of the input graph comes with a list of parts in which it is allowed to be placed. More formally, the input to the problem is a graph $G = (V, E)$ and a function $L\colon V \to \mathcal{P}(D)$, where $\mathcal{P}(\cdot)$ denotes the powerset. An M-partition σ of G *respects* the function L if $\sigma(v) \in L(v)$ for all vertices $v \in V$. The counting list M-partitions problem is defined as follows.

Name. $\#$LIST-M-PARTITIONS.

Instance. A graph G and a function $L\colon V(G) \to \mathcal{P}(D)$.

Output. The number of M-partitions of G that respect L.

The complexity of #List-M-partitions for all symmetric, square $\{0, 1, *\}$-matrices was recently determined by Göbel, Goldberg, McQuillan, Richerby and Yamakami [7]: depending on the structure of M, it is either in FP or is #P-complete.

The #M-partitions problem without lists is the special case of #List-M-partitions where $L(v) = D$ for every vertex v. Thus, there is a trivial polynomial-time Turing reduction from #M-partitions to #List-M-partitions. It is not known whether there is a polynomial-time Turing reduction in the other direction. As such, the dichotomy for counting list M-partitions does not necessarily translate into a dichotomy for counting M-partitions without lists.

M-partitions are also known as trigraph homomorphisms. Trigraphs are a generalisation of graphs, introduced by Chudnovsky [3], which allow $*$-edges. Thus trigraph homomorphisms are a generalisation of the well-known graph homomorphism problem [10]. Dyer and Greenhill [4] showed that, for any fixed graph H, the problem of counting homomorphisms from an input graph G to H is either in FP or is #P-complete, depending on the structure of H. The only polynomial-time cases are those where every component of H is either a complete graph with a self-loop on every vertex or a complete bipartite graph with no self-loops. The algorithm for the polynomial-time graph homomorphism cases is easily adapted to respect lists so, for any graph H, the problems of counting homomorphisms to H with and without lists have the same complexity [9].

We explain the criterion for the #List-M-partitions dichotomy from [7] in the following section. It is more complex than the criterion for graph homomorphisms, and so are the algorithms for the polynomial-time cases. Nonetheless, for every symmetric matrix M of size up to 4×4, it is true that #M-partitions and #List-M-partitions have the same complexity. We conjecture that this holds in general.

Conjecture 2. *Let M be a symmetric matrix in $\{0, 1, *\}^{D \times D}$. Then #M-partitions and #List-M-partitions have the same complexity.*

Proving this conjecture appears considerably more difficult than routinely extending the methods of Dyer and Greenhill [4], or even those of Bulatov's far-reaching generalisation [1]. The difficulty arises from the fact that some of the most powerful techniques used in proving those dichotomies do not seem to be applicable to the M-partitions problem.

1.1 The #List-M-partitions dichotomy

We now describe the complexity dichotomy for the #List-M-partitions problem, from [7]. The definitions and observation in this section are taken from that paper.

Definition 3. For any symmetric $M \in \{0, 1, *\}^{D \times D}$ and any sets $X, Y \in \mathcal{P}(D)$, define the binary relation
$$H^M_{X,Y} = \{(i, j) \in X \times Y \mid M_{i,j} = *\}\,.$$

The following notion of rectangularity was introduced by Bulatov and Dalmau [2].

Definition 4. A relation $R \subseteq D \times D'$ is *rectangular* if, for all $i, j \in D$, and $i', j' \in D'$,
$$(i, i'), (i, j'), (j, i') \in R \implies (j, j') \in R\,.$$

Definition 5. Given index sets X and Y, a matrix $M \in \{0, 1, *\}^{X \times Y}$ is *pure* if it has no 0s or has no 1s. M is *$*$-rectangular* if $H^M_{X,Y}$ is rectangular.

3

If M is a pure matrix with no 1s, then $Z_M(G)$ is the number of homomorphisms from the graph G to the graph whose adjacency matrix is obtained from M by changing all $*$s to 1s. If M is pure with no 0s, $Z_M(G)$ is the number of homomorphisms of the complement of G to the graph whose adjacency matrix is obtained from M by changing all 1s to 0s and then changing all $*$s to 1s. Thus, we sometimes refer to pure matrices as *homomorphism matrices*.

Definition 6. For any symmetric matrix $M \in \{0, 1, *\}^{D \times D}$, a set $\mathcal{L} \subseteq \mathcal{P}(D)$ is M-*purifying* if, for all $X, Y \in \mathcal{L}$, $M|_{X \times Y}$ is pure, where $M|_{X \times Y}$ is the submatrix formed by restricting to rows in X and columns in Y.

Definition 7. An $\mathcal{L}$-M-*derectangularising sequence* of length k is a sequence $D_1, \ldots, D_k$ with each $D_i \in \mathcal{L}$ such that:

- $\{D_1, \ldots, D_k\}$ is M-purifying and

- the relation $H^M_{D_1, D_2} \circ H^M_{D_2, D_3} \circ \cdots \circ H^M_{D_{k-1}, D_k}$ is not rectangular.

For brevity, we refer to a $\mathcal{P}(D)$-M-derectangularising sequence as an M-*derectangularising sequence* or as a derectangularising sequence of M.

Observation 8. If there is an $i \in \{1, \ldots, k\}$ such that $D_i = \emptyset$ then the relation $H = H^M_{D_1, D_2} \circ H^M_{D_2, D_3} \circ \cdots \circ H^M_{D_{k-1}, D_k}$ is the empty relation, which is trivially rectangular. If there is an i such that $|D_i| = 1$ then H is a Cartesian product, and is therefore rectangular. It follows that $|D_i| \geq 2$ for each i in a derectangularising sequence.

The complexity of #LIST-M-PARTITIONS is determined by the presence or absence of derectangularising sequences. The following is [7, Theorem 9].

Theorem 9. *Let M be a symmetric matrix in $\{0, 1, *\}^{D \times D}$. If there is an M-derectangularising sequence, then the problem #LIST-M-PARTITIONS is #P-complete. Otherwise, it is in* FP.

Thus, our conjecture that counting M-partitions has the same complexity as counting list M-partitions is the same as the following.

Conjecture 10. #M-PARTITIONS *is #P-complete if M has a derectangularising sequence, and is in* FP, *otherwise.*

1.2 Our contribution

Our main contribution is a computer-assisted proof of Theorem 1. This establishes a dichotomy for #M-PARTITIONS for 4×4 matrices that is consistent with Conjecture 2. We also show that Hell, Hermann and Nevisi's dichotomy for 2×2 and 3×3 matrices is consistent with our conjecture.

There are sufficiently few 2×2 and 3×3 $\{0, 1, *\}$-matrices that Hell, Hermann and Nevisi were able to determine the complexity of #M-PARTITIONS for all such matrices by case analysis. However, this approach does not seem feasible for larger matrices.

Recall that, for any symmetric matrix $M \in \{0, 1, *\}^{D \times D}$, #$M$-PARTITIONS is the special case of #LIST-M-PARTITIONS in which every vertex of the input graph is given list D. So, if #LIST-M-PARTITIONS is in FP, so is #M-PARTITIONS. By Theorem 9, this occurs precisely

when there is no M-derectangularising sequence. In Section 4, we give a method that can be used to show that some 4×4 matrices do not have M-derectangularising sequences.

In Section 5, we develop gadget-based techniques for showing #P-completeness of #M-PARTITIONS for symmetric $D \times D$ matrices M. Given an input graph G, we attach a gadget Γ to G. The parts of D into which the vertices of the gadget are placed determine the parts into which the vertices of G can be placed. If we could restrict to favourable partitions of the gadget, this would, in many cases, restrict G to be partitioned according to some proper submatrix M' for which #M'-PARTITIONS is known to be #P-complete by the work of Hell et al. [8].

We do not know how to restrict to specific partitions of the gadget. However, by varying the size of the gadget and using interpolation as follows, we are able to restrict to certain classes of partitions. This is enough to prove hardness in all but a few cases, by showing that we can use an oracle for #M-PARTITIONS to compute #M'-PARTITIONS for some hard submatrix M' of M. In more detail, let $J(\Gamma, G)$ be the graph that results from attaching the gadget Γ to the graph G. (In fact, we have two different ways of attaching the gadget, which are described in Section 5; we do not need the details, here.) For a set $S \subseteq D$, let $Z_M^S(\Gamma)$ be the number of M-partitions of the gadget Γ where exactly the parts in S are non-empty. In M-partitions of $J(\Gamma, G)$, placing Γ in the parts in S restricts the vertices of G to being placed in some set $E(S) \subseteq D$ of the parts. We can write

$$ Z_M(J(\Gamma, G)) = \sum_{S \subseteq D} Z_M^S(\Gamma) \, Z_{M|_{E(S)}}(G) \, , $$

where $M|_{E(S)}$ is the principal submatrix of M containing exactly the rows and columns with indices in $E(S)$.

The gadget Γ is just a clique or independent set of size k so $Z_M^S(\Gamma)$ is a polynomial-time computable function of M and k. Having computed these values, and also used the oracle to compute $Z_M(J(\Gamma, G))$, we can view the above equation as a linear equation in the "variables" $Z_{M|_{E(S)}}(G)$. By varying the size of the gadget, we can obtain a system of equations of this form, which we would hope to be able to solve. However, it is usually the case that there are distinct subsets $S_1, \ldots, S_r$ of D for which the functions $Z_M^{S_i}(\Gamma)$ for $1 \leq i \leq r$ are identical. In this case, we cannot solve for the variables $Z_{M|_{E(S_i)}}(G)$ individually but we can compute a weighted sum of them. In most cases, it turns out that only one of these variables is a #P-complete function. We can compute the weighted sum in polynomial time from the system of equations, and then compute all but one of the terms of that sum in polynomial time (with the assistance of the oracle, if needed), which allows us to compute a #P-complete function, completing the reduction from the problem of computing that function to #M-PARTITIONS.

We prove Theorem 1 with the aid of a computer program that, for each symmetric matrix $M \in \{0, 1, *\}^{4 \times 4}$ attempts to use the techniques of Section 4 to prove tractability and the interpolation technique of Section 5 to prove intractability. This is described in Section 6. The program resolves nearly all cases; the six exceptions (up to symmetries of the problem) are dealt with separately in Section 7. Finally, in Section 8, we show that our dichotomy for 4×4 matrices is consistent with our conjecture for the general case, Conjecture 10.

A similar computer-assisted proof could, in principle, be applied to 5×5 matrices, the number of which is not excessive (at most $3^{15} < 14,400,000$, even before symmetries are considered). Doing so requires automating more sophisticated handling of the sets of simultaneous linear equations and seems likely to result in a larger number of exceptional matrices than the six 4×4 matrices.

2 Preliminaries

Sets. We write $\mathcal{P}(D)$ for the powerset of D and $D^{(k)}$ for the set of k-element subsets of D. For convenience, we often list the elements of small sets as tuples (e.g., ac for $\{a, c\}$). For any natural number k, $[k]$ denotes the set $\{1, \ldots, k\}$.

Graphs. Since self-loops and parallel edges play no role in matrix partitions, we will assume that input graphs do not have self-loops or parallel edges. Let Γ_k^1 be the k-vertex complete graph and let Γ_k^0 be the k-vertex empty graph.[1] Let $\#\mathrm{IS}(G)$ and $\#\mathrm{Clique}(G)$ be the problems of determining, respectively, the number of independent sets and complete subgraphs of G.

Combinatorics. We write $(n)_k$ for the falling factorial $n(n-1)\cdots(n-k+1)$, taking $(n)_0 = 1$.

$$\left\{ {n \atop k} \right\} = \frac{1}{k!} \sum_{j=0}^{k} (-1)^{k-j} \binom{k}{j} j^n$$

denotes a Stirling number of the second kind. The number of surjective functions from a set of size n to a set of size k is $k! \left\{ {n \atop k} \right\}$. We will use the following bounds on $\left\{ {n \atop k} \right\}$:

$$\text{For } n \geq k \ln 2k, \quad \tfrac{1}{2} k^n/k! \leq \left\{ {n \atop k} \right\} \leq k^n/k!. \tag{1}$$

To see this, consider

$$\left\{ {n \atop k} \right\} = \sum_{j=0}^{k} (-1)^{k-j} \frac{j^n}{j!(k-j)!} = \sum_{j=0}^{k} (-1)^{k-j} s_j = S,$$

say. Now, $s_0 = 0$, $s_1 > 0$ and, for $j > 1$,

$$\frac{s_j}{s_{j-1}} = \frac{(j-1)!(k-j+1)!j^n}{j!(k-j)!(j-1)^n} = \frac{k-j+1}{j} \left(\frac{j}{j-1} \right)^n \geq 2,$$

if $(1-1/j)^n \leq (k-j+1)/(2j)$. Now $(1-1/j)^n \leq e^{-n/k}$, using $1-x \leq e^{-x}$ and $1 < j \leq k$. Also, $(k-j+1)/j \geq 1/k$ for $j \leq k$. Thus $s_j/s_{j-1} \geq 2$ if $e^{-n/k} \leq 1/(2k)$, i.e. $n \geq k \ln 2k$.

Thus, for $n \geq k \ln 2k$, S is an alternating series with strictly increasing terms. It follows that $s_k - s_{k-1} \leq S \leq s_k$. Equation (1) now follows, since $s_{k-1} \leq \tfrac{1}{2} s_k$ and $s_k = k^n/k!$.

Matrices. Let M be a symmetric $\{0, 1, *\}$-matrix with rows and columns indexed by a finite set D. For the 4×4 case, we adopt the convention that $D = \{a, b, c, d\}$ and we index the rows (and columns) a, b, c and d from top to bottom (left to right).

For sets $S, T \subseteq D$, we write $M|_{S \times T}$ for the submatrix of M obtained by restricting to the rows in S and the columns in T. $M|_S$ denotes the principal submatrix $M|_{S \times S}$.

Given a symmetric $D \times D$ matrix M and another symmetric $D' \times D'$ matrix M' with $|D| = |D'|$ we write $M \equiv M'$ if there is a bijection $\rho \colon D \to D'$ such that $M_{i,j} = M'_{\rho(i),\rho(j)}$ for all $i, j \in D$. It is clear that, if $M \equiv M'$, then $\#M$-PARTITIONS and $\#M'$-PARTITIONS have the same computational complexity.

---

[1] This nonstandard notation allows us to talk about a graph Γ_k^τ for $\tau \in \{0, 1\}$, simplifying the description of our gadget construction.

We write $\overline{M}$ for the matrix obtained from M by swapping all 0s and 1s. Note that the M-partitions of any graph G correspond directly to $\overline{M}$-partitions of the complement of G. Write $M \approx M'$ if $M \equiv M'$ or $M \equiv \overline{M'}$. Again, if $M \approx M'$, then $\#M$-PARTITIONS and $\#M'$-PARTITIONS have the same computational complexity.

We say that a matrix M is *easy* if the problem $\#M$-PARTITIONS is in FP and *hard* if it is $\#$P-complete.

3 2×2 and 3×3 matrices

Conjecture 2 is already known to hold for pure matrices. As we noted earlier, in this case $Z_M(G)$ is the number of homomorphisms from G (or its complement) to a graph whose edges correspond to the stars in M. The tractability criterion of Dyer and Greenhill [4, Theorem 1.1] for graph-homomorphism counting problems coincides with the tractability criterion for the problem with lists [9, Theorem 4]. The condition stated in these works concerns the graph H whose vertices are elements of D and whose edges (including self-loops) correspond to the stars in M. The tractability condition is that each component of H is either a complete graph in which every vertex has a self-loop or a complete bipartite graph in which no vertices have self-loops. Bulatov and Dalmau [2, Theorem 12] showed that this condition is equivalent to the condition that the relation $H_{D,D}^M$ is rectangular, which, in turn, is equivalent to the condition that M does not have $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ or $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$ or any permutation of these as a submatrix.

Conjecture 2 is also known to hold for impure 2×2 matrices. In particular, Hell, Hermann and Nevisi [8, Theorem 1] showed that for every impure symmetric 2×2 matrix M, $\#$LIST-M-PARTITIONS is in FP, hence so is $\#M$-PARTITIONS.

Hell, Hermann and Nevisi's dichotomy [8, Theorem 10] shows that if M is a symmetric impure 3×3 matrix then $\#M$-PARTITIONS is $\#$P-hard if M contains $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ or $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$ (or any permutation of these) as a principal submatrix. Otherwise, $\#M$-PARTITIONS is in FP. We will now show that this result is consistent with Conjecture 10, which we have already shown to be equivalent to Conjecture 2. In one direction, if M contains one of these hard principal submatrices then the rows and columns of this hard principal submatrix are an M-derectangularising sequence, so Conjecture 10 also says that M is hard. In the other direction, if M does not contain one of these hard principal submatrices then the following lemma shows that M has no derectangularising sequence, so Conjecture 10 also says that M is easy.

Lemma 11. *Let M be an impure 3×3 symmetric $\{0, 1, *\}$-matrix M with no principal hard 2×2 submatrix. Then M has no derectangularising sequence.*

Proof. Let $D_1, \ldots, D_k$ be a sequence of subsets of $D = \{a, b, c\}$. By Observation 8, if $|D_i| < 2$ for any i, the sequence cannot be derectangularising; if $|D_i| = 3$ for any i, the sequence is not derectangularising, since $M|_{D_i \times D_i} = M$ is not pure. Thus, $|D_i| = 2$ for all i.

Case 1. First, suppose that M has a non-principal hard 2×2 submatrix: without loss of generality, we may assume that $M|_{ab \times bc}$ contains three $*$s and one 0. Since M is impure, at least one of $M_{a,a}$ and $M_{c,c}$ must be 1: without loss of generality, assume that $M_{a,a} = 1$. In fact, we must have $M|_{ab \times bc} = \left(\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right)$ as, otherwise, every choice of $M_{c,c}$ would leave M containing a hard principal 2×2 submatrix. Therefore, $M = \left(\begin{smallmatrix} 1 & * & * \\ * & 0 & * \\ * & * & x \end{smallmatrix}\right)$ and $x \in \{0, 1\}$ since otherwise $M|_{ac}$ would be hard. The two choices for x lead to matrices that are $\approx$-equivalent, so we may assume that $x = 0$.

No derectangularising sequence can include $\{a, b\}$ or $\{a, c\}$ since $M|_{ab}$ and $M|_{ac}$ are impure. This leaves only $\{b, c\}$, but $H_{\{b,c\},\{b,c\}}^{M}$ is the disequality relation on the set $\{b, c\}$. Composing this with itself any number of times results in either equality or disequality, both of which are $*$-rectangular. Thus, M has no derectangularising sequence.

Case 2. Finally, suppose that M has no non-principal hard 2×2 submatrix. Let M' be the pure matrix formed from M by replacing every 1 with a 0. M' does not have $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ or any permutation of this as a submatrix. Equivalently, $H_{D,D}^{M'}$ is rectangular and the graph whose edges correspond to stars in M' has the property that every component is a complete graph in which every vertex has a self-loop or a complete bipartite graph in which no vertices have self-loops. There are only three elements in D so it is easy to see that M' has no derectangularising sequence. Since any M-derectangularising sequence is also an M'-derectangularising sequence, it follows that there is no M-derectangularising sequence. □

4 Tractability via #List-M-partitions

For any symmetric $D \times D$ matrix M, recall that #M-partitions is the special case of #List-M-partitions where the list of allowable parts for every vertex is D. Thus, if there is a polynomial-time algorithm for #List-M-partitions, a polynomial-time algorithm for #M-partitions is immediate.

By Theorem 9, #List-M-partitions is in FP if M has no derectangularising sequence. Determining that a general symmetric matrix has no derectangularising sequence is co-NP-complete [7, Theorem 10]. However, there are only finitely many 4×4 $\{0, 1, *\}$-matrices, so hardness of the general problem is moot. By [7, Lemma 27], any matrix in $\{0, 1, *\}^{4 \times 4}$ that has a derectangularising sequence has one of length at most 33,280 but it is not feasible to try all such sequences. In this section, we show that, in some cases, it is simple to determine that a 4×4 matrix has no derectangularising sequence.

Lemma 12. *Let M be a symmetric matrix in $\{0, 1, *\}^{D \times D}$ such that, for every $W \subseteq D^{(2)}$, at least one of the following holds:*

1. *There are $S, T \in W$ (not necessarily distinct) such that $M|_{S \times T}$ is not pure,*

2. *$W = \{S, T\}$, $S \cap T = \emptyset$ and $M|_{S \times T}$ is pure and $*$-rectangular, or*

3. *$M|_{\bigcup W}$ is pure and has no derectangularising sequence.*

Then #M-partitions is in FP.

Proof. If FP = #P, then #M-partitions is in FP for any matrix M. So we may assume that FP $\neq$ #P for the rest of the proof.

We prove the contrapositive. If #M-partitions is not in FP, then, by Theorem 9 and the assumption that FP $\neq$ #P, M has a derectangularising sequence. Choose such a sequence $D_1, \ldots, D_\ell$ that contains the least possible number of distinct sets among the D_i (i.e., a sequence that minimises $|\{D_1, \ldots, D_\ell\}|$). We show that none of the three properties holds for $W = \bigcup_i D_i^{(2)}$. By Observation 8, $|D_i| \geq 2$ for each $i \in [\ell]$.

For property 1, consider any $S \subseteq D_i$ and $T \subseteq D_j$ for any $i, j \in [\ell]$. $M|_{D_i \times D_j}$ is pure because $D_1, \ldots, D_\ell$ is M-purifying, so $M|_{S \times T}$ is pure. For property 3, suppose that $M|_{\bigcup W}$ is pure, since there is nothing more to prove if it is not. Since $|D_i| \geq 2$ for each i, $\bigcup W = \bigcup_i D_i$. Therefore, $D_1, \ldots, D_\ell$ is a derectangularising sequence of $M|_{\bigcup W}$.

It remains to show that W does not have property 2. Suppose that $W = \{S, T\}$ and $S \cap T = \emptyset$. If there were a D_i with $|D_i| > 2$, we would have $|W| > 2$, contradicting the assumption that $W = \{S, T\}$. Thus, $D_i \in \{S, T\}$ for each $i \in [\ell]$. By the definition of derectangularising sequence, $M|_{S \times S}$, $M|_{S \times T}$ and $M|_{T \times T}$ are all pure.

$M|_{S \times S}$ and $M|_{T \times T}$ must both be $*$-rectangular since, otherwise, S, S or T, T would be a derectangularising sequence, contradicting the choice of $D_1, \ldots, D_\ell$. If there is some $i \in [\ell - 1]$ such that $D_i = D_{i+1} = S$, then $H^M_{S,S}$ must be either the equality or disequality relation on S: any other relation would either not be rectangular or would prevent the sequence $D_1, \ldots, D_\ell$ from being derectangularising. Similarly, if we have $D_i = D_{i+1} = T$ for some i, then $H^M_{T,T}$ must be equality or disequality on T.

There must be some $i \in [\ell - 1]$ such that $D_i \neq D_{i+1}$. Without loss of generality, we may assume that $D_i = S$ and $D_{i+1} = T$. Consider $H^M_{S,T}$. If this were a matching or the complete relation $S \times T$, or if the projection onto its first and second columns were not S and T, respectively, then $D_1, \ldots, D_\ell$ would not be derectangularising. The only remaining possibility is that $H^M_{S,T}$ is not rectangular, i.e., $M|_{S \times T}$ is not $*$-rectangular. $\qquad\square$

Given a 4×4 matrix M, it is easy to check whether, for each of the 64 subsets of $D^{(2)}$, at least one of the three properties of Lemma 12 holds. If this is the case, we may deduce that M has no derectangularising sequence so is easy, even with lists.

5 Identifying hard matrices

For matrices M that are impure and, thus, not homomorphism matrices, we use a gadget construction and interpolation to "pick out" principal submatrices M' for which $\#M'$-PARTITIONS is $\#$P-complete. While we will be concerned with 4-element domains, the techniques in this section could potentially also be applied to arbitrary domains D, perhaps as part of a proof of a complexity dichotomy for all $\#M$-PARTITIONS problems, by induction on the size of the domain.

Given a Boolean value $\tau \in \{1, 0\}$, a graph G and a positive integer k, let $J^{0,\tau}(k, G)$ be the disjoint union of G and Γ^τ_k. The "0" in the notation is to remind us that there are no edges between G and the "gadget" Γ^τ_k (which is a complete graph if $\tau = 1$ and a graph with no edges if instead $\tau = 0$). Also, let $J^{1,\tau}(k, G)$ be the graph with vertex set $V(G) \cup V(\Gamma^\tau_k)$ and edge set $E(G) \cup E(\Gamma^\tau_k) \cup (V(G) \times V(\Gamma^\tau_k))$. The "1" in the notation is to remind us that all edges are present between G and the gadget Γ^τ_k.

The set of M-partitions of $J^{\pi,\tau}(k, G)$ can be broken down according to the set of parts $S \subseteq D$ in which vertices of the gadget Γ^τ_k are placed. For example, consider the matrix

$$M = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{c} a \quad b \quad c \quad d \\ \begin{pmatrix} 0 & 0 & 1 & * \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ * & 1 & 1 & * \end{pmatrix} \end{array} \qquad (2)$$

and take $\pi = \tau = 0$. In an M-partition of $J^{0,0}(k, G)$ in which the vertices of the Γ^0_k are all in part d, the vertices of G must be placed in parts a and d. Thus, the number of M-partitions of $J^{0,0}(k, G)$ in which the Γ^0_k is entirely within part d is equal to the number of $M|_{ad}$-partitions of G, which is the number of independent sets in G. If we could restrict

attention to only the M-partitions of $J^{0,0}(k, G)$ in which the Γ_k^0 is in part d, we could prove #P-completeness of #M-PARTITIONS by reduction from counting independent sets which, in the guise of monotone 2-SAT, was shown to be #P-complete by Valiant [12]. Unfortunately, we do not know how to restrict partitions in this way but, in this section, we set up machinery that nonetheless allows us to develop this idea into a method for proving hardness.

Definition 13. Let M be a symmetric matrix in $\{0, 1, *\}^{D \times D}$ and let $S \subseteq D$. An M-partition σ of a graph G is S-*surjective* if the image of σ is S. We write $Z_M^S(G)$ for the number of S-surjective M-partitions of G.

Given a set $S \subseteq D$, and a Boolean value $\pi \in \{0, 1\}$, let

$$E^\pi(S) = \{ j \in D \mid \forall i \in S, M_{i,j} \in \{\pi, *\} \}.$$

$E^1(S)$ is the set of parts in D that can be adjacent to every part in S; $E^0(S)$ is the set of parts that can be non-adjacent to every part in S. These will be interesting to us because we will proceed as follows in our reductions. Suppose that $M|_{E^\pi(S)}$ is a hard matrix and that we want to show that M is hard by reducing #$M|_{E^\pi(S)}$-PARTITIONS to #M-PARTITIONS. Then we can take an instance G of #$M|_{E^\pi(S)}$-PARTITIONS and form the gadget $J^{\pi,\tau}(k, G)$ for some value of k. Then, if we can choose τ so that the gadget Γ_k^τ is always partitioned surjectively into parts in S, we will have reduced #$M|_{E^\pi(S)}$-PARTITIONS to #M-PARTITIONS. Typically, we cannot do this, but we will be able to do is to compute the number of M-partitions of $J^{\pi,\tau}(k, G)$ for lots of values of k. Using polynomial interpolation, we will be able to work out the number of M-partitions of G which are consistent with an S-surjective partition of Γ_k^τ so this will enable us to count the $M|_{E^\pi(S)}$-partitions of G (solving a hard problem) by using an oracle for counting M-partitions. Thus, we will have proved that M is a hard matrix.

For $\pi \in \{0, 1\}$, we say that a principal submatrix M' of M is (M, π)-*accessed by* S if $M' \equiv M|_{E^\pi(S)}$. Note the equivalence — M' only has to be equivalent to $M|_{E^\pi(S)}$ — it doesn't have to *be* $M|_{E^\pi(S)}$. It is useful to define things this way because equivalent matrices correspond to matrix partition problems of equivalent difficulty. Also, we will not be able to separate them by interpolation, so we will have to consider them together.

To illustrate these definitions, consider the matrix M in Equation (2). Then $E^1(\{b, d\}) = \{c, d\}$ and $E^0(\{b, d\}) = \{a\}$. Thus, $M|_{cd}$ is $(M, 1)$-accessed by $\{b, d\}$ and $M|_a$ is $(M, 0)$-accessed by $\{b, d\}$. The matrix $M|_b$ is also $(M, 0)$-accessed by $\{b, d\}$ since $M|_b \equiv M|_a$. Also, $E^1(\{d\}) = \{a, b, c, d\}$. Thus, M itself is $(M, 1)$-accessed by $\{d\}$.

We say that a principal submatrix M' of M is *accessible* in the graph $J^{\pi,\tau}(k, G)$ if there is a set $S \subseteq D$ such that $Z_M^S(\Gamma_k^\tau) > 0$ and M' is (M, π)-accessed by S.

Continuing our example with $S = \{b, d\}$ and M as in Equation (2), note that for any $k > 1$, $Z_M^S(\Gamma_k^1) > 0$ since an S-surjective M-partition of Γ_k^1 may place one vertex in part b and the remaining vertices in part d. Thus, $M|_{cd}$ is accessible in $J^{1,1}(k, G)$ and $M|_a$ and $M|_b$ are accessible in $J^{0,1}(k, G)$. Note that accessibility in $J^{\pi,\tau}(k, G)$ depends on M, π, τ and possibly k but it does not depend on G. Because of this, we may talk about accessibility in $J^{\pi,\tau}(k, \cdot)$. In fact, we will see later in Theorem 18 that accessibility will not actually depend on k, provided that $k > |D|$ (this is not obvious at this point but will be important).

We now begin to decompose $Z_M(J^{\pi,\tau}(k, G))$ into more manageable units. The first step is to break the sum up over the set S which is used to surjectively partition the gadget Γ_k^τ:

$$Z_M(J^{\pi,\tau}(k, G)) = \sum_{S \subseteq D} Z_M^S(\Gamma_k^\tau) \, Z_{M|_{E^\pi(S)}}(G) \,. \tag{3}$$

Now let
$$\Psi_\pi = \{S \subseteq D \mid M \text{ itself is } (M, \pi)\text{-accessed by } S\}.$$

The set Ψ_π may be empty, depending on M. The reason that we have defined Ψ_π is that wish to use Equation (3) to show that M is a hard matrix — so we will use an oracle for M-parititons to compute the left-hand side and we will hope to discover the solution to some hard problem on the right-hand side. For this reason we don't want M itself to be one of the matrices $M|_{E^\pi(S)}$ appearing on the right-hand side. To ease the notation, let $\overline{\Psi_\pi} = \mathcal{P}(D) \setminus \Psi_\pi$; $\overline{\Psi_\pi}$ consists of all subsets S of D apart from those with $M|_{E^\pi(S)} = M$. From (3), we have

$$Z_M(J^{\pi,\tau}(k, G)) - \sum_{S \in \Psi_\pi} Z_M^S(\Gamma_k^\tau) Z_M(G) = \sum_{S \in \overline{\Psi_\pi}} Z_M^S(\Gamma_k^\tau) Z_{M|_{E^\pi(S)}}(G). \tag{4}$$

Now we would like to collect the terms on the right-hand side of Equation (4), gathering all terms with the same matrix $M|_{E^\pi(S)}$, and taking these together. So, for any principal submatrix M' of M, let

$$C_{M'}^{\pi,\tau}(k) = \sum_S Z_M^S(\Gamma_k^\tau),$$

Where the sum is over sets $S \subseteq D$ such that M' is (M, π)-accessed by S. Thus, M' is accessible in $J^{\pi,\tau}(k, \cdot)$ precisely when $C_{M'}^{\pi,\tau}(k)$ is positive. The quantity $C_{M'}^{\pi,\tau}(k)$ corresponds roughly to the coefficient of $Z_{M'}(G)$ in (4) though we will have to be careful about over-counting. As a first step, we can immediately rewrite the left-hand side of (4), combining the terms for all $S \in \Psi_\pi$, since these terms have a common factor of $Z_M(G)$.

$$Z_M(J^{\pi,\tau}(k, G)) - C_M^{\pi,\tau}(k) Z_M(G) = \sum_{S \in \overline{\Psi_\pi}} Z_M^S(\Gamma_k) Z_{M|_{E^\pi(S)}}(G). \tag{5}$$

Now, all of the matrices $M|_{E^\pi(S)}$ such that $Z_{M|_{E^\pi(S)}}(G)$ arises on the right-hand side of (5) are proper principal sub-matrices of M. Since a proper principal sub-matrix M' is (M, π)-accessed by S when $M' \equiv M|_{E^\pi(S)}$, the coefficient $C_{M'}^\pi$ captures the contribution of the entire equivalence class. Thus, we have

$$Z_M(J^{\pi,\tau}(k, G)) - C_M^{\pi,\tau}(k) Z_M(G) = \sum_{M'} C_{M'}^{\pi,\tau}(k) Z_{M'}(G), \tag{6}$$

where the sum is over one element from each $\equiv$-equivalence class of proper principal sub-matrices M' of M.

We now explain the point of Equation (6). Corollary 19 will show that all of the coefficients $C_{M'}^{\pi,\tau}(k)$ can be computed in polynomial time (as a function of k). Also, the left side of (6) can be computed in polynomial time with an oracle for computing Z_M — we just use the oracle twice to compute $Z_M(J^{\pi,\tau}(k, G))$ and $Z_M(G)$. So if we can show that it is hard to compute the right side of (6), then we can conclude that computing Z_M is hard.

Since each M' is a proper principal submatrix of M, the complexity of computing each $Z_{M'}$ is known from the dichotomy of Hell, Hermann and Nevisi [8] and is either in FP or is #P-complete.

We begin with two straightforward cases in Lemmas 14 and 15. These cases do not require interpolation, but we will handle these cases first and then explain the interpolation.

Lemma 14. *Suppose that M is a symmetric matrix in $\{0, 1, *\}^{D \times D}$, that π and τ are Boolean values in $\{0, 1\}$, and that k is some positive integer. If there is at least one proper hard submatrix of M that is accessible in $J^{\pi,\tau}(k, \cdot)$ and all such proper hard submatrices are $\equiv$-equivalent, then M is hard.*

Proof. Suppose that, up to $\equiv$-equivalence, M'' is the only hard proper submatrix that is accessible in $J^{\pi,\tau}(k, \cdot)$. Rearranging (6), we obtain, for any graph G,

$$Z_{M''}(G) = \frac{1}{C_{M''}^{\pi,\tau}(k)} \left(Z_M(J^{\pi,\tau}(k, G)) - C_M^{\pi,\tau}(k) \, Z_M(G) \; - \sum_{M' \not\equiv M''} C_{M'}^{\pi,\tau}(k) \, Z_{M'}(G) \right).$$

Since all the quantities $C_{M'}^{\pi,\tau}(k)$ and $Z_{M'}(G)$ are computable in FP (which follows since all M' are easy by assumption, and the coefficients $C_{M'}^{\pi,\tau}(k)$ are constants) this gives a polynomial-time Turing reduction from #M''-PARTITIONS to #M-PARTITIONS. $\qquad \square$

Lemma 15. *Suppose that M is a symmetric matrix in $\{0, 1, *\}^{D \times D}$, that π and τ are Boolean values in $\{0, 1\}$, and that k is some positive integer. Suppose that there is at least one proper hard submatrix of M that is accessible in $J^{\pi,\tau}(k, \cdot)$ that is $\equiv$-equivalent to $M_0 = \left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix} \right)$ and that there is at least one proper hard submatrix of M that is accessible in $J^{\pi,\tau}(k, \cdot)$ that is $\equiv$-equivalent to $M_1 = \left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix} \right)$. Suppose that every proper hard submatrix that is accessible is either $\equiv$-equivalent to $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix} \right)$ or to $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix} \right)$. Then M is hard.*

Proof. Recall that #IS(G) and #Clique(G) are, respectively, the number of independent sets and complete subgraphs in a graph G. Computing each of these is #P-complete [12] and they correspond to #$\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix} \right)$-PARTITIONS and #$\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix} \right)$-PARTITIONS, respectively.

We first show that, for any fixed integers α and β, computing the function $\theta_{\alpha,\beta}(G) = \alpha$#IS$(G) + \beta$#Clique$(G)$ is also #P-complete unless $\alpha = \beta = 0$. Assume that α and β are both non-zero as the result is trivial, otherwise. Observe that, for any graph G,

$$\#\text{IS}(G + K_1) = 2\#\text{IS}(G)$$
$$\#\text{Clique}(G + K_1) = \#\text{Clique}(G) + 1 \,.$$

Therefore,

$$\theta_{\alpha,\beta}(G + K_1) - \theta_{\alpha,\beta}(G) = \alpha\#\text{IS}(G) + \beta \,,$$

which is #P-complete to compute since $\alpha \neq 0$. Thus, we have shown that computing $\theta_{\alpha,\beta}(\cdot)$ is #P-complete.

Now rearrange (6) as in the proof of Lemma 14.

$$C_{M_0}^{\pi,\tau}(G) \, Z_{M_0}(G) \; + \; C_{M_1}^{\pi,\tau}(G) \, Z_{M_1}(G)$$
$$= Z_M(J^{\pi,\tau}(k, G)) \; - \; C_M^{\pi,\tau}(k) \, Z_M(G) \; - \sum_{M' \notin \{M_0, M_1\}} C_{M'}^{\pi,\tau}(k) \, Z_{M'}(G) \,,$$

where the sum is over one element from each $\equiv$-equivalence class of proper principal submatrices M' of M other than the equivalence classes of M_0 and M_1. Writing #IS(G) for $Z_{M_0}(G)$ and #Clique(G) for $Z_{M_1}(G)$, and taking $\alpha = C_{M_0}^{\pi,\tau}(k)$ and $\beta = C_{M_1}^{\pi,\tau}(k)$, we get

$$\theta_{\alpha,\beta}(G) \; = \; Z_M(J^{\pi,\tau}(k, G)) \; - \; C_M^{\pi,\tau}(k) \, Z_M(G) \; - \sum_{M' \notin \{M_0, M_1\}} C_{M'}^{\pi,\tau}(k) \, Z_{M'}(G) \,.$$

Thus, we have reduced the #P-hard problem of computing $\theta_{\alpha,\beta}(\cdot)$ to the problem of evaluating the right-hand side, which can be done in polynomial time with an oracle for #M-PARTITIONS. We conclude that #M-PARTITIONS is #P-complete. $\square$

Lemmas 14 and 15 give us a tool for identifying some hard matrices M. However, neither of these lemmas helps with our example matrix (2). To make progress, we will use interpolation. First, in Theorem 18, we will show that the value of $Z_M^S(\Gamma_k^\tau)$ is very constrained — there are only a few possible values, depending on k. Further, in Lemma 20 we will show that these values are linearly independent as functions of k. We will later use this fact to prove hardness by interpolation.

Definition 16. Let $f_{\ell,s}(k) = (k)_\ell \, (s-\ell)! \left\{ {k-\ell \atop s-\ell} \right\}$.

$f_{\ell,s}(k)$ is the number of ways that a set of size k can be partitioned into s parts, the first ℓ of which have size exactly 1 and the remaining $s-\ell$ of which have size at least 1.

Definition 17. Let M be any symmetric matrix in $\{0,1,*\}^{D \times D}$. Let $\tau \in \{0,1\}$ be a Boolean value. For $S \subseteq D$, let $\ell(M,S,\tau) = |\{i \in S \mid M_{i,i} = \tau \oplus 1\}|$. Let

$$\mathcal{E}(M,\tau) = \left\{ S \mid \ell(M,S,\tau) = |S| \text{ or there are distinct } i,j \in S \text{ with } M_{i,j} = \tau \oplus 1 \right\}.$$

Intuitively, $\mathcal{E}(M,\tau)$ is the set of subsets S of D that will not be useful for S-surjectively partitioning the gadget K_k^τ (as long as $k > |D|$). For example, if there are distinct $i,j \in S$ with $M_{i,j} = \tau \oplus 1$ then we can't simultaneously use parts i and j, so an S-surjective partition is impossible. We will see below that an S-surjective partition is also impossible if $\ell(M,S,\tau) = |S|$. The following theorem shows that as long as $S \notin \mathcal{E}(M,\tau)$ the number of S-surjective M-partitions of Γ_k^τ is a simple function of k.

Theorem 18. *Let M be any symmetric matrix in $\{0,1,*\}^{D \times D}$ and suppose $S \subseteq D$ and $\tau \in \{0,1\}$. If $S \in \mathcal{E}(M,\tau)$ then, for all $k > |D|$, $Z_M^S(\Gamma_k^\tau) = 0$. Otherwise, for all $k > |D|$, $Z_M^S(\Gamma_k^\tau) = f_{\ell(M,S,\tau),|S|}(k)$.*

Proof. **Case 1.** Suppose there are distinct $i,j \in S$ with $M_{i,j} = \tau \oplus 1$. Then no M-partition of any Γ_k^τ can place elements in both parts i and j. Thus, for any k, there are no S-surjective M-partitions of Γ_k^τ, so $Z_M^S(\Gamma_k^\tau) = 0$.

Case 2. Suppose we are not in Case 1. Let $S' = \{i \in S \mid M_{i,i} = \tau \oplus 1\}$ so $\ell(M,S,\tau) = |S'|$. In any S-surjective M-partition of any Γ_k^τ, every part in S' must contain exactly one vertex.

Case 2a. If $S \in \mathcal{E}(M,\tau)$ then $|S'| = |S|$, so for all $k > |D| \geq |S'|$, we have $Z_M^S(\Gamma_k^\tau) = 0$.

Case 2b. Otherwise, $S \notin \mathcal{E}(M,\tau)$. Let $\ell = \ell(M,S,\tau) < |S|$. Now, for any $k > |S|$, $Z_M^S(\Gamma_k^\tau) = f_{\ell,|S|}(k)$. To see this, note that there are $(k)_\ell$ ways to choose one vertex of Γ_k^τ to place in each part in S'. This leaves the remaining $k - \ell$ vertices to be surjectively placed in the $|S| - \ell$ parts in $S \setminus S'$. There are $(|S| - \ell)! \left\{ {k-\ell \atop |S|-\ell} \right\}$ ways of doing this. $\square$

Since $f_{\ell,s}(k)$ can be evaluated in polynomial time (as a function of k), we obtain the following corollary.

Corollary 19. *For any symmetric matrix M in $\{0,1,*\}^{D \times D}$ and any $S \subseteq D$, the S-surjective M-partitions of complete and empty graphs can be counted in polynomial time.*

Lemma 20. *Suppose $|D| \geq 2$. Then there is a full rank matrix F satisfying the following properties.*

- *The columns of F are indexed by the pairs (ℓ, s) with $0 \le \ell < s \le |D|$.*

- *The rows of F are indexed by $\binom{|D|+1}{2}$ distinct values $k_1 < k_2 < \ldots$, all of which are greater than $|D|$.*

- *For each row k_i and each column (ℓ, s), the corresponding entry in F is $f_{\ell,s}(k_i)$.*

Proof. Let $d = |D|$ and let $U = \{(\ell, m) \mid 0 \le \ell < d \text{ and } 1 \le m \le d - \ell\}$. For $(\ell, m) \in U$, let $\phi_{\ell,m}(k) = f_{\ell,\ell+m}(k)$. The stated properties of the matrix F indicate that the function $\phi_{\ell,m}$ maps every row index k to the entry in row k and column $(\ell, \ell + m)$ of F. Let $\Phi = \{\phi_{\ell,m} \mid (\ell, m) \in U\}$.

We will show that the functions in Φ (which correspond to the columns of F) are linearly independent (as functions of k). To do this, we define a strict ordering $<$ on functions in Φ. Then we will show that for any $\phi \in \Phi$, the function ϕ cannot be expressed as a linear combination of the functions in $\{\phi' \in \Phi \mid \phi' < \phi\}$, because it grows too fast as k increases. Then we will also be able to conclude that $\binom{d+1}{2}$ row indices can be chosen so that the matrix F has full rank, and the other properties in the statement of the lemma are satisfied.

We first define the ordering on the $\binom{d+1}{2}$ functions in Φ. We do this by defining a lexicographic ordering on the set U of column indices, and then ordering the functions in Φ accordingly. For (ℓ', m') and (ℓ, m) in U, we say that $(\ell', m') < (\ell, m)$ if one of the following is true:

- $m' < m$, or

- $m' = m$ and $\ell' < \ell$.

We use the natural induced order on functions: $\phi_{\ell',m'} < \phi_{\ell,m}$ if and only if $(\ell', m') < (\ell, m)$.

For convenience, let $\Phi_{\ell,m} = \{\phi \in \Phi \mid \phi < \phi_{\ell,m}\}$. We will show that $\phi_{\ell,m}$ is not in the span of $\Phi_{\ell,m}$, for all $(\ell, m) \in U$. We start by deriving bounds on $\phi_{\ell,m}(k)$. If k is an integer that is at least $\ell + m \ln 2m$, then, from Equation (1), we have

$$\phi_{\ell,m}(k) = (k)_\ell \, m! \begin{Bmatrix} k - \ell \\ m \end{Bmatrix} = \begin{cases} \le (k)_\ell \, m! \, m^{k-\ell}/m! & = (k)_\ell \, m^{k-\ell}, \\ \ge (k)_\ell \, m! \, \frac{1}{2} m^{k-\ell}/m! & = \frac{1}{2}(k)_\ell \, m^{k-\ell}. \end{cases}$$

Now $k^\ell \ge (k)_\ell \ge (k - \ell)^\ell = k^\ell (1 - \frac{\ell}{k})^\ell \ge k^\ell (1 - \ell^2/k) \ge \frac{1}{2}k^\ell$ if $k \ge 2\ell^2$. So, if $k \ge 2\ell^2 + m \ln 2m$, then

$$\phi_{\ell,m}(k) = (k)_\ell \, m! \begin{Bmatrix} k - \ell \\ m \end{Bmatrix} = \begin{cases} \le (k)_\ell \, m^{k-\ell} & \le k^\ell \, m^{k-\ell}, \\ \ge \frac{1}{2}(k)_\ell \, m^{k-\ell} & \ge \frac{1}{4}k^\ell \, m^{k-\ell}. \end{cases} \tag{7}$$

Now, we wish to show that $\phi_{\ell,m}$ is not in the span of $\Phi_{\ell,m}$. The claim is trivial if $\ell = 0$ and $m = 1$ since $\Phi_{0,1} = \emptyset$, so suppose otherwise. Consider any function ψ in the linear span of $\Phi_{\ell,m}$. We will show that ψ is not equal to $\phi_{\ell,m}$. Clearly, we can assume that ψ is not identically 0 since $\phi_{\ell,m}$ is not identically zero. By the definition of linear span, there are real numbers β_ϕ, not depending on k, so that $\psi(k) = \sum_{\phi \in \Phi_{\ell,m}} \beta_\phi \phi(k)$. First suppose $m' \le m - 1$ for all $\phi_{\ell',m'} \in \Phi$. Plugging in (7), we will show that, if k is sufficiently large, then

$$\psi(k) \le \sum_{\phi \in \Phi_{\ell,m}} \beta_\phi k^d (m-1)^k \le \beta_\Phi k^d (m-1)^k < \frac{1}{8} k^\ell m^{k-\ell} \le \frac{1}{2} \phi_{\ell,m}(k), \tag{8}$$

14

where $\beta_\Phi = \sum_{\phi \in \Phi_{\ell,m}} |\beta_\phi| > 0$. Note that β_Φ depends on ψ, ℓ and m but not on k. Now (8) holds if $k \geq 2\ell^2 + m \ln 2m$ (for the final inequality) and $8\beta_\Phi m^\ell k^d (1 - 1/m)^k < 1$ (for the strict inequality). The latter inequality is true if $k^d e^{-k/m} < 1/(8\beta_\Phi m^\ell)$. Now $k^d \leq e^{k/2m}$ if $k/(\ln k) > 2dm$, which is true if $k > 4m^2 d^2$ (since $\ln k < \sqrt{k}$ for all $k \geq 1$). So, if $k > \max(2\ell^2 + m \ln 2m, 4m^2 d^2)$, the condition becomes $e^{k/2m} > 8\beta_\Phi m^\ell$, i.e. $k > 2m \ln(8\beta_\Phi m^\ell)$. So, if $k > \max(2\ell^2 + m \ln 2m, 4m^2 d^2, 2m \ln(8\beta_\Phi m^\ell))$, then $\psi(k) < \frac{1}{2}\phi_{\ell,m}(k)$, so $\psi \neq \phi_{\ell,m}$.

In the general case, let $\Phi' = \{\phi \in \Phi_{\ell,m} \mid \phi < \phi_{d,m-1}\}$ and $\Phi'' = \{\phi_{\ell',m} \in \Phi_{\ell,m} \mid \ell' < \ell\}$. Thus

$$\psi(k) = \sum_{\phi \in \Phi_{\ell,m}} \beta_\phi \phi(k) = \sum_{\phi \in \Phi'} \beta_\phi \phi(k) + \sum_{\phi \in \Phi''} \beta_\phi \phi(k).$$

Now, using the proof of (8) above,

$$\sum_{\phi \in \Phi'} \beta_\phi \phi(k) \;\; < \;\; \tfrac{1}{2}\phi_{\ell,m}(k),$$

if $k > \max(2\ell^2 + m \ln 2m, 4m^2 d^2, 2m \ln(8\beta_{\Phi'} m^\ell))$, where $\beta_{\Phi'} = \sum_{\phi \in \Phi'} |\beta_\phi|$. Also, using (7) again,

$$\sum_{\phi \in \Phi''} \beta_\phi \phi(k) < \beta_{\Phi''} k^{\ell-1} m^{k-\ell} < \tfrac{1}{8} k^\ell m^{k-\ell} \leq \tfrac{1}{2}\phi_{\ell,m}(k),$$

provided that we also have $k > 8\beta_{\Phi''}$, where $\beta_{\Phi''} = \sum_{\phi \in \Phi''} |\beta_\phi|$. Thus if

$$k > k' = \max(2\ell^2 + m \ln 2m, 4m^2 d^2, 2m \ln(8\beta_{\Phi'} m^\ell)), 8\beta_{\Phi''}), \tag{9}$$

we have $\psi(k) < \phi_{\ell,m}(k)$, and so $\psi \neq \phi_{\ell,m}$.

Now we will show how to choose $\binom{d+1}{2}$ row indices $k_1, k_2, \ldots$, so that F has full rank, and the other properties in the statement of the lemma are satisfied. Order the columns of F according to the ordering $<$ defined above. We will choose the row-indices k_1, k_2 inductively, using the invariant that F^i, which the sub-matrix defined by the row-indices $k_1, \ldots, k_i$ and the first i columns in U, has full rank. The base case, $i = 1$, is trivial — for concreteness, take $k_1 = d + 1$. Now consider the inductive step, and the choice of k_{i+1}. Let (ℓ, m) denote the $(i+1)$st pair in U. Since F^i has full rank, there is exactly one linear combination of the first i columns of F^i that agrees with the $(i+1)$st column on the rows with indices $k_1, \ldots, k_i$. Thus, there is only one possible linear combination ψ in the linear span of $\Phi_{\ell,m}$ that that agrees with $\phi_{\ell,m}$ on $k_1, \ldots, k_i$. Now, use (9) to choose k' so that $\phi_{\ell,m}(k) > \psi(k)$ for $k > k'$, and set $k_{i+1} = \min(k_i, \lceil k' \rceil) + 1$. This completes the inductive step, and the proof. $\square$

At this point is helpful to recall our construction of the graph $J^{\pi,\tau}(k, G)$ from G. It also helps to recall Equation (3).

$$Z_M(J^{\pi,\tau}(k, G)) = \sum_{S \subseteq D} Z_M^S(\Gamma_k^\tau) \, Z_{M|_{E^\pi(S)}}(G) \,.$$

We know from Theorem 18 that, for any matrix $M|_{E^\pi(S)}$ corresponding to an element S of the sum, either $S \in \mathcal{E}(M, \tau)$ in which case the function $Z_M^S(\Gamma_k^\tau)$ is identically zero (assuming $k > |D|$) or $S \notin \mathcal{E}(M, \tau)$ in which case it is identically the function $f_{\ell(M,S,\tau),|S|}(k)$ (as a function of k). Let

$$\mathcal{S}(\ell, s, M, \tau) = \{S \in \mathcal{P}(D) \setminus \mathcal{E}(M, \tau) \text{ such that } |S| = s \text{ and } \ell(M, s, \tau) = \ell\} \,.$$

15

$\mathcal{S}(\ell, s, M, \tau)$ is the set of sets $S \subseteq D$ such that $Z_M^S(\Gamma_k^\tau) = f_{\ell,s}(k)$. Thus, we can rewrite Equation (3) for $k > |D|$ as

$$Z_M(J^{\pi,\tau}(k, G)) \quad = \sum_{0 \le \ell < s \le |D|} f_{\ell,s}(k) \sum_{S \in \mathcal{S}(\ell, s, M, \tau)} Z_{M|_{E^\pi(S)}}(G) \,. \qquad (10)$$

Now the point is that the $f_{\ell,s}(k)$ entries are linearly independent functions of k by Lemma 20. We will see in the proof of Theorem 21 that we will be able to choose sufficiently many values of k, evaluate the left-hand side $Z_M(J^{\pi,\tau}(k, G))$ for each of these using an oracle for $\#M$-PARTITIONS and then interpolate to compute each "coefficient" of $f_{\ell,s}(k)$ on the right-hand side. That is, we show how to compute each value $\sum_{S \in \mathcal{S}(\ell, s, M, \tau)} Z_{M|_{E^\pi(S)}}(G)$. If computing one of these values (for an input G) is a hard problem, then we will have proved that $\#M$-PARTITIONS is also $\#$P-complete.

Before we proceed it will help to rewrite (10) one last time, splitting the sum over principal submatrices of M. For $0 \le \ell < s \le |D|$, let

$$A_M^{\pi;\tau}(\ell, s) = \{M|_{E^\pi(S)} \mid S \in \mathcal{S}(\ell, s, M, \tau)\} \,.$$

$A_M^{\pi;\tau}(\ell, s)$ is just the set of matrices M' such that the coefficient of $f_{\ell,s}(k)$ in (10) has a $Z_{M'}(G)$ term. As before, we will need to deal with equivalences between matrices. Let $A_M^{\pi;\tau}(\ell, s)/\equiv$ be the set containing one matrix from each $\equiv$-equivalence class of $A_M^{\pi;\tau}(\ell, s)$. For each matrix M' in $A_M^{\pi;\tau}(\ell, s)/\equiv$, let

$$n_{M'}(\ell, s) = |\{S \in \mathcal{S}(\ell, s, M, \tau) \mid M|_{E^\pi(S)} \equiv M'\}| \,.$$

$n_{M'}(\ell, s)$ is just the number of times that a term $Z_{M''}(G)$ arises in the coefficient of $f_{\ell,s}(k)$ where $M'' \equiv M'$. Now, for $k > |D|$ we can rewrite Equation (10) as

$$Z_M(J^{\pi,\tau}(k, G)) \quad = \sum_{0 \le \ell < s \le |D|} f_{\ell,s}(k) \, T_{M,\ell,s}^{\pi,\tau}(G) \,, \qquad (11)$$

where

$$T_{M,\ell,s}^{\pi,\tau}(G) \quad = \sum_{M' \in A_M^{\pi,\tau}(\ell,s)/\equiv} n_{M'}(\ell, s) \, Z_{M'}(G) \,. \qquad (12)$$

Theorem 21. *Let M be any symmetric matrix in $\{0, 1, *\}^{D \times D}$. Suppose that there are ℓ and s satisfying $0 \le \ell < s \le |D|$ and Boolean values π and τ in $\{0, 1\}$ such that, up to $\equiv$-equivalence, $A_M^{\pi,\tau}(\ell, s)$ contains either*

- *exactly one hard proper principal submatrix of M; or*

- *exactly two hard proper principal submatrices of M and these are $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$.*

Then $\#M$-PARTITIONS is $\#$P-complete.

Proof. First, let's go back to Equation (11). Let $\Delta = \binom{|D+1|}{2}$. Note that M, π and τ are all fixed. Consider a graph G. In the proof we will consider the quantities $T_{M,\ell,s}^{\pi,\tau}(G)$ to be a set of Δ "variables" indexed by the pairs (ℓ, s). We will compute the values of these variables by

16

making multiple evaluations of $Z_M(J^{\pi,\tau}(k,G))$ for different values of k (using an oracle for $\#M$-PARTITIONS).

It will help to have an enumeration of the Δ pairs (ℓ,s) with $0 \le \ell < s \le |D|$, so let (ℓ_j, s_j) be the j'th such pair (for $1 \le j \le \Delta$). Choose Δ distinct values $k_1, \ldots, k_\Delta$, which meet the requirements of Lemma 20. Let F be the $\Delta \times \Delta$ integer matrix whose (i,j)'th entry $F_{i,j}$ is $f_{\ell_j, s_j}(k_i)$.

Using an oracle for $\#M$-PARTITIONS, we can compute the entries of a length-Δ column vector $\overline{Z}$ whose i'th entry is $Z_M(J^{\pi,\tau}(k_i, G))$.

Let $\overline{T}$ be a length-Δ column vector whose j'th entry is the j'th variable $T^{\pi,\tau}_{M,\ell_j,s_j}(G)$. Then Equation (11) gives the system of equations $\overline{Z} = F\overline{T}$.

Lemma 20 shows that F has full rank so F can be inverted, and we can compute all of the variables $T^{\pi,\tau}_{M,\ell,s}(G)$ using $F^{-1}\overline{Z} = \overline{T}$ and using the $\#M$-PARTITIONS oracle to compute the values of $\overline{Z}$.

By Equation (12), each variable $T^{\pi,\tau}_{M,\ell,s}(G)$ is a sum of terms, each of which is a constant multiple of $Z_M(G)$ or of $Z_{M'}(G)$ for some proper principal submatrix M' of M. If exactly one of these submatrices M' is hard, we can use the polynomial-time algorithms for the other problems $Z_{M''}$ ($M'' \not\equiv M'$) to compute $Z_{M'}(G)$ in polynomial time. If exactly two of the submatrices M' are hard and these are $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$, we can similarly compute $\alpha\#\mathrm{IS}(G) + \beta\#\mathrm{Clique}(G)$ in polynomial time for constants $\alpha, \beta \ge 1$, which is $\#P$-complete by Lemma 15. In both cases, we conclude that $\#M$-PARTITIONS is $\#P$-complete. $\qquad\square$

We could, in fact, go further and consider the equations (12) for different values for ℓ and s as a system of linear equations in variables $Z_{M'}(G)$ for principal submatrices M' of M. This system may be underdetermined so it might not be possible to solve for all the terms $Z_{M'}(G)$ that appear; however, we do not necessarily need to. We can still deduce $\#P$-completeness for any matrix M for which we can solve the equations for at least one variable $Z_{M'}(G)$ where M' is a hard proper principal submatrix. Similarly, we can still deduce $\#P$-completeness for any matrix M for which we can solve the equations for a linear combination of $Z_{M'}(G)$ and $Z_{M''}(G)$ where M' and M'' are equivalent to $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$.

It turns out that this extension of our technique is not necessary for 4×4 matrices, apart from one exceptional case which we resolve by hand; but this extension would be required to extend the technique to larger matrices.

Theorem 21 allows us to show that our example matrix is hard. Recall that the matrix is

$$
M = \begin{array}{c c} & \begin{array}{c c c c} a & b & c & d \end{array} \\ \begin{array}{c} a \\ b \\ c \\ d \end{array} & \left(\begin{array}{c c c c} 0 & 0 & 1 & * \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ * & 1 & 1 & * \end{array}\right) \end{array}
$$

and consider again the graph $J^{0,0}(k,G)$ for some $k > 4$ and some G. For $S \in \{a,b,d,ab,ad\}$ we find that $S \notin \mathcal{E}(M,0)$ so there are S-surjective M-partitions of Γ^0_k. Thus, we have

S	a	b	d	ab	ad
$Z^S_M(\Gamma^0_k)$	$f_{0,1}$	$f_{0,1}$	$f_{0,1}$	$f_{0,2}$	$f_{0,2}$
$E^0(S)$	abd	ab	ad	ab	ad
$M\vert_{E^0(S)}$	hard	easy	hard	easy	hard

Equation (11) gives

$$Z_M(J^{0,0}(k,G)) = f_{0,1}(k)\, T^{\pi,\tau}_{M,0,1}(G) + f_{0,2}(k)\, T^{\pi,\tau}_{M,0,2}(G)\,,$$

where

$$T^{\pi,\tau}_{M,0,1}(G) = Z_{M|_{abd}}(G) + Z_{M|_{ab}}(G) + Z_{M|_{ad}}(G)$$
$$T^{\pi,\tau}_{M,0,2}(G) = Z_{M|_{ab}}(G) + Z_{M|_{ad}}(G)\,.$$

$T^{\pi,\tau}_{M,0,1}(G)$ contains two terms that are partition functions of hard matrices so is not useful to us but $T^{\pi,\tau}_{M,0,2}(G)$ contains only one ($Z_{M|_{ad}}$, which counts independent sets). Therefore, by Theorem 21, $\#M$-PARTITIONS is #P-complete. Given an oracle for Z_M, we could obtain the value of $T^{\pi,\tau}_{M,0,2}(G)$ by interpolation and, from that, we could compute $Z_{M|_{ad}}$.

6 The computer-assisted dichotomy

So far, we have seen three techniques for determining the computational complexity of the $\#M$-PARTITIONS problem for a given matrix M. If M is pure, $\#M$-PARTITIONS is a graph homomorphism problem, so M is hard if, and only if, it has a 2×2 submatrix containing exactly three $*$s. For impure M, Lemma 12 allows us to identify a class of tractable matrices and the techniques of Section 5 allow us to identify a class of hard matrices. We were unable to prove that the last two cases cover all impure 4×4 matrices, so we wrote a computer program to check all such matrices, as follows.

The number of distinct symmetric 4×4 $\{0, 1, *\}$-matrices is modest: at most $3^{10} = 59,049$. Thus, from a computational point of view it is not necessary to do anything to reduce the search space. However, it turns out that the methods described above are not enough to determine the complexity of $\#M$-PARTITIONS for all symmetric 4×4 matrices. Recall that $M_1 \approx M_2$ if $M_1 \equiv M_2$ or $\overline{M_1} \equiv M_2$ (i.e., M_1 can be transformed into M_2 by permuting D and possibly exchanging 0s and 1s). Since $\#M_1$-PARTITIONS and $\#M_2$-PARTITIONS are computationally equivalent when $M_1 \approx M_2$, it suffices to consider only one matrix from each $\approx$-equivalence class. This minimises the set of matrices that the program fails to resolve.

To do this, we associate each 4×4 symmetric matrix M with the string

$$w(M) = M_{a,a}M_{b,b}M_{c,c}M_{d,d}M_{a,b}M_{b,c}M_{c,d}M_{a,c}M_{b,d}M_{a,d} \in \{0, 1, *\}^{10}\,.$$

The program generates 4×4 matrices in the lexicographic order induced by taking $0 < 1 < *$. For each matrix M, we check whether $w(M') < w(M)$ for any matrix $M' \approx M$. If there is such an M', we have already considered a matrix equivalent to M so we do not need to consider it again.

For each matrix M that survives (i.e., for the lexicographically first member of every $\approx$-equivalence class), we apply the following tests. The correctness of these tests will be explained below.

1. If M is pure (contains no 0's or no 1's)

   (a) If M contains a 2×2 submatrix with exactly three $*$s then $\#M$-PARTITIONS is #P-complete.

(b) Otherwise, $\#M$-PARTITIONS is in FP.

2. Otherwise, if the test of Lemma 12 shows that M has no derectangularising sequence then $\#M$-PARTITIONS is in FP.

3. Otherwise, for each proper principal submatrix M' of M, we can determine whether M' is easy or hard using the characterisations of Hell, Hermann and Nevisi [8] and Dyer and Greenhill [4]. The program now does the following for each $\pi, \tau \in \{0,1\}$, and each $0 \le \ell < s \le |D|$, using the notation of Section 5. It computes the elements of $A_M^{\pi,\tau}(\ell, s)$, up to $\equiv$-equivalence and makes the following conclusions.

   (a) If this set contains exactly one hard proper principal submatrix of M then $\#M$-PARTITIONS is $\#$P-complete.
   (b) If this set contains exactly two hard proper submatrices of M and these are $\left(\begin{smallmatrix} * & * \\ * & 0 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} * & * \\ * & 1 \end{smallmatrix}\right)$ then $\#M$-PARTITIONS is $\#$P-complete.

4. If none of the above tests resolves the complexity of $\#M$-PARTITIONS, output the matrix as having unknown complexity.

The program resolves the complexity of $\#M$-PARTITIONS for all but six $\approx$-equivalence classes of matrices. These six are handled in the next section; all turn out to be hard.

We conclude this section by justifying the correctness of the program. If M is pure then $\#M$-PARTITIONS is equivalent to a homomorphism-counting problem so the correctness of Step 1 follows from the dichotomy theorem of Dyer and Greenhill [4]. Now consider Step 2. If M has no derectangularising sequence then $\#$LIST-M-PARTITIONS is in FP by Theorem 9. Since $\#M$-PARTITIONS is just the special case where every vertex has list D, $\#M$-PARTITIONS is also in FP. Finally, the correctness of Step 3 follows from Theorem 21.

7 The last six matrices

In this section, we despatch the six matrices that our program could not resolve.

7.1 Bipartite problems

Let $G = (U, V, E)$ be a bipartite graph and let its *bipartite complement* be the graph $(U, V, (U \times V) \setminus E)$. Note that the bipartite complement of G depends on the partition (U, V) and not just on the vertices and edges of G. A *bipartite clique* in G is a set $S \subseteq U \cup V$ such that G contains an edge between every vertex of $S \cap U$ and every vertex of $S \cap V$. Note the trivial case that S is a bipartite clique in G if $S \subseteq U$ or $S \subseteq V$.

Counting bipartite cliques in a bipartite graph is $\#$P-complete. This is because a bipartite clique in G is an independent set in G's bipartite complement and counting independent sets in a bipartite graph is $\#$P-complete [11]. The problem of counting bipartite cliques remains $\#$P-complete when the input is restricted to be a connected bipartite graph. To see this, note that counting non-trivial bipartite cliques (with at least one edge) is inter-reducible with the problem of counting all bipartite cliques (since the number of trivial ones is easy to compute). But the number of non-trivial bipartite cliques in a graph is the sum of the numbers of non-trivial bipartite cliques in each component.

Lemma 22. #M-PARTITIONS *is* #P-*complete for*

$$M = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{cccc} a & b & c & d \\ \begin{pmatrix} 0 & 0 & * & * \\ 0 & 0 & 1 & * \\ * & 1 & 0 & 0 \\ * & * & 0 & 0 \end{pmatrix} \end{array}.$$

Proof. The problem of counting bipartite cliques in a connected bipartite graph reduces immediately to counting M-partitions. Consider a connected bipartite graph G with vertex bipartition (U, V). Since G is connected, any M-partition of G either

- assigns vertices in U to parts a and b and assigns vertices in V to parts c and d, or

- assigns vertices in U to parts c and d and assigns vertices in V to parts a and b.

In each case, the vertices in parts b and c form a bipartite clique because $M_{b,c} = 1$ whereas the other relevant entries of M are all stars.

So the M-partitions in each case are in one-to-one correspondence with the bipartite cliques of G. Therefore, $Z_M(G)$ is twice the number of bipartite cliques in G. □

Lemma 23. #M-PARTITIONS *is* #P-*complete for* $M \in \{M_1, M_2, M_3\}$, *where*

$$M_1 = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{cccc} a & b & c & d \\ \begin{pmatrix} 0 & 0 & * & * \\ 0 & 0 & 0 & * \\ * & 0 & 1 & 1 \\ * & * & 1 & 1 \end{pmatrix} \end{array} \quad M_2 = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{cccc} a & b & c & d \\ \begin{pmatrix} 0 & 0 & * & * \\ 0 & 0 & 0 & * \\ * & 0 & 1 & * \\ * & * & * & 1 \end{pmatrix} \end{array} \quad M_3 = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{cccc} a & b & c & d \\ \begin{pmatrix} 0 & * & * & * \\ * & 0 & 0 & * \\ * & 0 & 1 & * \\ * & * & * & 1 \end{pmatrix} \end{array}.$$

Proof. In all three cases $M \in \{M_1, M_2, M_3\}$ we will show how to reduce from the #P-complete problem of counting independent sets in a bipartite graph to counting M-partitions.

Let G be a bipartite graph with vertex bipartition (U, V). For an integer $k > 4$, construct G_k from G by adding a set W of k new vertices and adding all edges between distinct vertices w and v where $w \in W$ and $v \in V \cup W$. Note that G_k is not bipartite because the vertices of W form a complete subgraph.

This complete subgraph is the same as the gadget Γ_k^1 that we have already considered, so it will be useful to apply Theorem 18 to all 3-element sets $S \subseteq D$. The outcomes for $k > 4$ are:

S	abc	abd	acd	bcd
$Z_{M_1}^S(\Gamma_k^1)$	0	0	$f_{1,3}(k)$	0
$Z_{M_2}^S(\Gamma_k^1)$	0	0	$f_{1,3}(k)$	0
$Z_{M_3}^S(\Gamma_k^1)$	0	$f_{2,3}(k)$	$f_{1,3}(k)$	0

By Theorem 18, no set S with $|S| \neq 3$ has $Z_M^S(\Gamma_k^1) = f_{1,3}(k)$. The exact value of $Z_M^S(\Gamma_k^1)$ as a function of k for such a set S will not be important in the following interpolation argument.

Now consider $S \subseteq D$ so that $Z_M^S(\Gamma_k^1) = f_{\ell,|S|}(k)$. Let $Z_M^{W \mapsto S}(G)$ denote the number of M-partitions of G in which every vertex of U is assigned a part in $E^0(S)$ and every vertex of V is assigned a part in $E^1(S)$. This is the number of M-partitions of G that can be combined

with an S-surjective M-partition of the k-clique on W to get a valid M-partition of G_k. We will use interpolation as in the proof of Theorem 21. Suppose $k > 4$. Using the table above and noting the value $f_{1,3}(k)$ in the acd column, we can write

$$Z_M(G_k) = f_{1,3}(k)\, Z_M^{W \mapsto \{a,c,d\}}(G) \ + \sum_{\substack{0 \le \ell < s \le 4, \\ (\ell,s) \ne (1,3)}} \sum_{\substack{S \subseteq D, \\ |S| = s}} \mathbf{1}_{Z_M^S(\Gamma_k^1) = f_{\ell,s}(k)}\, f_{\ell,s}(k)\, Z_M^{W \mapsto S}(G)\,,$$

where $\mathbf{1}_{Z_M^S(\Gamma_k^1) = f_{\ell,s}(k)}$ is the indicator for the event that $Z_M^S(\Gamma_k^1) = f_{\ell,s}(k)$ – we know from Theorem 18 that, if this event does not hold, then $Z_M^S(\Gamma_k^1) = 0$.

As in the proof of Theorem 21, Lemma 20 guarantees that the $f_{\ell,s}(k)$ values are linearly independent. So, by varying k and using an oracle for #M-PARTITIONS to compute the left-hand side, we can compute the coefficient of $f_{1,3}(k)$, which is $Z_M^{W \mapsto \{a,c,d\}}(G)$. So, to finish the proof, we just need to show that computing $Z_M^{W \mapsto \{a,c,d\}}(G)$ is #P-hard.

$Z_M^{W \mapsto \{a,c,d\}}(G)$ is the number of M-partitions of G in which every vertex of U is assigned to a part in $E^0(\{a,c,d\}) = \{a,b\}$ and every vertex of V is assigned to a part in $E^1(\{a,c,d\}) = \{c,d\}$. But note that edges are forbidden between part b and part c so there is a one-to-one correspondence between these partitions of G and the independent sets of G. (Vertices assigned to these parts are in the corresponding independent set.) The result follows, since computing independent sets of a bipartite graph G is #P-hard. □

The proof of the following lemma is similar in spirit but with more details to track.

Lemma 24. #M-PARTITIONS *is* #P-*complete for*

$$M \;=\; \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{pmatrix} \begin{array}{cccc} a & b & c & d \end{array} \\ \begin{array}{cccc} 0 & 0 & * & * \\ 0 & 0 & 1 & * \\ * & 1 & 1 & * \\ * & * & * & 1 \end{array} \end{pmatrix}.$$

Proof. Let G be a bipartite graph with vertex bipartition (U, V). For any integer $k > 4$, construct G_k from G as follows. The vertices of G_k are $U \cup V \cup W \cup \{x_c, x_d\}$, where $|W| = k$ and W, x_c and x_d are new vertices. The edges are as follows (see Figure 1):

- every edge (x, y) for $x \in \{x_c, x_d\}$, $y \in V \cup W$;

- every edge (v, w) for $v \in V$, $w \in W$;

- every edge (v, v') for distinct $v, v' \in V$;

- every edge (x_c, u) for $u \in U$;

- every edge (u, v) where $u \in U$, $v \in V$ and $(u, v) \notin E(G)$.

The subgraph induced on W is an independent set, so it is the same as Γ_k^0. We now apply Theorem 18 to all 2-element sets $S \subseteq D$. The outcomes for $k > 4$ are:

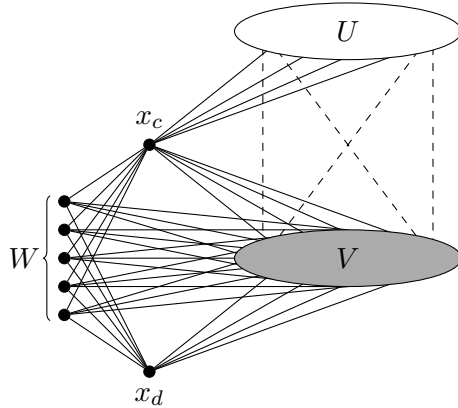| S	ab	ac	ad	bc	bd	cd
$Z_M^S(\Gamma_k^0)$	$f_{0,2}(k)$	$f_{1,2}(k)$	$f_{1,2}(k)$	0	$f_{1,2}(k)$	0

21

Figure 1: The construction of G_k, used in the proof of Lemma 24, shown with $k = 5$. The dotted lines U–V denote the complement of G's edge relation between U and V; the shading of V indicates a clique on those vertices.

By Theorem 18, no set S with $|S| \neq 2$ has $Z_M^S(\Gamma_k^0) = f_{0,2}(k)$ as a function of k.

Now consider $S \subseteq D$ so that $Z_M^S(\Gamma_k^0) = f_{\ell,|S|}(k)$ for some $\ell < |S|$. Let $G_k - W$ denote the subgraph of G_k induced by all vertices other than those in W. Let $Z_M^{W \mapsto S}(G)$ denote the number of M-partitions of $G_k - W$ in which every vertex of U is assigned a part in $E^0(S)$ and every other vertex is assigned a part in $E^1(S)$. As in the proof of Lemma 23, each such M-partition of $G_k - W$ extends to $Z_M^S(\Gamma_k^0)$ M-partitions of G_k, so we can write

$$Z_M(G_k) = f_{0,2}(k)\, Z_M^{W \mapsto \{a,b\}}(G) \;\; + \sum_{\substack{0 \leq \ell < s \leq 4, \\ (\ell,s) \neq (0,2)}} \sum_{\substack{S \subseteq D, \\ |S| = s}} \mathbf{1}_{Z_M^S(\Gamma_k^0) = f_{\ell,s}(k)}\, f_{\ell,s}(k)\, Z_M^{W \mapsto S}(G) \,.$$

As in the proof of Theorem 21, Lemma 20 guarantees that the $f_{\ell,s}(k)$ values are linearly independent. So, by varying k and using an oracle for #M-PARTITIONS to compute the left-hand side, we can compute the coefficient of $f_{0,2}(k)$, which is $Z_M^{W \mapsto \{a,b\}}(G)$. So to finish the proof, we just need to show that computing $Z_M^{W \mapsto \{a,b\}}(G)$ is #P-hard.

$Z_M^{W \mapsto \{a,b\}}(G)$ is the number of M-partitions of $G_k - W$ in which every vertex of U is assigned to a part in $E^0(\{a,b\}) = \{a,b,d\}$ and every vertex in $V \cup \{x_c, x_d\}$ is assigned to a part in $E^1(\{a,b\}) = \{c,d\}$.

Note that the edges between vertices in V add no further restriction on the parts assigned to vertices in V, since $M|_{cd}$ contains no zeroes. Vertices x_c and x_d are not adjacent, so one of them is assigned to part c and the other to part d.

In the first case, x_c is assigned to part c and x_d is assigned to part d. Vertices in U are adjacent to part c and not to part d. Since they are not adjacent to part d, and we already know (from above) that they are not assigned to part c, each must be assigned to part a or b. So we have selected M-partitions in which vertices in U are assigned to parts a or b and vertices in V are assigned to parts c or d. This counts independent sets in G, with parts b and c corresponding to being in the independent set (all edges between these parts must exist in G_k, which corresponds to an independent set in G).

In the second case, x_c is assigned to part d and x_d is in part c. Vertices in U are adjacent to part d and not to part c so they can only be in parts a and d. Since U is an independent set

and $M_{d,d} = 1$, at most one of its vertices is in part d. We can count all such M-partitions in polynomial time by considering each possible vertex $u \in U$ that might be assigned to part d and assigning the rest to part a. The vertex in part d restricts its non-neighbours in V to be assigned part c. The rest of the vertices in V can be assigned to either c or d.

In conclusion, computing $Z_M(\cdot)$ enables us to compute $Z_M^{W \mapsto \{a,b\}}(G)$. But computing $Z_M^{W \mapsto \{a,b\}}(G)$ enables us to count independent sets of G. Since counting independent sets of a bipartite graph is #P-hard, so is counting M-partitions. $\qquad\square$

7.2 A matrix proved hard by solving simultaneous linear equations

Recall the definition of $T_{M,\ell,s}^{\pi,\tau}(G)$ from (12). In Section 5, our gadgets were large cliques and independent sets and we used interpolation on the number of vertices in the gadget to compute $Z_{M'}(G)$ for some submatrix M' such that $\#M'$-PARTITIONS is #P-complete. Our final case is a matrix M where this technique only allows us to compute the linear combinations $T_{M,\ell,s}^{\pi,\tau}(G) = \sum_i \alpha_i Z_{M_i}(G)$ where, although each subproblem $\#M_i$-PARTITIONS is hard, we do not have enough independent linear equations to compute any single term $Z_{M_i}(G)$. The solution is to use a similar gadget to generate an extra linear equation that allows us to solve for a hard $Z_{M'}$.

Lemma 25. $\#M$-PARTITIONS *is* #P-*complete for*

$$
M = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array}
\begin{array}{c} a \quad b \quad c \quad d \\
\begin{pmatrix} 0 & * & * & * \\ * & * & 0 & * \\ * & 0 & * & 1 \\ * & * & 1 & * \end{pmatrix} \end{array}.
$$

Proof. We show how to reduce $\#M|_{abd}$-PARTITIONS to $\#M$-PARTITIONS. The matrix $M|_{abd}$ is hard by Hell, Hermann and Nevisi's characterisation of the hard 3×3 matrices [8]: the principal submatrix $M|_{ad} = \left(\begin{smallmatrix} 0 & * \\ * & * \end{smallmatrix} \right)$ is hard.

First, consider $J^{1,0}(k,G)$ and the set of M-partitions in which the vertices of the Γ_k^0 appear in exactly two parts, corresponding to the term $T_{M,0,2}^{1,0}(G)$. Note that every pair of parts is possible, except for $\{c,d\}$:

S	ab	ac	ad	bc	bd
$E^1(S)$	bd	cd	bcd	ad	abd
$M|_{E^1(S)}$	easy	easy	easy	hard	hard

Thus, noting that $M|_{ab} = M|_{ad} = \left(\begin{smallmatrix} 0 & * \\ * & * \end{smallmatrix} \right)$, there is a polynomial-time-computable function $p(G)$ (corresponding to the "easy" entries of the above table) such that

$$
T_{M,0,2}^{1,0}(G) = p(G) + Z_{M|_{abd}}(G) + Z_{M|_{ad}}(G). \tag{13}
$$

Second, consider $J^{1,0}(k, G+x)$ where $G+x$ denotes the union of G and a new isolated vertex x. We again consider M-partitions of this graph in which vertices of the Γ_k^0 appear in exactly two parts and we divide up these partitions according to the part in which the new vertex x appears. If the vertices of the Γ_k^0 are in parts $S \subset D$ and x is in part i, then the vertices of G must be in some subset of the parts $P(i,S) = E^0(\{i\}) \cap E^1(S)$. The possible combinations are as follows.

i	a		b			c		d				
S	bc	bd	ab	ad	bd	ac	ad	ab	ac	ad	bc	bd
$P(i,S)$	ad	abd	bd	bcd	abd	c	bc	bd	d	bd	ad	abd
$M\|_{P(i,S)}$	hard	hard	easy	easy	hard	easy	easy	easy	easy	easy	hard	hard

This gives a second equation,

$$T_{M,0,2}^{1,0}(G + x) = p'(G) + 3Z_{M|_{abd}}(G) + 2Z_{M|_{ad}}(G)\,, \tag{14}$$

where, again, $p'(G)$ is a polynomial-time computable function.

As in the proof of Theorem 21, we can compute $T_{M,0,2}^{1,0}(G)$ and $T_{M,0,2}^{1,0}(G + x)$ by interpolations on k, using an oracle for $\#M$-PARTITIONS. Thus, we can solve (13) and (14) for $Z_{M|_{abd}}(G)$ (and $Z_{M|_{ad}}(G)$, which is also #P-hard), completing the reduction. □

8   The dichotomy for 4×4 matrices

Finally, we establish Theorem 1 and show that Conjecture 10 holds for 4×4 matrices.

Theorem 26. *Let M be a symmetric matrix in $\{0, 1, *\}^{4 \times 4}$. Then $\#M$-PARTITIONS is #P-complete if M has a derectangularising sequence, and is in* FP*, otherwise.*

Proof. The conjecture is already know to hold for pure matrices (see Section 3).

The impure matrices covered by Lemma 12 have no derectangularising sequence, so are easy by Theorem 9.

For the matrices proved hard via Theorem 21, the computer program finds a hard principal submatrix of size either 2×2 or 3×3. If the 2×2 submatrix $M|_S$ is hard, then S, S is a derectangularising sequence; if the 3×3 submatrix $M|_T$ is hard then, by Lemma 11, $M|_T$ has a derectangularising sequence, and this is also derectangularising for M.

For each of the six matrices proved hard in Section 7, it is easy to check that $\{a, b\}, \{c, d\}$ is a derectangularising sequence. □

9   Acknowledgements

References

[1] A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34:1–34:41, 2013.

[2] A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Inform. Comput.*, 205(5):651–678, 2007.

[3] M. Chudnovsky. *Berge trigraphs and their applications*. PhD thesis, Princeton University, 2003.

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

[5] T. Feder, P. Hell, S. Klein, and R. Motwani. Complexity of graph partition problems. In *Proc. 31st ACM Symposium on Theory of Computing (STOC 1999)*, pages 464–472. ACM, 1999.

[6] T. Feder, P. Hell, S. Klein, and R. Motwani. List partitions. *SIAM J. Discrete Math.*, 16(3):449–478, 2003.

[7] A. Göbel, L. A. Goldberg, C. McQuillan, D. Richerby, and T. Yamakami. Counting list matrix partitions of graphs. In *Proc. 29th Conference on Computational Complexity (CCC 2014)*, pages 56–65. IEEE, 2014. Full version: ArXiv CoRR abs/1306.5176.

[8] P. Hell, M. Hermann, and M. Nevisi. Counting partitions of graphs. In *Proc. 23rd International Symposium on Algorithms and Computation (ISAAC 2012)*, volume 7676 of *LNCS*, pages 227–236. Springer, 2012.

[9] P. Hell and J. Nešetřil. Counting list homomorphisms and graphs with bounded degrees. In J. Nešetřil and P. Winkler, editors, *Graphs, Morphisms and Statistical Physics*, volume 63 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 105–112, 2004.

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

[11] J. S. Provan and M. O. Ball. The complexity of counting cuts and computing the probability that a graph is connected. *SIAM J. Comput.*, 12(4):777–788, 1983.

[12] L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979.