

Matrix Partitions with Finitely Many Obstructions

Tomás Feder*, Pavol Hell† and Wing Xie‡

Submitted: Jan 29, 2007; Accepted: Aug 10, 2007; Published: Aug 20, 2007

Mathematics Subject Classification: 05C75

Abstract

Each m by m symmetric matrix M over $0, 1, *$, defines a partition problem, in which an input graph G is to be partitioned into m parts with adjacencies governed by M , in the sense that two distinct vertices in (possibly equal) parts i and j are adjacent if $M(i, j) = 1$, and nonadjacent if $M(i, j) = 0$. (The entry $*$ implies no restriction.)

We ask which matrix partition problems admit a characterization by a finite set of forbidden induced subgraphs. We prove that matrices containing a certain two by two diagonal submatrix S never have such characterizations. We then develop a recursive technique that allows us (with some extra effort) to verify that matrices without S of size five or less always have a finite forbidden induced subgraph characterization. However, we exhibit a six by six matrix without S which cannot be characterized by finitely many induced subgraphs. We also explore the connection between finite forbidden subgraph characterizations and related questions on the descriptive and computational complexity of matrix partition problems.

1 Introduction

Many graph partition problems (especially those arising from the study of perfect graphs [6, 7, 16]) can be formulated in the following terms. Let M be a symmetric m by m matrix over $\{0, 1, *\}$. An M -partition of a graph G is a partition of $V(G)$ into parts V_1, V_2, \dots, V_m such that for distinct vertices $u \in V_i, v \in V_j$, we have $uv \in E(G)$ if $M(i, j) = 1$, and $uv \notin E(G)$ if $M(i, j) = 0$. Note that we admit $i = j$; in particular, if $M(i, i) = 0$, the set V_i is independent in G , and if $M(i, i) = 1$, it is a clique. Also note that $*$ means no restriction. For each fixed matrix M we obtain the M -partition problem - to decide whether or not an input graph G admits an M -partition. For instance, for the identity

*268 Waverley St., Palo Alto, CA 94301, USA; tomas@theory.stanford.edu

†School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada, V5A 1S6; pavol@cs.sfu.ca

‡School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada, V8W 3P4; wingx@cs.sfu.ca

matrix, the I_m -partition problem asks whether or not G is a union of at most m disjoint cliques with no edges joining them. If C_m is the matrix in which the diagonal entries are 0 and all other entries are *, then the C_m -partition problem asks whether or not the graph G is m -colourable, in the usual sense. Many other examples are described in [16, 25, 27]; they include problems such as H -colourability (also known as the *homomorphism* problem) [26], the existence of a clique cutset or a skew cutset [7, 29], or being a split graph [23]. The homomorphism language is particularly appropriate here, and the M -partition problem can be cast as a homomorphism problem to a suitable structure H , called a *trigraph* [19, 25, 27, 31].

We note in passing that there are many variants of the basic M -partition problem, including partitioning digraphs (M is not necessarily symmetric) [19], equipping the vertices of G with lists (of parts in which the vertex is allowed to be placed) [4, 16, 12, 13, 14, 22, 25, 27], requiring all parts to be nonempty [9, 30], generalizing to certain constraint satisfaction problems [10, 25], or restricting the input graphs to have special structure [11, 17, 20, 24, 25, 27].

The I_m -partition problem is clearly solvable in polynomial time - it suffices to compute the connected components of G and check whether each is a clique, and whether there are at most m of them. On the other hand, for $m > 2$, the C_m -partition problem is well known to be NP-complete. There are many other matrices M for which the M -partition problem has been shown to be polynomial or NP-complete [4, 16, 18, 22, 26, 27], including for instance all matrices of size $m < 5$. However, in general, we do not know how to tell, for a given matrix M , what is the complexity of the M -partition problem. In fact, we do not know whether each M -partition problem is polynomial or NP-complete [17]. If this were the case, it would imply the validity of the so-called dichotomy conjecture of Feder and Vardi [21, 17]. Moreover, we cannot decide the complexity of some concrete M -partition problems with small matrices M ; for instance if we consider M -partitions with lists, one such problem with $m = 4$ is described in [4]; see also [10, 25].

Each M -partition problem can be described in monadic second-order logic [8] (and hence solved efficiently on graphs on bounded treewidth or cliquewidth). However, certain M -partition problems can be described in first-order logic (and hence solved efficiently on all graphs). For instance, it is easy to see that a graph G admits an I_m -partition if and only if it does not contain an induced path with three vertices, P_3 , or an induced union of $m + 1$ isolated vertices, $(m + 1)K_1$. Since m is fixed, having such an induced subgraph can be described by a first-order sentence, and hence so can being I_m -partitionable.

Thus there are three basic questions we may ask, for a given matrix M :

- Can M -partitionable graphs be recognized in polynomial time?
- Can M -partitionable graphs be described by a first-order sentence?
- Can M -partitionable graphs be characterized by a finite set of forbidden induced subgraphs?

Note that the questions have been ordered so that a positive answer to a later question implies a positive answer to an earlier question. As we have argued above, the first

question seems hard to answer in full generality. If M has no 1's, it is known that the last two questions have the same answer (even in the more general context of constraint satisfaction problems) [1, 28]. In this note we focus on the last question.

The *complement* \overline{M} of a matrix M has the entries 0 and 1 interchanged, i.e., $\overline{M}(i, j) = 1 - M(i, j)$, where $1 - *$ is defined to be $*$. Clearly a graph G admits an M -partition if and only if its complement \overline{G} admits an \overline{M} -partition.

In the remainder of the note we shall assume that all $M(i, i) \neq *$. Indeed, if some $M(i, i) = *$, then every graph G would be M -partitionable, and all three questions would be trivial. By simultaneously permuting the rows and columns of M , we may assume that $M(1, 1) = \dots = M(k, k) = 0$, $M(k + 1, k + 1) = \dots = M(m, m) = 1$. Let A denote the submatrix of M with rows $1, \dots, k$ and columns $1, \dots, k$; let B denote the submatrix with rows $k + 1, \dots, m$ and columns $k + 1, \dots, m$; and let C denote the submatrix with rows $1, \dots, k$ and columns $k + 1, \dots, m$.

We define two vertices u, v in a graph G to be *similar*, if they have exactly the same neighbours other than u and v . Note that similar vertices may be adjacent or non-adjacent; in the former case they will be called *c-similar*, and in the latter case *i-similar*. (Other terms used in the literature are joined duplicates, joined twins and true twins for c-similar vertices, and unjoined duplicates, unjoined twins and false twins for i-similar vertices.) Clearly, both i-similarity and c-similarity are equivalence relations on $V(G)$. Moreover, it is easy to check that similarity itself is also an equivalence relation. Each similarity class is an independent set or a clique, while i-similarity classes are just independent sets, and c-similarity classes are just cliques.

2 Friendly and Unfriendly Matrices

We say that a matrix M is *friendly* if it has no $*$ entries in A and in B . A matrix is *unfriendly* if it is not friendly. Note that a matrix is unfriendly if and only if it contains a diagonal two by two submatrix S with $S(1, 1) = S(2, 2) \neq *$ and $S(1, 2) = S(2, 1) = *$.

We shall prove that if M is an unfriendly matrix, then M -partitionable graphs cannot be characterized by a finite set of forbidden induced subgraphs. We shall cast our discussion in the following terms.

A *minimal obstruction* to M -partition is a graph G which is not M -partitionable, and such that for each vertex $v \in V(G)$ the graph $G - v$ is M -partitionable. To prove the above claim for an unfriendly matrix M , we shall exhibit infinitely many non-isomorphic minimal obstructions to M -partition. For instance, the matrix C_2 is unfriendly, and there are infinitely many minimal obstructions to C_2 -partition, namely all odd cycles. (C_2 -partitionability is the same as two-colourability.) It is clear that this means that it is not possible to characterize the property by finitely many forbidden subgraphs. Note that C_2 -partitionable graphs can nevertheless be recognized in polynomial time.

Theorem 2.1 *If M is an unfriendly matrix, then there are infinitely many minimal obstructions for M -partitionability.*

Proof. Let M be an unfriendly matrix of size m , with k defined as above to be the number of 0's on the main diagonal. By taking the complement if necessary, we may assume without loss of generality that M contains the diagonal submatrix $S \begin{pmatrix} 0 & * \\ * & 0 \end{pmatrix}$.

We first consider the case that $m = k$, i.e., M has only zeros on the main diagonal. This means, in particular, that each M -partitionable graph is k -colourable. If $k = 2$, i.e., if $M = S$, then a graph is M -partitionable if and only if it is bipartite; hence each odd cycle is a minimal obstruction. In general, we shall appeal to the well-known fact [27] that there exist graphs with arbitrarily high chromatic number and odd girth. We let G_0 be any graph with chromatic number greater than k ; as noted above G_0 is not M -partitionable, and hence contains some minimal M -obstruction G'_0 . Suppose G'_0 has odd circumference c_0 . (The *odd circumference* of a graph is the maximum length of an odd cycle; since G'_0 is not M -partitionable it must contain an odd cycle.) We proceed recursively, assuming that we have already constructed a minimal M -obstruction G'_i with odd circumference c_i . Let G_{i+1} be a graph with chromatic number greater than k and odd girth greater than c_i . It is again the case that G_{i+1} is not M -partitionable, and hence contains a minimal M -obstruction G'_{i+1} ; of course G'_{i+1} must again contain an odd cycle, and hence an odd cycle of length greater than any G'_j with $j < i$. Therefore we have infinitely many minimal M -obstructions $G'_i, i = 0, 1, \dots$; moreover, our graphs G'_i have the additional property that the odd girth of each G'_i is greater than the odd circumference of any $G'_j, j < i$. By complementation, this proof covers also the case when $k = 0$, i.e., when M has only ones on the main diagonal.

If $m > k > 0$, we proceed recursively, letting M' be obtained from M by deleting the m -th row and column, and assuming we have already constructed infinitely many minimal M' -obstructions $G'_i, i = 0, 1, \dots$ with the property that the odd girth of G'_i is greater than the odd circumference of any $G'_j, j < i$. We note that the disjoint union of two copies of any G'_i cannot be M -partitionable: at most one copy can use the m -th part (the set V_m from the definition of M -partition), since $M(m, m) = 1$ implies that all vertices in this part are adjacent to each other. Thus the disjoint union of two copies of G'_i contains a minimal M -obstruction G''_i . Now we observe that the odd girth of G''_i is at least the odd girth of G'_i , and the odd circumference of G''_i is at most the odd circumference of G'_i ; thus the additional property is maintained, and the graphs G''_i are not isomorphic. \square

We have several classes of friendly matrices M for which M -partitionable graphs are known to have a characterization by finitely many forbidden subgraphs. The simplest case occurs when M has no $*$ entries at all. In [15] we have shown the following fact.

Theorem 2.2 [15] *If M has no $*$ entries, then a minimal M -obstruction has at most $(k+1)(m-k+1)$ vertices; moreover, there are at most two minimal M -obstructions with precisely $(k+1)(m-k+1)$ vertices.*

The proof of the above theorem is quite involved [15]. However, it is not difficult to prove directly that the number of minimal M -obstructions is finite [31]. (This fact also follows from a more general result for constraint satisfaction problems [2].)

Proposition 2.1 *If M has no $*$ entries, then a minimal M -obstruction has at most $m(2k' + 2) + 1$ vertices, where $k' = \max(k, m - k)$.*

Proof. Suppose a minimal M -obstruction G has at least $m(2k' + 2) + 2$ vertices. Then for any vertex v , the graph $G - v$ admits an M -partition; since $G - v$ has at least $m(2k' + 2) + 1$ vertices, some set S of at least $2k' + 3$ vertices belongs to the same part - either a clique or an independent set. In either case, all these vertices are similar to each other in $G - v$.

Suppose first S is an independent set: then any two vertices of S have exactly the same neighbours in $G - v$, and hence there is a subset T of at least $k' + 2$ independent vertices that have exactly the same neighbours in G . Let $t \in T$: as before, $G - t$ has an M -partition. Note that exactly $m - k \leq k'$ of the parts are cliques; thus of the $k' + 1$ vertices of $T - t$, at least one must be placed into a part that is an independent set; clearly, t can be placed into the same part, as it has the same neighbours.

If S is a clique, then the argument is analogous: any two vertices of S have the same neighbours in $\overline{G} - v$, hence a subset T of at least $k' + 2$ vertices have the same neighbours in \overline{G} . In any M -partition of $G - t$, $t \in T$, some vertex of $T - t$ must be placed into a part that is a clique, and t can be placed into the same part. \square

We can extend the validity of this result by applying the so-called sparse-dense technique from [11, 16]. We formulate it here specifically for the application at hand; its derivation from the general sparse-dense technique is easy to see. (In the notation of [11], we set \mathcal{S} to be all A -partitionable graphs and \mathcal{D} to be all B -partitionable graphs.)

We claim that for any fixed M there exists an integer r such that any graph G that is both A -partitionable and B -partitionable has at most r vertices. Indeed, such a graph G is, in particular, k -colourable and hence contains no clique of size $k + 1$, and its complement \overline{G} is $(m - k)$ -colourable and hence contains no clique of size $m - k + 1$. Thus the existence of such an integer r follows from Ramsey's theorem. This integer r will be used in all applications of the sparse-dense technique - both in Theorem 2.3 below, and also in Theorem 5.2 later on.

A *labeled* graph G is a graph in which each vertex has a label, either A or B . Label A means that the vertex has to be placed into parts V_1, \dots, V_k , label B means the vertex has to be placed into parts V_{k+1}, \dots, V_m . (Thus the labels are lists of a restricted kind.) A *labeled M -partition* of G is an M -partition of G that satisfies these constraints. A *minimal labeled M -obstruction* is a labeled graph G which has no labeled M -partition, such that for each vertex $v \in V(G)$ the graph $G - v$ (with the inherited labels) has a labeled M -partition.

Note that we have used A and B before, to denote the two diagonal submatrices of M . Since the label A (respectively B) actually restricts the vertex to be placed into a part from the submatrix A (respectively B), this ambiguity will not cause a problem.

For the number r derived above, Theorem 3.1 from [11] (proved using the sparse-dense technique) allows us to make the following conclusion.

Theorem 2.3 [11] *Suppose each minimal labeled M -obstruction has at most p vertices. Then each minimal M -obstruction to M -partition has at most $2p^{2^{r+1}}$ vertices. \square*

Now suppose M is a matrix in which C has only $*$ entries. In this case a labeled graph G is M -partitionable if and only if its subgraph on the vertices labeled A is A -partitionable and its subgraph on the vertices labeled B is B -partitionable. Since the submatrices A and B of a friendly matrix have no $*$ entries, Proposition 2.1 (or Theorem 2.2) implies that for friendly matrices M , there are only finitely many minimal labeled M -obstructions. Thus we conclude from Theorem 2.3 that there are only finitely many minimal M -obstructions.

Corollary 2.4 *If M is a friendly matrix in which the submatrix C has only $*$ entries, then M -partitionable graphs can be characterized by a finite set of forbidden induced subgraphs.*

□

Note that the same conclusion applies when C has no $*$ entries, by Proposition 2.1. We single out a few example consequences of Corollary 2.4.

The following classes of graphs can be characterized by a finite set of forbidden induced subgraphs.

- Graphs partitionable into an independent set and a clique. These are known as *split graphs*, and the (exactly three) forbidden induced subgraphs are known [23].
- Graphs partitionable into a complete k -partite graph and a graph which is the complement of a complete ℓ -partite graph (i.e., is a union of ℓ disjoint cliques with no other edges). These are known as *polar graphs* [5]; we have parametrized them by the number of independent sets and cliques. Recognizing polar graphs without fixing these numbers is NP-complete [5].
- Graphs partitionable into one clique and one graph which is the disjoint union of k complete bipartite graphs with no other edges. These are a variant of k -bisplit graphs [3, 18].

Theorem 2.3 shows that if there are only finitely many minimal labeled M -obstructions, then there are only finitely many minimal M -obstructions. This statement has an easy converse.

Proposition 2.2 *If there are only finitely many minimal M -obstructions, then there are also only finitely many labeled minimal M -obstructions.*

Proof. We shall prove, specifically, that if every minimal M -obstruction has at most p vertices, then every minimal labeled M -obstruction has also at most p vertices. Otherwise, some minimal labeled M -obstruction G has more than p vertices. Let G' be obtained from G by replacing each vertex labeled A by an independent set of $m - k + 1$ vertices, and each vertex labeled B by a clique of $k + 1$ vertices. It is easy to see that the unlabeled graph G' does not admit an M -partition, as at least one vertex from the independent set replacing a vertex of G labeled A must be placed into a part that is an independent set

(there are only $m - k$ parts that are cliques), and at least one vertex from each clique replacing a vertex of G labeled B must be placed into a part that is a clique. (Hence if G' admitted an M -partition, then so would the labeled graph G .) Therefore G' contains a minimal M -obstruction G'' with at most p vertices. Let G^* be the subgraph of G obtained by taking all the (at most p) vertices of G which correspond to vertices in G'' . Clearly, if the labeled graph G^* admitted an M -partition, then so would the graph G'' ; therefore G^* is a proper subgraph of G which does not admit an M -partition, contradicting the fact that G is a minimal M -obstruction. \square

3 A Recursive Technique

For a matrix M , we denote by $M(i)$ the submatrix obtained from M by deleting the i -th row and i -th column. Note that if M is a friendly matrix, then so is each matrix $M(i), i = 1, 2, \dots, m$.

Our main technique is the following recursive method.

Theorem 3.1 *Suppose M is a friendly matrix such that all rows of its submatrix A are distinct, or such that all rows of its submatrix B are distinct.*

If there are only finitely many minimal labeled $M(i)$ -obstructions for each $i = 1, \dots, m$, then there are also only finitely many minimal labeled M -obstructions.

Proof. Without loss of generality, we shall focus on the case when the rows of the submatrix A are distinct. Some rows of the submatrix B may be the same - we partition the $m - k$ rows into $n_1 \leq m - k$ different groups of equal rows. For symmetry, we also imagine A partitioned into $n_0 = k$ groups of equal rows (in this case each group has just one row).

Suppose all minimal labeled $M(i)$ -obstructions ($i = 1, 2, \dots, m$) have at most p vertices, and consider a minimal labeled M -obstruction G . Let G_A be the subgraph of G induced by the vertices labeled A , and let G_B the subgraph induced by vertices labeled B . Let n_A denote the number of i -similarity classes in G_A , and let n_B denote the number of c -similarity classes in G_B .

We first prove that G cannot be too big if it has $n_A > n_0$. Indeed, in this case, consider the labeled subgraph G' of G induced by taking just one vertex each from $n_0 + 1$ i -similarity classes of G_A (with the inherited labels A). We claim that the labeled graph G' is itself not M -partitionable, since it has $n_0 + 1$ non- i -similar vertices labeled A : placing such vertices would require $n_0 + 1$ parts corresponding to distinct rows of the matrix A . Since G is a minimal labeled M -obstruction, we must have $G' = G$, i.e., G has at most $n_0 + 1$ vertices. A similar argument shows that G cannot be too big if it has $n_B > n_1$; in that case G has at most $n_1 + 1$ vertices.

Next we focus on the case when $n_A < n_0$. In this situation, G_A cannot admit an A -partition in which each of the n_0 groups of parts contains a nonempty part. Otherwise, two i -similar vertices v, v' of G_A would be placed into two parts corresponding to different rows i, i' of A : if the two rows differ in column j , then the vertex of G_A placed into the

j -th part (or a part corresponding to the same group) would have different adjacency to the two supposedly i -similar vertices v, v' . This contradiction shows that G must contain a minimal labeled $M(i)$ -obstruction for every $i = 1, 2, \dots, k$. Moreover, if G' is a labeled subgraph of G containing a minimal labeled $M(i)$ -obstruction for every $i = 1, 2, \dots, k$, then G' already cannot be M -partitionable, by the same argument; since G is a minimal labeled M -obstruction, we would again have to have $G = G'$. It now follows that G can have at most kp vertices. A similar argument applies in the case when $n_B < n_1$; in this case G can have at most $(m - k)p$ vertices.

It remains to consider minimal labeled M -obstructions G that have $n_A = n_0$ and $n_B = n_1$. According to the previous arguments, any M -partition of G must place all vertices of an i -similarity class C of G_A into one part P corresponding to a row of A , and place all vertices of a c -similarity class D of G_B to a set of parts corresponding to a group Q of equal rows into B . Let W be the set of such assignments, in which the placement of the vertices of G_A is an A -partition and the placement of the vertices of G_B corresponds to a B -partition (by selecting, for each $d \in D$ one of the parts in the group to which it was assigned). Clearly, W must be nonempty: otherwise, the labeled graph obtained from G_A by selecting one vertex in each i -similarity class has no M -partition, or the labeled graph obtained from G_B by selecting one vertex from each c -similarity class has no M -partition, implying that the minimal labeled M -obstruction G is actually equal to one of these two labeled graphs, and hence has at most n_0 , or n_1 , vertices. On the other hand, W has at most $n_0!n_1!$ assignments. We now consider how many vertices of G are necessary in order to ensure that none of the assignments $w \in W$ arises from an actual M -partition of G . Suppose that $w \in W$ assigns each i -similarity class C_x of G_A to the part corresponding to a row a_x in A , and each c -similarity class D_y of G_B to a set of parts corresponding to a group Q_y of equal rows of B . Note that the assignment w completely determines the placements of all vertices of G_A . If w does not arise from an actual M -partition, then it must be impossible to place the vertices of some c -similarity class D_y . In other words, for some vertex $d \in D_y$, no row of Q_y has the right entries in M in the columns corresponding to A . Since Q_y has at most $m - k$ rows, and for each of them we need only one vertex in G_A placed by w into a part corresponding to a row in A with the wrong entry, the assignment w is disqualified on the basis of the vertex d plus at most $m - k$ vertices of G_A . Let the labeled graph G' be the induced subgraph of G (with the inherited labels) on the set of vertices used for disqualifying any assignment $w \in W$. Then G' has at most $n_0!n_1!(1 + m - k)$ vertices, and does not admit an M -partition. Since G is a minimal labeled M -obstruction, we must have $G = G'$, and hence G has at most $n_0!n_1!(1 + m - k)$ vertices. Since in every case G has a number of vertices bounded by a function of the fixed matrix M , the number of minimal labeled M -obstructions is finite. \square

Proposition 2.2 allows us to state Theorem 3.1 in a simpler form.

Corollary 3.2 *Suppose M is a friendly matrix such that all rows of its submatrix A are distinct, or such that all rows of its submatrix B are distinct.*

If there are only finitely many minimal $M(i)$ -obstructions for each $i = 1, \dots, m$, then there are also only finitely many minimal M -obstructions. \square

4 Small Friendly Matrices

In this section, a matrix M is *small* if it has size $m, m \leq 5$. We prove that all small friendly matrices M have finitely many minimal M -obstructions, and hence admit a characterization of M -partitionability by a finite set of forbidden induced subgraphs.

Theorem 4.1 *If M is a small friendly matrix then the number of minimal M -obstructions is finite.*

Proof. We first consider the case that $n_0 = 1$ and $n_1 = 1$. This means that A is the all-zero matrix and B is the all-one matrix. In this case, if the matrix M contains a $*$, then a graph is M -partitionable if and only if it is a *split graph*, i.e., can be partitioned into an independent set and a clique. Such graphs have a finite forbidden subgraph characterization by [23] (see also Corollary 2.4 and its consequences). Otherwise, the matrix M has no $*$, and then M -partitionability has a finite forbidden subgraph characterization by Proposition 2.1 (or Theorem 2.2).

Next consider the case that $k = 0$ or $k = m$: then the fact that M is friendly implies that M has no $*$, and hence only finitely many minimal M -obstructions by Proposition 2.1 (or Theorem 2.2).

Thus we may now assume that $1 \leq k \leq m - 1$ and $3 \leq n_0 + n_1 \leq m$. A matrix M with $m = 3$ must have either $k = 1$ or $m - k = 1$, and Theorem 3.1 and Theorem 2.3 imply there are only finitely many minimal M -obstructions. The same argument applies to a matrix M with $m = 4$, since in the case $k = 2$ we can use the fact that $n_0 + n_1 \geq 3$.

In case $m = 5$, we can again make the same argument, unless we have (up to complementation) $k = 2$ and $n_0 = 1, n_1 = 2$. This means that $A = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$ and

$$B = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

In this case we make a separate argument, akin to the proof of Theorem 3.1. Note that the matrix A has rows 1 and 2 equal, forming one group of rows, while matrix B has two groups - the equal rows 3, 4, and the separate group containing just row 5. As in the proof of Theorem 3.1, we may assume that the minimal labeled M -obstruction G has one i -similarity class C in G_A , and two c -similarity classes D, D' in G_B . There are just two possible assignments of the classes of G to the groups of M since C must be assigned to rows 1, 2, and either D to rows 3, 4 and D' to row 5, or conversely. We shall estimate how many vertices of G are necessary to prevent one of these assignments from being a labeled M -partition. Without loss of generality, consider D being assigned to parts (rows) 3, 4 and D' to part (row) 5.

Suppose first that $M(1, 5) = M(2, 5) = 0$. If there is a vertex $v \in C$ and a vertex $w \in D'$ which are adjacent, then these two vertices already prevent the assignment, so only two vertices of G are needed for it. Otherwise, all vertices of D' can be placed into part 5, and so the only vertices of G that can prevent the assignment are vertices ensuring that $G - D'$ does not admit a labeled M_5 -partition. (Recall that M_5 is obtained from M by eliminating row and column 5.) This means that $G - D'$ is a minimal labeled M_5 -obstruction, and hence has a bounded number of vertices; these are the only vertices G needs to prevent the assignment. If $M(1, 5) = M(2, 5) = 1$, the argument is analogous.

If one of $M(1, 5), M(2, 5)$ is 0 and the other 1, say if $M(1, 5) = 0, M(2, 5) = 1$, we argue as follows. If some vertex $x \in C$ is adjacent to some vertex of D' and nonadjacent to another, then these three vertices are all that G needs to prevent the assignment. Otherwise, the vertices of C are partitioned into C_1 , the set of those $x \in C$ that are nonadjacent to all vertices of D' , and C_2 , the set of those $x \in C$ that are adjacent to each vertex of D' . We have now effectively assigned C_1 to 1, C_2 to 2, D to 3, 4, and D' to 5, and can proceed as in the proof of Theorem 3.1. Specifically, if each $v \in D$ can be placed into 3 or in 4, we would have an M -partition of G ; thus there must exist some vertex $v \in D$ which cannot be placed into 3 or in 4. The former is ensured by a vertex u in C having the wrong kind of connection to v (there are four ways to have such a wrong connection, for instance $u \in C_1$ adjacent to v while $M(1, 3) = 0$), and the latter is similarly ensured by another vertex $w \in C$ having the wrong kind of connection to v . It is now clear that G only needs u, v, w to prevent the assignment.

If $M(1, 5) = M(2, 5) = *$, then $G - D'$ must be a minimal labeled M_5 -obstruction as above, and hence has a bounded number of vertices - and G only needs the vertices of $G - D'$ to prevent the assignment. If only one of $M(1, 5), M(2, 5)$ is $*$, we argue as follows. Without loss of generality, assume $M(1, 5) = *, M(2, 5) = 1$. Note that C is an independent set, and D, D' are cliques with no edges joining them; the assumption $M(2, 5) = 1$ implies that any vertex of C that is nonadjacent to a vertex of D' must be placed into part 1. If such a vertex does not exist, then C has all possible edges to D' , and hence G only needs the vertices of the minimal labeled M_5 -obstruction $G - D'$ to prevent the assignment, whence G has a bounded number of vertices. Otherwise, let C_1 be the set of all vertices of C that have a nonneighbour in D' . If $M(1, 3) = *$ or $M(1, 4) = *$, then G would admit an M -partition - placing all vertices of C to part 1, all vertices of D to part 3 (respectively 4), and all vertices of D' to part 5.

It remains to consider the following cases.

1. $M(1, 3) = M(1, 4) = 0$
2. $M(1, 3) = M(1, 4) = 1$
3. $M(1, 3) = 1, M(1, 4) = 0$
4. $M(1, 3) = 0, M(1, 4) = 1$

Consider case 1. If some $u \in C_1$ is adjacent to some $v \in D$, then u, v , and any non-neighbour $w \in D'$ of u , already prevent the assignment. Otherwise, $G - C_1$ must

not admit a labeled M -partition according to this assignment; since $C_1 \neq \emptyset$ and G is a minimal labeled M -obstruction, we have a contradiction. Case 2 is analogous.

In case 3, any labeled M -partition of G must partition the vertices of D into those adjacent to all vertices of C_1 , forming a set D_1 (and being placed into part 3), and those not adjacent to any vertices of C_1 , forming a set D_2 (and being placed into part 4). Thus if there is a vertex $v \in D$ adjacent to some $v \in C_1$ and nonadjacent to some $v' \in C_1$, then this assignment of vertices, together with $w \in D'$ nonadjacent to v , and $w' \in D'$ nonadjacent to v' , is impossible. We have now placed all vertices of D' (into part 5), all vertices of D (those in D_1 to part 3 and those in D_2 to part 4), and all vertices of C_1 (into part 1). The remaining vertices, in $C - C_1$, should be placed into part 3 or part 4, but a placement is not possible. Thus some vertex $u \in C - C_1$ cannot be placed into part 1 because of its connection to some $v \in D$ (this means u is adjacent to v and $v \in D_2$ or u is nonadjacent to v and $v \in D_1$), and cannot be placed into part 2 because of some $v' \in D$ (the reasons for this depend on $M(2, 3)$ and $M(2, 4)$, but in any case they concern just one vertex v'). Then u, v, v' , and any one vertex of C_1 form a labeled M -obstruction. Case 4 is again analogous.

Since a bounded number of vertices is needed to prevent this assignment, and there is just one more analogous assignment, the minimal labeled M -obstruction G has a bounded number of vertices. \square

5 Larger Friendly Matrices

It may seem that all friendly matrices admit characterizations by finite sets of forbidden induced subgraphs. However, for $m > 5$ we have a counterexample.

Let M_6 be the matrix

$$\begin{pmatrix} 0 & 1 & 0 & * & 0 & 0 \\ 1 & 0 & 1 & * & 0 & 0 \\ 0 & 1 & 0 & 0 & * & * \\ * & * & 0 & 1 & 0 & 1 \\ 0 & 0 & * & 0 & 1 & 0 \\ 0 & 0 & * & 1 & 0 & 1 \end{pmatrix}$$

Theorem 5.1 *We have the following facts.*

- *The matrix M_6 is friendly.*
- *The M_6 -partition problem cannot be characterized by a finite set of forbidden induced subgraphs.*

Proof. The first statement is obvious. The second statement will be proved by exhibiting infinitely many non-isomorphic minimal labeled M_6 -obstructions G_n , $n = 3, 4, \dots$. These

are described in Figure 1: there is a path of $2n$ vertices v_1, v_2, \dots, v_{2n} , alternately labeled A, B . The vertex v_1 is moreover adjacent to all v_i with odd $i > 1$, and all vertices v_j with even $j < 2n$ are adjacent to each other. Then Proposition 2.2 implies that there are also infinitely many minimal M_6 -obstructions, and hence the M_6 -partition problem cannot be characterized by a finite set of forbidden induced subgraphs. (With a bit more effort, it can be shown that the graphs G_n in the figure are themselves minimal M_6 -obstructions even without the labels [31].)

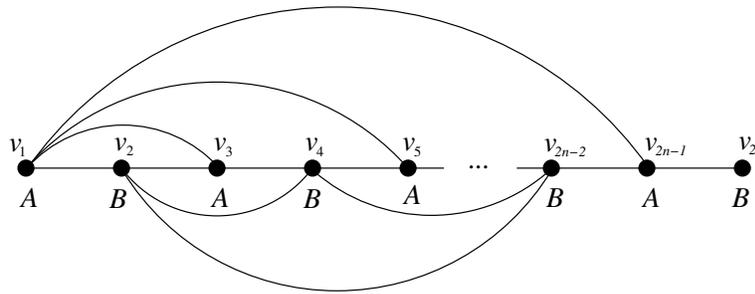


Figure 1: The minimal labeled obstruction G_n

An M_6 -partition of a graph G has six parts, corresponding to the six rows/columns of M_6 . Since rows 1 and 3 coincide in A , the subgraph G_A induced by A has two i -similarity classes, the first one, A_1 , consisting of parts 1 and 3, and the second one, A_2 , consisting of part 2. Since rows 4 and 6 coincide in B , we also have two c -similarity classes in the subgraph G_B induced by B , the first one, B_1 , consisting of parts 4 and 6 and the second one, B_2 , consisting of part 5. Amongst the A -labeled vertices of the graph G_n , the vertex v_1 is adjacent to all the others, which are not adjacent to each other. This means that

- either v_1 is placed into A_1 and all the others into A_2 , or
- v_1 is placed into A_2 and all the others into A_1 .

Similarly, amongst two B -labeled vertices of G_n , the vertex v_{2n} is non-adjacent to all the others, and the others are all adjacent to each other. Thus

- either v_{2n} is placed into B_1 and all the others in B_2 , or
- v_{2n} is placed into B_2 and all the others in B_1 .

Note that $M(2, 5) = 0$: this means that of the four possibilities implicit in the above choices, v_1 must be placed into A_2 (i.e., in part 2), and v_{2n} must be placed into B_2 (i.e., in part 5). (Otherwise some adjacent vertices v_i, v_j of G_n would be placed into parts 2, 5 respectively, which is impossible.) On the other hand, if v_1 is in part 2, then vertex v_2 must be placed into part 4, vertex v_3 into part 1, vertex v_4 into part 4, and so on, implying that v_{2n-1} is in part 1, contrary to the fact that its neighbour v_{2n} is in part 5 (and $M(1, 5) = 0$). Thus the labeled graphs G_n do not admit a labeled M_6 -partition.

Consider now the labeled graph $G_n - v_j, j = 1, 2, \dots, 2n$. This graph does admit a labeled M_6 -partition as follows: vertices v_1, v_2, \dots, v_{j-1} can be placed into parts 2, 4, 1, 4, 1, \dots respectively, while vertices $v_{2n}, v_{2n-1}, v_{2n-2}, \dots, v_{j+1}$ can be placed into parts 5, 3, 6, 3, 6 \dots respectively. Thus each graph G_n is a minimal labeled M_6 -obstruction. \square

It turns out that the M_6 -partition problem can be solved in polynomial time. In fact, we have the following observation.

Theorem 5.2 *If M is a friendly matrix in which neither A nor B has three identical rows, then the M -partition problem is polynomial time solvable.*

Proof. First we choose which parts of A and B will actually have vertices of the input graph G placed into them. There are only finitely many possible choices, and we shall test each. Clearly, it suffices to explain how to test the case when all parts of A and all parts of B are nonempty. We next apply the sparse-dense method from [16], using the same integer r derived before Theorem 2.3. Theorem 3.1 from [16] implies that a graph G on n vertices has at most n^2r possible partitions into induced subgraphs G_A, G_B , such that G_A is A -partitionable and G_B is B -partitionable. Moreover, these partitions can be found in polynomial time, since A and B are matrices without $*$ (see [16]). We shall test each such partition separately, viewing all vertices of G_A as labeled by A , and all vertices of G_B as labeled by B . Thus we may assume that the graph G is labeled. We now find the i -similarity classes of G_A and the c -similarity classes of G_B . Consider the groups of identical rows in A and in B . Since we are assuming that no parts are empty, the number of groups of rows in A must equal the number of i -similarity classes of G_A , and the number of groups of rows in B must equal the number of c -similarity classes of G_B . We shall consider separately each of the bijective assignments between the i -similarity classes of G_A and c -similarity classes of G_B to the groups of identical rows of A and groups of identical rows of B . For each such assignment, we shall decide whether a corresponding M -partition of G is possible. Next we use the fact that A and B don't have three identical rows. If v is a vertex assigned to a group of rows (parts), the only decision left is to choose one of those rows. Since there are at most two possibilities, we can associate to each vertex v a boolean variable x_v . It is clear that the resulting constraints (between the vertices of G_A and the vertices of G_B) can be expressed by clauses with at most two literals each. Thus this problem can be solved in polynomial time using a 2-satisfiability algorithm. \square

There are, however, friendly matrices (in which A or B have three identical rows), such that the M -partition problem is NP-complete.

Let M_9 be the matrix

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & * & 0 & * \\ 0 & 0 & 0 & 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & * & * \\ 0 & 0 & 0 & 0 & 1 & 1 & * & 0 & * \\ 0 & 0 & 0 & 1 & 0 & 1 & * & * & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & * & * \\ * & * & 0 & * & * & 0 & 1 & 1 & 1 \\ 0 & * & * & 0 & * & * & 1 & 1 & 1 \\ * & 0 & * & * & 0 & * & 1 & 1 & 1 \end{pmatrix}$$

Theorem 5.3 *We have the following facts.*

- *The matrix M_9 is friendly.*
- *The M_9 -partition problem is NP-complete.*

Proof. The first statement is obvious. To prove the NP-completeness, we observe that $M(1, 7) = M(7, 2) = M(2, 8) = M(8, 3) = M(3, 9) = M(9, 1) = *$ while all other entries in rows 1, 2, 3 and columns 7, 8, 9 are 0; thus parts 1, 7, 2, 8, 3, 9 are connected in a hexagonal fashion (and similarly for parts 4, 7, 5, 8, 6, 9).

The following problem has been proved NP-complete in [13]. Let H be the hexagon a, b, c, d, e, f (with edges ab, bc, cd, de, ef, fa). Given a bipartite graph G containing H as an induced subgraph, can the vertices of G be mapped to the vertices of H by an edge-preserving mapping (a homomorphism) in which each vertex of H is mapped to itself (a retraction)? We reduce this problem to the M_9 -partition problem. First of all, we rename the vertices of the hexagon H to 1, 7, 2, 8, 3, 9, instead of a, b, c, d, e, f , in this order (to better correspond to the parts of M_9 -partitions). We may assume that the bipartite graph G is two-coloured so that vertices 1, 2, 3 are white, and vertices 7, 8, 9 are black. It is then easy to see that the desired mapping must take white vertices of G to 1, 2, 3 and black vertices of G to 7, 8, 9.

We first join all the black vertices to each other, to form a large clique. Additionally, we also add the three edges joining the white vertices 1, 2, 3 to each other. Finally, we replace each white vertex by four independent vertices, and each black vertex by four vertices in a clique. (The replacing vertices have the same adjacencies as the original vertices they replaced.) We call the resulting graph G' . It is easy to check that G admits a retraction to H if and only if G' has an M_9 -partition. Indeed, if f is an edge-preserving mapping of $V(G)$ to $V(H)$ in which $f(i) = i$ for $i = 1, 2, 3, 7, 8, 9$, then G' admits an M_9 -partition in which all vertices of G' replacing a black vertex x of G or a white vertex x not in H are placed in part $f(x)$, while all vertices of G' , replacing a white vertex y of H are placed in part $f(y) + 3$. Conversely, if G' has an M_9 -partition, then at least one of any four vertices replacing a white vertex x must be placed in parts 1, 2, 3, 4, 5, 6, and

at least one of any four vertices replacing a black vertex y must be placed in parts 7, 8, 9. Moreover, any vertex replacing a white vertex x of H must be placed in parts 1, 2, 3. It is now easy to conclude that this implies that G admits a retraction to H . \square

Note that the matrix M_9 has three identical rows in both A and B .

6 Conclusions

Unfriendly matrices lead to problems that cannot be characterized by finitely many forbidden induced subgraphs. While many friendly matrices allow such characterizations, including all small matrices (up to size five), we have shown that there are friendly matrices for which such characterizations are not possible, and friendly matrices for which the partition problems are NP-complete.

All three of our basic questions remain open. The problem of whether the last two questions are equivalent for all matrices M is also open (for matrices without 1's equivalence follows from [1, 28]).

References

- [1] A. Atserias, “On digraph coloring problems and treewidth duality”, Proceedings of the 20th Annual IEEE Symposium on Logic in Computer Science (LICS 2005) 106 - 115.
- [2] R. N. Ball, J. Nešetřil, and A. Pultr, “Dualities in full homomorphisms”, manuscript 2006.
- [3] A. Brandstädt, P.L. Hammer, V. Bang Le, and V.V. Lozin, “Bisplit graphs”, Discrete Math. 299 (2005) 11 - 32.
- [4] K. Cameron, E. E. Eschen, C. T. Hoang and R. Sritharan, “The list partition problem for graphs,” Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2004, 391–399.
- [5] Z.A. Chernyak and A.A. Chernyak, ”About recognizing (a, b) -classes of polar graphs”, Discrete Math. 62 (1986) 133 - 138.
- [6] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas, “The strong perfect graph theorem”, Annals of Math. 164 (2006) 51 - 229.
- [7] V. Chvátal, “Star-cutsets and perfect graphs”, *J. Combinatorial Theory B* 39 (1985) 189–199.
- [8] Bruno Courcelle, “Graph Operations and Monadic Second-Order Logic: A Survey”, LPAR 2000, 20–24.
- [9] S. Dantas, C.M.H. de Figueiredo, S. Gravier and S. Klein, “Finding H-partitions efficiently”, Theoret. Informatics Appl. 39 (2005) 133–144.

- [10] T. Feder and P. Hell, “Full constraint satisfaction problems”, *SIAM J. on Computing* 36 (2006) 230–246.
- [11] T. Feder and P. Hell, “Matrix partitions of perfect graphs”, *Discrete Applied Math.* 306 (2006) 2450 – 2460.
- [12] T. Feder and P. Hell, “List homomorphisms to reflexive graphs”, *J. Combinatorial Theory B* 72 (1998) 236–250.
- [13] T. Feder, P. Hell, and J. Huang, “List homomorphisms and circular arc graphs”, *Combinatorica* 19 (1999) 487–505.
- [14] T. Feder, P. Hell, and J. Huang, “Bi-arc graphs and the complexity of list homomorphisms”, *J. Graph Theory* 42 (2003) 61–80.
- [15] T. Feder and P. Hell, “On realizations of point determining graphs, and obstructions to full homomorphisms”, *Discrete Math.*, in press doi:10.1016/j.disc.2006.11.026.
- [16] T. Feder, P. Hell, S. Klein, and R. Motwani, “Complexity of list partitions,” *SIAM J. Discrete Mathematics* 16 (2003) 449–478.
- [17] T. Feder, P. Hell, S. Klein, L. Nogueira, and F. Protti, “List matrix partitions of chordal graphs”, *Theoretical Computer Science* 349 (2005) 52–66.
- [18] T. Feder, P. Hell, D. Král’, and J. Sgall, “Two algorithms for list matrix partitions.” Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2005, 870–876.
- [19] T. Feder, P. Hell, and K. Tucker-Nally, “Digraph matrix partitions and trigraph homomorphisms”, *Discrete Applied Math.* 154 (2006) 2458–2469.
- [20] T. Feder, P. Hell, and W. Hochstättler, “Generalized colourings (matrix partitions) of cographs”, in **Graph Theory in Paris** (A. Bondy et al. editors), Trend in Mathematics, Birkhäuser Verlag, Basel (2006) 149–167.
- [21] T. Feder and M.Y. Vardi, “The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory,” *SIAM J. Comput.* 28 (1998) 236–250.
- [22] C. M. H. de Figueiredo, S. Klein, Y. Kohayakawa and B. Reed, “Finding skew partitions efficiently”, *J. Algorithms* 37 (2000) 505 – 521.
- [23] S. Foldes and P. Hammer, “Split graphs”, *Congr. Numer.* 19 (1977) 311–315.
- [24] P. Hell, S. Klein, L. Tito Nogueira, and F. Protti, “Partitioning chordal graphs into independent sets and cliques”, *Discrete Applied Math.* 141 (2004) 185-194.
- [25] P. Hell, “From graph colouring to constraint satisfaction: there and back again”, in **Topics in Discrete Mathematics**, Springer Verlag Algorithms and Combinatorics Series 26 (2006) 407 – 432.
- [26] P. Hell and J. Nešetřil, “On the complexity of H -coloring,” *J. Combinatorial Theory, Series B* 48 (1990), 92–110.
- [27] P. Hell and J. Nešetřil, **Graphs and Homomorphisms**, Oxford University Press 2004.

- [28] B. Rossman, “Existential positive types and preservation under homomorphisms”, 20th IEEE LICS (2005).
- [29] R. E. Tarjan, “Decomposition by clique separators”, *Discrete Math.* 55 (1985) 221–232.
- [30] N. Vikas, “Computational complexity of compaction to reflexive cycles,” *SIAM J. on Computing* 32 (2003) 253 – 280.
- [31] W. Xie, “Obstructions to trigraph homomorphisms”, M.Sc. Thesis, School of Computing Science, Simon Fraser University 2006.