

The Minimum Feasible Tileset problem*

Yann Disser¹, Stefan Kratsch², and Manuel Sorge³

¹Institut für Mathematik, Graduate School CE, TU Darmstadt, Germany,
disser@mathematik.tu-darmstadt.de

²Institut für Informatik, Humboldt-Universität zu Berlin, Berlin, Germany,
kratsch@informatik.hu-berlin.de

³Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Germany and Department of
Industrial Engineering and Management, Ben-Gurion University of the Negev, Beer Sheva, Israel,
sorge@post.bgu.ac.il

July 26, 2018

Abstract

We introduce and study the MINIMUM FEASIBLE TILESET problem: Given a set of symbols and subsets of these symbols (scenarios), find a smallest possible number of pairs of symbols (tiles) such that each scenario can be formed by selecting at most one symbol from each tile. We show that this problem is APX-hard and that it is NP-hard even if each scenario contains at most three symbols. Our main result is a $4/3$ -approximation algorithm for the general case. In addition, we show that the MINIMUM FEASIBLE TILESET problem is fixed-parameter tractable both when parameterized with the number of scenarios and with the number of symbols.

1 Introduction

Consider the general assignment problem where several devices (e.g., workers, robots, microchips, ...) each can be used in one of k functions/modes at a time (e.g., employing different skills, tools, instruction sets, ...). Given a set of scenarios, the goal is to assign k different functions to each device, such that, for each scenario, all functions requested by the scenario are available simultaneously. In this paper, we initiate the study of this problem for $k = 2$ and the case that each function is requested at most once by each scenario. Formally, we study the following problem (we use “tile” instead of “device” to intuitively capture the fact that a device/tile has two modes/sides).

MINIMUM FEASIBLE TILESET

Input: A universe of symbols F , scenarios $\mathcal{S} \subseteq 2^F \setminus \{F\}$.

Problem: Find a minimum-size tileset \mathcal{T} that is feasible for all scenarios in \mathcal{S} .

*An extended abstract of this article appeared at the *12th Workshop on Approximation and Online Algorithms*, Wrocław, September 2014 [9]. In comparison, apart from full proof details, the present article additionally contains examples and an APX-hardness proof. The authors gratefully acknowledge support by the Alexander von Humboldt-Foundation (Yann Disser), the ‘Excellence Initiative’ of the German Federal and State Governments and the Graduate School CE at TU Darmstadt (Yann Disser), the German Research Foundation (DFG), projects KR 4286/1 (Stefan Kratsch) and NI 369/12 (Manuel Sorge), the Israel Science Foundation, grant no. 551145/14 (Manuel Sorge), and the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 631163.11 (Manuel Sorge).

In the above, a *tile* is a two-element subset of F and we refer to (multi-)sets of tiles as *tilesets*. A tileset \mathcal{T} is *feasible* for scenario S if we can produce all symbols in S by taking at most one symbol from each tile in \mathcal{T} . Formally, a tileset \mathcal{T} is feasible for a scenario $S \subset F$ if there is a mapping $\phi: \mathcal{T} \rightarrow F$, such that $\phi(T) \in T$ for all $T \in \mathcal{T}$, and $S \subseteq \phi[\mathcal{T}] := \{\phi(T) \mid T \in \mathcal{T}\}$. By definition, no scenario contains all symbols of F . Note that such a scenario would require $|F|$ tiles, making the problem trivial. Similarly, we may assume that all symbols in F appear in at least one scenario, otherwise we can simply remove each symbol that does not occur in any scenario. Finally, the requirement that tiles contain no less than two symbols can be met by arbitrarily assigning a second symbol to all tiles of cardinality one.

Example 1. Let us illustrate the problem with two instances of MINIMUM FEASIBLE TILESET:

If our set of symbols is $F = \{A, B, C, 1, 2, 3\}$ and our set \mathcal{S} of scenarios consists of $\{A, B, 1, 2\}$, $\{A, C, 1, 3\}$, and $\{B, C, 2, 3\}$, then it is not hard to check that a feasible tileset is

$$\begin{array}{|c|} \hline A \\ \hline B \\ \hline \end{array}, \begin{array}{|c|} \hline B \\ \hline C \\ \hline \end{array}, \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline \end{array}, \begin{array}{|c|} \hline 2 \\ \hline 3 \\ \hline \end{array}.$$

Herein, each tile is represented by two adjoined boxes which correspond to the two modes in which we can use the tile. Clearly, the feasible tileset above is also of minimum size since each scenario requires at least four tiles.

If we have $n \in \mathbb{N}$, $F = \{1, \dots, 3n\}$, and our scenarios are all size-2 subsets of F , that is, $\mathcal{S} = \binom{F}{2}$, then a feasible tileset with $2n$ tiles is

$$\begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline \end{array}, \begin{array}{|c|} \hline 2 \\ \hline 3 \\ \hline \end{array}, \begin{array}{|c|} \hline 4 \\ \hline 5 \\ \hline \end{array}, \begin{array}{|c|} \hline 5 \\ \hline 6 \\ \hline \end{array}, \dots, \begin{array}{|c|} \hline 3i-2 \\ \hline 3i-1 \\ \hline \end{array}, \begin{array}{|c|} \hline 3i-1 \\ \hline 3i \\ \hline \end{array}, \dots, \begin{array}{|c|} \hline 3n-2 \\ \hline 3n-1 \\ \hline \end{array}, \begin{array}{|c|} \hline 3n-1 \\ \hline 3n \\ \hline \end{array}.$$

To see that the above tileset is feasible for $\binom{F}{2}$, pick any integer $a \in \{1, \dots, 3n\}$. If $a \neq 3i-1$ for each $i \in \{1, \dots, n\}$, then, whichever tile we pick to produce a , each other integer from 1 to $3n$ is available on some other tile. If $a = 3i-1$ for some $i \in \{1, \dots, n\}$, then, to produce a and any other integer $b \in \{1, \dots, 3n\} \setminus \{a\}$, we can pick an arbitrary tile for b and a will be available on another tile. Hence, we can produce any scenario containing two distinct integers from 1 to $3n$ by two distinct tiles. It will become clear later, that each feasible tileset for the instance $(F, \binom{F}{2})$ of MINIMUM FEASIBLE TILESET contains at least $2n$ tiles.

Apart from practical motivations, MINIMUM FEASIBLE TILESET is appealing from a structural point of view. In this work we exhibit equivalent definitions for the problem which are interesting in their own right. At first glance, MINIMUM FEASIBLE TILESET is a covering problem since we must cover all scenarios using tiles that can each cover one of the tile's two symbols in each scenario. It turns out that the problem can also be phrased as a packing/partitioning problem, but with an objective function different from the classical one in terms of number of packed objects or sets (see Section 3). In addition, having tiles be symbol sets of size two suggests a graph interpretation where we are asked to find a minimum set of edges such that for each scenario there is an orientation where each vertex has indegree at least one. For our presentation however, we favor the tileset formulation, since it most naturally generalizes to the original assignment problem with tiles of larger sizes and scenarios which contain multiple copies of the same symbols. Also, the MINIMUM FEASIBLE TILESET interpretation appears suitable for studying the effect of parameters, such as the number of symbols/scenarios, on the complexity.

Results and Outline. We analyze the structure of the graph that has the tiles of a minimum cardinality tiling as its edges, and show that this graph is always (wlog.) a forest. In fact, only the component structure of this forest matters: We may replace trees by arbitrary trees spanning the same components without affecting the feasibility of the corresponding tiling (Section 2). This lets us view MINIMUM FEASIBLE TILING as a partitioning problem, which in turn allows us to prove NP-completeness even when scenarios have size at most three and lets us show APX-hardness for the general case (Section 3). As our main result, we complement the hardness with a $4/3$ -approximation algorithm (for scenarios of arbitrary sizes) inspired by the component structure of the optimum solution (Section 4). We show that the problem is fixed-parameter tractable with respect to the number of scenarios (Section 5) and the number of symbols (Section 6), respectively. We also observe that, when each scenario has size at most d , a polynomial-time compression of an arbitrary instance to $O(|F|^d)$ bits is possible without losing the information about the size of the optimum solution and such a compression to $O(|F|^{d-\epsilon})$ bits is unlikely (Section 6). Finally, we provide a preliminary result on the relevant variant of MINIMUM FEASIBLE TILING where the scenarios are multisets rather than sets and show that also this case is fixed-parameter tractable with respect to the number of symbols (Section 6).

Related Work. The problem most closely related to MINIMUM FEASIBLE TILING is arguably SET PACKING, as 3-SET PACKING appears as a subproblem in our approximation algorithm and also as the source problem for our NP-hardness reduction (in the form of MAXIMUM THREE-DIMENSIONAL MATCHING). SET PACKING has been extensively studied for both approximability and parameterized complexity (see, e.g., [1, 6, 27] and [7, 21] for some recent results). The main difference between the two problems is that SET PACKING is a maximization problem whereas MINIMUM FEASIBLE TILING seeks to minimize the size of a feasible tiling — a measure that is only indirectly related to the number of sets (scenarios). In particular, SET PACKING becomes trivial for a bounded number of sets, whereas for MINIMUM FEASIBLE TILING we get a nontrivial polynomial-time algorithm via integer linear programming (see Section 5).

As alluded to above, the MINIMUM FEASIBLE TILING problem can equivalently be seen as designing an edge-minimal graph on the set of symbols such that, for each scenario, the edges (tiles) can be oriented in such a way that all symbols in the scenario have indegree at least one. The question whether a *given* graph admits an orientation with certain properties has been studied in various settings. For example, Biedl et al. [2] proposed an approximation algorithm for finding a balanced acyclic orientation. Another natural constraint on an orientation that has been studied is to prescribe degrees for each vertex [10, 13, 17].

More abstractly, we are looking for a graph on the set of symbols that fulfills a certain constraint for each scenario. The case where the subgraph induced by each scenario has to be connected is well-studied [3, 4, 11, 16, 18, 28]. In particular, it is NP-hard to find the minimum number of edges needed [11] and to decide whether a planar solution [3, 18] or a solution of treewidth at most three [16] exists.

Preliminaries. For some positive integer $\ell \in \mathbb{N}$, we denote $[\ell] := \{1, 2, \dots, \ell\}$. For a set family \mathcal{S} we use $\bigcup \mathcal{S}$ as a shorthand for $\bigcup_{S \in \mathcal{S}} S$. Apart from standard Landau notation for running times, we also use the O^* notation, which disregards factors that are polynomial in the input size. We use standard graph notation, see the book by Diestel [24], for example. For the relevant notions from parameterized complexity and approximation complexity we refer to textbooks [12, 23, 25] and refs. [5, 26], respectively.

2 Graph structure of tilesets

The tiles in a tileset \mathcal{T} over a universe of symbols F can be viewed as the edges of the undirected (multi-) graph $G(\mathcal{T}) := (F, \mathcal{T})$. In this section, we establish that there always exist optimal tilesets with a simple graph structure. This is made formal in the following lemma which will be useful in later sections.

Lemma 1. *Let F be a universe of symbols, S a family of scenarios over F , and \mathcal{T} a tileset feasible for S . There is a tileset $\mathcal{T}' \subseteq \binom{F}{2}$ feasible for S such that $|\mathcal{T}'| \leq |\mathcal{T}|$ and $G(\mathcal{T}')$ is a forest.*

Note that each connected component of $G(\mathcal{T}')$ has size at least two because each symbol occurs in at least one scenario and hence is incident with at least one edge.

In the proof of Lemma 1 it is convenient to think of feasibility of \mathcal{T} via orientations of the graph $G(\mathcal{T})$. Let us say that an orientation of $G(\mathcal{T})$ is *feasible* for the scenario S if each vertex in S has indegree at least one. It is easy to see that deciding whether \mathcal{T} is feasible for some scenario $S \subseteq F$ is equivalent to deciding whether there is a *feasible* orientation of the edges of $G(\mathcal{T})$ for S . We obtain the following lemma.

Lemma 2. *For every tileset \mathcal{T} and scenario S over a universe of symbols F the following are equivalent.*

- (i) \mathcal{T} is feasible for S ,
- (ii) there is a feasible orientation of $G(\mathcal{T})$ for S ,
- (iii) for every connected component C of $G(\mathcal{T})$, there is a feasible orientation of $G(\mathcal{T})$ for $S \cap C$,
- (iv) for every connected component C of $G(\mathcal{T})$, tileset \mathcal{T} is feasible for $S \cap C$.

Proof. Note that it suffices to prove the equivalence of the first three statements, since equivalence of (i) and (ii) implies equivalence of (iii) and (iv).

(i) \Rightarrow (ii): Assume that \mathcal{T} is feasible for S and let $\phi : \mathcal{T} \rightarrow F$ be the corresponding mapping with $\phi(T) \in T$ for all $T \in \mathcal{T}$, and $S \subseteq \phi[\mathcal{T}]$. We obtain an orientation of $G(\mathcal{T})$ by orienting each edge $T \in \mathcal{T}$ towards $\phi(T)$. Since $S \subseteq \phi[\mathcal{T}]$, each symbol in S has indegree at least one and we have a feasible orientation of $G(\mathcal{T})$ for S .

(ii) \Rightarrow (iii): Clearly, a feasible orientation of $G(\mathcal{T})$ for S is, in particular, a feasible orientation for $S \cap C$ for every connected component C of $G(\mathcal{T})$.

(iii) \Rightarrow (i): Let C_1, \dots, C_k denote the connected components of $G(\mathcal{T})$ and assume that there are feasible orientations $\vec{G}_1, \dots, \vec{G}_k$ for $S \cap C_1, \dots, S \cap C_k$, respectively. Since $S \subseteq \bigcup_i C_i$, we obtain a feasible orientation \vec{G} of $G(\mathcal{T})$ for S as $\vec{G}_1[C_1] \uplus \dots \uplus \vec{G}_k[C_k]$. We define the mapping $\phi : \mathcal{T} \rightarrow F$ by setting $\phi(T) = s$ for each $T \in \mathcal{T}$, where s is the symbol towards which the edge T is oriented in \vec{G} . By definition, $\phi(T) \in T$ for all $T \in \mathcal{T}$, and, since \vec{G} is feasible for S , we have $S \subseteq \phi[\mathcal{T}]$. The existence of the mapping ϕ hence proves that \mathcal{T} is feasible for S . \square

Using the notion of feasible orientations we now observe that connected components in $G(\mathcal{T})$ yield feasibility for each of their strict subsets.

Lemma 3. *Let \mathcal{T} be a tileset, C a connected component of $G(\mathcal{T})$ and $C' \subsetneq C$. Then, \mathcal{T} is feasible for C' .*

Proof. The proof is by induction over the size of C' . If C' contains a single symbol, that is, $C' = \{s\}$, then we obtain a feasible orientation by orienting an arbitrary edge towards s ; such an edge exists because C' is part of the (larger) connected component C . Consider the case $|C'| > 1$. First assume that $G(\mathcal{T})[C']$ contains no edges, i.e., C' is an independent set. Then, there is an edge in $G(\mathcal{T})[C]$ for each symbol $s \in C'$, connecting s to $C \setminus C'$. A feasible

orientation for C' can simply be obtained by orienting these edges towards C' . Now, assume there is an edge $\{s, s'\}$ in $G(\mathcal{T})[C']$ and consider the graph G' obtained by contracting $\{s, s'\}$. By induction, there is a feasible orientation of G' for C'' , where C'' is obtained from C' by identifying s and s' . Hence, there is an orientation of $G(\mathcal{T})$ such that all vertices in C' except one of $\{s, s'\}$ have indegree at least one. We orient the edge $\{s, s'\}$ towards the vertex of smaller indegree to obtain the desired feasible orientation of $G(\mathcal{T})$. \square

We are ready for a proof of Lemma 1. Intuitively, we observe that cycle components in $G(\mathcal{T})$ yield feasibility for *any* of their subsets and hence are a safe replacement for every component with a large number of edges. Then we show how to break cycle components into trees.

Proof of Lemma 1. We replace connected components in $G(\mathcal{T})$, maintaining feasibility of \mathcal{T} and without increasing the cardinality of \mathcal{T} .

Let C_1, \dots, C_k be the connected components of $G(\mathcal{T})$ that contain a cycle, and let $C = \bigcup_{i=1}^k C_i$. We obtain a new tiling \mathcal{T}' by replacing the edges in $G(\mathcal{T})[C]$ with a cycle on C . (If $|C| = 1$, we introduce a self-loop and if $|C| = 2$, we introduce two parallel edges.) Since each component C_i , $i \in [k]$, originally contained at least $|C_i|$ edges, and we only introduced $|C|$ new edges, the cardinality of \mathcal{T}' is not larger than the cardinality of \mathcal{T} . We observe that \mathcal{T}' is still feasible. Consider an arbitrary scenario $S \in \mathcal{S}$. Clearly, if $S \cap C = \emptyset$, then \mathcal{T}' is feasible for S . Hence, assume $S \cap C \neq \emptyset$. By Lemma 2 there is a feasible orientation of $G(\mathcal{T})$ for $S \setminus C$. This implies that there is still a feasible orientation of $G(\mathcal{T}')$ for $S \setminus C$. By Lemma 2 it suffices to prove that there is a feasible orientation of $G(\mathcal{T}')$ for $S \cap C$. Indeed, since $G(\mathcal{T}')[C]$ is a cycle, orienting the edges in one direction along the cycle yields a feasible orientation for any subset of C , and, in particular, for $S \cap C$. Hence, \mathcal{T}' is still feasible for every $S \in \mathcal{S}$.

By definition, every connected component C' of $G(\mathcal{T})$ outside of C is a tree. Hence, the connected components of $G(\mathcal{T}')$ are trees, with the exception of at most one component C that is a cycle. We now modify C in order to obtain our final feasible tiling \mathcal{T}'' with the desired structure.

First, consider the case that C is the only connected component of $G(\mathcal{T}')$. Then, we can remove an arbitrary tile from \mathcal{T}' to obtain \mathcal{T}'' : Since, by definition, \mathcal{S} does not contain F as a scenario, by Lemma 3, \mathcal{T}'' is feasible for all scenarios $S \in \mathcal{S}$. Clearly, $G(\mathcal{T}'')$ is a tree, as required.

Now assume that there is at least one tree component C' in $G(\mathcal{T}')$ along with C . Consider an arbitrary edge $\{s, s'\}$ in C and an arbitrary vertex $s'' \in C'$. We remove $\{s, s'\}$ from \mathcal{T}' and instead add the edge $\{s, s''\}$ to obtain the tiling \mathcal{T}'' . Clearly, $G(\mathcal{T}'')$ is a forest. It remains to prove that \mathcal{T}'' is feasible for every scenario $S \in \mathcal{S}$. By Lemma 2, \mathcal{T}' is feasible for $S \setminus C$ and, hence, so is \mathcal{T}'' . Because $C \cup C'$ is a connected component in $G(\mathcal{T}'')$, Lemma 3 guarantees that \mathcal{T}'' is feasible for every $S' \subsetneq (C \cup C')$ and, in particular, \mathcal{T}'' is feasible for $S \cap C$. Hence, as \mathcal{T}'' is feasible for $S \cap C$ and $S \setminus C$, applying Lemma 2 we obtain that \mathcal{T}'' is feasible for S . Clearly, $G[\mathcal{T}'']$ is a forest, as required. \square

Intuitively, Lemmas 2 and 3 imply that only the partition of the symbols induced by the component structure of the graph of a tiling matters, but not the exact topology of each of the trees. This leads to the following.

Theorem 1. *Let \mathcal{S} be a family of scenarios and \mathcal{T} be a tiling over symbols F . If $G(\mathcal{T})$ is a forest, then \mathcal{T} is feasible for \mathcal{S} if and only if no connected component C of $G(\mathcal{T})$ is fully contained in any scenario $S \in \mathcal{S}$, i.e., $C \not\subseteq S$ for all scenarios $S \in \mathcal{S}$ and all connected components C of $G(\mathcal{T})$.*

Proof. “ \Rightarrow ”: Assume towards a contradiction that \mathcal{T} is feasible, $G(\mathcal{T})$ is a forest, and there is a scenario $S \in \mathcal{S}$ and a component C of $G(\mathcal{T})$ such that $C \subseteq S$. By Lemma 2 there is a feasible orientation of $G(\mathcal{T})$ for $S \cap C = C$. But this is absurd, because $G(\mathcal{T})$ is a forest and hence $G(\mathcal{T})[C]$ contains only $|C| - 1$ edges.

“ \Leftarrow ”: For each component C of $G(\mathcal{T})$ and every scenario $S \in \mathcal{S}$, we have that $C \cap S \subsetneq C$, since $C \not\subseteq S$. By Lemma 3, we get that \mathcal{T} is feasible for $C \cap S$. Since this is true for all choices of C and S , Lemma 2 implies that \mathcal{T} is feasible for \mathcal{S} . \square

Example 2. We can now observe that for the instance $(F = \{1, \dots, 3n\}, \mathcal{S} = \binom{F}{2})$ from Example 1 each feasible tileset contains at least $2n$ tiles: Let \mathcal{T} be a feasible tileset for \mathcal{S} . No connected component of $G(\mathcal{T})$ can be contained in any set in \mathcal{S} . Thus, each connected component has size at least three, meaning that there are at least n connected components. Since each such connected component induces at least two tiles, \mathcal{T} contains at least $2n$ tiles.

3 NP-hardness and APX-Hardness of Minimum Feasible Tileset

In this section we establish the following result.

Theorem 2. MINIMUM FEASIBLE TILESET is APX-hard. MINIMUM FEASIBLE TILESET is NP-hard even if each scenario has size at most three.

Before proving Theorem 2, let us check that the decision variant of MINIMUM FEASIBLE TILESET, in which we want to check for feasible tilesets of size at most a given integer, is contained in NP: A feasible tileset can be encoded using polynomially many bits with respect to $|F|$. Verifying feasibility comes down to solving one bipartite matching problem for each scenario on an auxiliary graph that has an edge between each symbol in the scenario and every tile containing that symbol, which is possible in polynomial time. Thus we can infer from Theorem 2 that the decision variant of MINIMUM FEASIBLE TILESET is NP-complete.

We now prove NP- and APX-hardness of MINIMUM FEASIBLE TILESET. For this, we first give a relation of MINIMUM FEASIBLE TILESET to a partitioning problem. Let us say that, for a finite set of symbols F and a family of scenarios $\mathcal{S} \subseteq 2^F$, a partition \mathcal{P} of F is *admissible*, if for every $P \in \mathcal{P}$ and every $S \in \mathcal{S}$ we have $P \not\subseteq S$. We obtain the following.

Lemma 4. Let F be a set of symbols and $\mathcal{S} \subseteq 2^F \setminus \{F\}$ a family of scenarios. There is a feasible tileset of size ℓ for \mathcal{S} if and only if there is a partition of F which is admissible for \mathcal{S} and comprises $|F| - \ell$ parts.

Proof. “ \Rightarrow ”: By Lemma 1 there is a feasible tileset \mathcal{T}' for \mathcal{S} of cardinality ℓ such that $G(\mathcal{T}')$ is a forest. The connected components C_1, \dots, C_k of $G(\mathcal{T}')$ induce a partition \mathcal{P}' which we claim to be admissible: Indeed, by Theorem 1 we have $C_i \not\subseteq S$ for all connected components C_i , $i \in [k]$, and scenarios $S \in \mathcal{S}$. Furthermore, since there are exactly ℓ edges in $G(\mathcal{T}')$ and each connected component is a tree, we have $\ell = \sum_{i=1}^k |C_i| - 1 = |F| - k$. Hence, our partition has $k = |F| - \ell$ parts, as required.

“ \Leftarrow ”: Let $\mathcal{P} = \{P_1, \dots, P_p\}$ be an admissible partition with $|F| - \ell$ parts. We construct a tileset \mathcal{T} by setting $G(\mathcal{T})[P_i]$ to an arbitrary spanning tree for each $i \in [p]$. Since $P_i \not\subseteq S$ for each $S \in \mathcal{S}$ and each $i \in [p]$, by Theorem 1, \mathcal{T} is feasible for \mathcal{S} . The number of tiles in \mathcal{T} is $\sum_{i=1}^p |P_i| - 1 = |F| - p = |F| - (|F| - \ell) = \ell$, as required. \square

Thus, MINIMUM FEASIBLE TILESET is equivalent to finding a finest-possible partition, i.e. with maximum number of parts, of the symbols such that no part in the partition is contained in any scenario.

We now give a reduction from MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING which is both NP-hard and APX-hard [19]:

MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING

Input: Three pairwise disjoint sets X, Y, Z , and a set $D \subseteq X \times Y \times Z$ of triples such that each element in $X \cup Y \cup Z$ occurs in at most three triples in D .

Problem: Find a maximum-size *three-dimensional matching* for D , i.e., a maximum-size subset $D' \subseteq D$ such that no element of $X \cup Y \cup Z$ occurs in two triples in D' .

Proof of Theorem 2. We give a PTAS-reduction from MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING [5]. More precisely, given an instance (X, Y, Z, D) and a desired approximation ratio r , we construct an instance (F, \mathcal{S}) of MINIMUM FEASIBLE TILESET in time polynomial in the instance size for every fixed r , subject to the following condition. There is a function $f: (0, 1) \rightarrow \mathbb{Q}$ such that for every $r \in (0, 1)$ we have $f(r) > 1$ and for a arbitrary given $f(r)$ -approximate feasible tiling \mathcal{T} for (F, \mathcal{S}) we can construct an r -approximate three-dimensional matching M for (X, Y, Z, D) in polynomial time. We specify the function f below (while ensuring that $f(r) > 1$).

We first describe how to construct the MINIMUM FEASIBLE TILESET instance (F, \mathcal{S}) from a MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING instance (X, Y, Z, D) . Set the universe $F := X \cup Y \cup Z$. We choose a function $g: (0, 1) \rightarrow \mathbb{N}$ with $g(r) \geq 3$ for all $r \in (0, 1)$. The precise function is given below. The scenarios \mathcal{S} consist of all subsets of F that have size at most $g(r)$ and do not contain any triple in D as a subset (we interpret D as a family of three-element sets). Formally, $\mathcal{S} := \{S \subseteq F \mid |S| \leq g(r) \wedge \forall E \in D: E \setminus S \neq \emptyset\}$. This concludes the construction. Let $n := |X \cup Y \cup Z| = |F|$. Clearly, for any fixed r , we can carry out the construction in time $\mathcal{O}(n^{g(r)+1})$, that is, in polynomial time in the instance size.

Before we show how to compute an approximate three-dimensional matching from an approximate feasible tiling, we find a relation between the optimal solution sizes of the two instances. Note that, from each three-dimensional matching N we can construct an admissible partition \mathcal{R} of F for \mathcal{S} satisfying $|\mathcal{R}| = |N|$ as follows. Initially, take $\mathcal{R} = N$ (where N is interpreted as a family of three-element sets). Then, replace an arbitrary part $P \in \mathcal{R}$ with $P \cup ((X \cup Y \cup Z) \setminus \bigcup \mathcal{R})$. That is, add to P all elements not covered by N . By definition of \mathcal{S} , there is no set $S \in \mathcal{S}$ that contains any $P \in \mathcal{R}$, and hence \mathcal{R} is admissible. Letting $\text{opt}_{3\text{DM}}$ denote the size of an optimal solution to the MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING instance, we thus have $|\mathcal{P}^*| \geq \text{opt}_{3\text{DM}}$ for an admissible partition \mathcal{P}^* containing the maximum number of parts. Lemma 4 implies that $|\mathcal{P}^*| = n - \text{opt}_{\text{FT}}$, where opt_{FT} denotes the size of an optimal solution for the MINIMUM FEASIBLE TILESET instance. Rearranging terms hence yields $\text{opt}_{\text{FT}} = n - |\mathcal{P}^*|$. Because of $|\mathcal{P}^*| \geq \text{opt}_{3\text{DM}}$, we have

$$\text{opt}_{\text{FT}} \leq n - \text{opt}_{3\text{DM}}. \tag{1}$$

Now let \mathcal{T} be an arbitrary $f(r)$ -approximate feasible tiling \mathcal{T} for (F, \mathcal{S}) . We construct an r -approximate three-dimensional matching M for (X, Y, Z, D) in polynomial time as follows. Along the way, we gather observations that allow us to prove that M is r -approximate in the end.

First, by Lemma 4 there is a partition \mathcal{P} of F which is admissible for \mathcal{S} and has $n - |\mathcal{T}|$ parts. In other words, \mathcal{T} has $n - |\mathcal{P}|$ tiles. (As the proof of Lemma 4 is constructive, it is not hard to check that \mathcal{P} can be computed in polynomial time.) As \mathcal{T} is $f(r)$ -approximate, $|\mathcal{T}| \leq f(r) \text{opt}_{\text{FT}}$ and hence, $n - |\mathcal{P}| \leq f(r) \text{opt}_{\text{FT}}$. Applying Inequality (1) we thus obtain

$$n - |\mathcal{P}| \leq f(r)(n - \text{opt}_{3\text{DM}}). \tag{2}$$

We create a partition $\mathcal{P}^1 = \mathcal{P}_3^1 \cup \mathcal{P}_{g(r)+1}^1$ from \mathcal{P} as follows. Obtain \mathcal{P}_3^1 by picking, for each part $P \in \mathcal{P}$ which contains a triple of D , one triple $E \in D$ with $E \subseteq P$ and putting E (as a set) into \mathcal{P}_3^1 . To create $\mathcal{P}_{g(r)+1}^1$, if $|F \setminus \bigcup_{P \in \mathcal{P}_3^1} P| < g(r) + 1$, then put $F \setminus \bigcup_{P \in \mathcal{P}_3^1} P$ into $\mathcal{P}_{g(r)+1}^1$ as a single set of size at most $g(r)$. In this case we call $\mathcal{P}_{g(r)+1}^1$ *degenerate*. Otherwise, put $\mathcal{P}_{g(r)+1}^1$ to be an arbitrary partition of $F \setminus \bigcup_{P \in \mathcal{P}_3^1} P$ into parts of size $g(r) + 1$ and, perhaps, one part of size at least $g(r) + 2$ and at most $2g(r) + 1$. Note that \mathcal{P}_3^1 is admissible for \mathcal{S} because each triple in \mathcal{P}_3^1 is not contained in any set in \mathcal{S} by definition of \mathcal{S} . We claim that $|\mathcal{P}^1| \geq |\mathcal{P}|$. Clearly, for each part in \mathcal{P} that contains a triple of D there is at least one part also in \mathcal{P}^1 . Furthermore, since \mathcal{S} contains all $g(r)$ -element sets which do not contain any triple of D , each set $P \in \mathcal{P}$ that does not contain a triple from D must contain at least $g(r) + 1$ elements because \mathcal{P} is admissible for \mathcal{S} . Hence, $|\mathcal{P}^1| \geq |\mathcal{P}|$. From Inequality (2) it follows that

$$n - |\mathcal{P}^1| \leq f(r)(n - \text{opt}_{3\text{DM}}). \quad (3)$$

Note that the three-element sets in \mathcal{P}^1 , i.e., \mathcal{P}_3^1 , form a three-dimensional matching for the instance (X, Y, Z, D) . The sets in \mathcal{P}_3^1 will form our r -approximate three-dimensional matching M after one further augmentation step. The aim of this augmentation is to bound $|\mathcal{P}_3^1|$ by a function of $|\mathcal{P}_{g(r)+1}^1|$. This enables us to give a lower bound on $|M|$ via the size of \mathcal{P}^1 .

Consider the following modification of \mathcal{P}^1 . If there is a triple in D that is disjoint from $\bigcup_{P \in \mathcal{P}_3^1} P$, then add this triple to \mathcal{P}_3^1 . If we now have $|F \setminus \bigcup_{P \in \mathcal{P}_3^1} P| < g(r) + 1$, then replace the sets in $\mathcal{P}_{g(r)+1}^1$ by the single set $F \setminus \bigcup_{P \in \mathcal{P}_3^1} P$. Otherwise, replace $\mathcal{P}_{g(r)+1}^1$ by an arbitrary partition of $F \setminus \bigcup_{P \in \mathcal{P}_3^1} P$ into parts of size $g(r) + 1$ and, perhaps, one part of size at least $g(r) + 2$ and at most $2g(r) + 1$.

We claim that the above two modification steps do not decrease the number of sets in \mathcal{P}^1 . This is clear if $\mathcal{P}_{g(r)+1}^1$ was degenerate before applying them. Otherwise, we have $|F \setminus \bigcup_{P \in \mathcal{P}_3^1} P| > g(r) + 3$ before applying the modification. Hence, each set in $\mathcal{P}_{g(r)+1}^1$ had size at least $g(r) + 1 \geq 4$. Since we moved only three elements from these sets to \mathcal{P}_3^1 and afterwards repartition the remaining elements with sets of size at least 4, the total number of sets cannot decrease.

As before, \mathcal{P}_3^1 remains admissible for \mathcal{S} . Let \mathcal{P}^2 be the partition obtained by exhaustively applying the above modification, let \mathcal{P}_3^2 equal the resulting set \mathcal{P}_3^1 and let $\mathcal{P}_{g(r)+1}^2$ be the resulting set $\mathcal{P}_{g(r)+1}^1$. We define the three-dimensional matching M as the family of three-element parts in $\mathcal{P}^2 \cap D$, that is $M = \mathcal{P}_3^2$. As mentioned, we have $|\mathcal{P}^2| \geq |\mathcal{P}^1|$. From Inequality (3) it thus follows that

$$n - |\mathcal{P}^2| \leq f(r)(n - \text{opt}_{3\text{DM}}). \quad (4)$$

It remains to show that M is r -approximate (for appropriate functions $f(r)$ and $g(r)$).

We now claim that $12|\mathcal{P}_3^2|/(g(r) + 1) \geq |\mathcal{P}_{g(r)+1}^2|$. If $\mathcal{P}_{g(r)+1}^1$ is degenerate and if this relation does not hold, then $|F|$ is upper bounded by $3|\mathcal{P}_3^2| + g(r) - 1 < 3(g(r) + 1)/12 + g(r) - 1$, that is, $|F|$ is upper bounded by a constant. Hence, we may compute the optimal solution in constant time in this case. Thus, we may, without loss of generality, assume that the relation holds if $\mathcal{P}_{g(r)+1}^1$ is degenerate.

If $\mathcal{P}_{g(r)+1}^1$ is not degenerate, we claim that the above modification of \mathcal{P}^1 is applicable as long as $12|\mathcal{P}_3^1| < (g(r) + 1)|\mathcal{P}_{g(r)+1}^1|$, where the right hand side bounds the number of unmatched elements. Indeed, since each element in $F = X \cup Y \cup Z$ is contained in at most three triples in D , for each triple E in \mathcal{P}_3^1 , there are at most twelve elements of F whose incident triples in D cannot be added to \mathcal{P}_3^1 , because they overlap with E . Hence, if $12|\mathcal{P}_3^1| < (g(r) + 1)|\mathcal{P}_{g(r)+1}^1|$, then there exists at least one element of F whose incident triples do not overlap with any triple in \mathcal{P}_3^1 . This means that at least one triple will be added to \mathcal{P}_3^1 in the above modification step,

because, without loss of generality, each element is in at least one triple. This indeed implies for partition \mathcal{P}^2 (after exhaustive modification) that $12|\mathcal{P}_3^2|/(g(r)+1) \geq |\mathcal{P}_{g(r)+1}^2|$. We thus have

$$\begin{aligned} n - |\mathcal{P}^2| &= n - (|\mathcal{P}_3^2| + |\mathcal{P}_{g(r)+1}^2|) \\ &\geq n - \left(1 + \frac{12}{g(r)+1}\right) |\mathcal{P}_3^2|, \end{aligned}$$

and since $\mathcal{P}_3^2 = M$, in combination with Inequality (4) we have

$$n - \left(1 + \frac{12}{g(r)+1}\right) |M| \leq f(r)(n - \text{opt}_{3\text{DM}}).$$

Thus,

$$\frac{g(r)+1}{g(r)+13} \cdot ((1-f(r))n + f(r)\text{opt}_{3\text{DM}}) \leq |M|. \quad (5)$$

The same modification that we applied above to (the three-element sets of) \mathcal{P}^1 works for each three-dimensional matching. As we cannot improve a maximum three-dimensional matching, each element of F is either matched — of this type there are $3\text{opt}_{3\text{DM}}$ elements — or it is in a triple together with a matched element — of this type there are at most $12\text{opt}_{3\text{DM}}$ elements because each element is in at most three triples in D . Hence, $n \leq 3\text{opt}_{3\text{DM}} + 12\text{opt}_{3\text{DM}} = 15\text{opt}_{3\text{DM}}$. Note that $1 - f(r) < 0$. Inequality (5) thus implies

$$\frac{g(r)+1}{g(r)+13} \cdot (15 - 14f(r)) \cdot \text{opt}_{3\text{DM}} \leq |M|. \quad (6)$$

We now define f and g by setting $g(r) := \max\{3, \lceil 13r/(1-r) \rceil\}$ and

$$f(r) := \frac{-rg(r) - 13r + 15g(r) + 15}{14g(r) + 14}.$$

Clearly, $g(r) \geq 3$ as required. We claim also that $f(r) > 1$. To see this, consider subtracting the denominator from the numerator in $f(r)$ to obtain x , that is,

$$x := -rg(r) - 13r + 15g(r) + 15 - 14g(r) - 14 = -rg(r) - 13r + g(r) + 1 = (1-r)g(r) - 13r + 1.$$

If $3 \leq 13r/(1-r)$, then $x \geq (1-r) \cdot \frac{13r}{1-r} - 13r + 1 > 0$, that is $f(r) > 1$. If $3 > 13r/(1-r)$ then $r < 3/16$. Thus, $x \geq (1 - 3/16) \cdot 3 - 13 \cdot 3/16 + 1 = 1 > 0$ and again $f(r) > 1$. Thus, these are suitable definitions. All that remains is to show that the approximation factor in Inequality (6) is at least r , that is,

$$\frac{g(r)+1}{g(r)+13} \cdot (15 - 14f(r)) \geq r. \quad (7)$$

Observe that

$$15 - 14f(r) = \frac{15g(r) + 15}{g(r) + 1} - \frac{15g(r) + 15 - rg(r) - 13r}{g(r) + 1} = \frac{rg(r) + 13r}{g(r) + 1}.$$

Hence,

$$\frac{g(r)+1}{g(r)+13} \cdot (15 - 14f(r)) = \frac{rg(r) + 13r}{g(r) + 13} = r.$$

This implies that Inequality (7) holds. Hence, there is a PTAS-reduction from MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING to MINIMUM FEASIBLE TILESET.

NP-hardness of the decision version of MINIMUM FEASIBLE TILESET follows from the following modification to the reduction above. Instead of MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING we reduce from the NP-hard decision problem which asks whether there is a three-dimensional matching with $n/3$ triples [15]. We use the reduction above and set $g(r) = 3$ and the desired feasible tileset size to $2n/3$. As mentioned, this can be done in polynomial time. For the correctness, each three-dimensional matching of size $n/3$ is also an admissible partition for \mathcal{S} . Hence, it implies a feasible tileset of size $n - n/3 = 2n/3$ by Lemma 4. In the reverse direction, each feasible tileset of size $2n/3$ implies an admissible partition with $n/3$ parts by Lemma 4. Each of these parts is of size three because \mathcal{S} contains all size-two subsets of F . Among sets of size three, the only sets not contained in \mathcal{S} are precisely the sets in D ; hence, each part of an admissible partition is in D . That is, any admissible partition is a three-dimensional matching as well. \square

4 A 4/3-approximation for Minimum Feasible Tileset

In this section, we propose an approximation algorithm for MINIMUM FEASIBLE TILESET with unbounded scenario size. Motivated by the structural insights of Section 2, we construct a tileset that induces a forest in the corresponding graph, with the property that none of its components are contained in a single scenario. Since a component of size k requires $k - 1$ tiles, we additionally aim for small components in order to keep the resulting tileset small.

We first take as many components of size two as possible among all disjoint sets of two symbols that are not both contained in the same scenario. This can easily be achieved by computing a maximum matching in the graph that has an edge for each candidate component. Similarly, among all remaining symbols, we try to form many (disjoint) components of size three, without creating components that are contained in a single scenario. For this, we employ a simple greedy strategy, that repeatedly takes any possible component until no possible candidates remain. (While there are better packing strategies available for sets of size three, we will see that improving the packing strategy alone does not improve our approximation ratio.) Finally, for each leftover symbol we add an individual tile (pairing that symbol in such a way as to prevent cycles).

We give a more formal listing in Algorithm A. We use $\bar{F}_i(F') = \{C \in \binom{F'}{i} \mid \forall S \in \mathcal{S}: C \not\subseteq S\}$ to denote the family of all sets of symbols in F' that are of size i and not fully contained in a single scenario. In the following, we identify connected components with their sets of vertices.

Algorithm A: 4/3-approximation for minimum feasible tilesets

Input: A set F of symbols and a set \mathcal{S} of scenarios, where $\mathcal{S} \subseteq 2^F \setminus \{F\}$.

Output: A set of tiles \mathcal{T} .

$\mathcal{T}_2 \leftarrow$ maximum matching in graph $G(\bar{F}_2(F))$.

$\mathcal{P} \leftarrow$ greedy set packing of $\bar{F}_3(F \setminus \bigcup_{t \in \mathcal{T}_2} t)$.

$\mathcal{T}_3 \leftarrow \bigcup_{\{f_1, f_2, f_3\} \in \mathcal{P}} \{\{f_1, f_2\}, \{f_2, f_3\}\}$.

if $\mathcal{T}_2 \cup \mathcal{T}_3 \neq \emptyset$ **then** take $f_{\text{root}} \in \bigcup_{t \in \mathcal{T}_2 \cup \mathcal{T}_3} t$

else take $f_{\text{root}} \in F$.

$\mathcal{T}_1 \leftarrow \{\{f, f_{\text{root}}\} \mid f \in F \setminus \bigcup_{t \in \mathcal{T}_2 \cup \mathcal{T}_3} t, f \neq f_{\text{root}}\}$.

return $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{T}_3$.

Theorem 3. *Algorithm A computes a 4/3-approximation for MINIMUM FEASIBLE TILESET.*

Proof. We first argue that the set of tiles $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{T}_3$ computed by Algorithm A is feasible for \mathcal{S} . First observe that $G(\mathcal{T})$ is a forest. This is true, because $G(\mathcal{T}_2 \cup \mathcal{T}_3)$ consists of trees of sizes 2 and 3, $G(\mathcal{T}_1)$ is a star, and $\mathcal{T}_1 \cap (\mathcal{T}_2 \cup \mathcal{T}_3)$ contains at most one node (f_{root}). Using Theorem 1 it only remains to show that no connected component C of $G(\mathcal{T})$ is contained in any scenario $S \in \mathcal{S}$, i.e. $C \cap S \subsetneq C$. By definition of Algorithm A this is true for all connected components of the graph $G(\mathcal{T}_2 \cup \mathcal{T}_3)$. If $\mathcal{T}_2 \cup \mathcal{T}_3 \neq \emptyset$, then each component of $G(\mathcal{T})$ is a superset of a component of $G(\mathcal{T}_2 \cup \mathcal{T}_3)$, and is thus not contained in any scenario. If $\mathcal{T}_2 \cup \mathcal{T}_3$ is empty, then $G(\mathcal{T}) = G(\mathcal{T}_1)$ consists of a single component that is not contained in any scenario, since, by definition, $F \notin \mathcal{S}$. Thus \mathcal{T} is feasible for \mathcal{S} .

We now bound the size of \mathcal{T} with respect to a minimum cardinality tiling \mathcal{T}^* . To do this we *distribute* virtual currency (*gold*) to the symbols in F , such that the total gold distributed is $4/3$ times the size of \mathcal{T}^* . We later use this gold to *pay* one unit of gold to certain symbols that these can in turn use to *provide for* (at most) one tile of \mathcal{T} that involves this symbol. To complete the proof, we establish that each tile of \mathcal{T} is provided for by one of its two symbols.

Let $G^* := G(\mathcal{T}^*)$ be the graph induced by \mathcal{T}^* and \bar{F}_i^* be the set of connected components of size $i \in \{2, \dots, |F|\}$ in G^* . By Lemma 1, we may assume that G^* is a forest. Furthermore, because each symbol appears in at least one scenario, graph G^* does not contain components of size 1. Since the symbols in a component of size $i > 1$ are part of exactly $i - 1$ tiles in \mathcal{T}^* , we may distribute all available gold by giving $4/3 \cdot \frac{i-1}{i}$ gold to each symbol in a component of \bar{F}_i^* , for all $i \in \{2, \dots, |F|\}$. This gold is used to pay symbols in what follows. We call a symbol $s \in F$ *sufficiently paid* if one of the following holds: (i) s is paid, (ii) s appears in a tile $T \in \mathcal{T}_2$ and the other symbol of T is paid, or (iii) s appears in a tile $T \in \mathcal{T}_3$ and the other two symbols in the same component of $G(\mathcal{T}_3)$ are paid. Below, we show how to sufficiently pay all symbols. This completes the proof, since then all tiles in $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{T}_3$ can be provided for (note that then each tile in \mathcal{T}_1 contains its own paid symbol). We call a component of G^* *sufficiently paid*, if all its symbols are sufficiently paid. Let $F_{\geq 4}^* := F \setminus \bigcup_{C \in \bar{F}_2^* \cup \bar{F}_3^*} C$ be the set of all symbols not in components of size two or three in G^* . In paying the symbols we will maintain the invariant that each element of $\bar{F}_2^* \cup \bar{F}_3^* \cup F_{\geq 4}^*$ is either sufficiently paid, or it still holds its gold (all its symbols still hold their gold, respectively).

We define a graph $H = (V, E)$ that has the components in $\bar{F}_2^* \cup \bar{F}_3^*$ as its vertices, as well as the symbols that are not part of these components, i.e., $V = \bar{F}_2^* \cup \bar{F}_3^* \cup F_{\geq 4}^*$ (cf. Figure 1). In this way, each vertex of H represents up to three symbols. For each tile $T \in \mathcal{T}_2$ we introduce an edge connecting the vertices of H representing the two symbols of T , possibly introducing self-loops. Since \mathcal{T}_2 is a matching, and since the vertices in H represent at most three symbols each, all vertices in H have degree at most 3. We partition the edges of H into paths, cycles, and self-loops, and show for each how to use the gold remaining at its vertices to pay all symbols in the components of G^* that are intersected by the path/cycle/self-loop. We will ensure that every symbol (except possibly f_{root}) on a tile in \mathcal{T}_1 is paid. Since each symbol on a tile of \mathcal{T}_2 appears only exactly on this and no other tile of $\mathcal{T}_2 \cup \mathcal{T}_3$, it is thus sufficient to pay only one of the two symbols on each tile of \mathcal{T}_2 .

Let \mathcal{P} be the set of all paths in H connecting (different) vertices of degree 1 or 3 with internal nodes of degree 2. Consider the paths in \mathcal{P} one by one. We use the gold available along path $P \in \mathcal{P}$ of length k as follows (cf. Figure 2). Let N_2, N_3 be the number of internal nodes of P that represent 2 and 3 symbols, respectively. Note that P has no inner nodes that represent a single symbol, since \mathcal{T}_2 is a matching, and hence $k = 1 + N_2 + N_3$. Also, P is the only path visiting these inner nodes and hence they all still hold their gold. Let $N_1^{\text{end}}, N_2^{\text{end}}, N_3^{\text{end}} \leq 2$ be the number of endpoints of P that still hold gold and represent 1, 2, and 3 symbols, respectively. Similarly, let N_0^{end} be the number of endpoints without gold. By our invariant, the symbols or components represented by the endpoints without gold left have already been sufficiently paid before. We

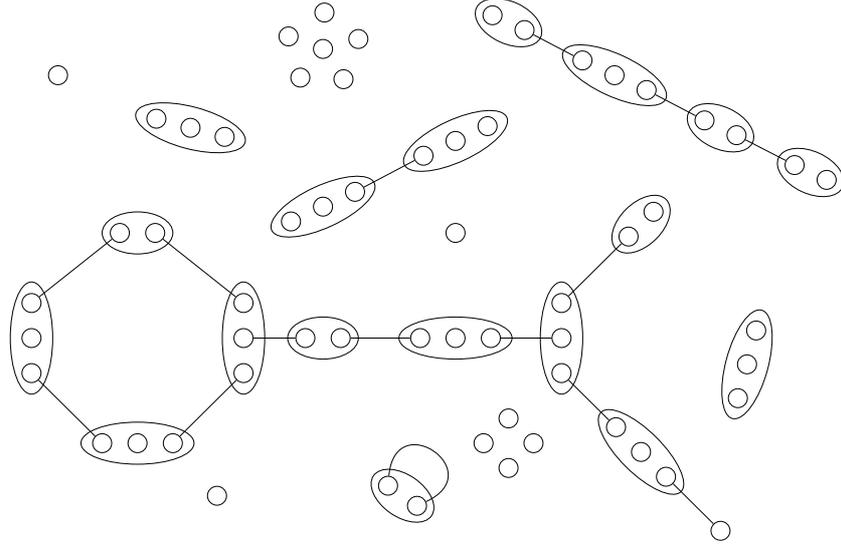


Figure 1: Illustration of the graph H that has as its nodes the components of sizes 2 and 3 in G^* and all symbols that appear in components of other sizes. An edge between symbols corresponds to a tile in \mathcal{T}_2 .

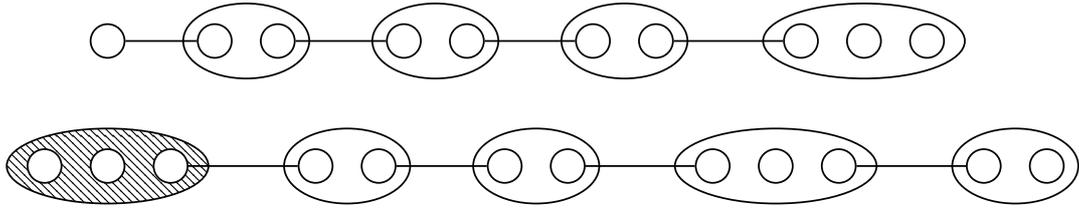


Figure 2: Illustration of our procedure for paying all symbols represented by the nodes along a path in graph H . Shaded components have been sufficiently paid for previously. For the top path we have $N_1^{\text{end}} = N_3^{\text{end}} = 1$, $N_2 = 3$, $N_0^{\text{end}} = N_2^{\text{end}} = N_3 = 0$. For the bottom path we have $N_0^{\text{end}} = N_2^{\text{end}} = 1$, $N_2 = 2$, $N_3 = 1$, $N_1^{\text{end}} = N_3^{\text{end}} = 0$.

make sure that all other nodes along P are sufficiently paid. We do this by, for all tiles that form the path P , paying one of the two corresponding symbols, and, in addition, paying *every* further symbol represented by nodes along P . Note that this preserves the invariant. The total cost is

$$C^- = k + N_2^{\text{end}} + 2N_3^{\text{end}} + N_3 - N_0^{\text{end}} = 1 + N_2^{\text{end}} + 2N_3^{\text{end}} + N_2 + 2N_3 - N_0^{\text{end}}. \quad (8)$$

Using that each endpoint of P that contributes to N_1^{end} represents a symbol that is part of a component in G^* of size $i \geq 4$, we get that the gold available at this symbol is at least $\frac{4}{3} \cdot \frac{i-1}{i} \geq 1$. Hence, the gold available to us is at least

$$C^+ = \frac{4}{3}(N_2^{\text{end}} + 2N_3^{\text{end}} + N_2 + 2N_3 + \frac{3}{4}N_1^{\text{end}}). \quad (9)$$

Since $N_0^{\text{end}} + N_1^{\text{end}} + N_2^{\text{end}} + N_3^{\text{end}} = 2$, we get

$$C^+ - C^- = 1 - \frac{2}{3}N_2^{\text{end}} - \frac{1}{3}N_3^{\text{end}} + \frac{1}{3}N_2 + \frac{2}{3}N_3.$$

Hence, we have $C^+ \geq C^-$, unless $N_2^{\text{end}} = 2$ and $N_0^{\text{end}} = N_1^{\text{end}} = N_3^{\text{end}} = N_2 = N_3 = 0$, i.e. P is of length one, connecting two tiles $p_1, p_2 \in \bar{F}_2^*$ by an edge which corresponds to a

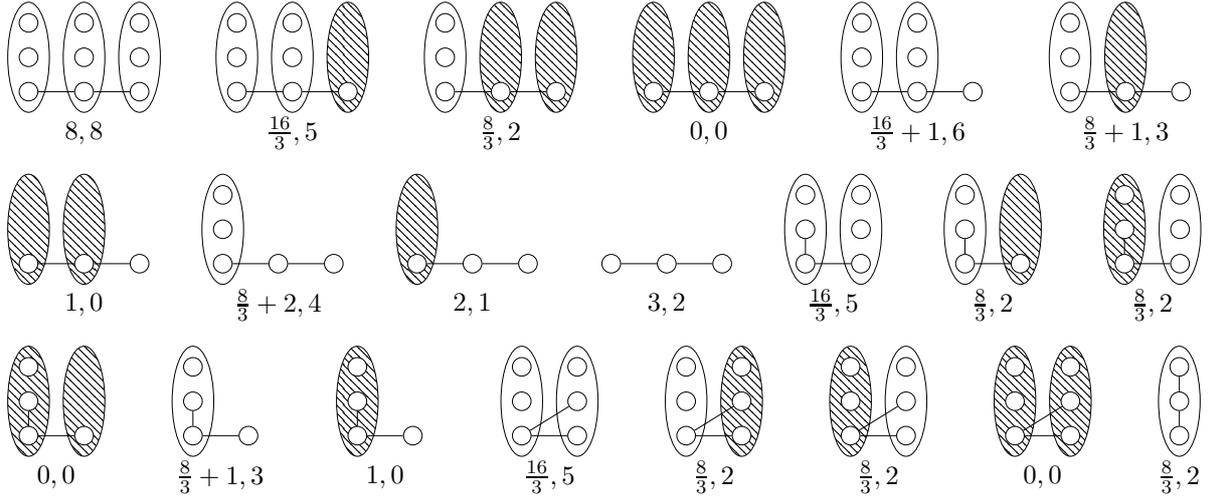


Figure 3: Possible intersections of components of $G(\mathcal{T}_3)$ (arcs) and G^* (ellipses). Shaded components have been sufficiently paid previously. Configurations are labeled by the available gold C^+ and the required gold C^- . Symmetrical configurations are omitted.

tile $t \in \mathcal{T}_2$. To see that this case cannot occur, observe that, first, p_1 and p_2 are of degree 1 in H . Second, since \mathcal{T}^* is feasible, no component of G^* is contained in a single scenario (Theorem 1), and thus $p_1, p_2 \in \bar{F}_2^* \subseteq \bar{F}_2(F)$. This is a contradiction to \mathcal{T}_2 being a maximum matching in graph $G(\bar{F}_2(F))$, as the matching can be augmented by removing t and adding p_1 and p_2 .

Similarly to the above, we can consider all cycles in H with at most one node of degree 3 one by one. (Note that cycles with at least two nodes of degree 3 contain a path as before.) If a cycle of length k does not contain a node of degree 3, or the node of degree 3 is not yet sufficiently paid (and thus still holds its gold), the cost for the cycle and its available gold are

$$C^- = k + N_3 = N_2 + 2N_3 = \frac{3}{4}C^+ < C^+,$$

where N_2, N_3 are the numbers of nodes of P that represent 2 and 3 symbols, respectively. If the node of degree 3 has no gold left, then it has already been sufficiently paid and $C^- = N_2 + 2N_3 - 3 < \frac{3}{4}C^+$. In either case, the available gold allows to sufficiently pay all nodes along the cycle. Finally, each self-loop in H connects two symbols in the same component C of size 2 or 3 in G^* . If $|C| = 2$, the gold available among the two symbols is $C^+ = \frac{4}{3}$, while we require only $C^- = 1$ unit of gold. If $|C| = 3$, we have $C^+ = \frac{8}{3}$ and $C^- = 2$.

After processing all paths, cycles, and self-loops all nodes of H intersecting a tile of \mathcal{T}_2 are sufficiently paid. In particular, since \mathcal{T}_2 is a maximum matching, all components in \bar{F}_2^* are sufficiently paid. In the next step we ensure that all components of \bar{F}_3^* are sufficiently paid. By construction, every element of \bar{F}_3^* , that is not sufficiently paid yet, intersects at least one tile of \mathcal{T}_3 . We can thus consider the components of $G(\mathcal{T}_3)$ one by one and make sure to sufficiently pay each element of \bar{F}_3^* that intersects the considered component of $G(\mathcal{T}_3)$.

Consider a component of $G(\mathcal{T}_3)$ involving the three symbols f_1, f_2, f_3 (cf. Figure 3 in the following). Let $\mathcal{C}_3 \subseteq \bar{F}_3^*$ be the set of components of size 3 in G^* that involve at least one of these symbols and have not yet been sufficiently paid (i.e., still hold their gold). Further, let N_n be the number of symbols among $\{f_1, f_2, f_3\} \cap F_{\geq 4}^*$ that are not yet sufficiently paid. Since all components in \bar{F}_2^* are sufficiently paid, the gold we have available is at least $C^+ \geq \frac{4}{3}(2|\mathcal{C}_3| + \frac{3}{4}N_n)$. We ensure that (at least) two symbols among f_1, f_2, f_3 are paid, as well as all other symbols appearing in \mathcal{C}_3 . In this way, each component in \mathcal{C}_3 is sufficiently paid. Note that this preserves our

invariant that each element of $\bar{F}_2^* \cup \bar{F}_3^* \cup F_{\geq 4}^*$ is either sufficiently paid, or still holds its gold. The cost for paying the symbols f_1, f_2, f_3 is at most 2. Since in addition to f_1, f_2, f_3 there are $3|\mathcal{C}_3| + N_n - 3$ symbols needing pay in $\bigcup_{C \in \mathcal{C}_3} C \cup \{f_1, f_2, f_3\}$, and because $|\mathcal{C}_3| \leq 3$, the total cost is

$$C^- \leq 3|\mathcal{C}_3| + N_n - 1 \leq \frac{8}{3}|\mathcal{C}_3| + N_n \leq C^+.$$

At this point, we have sufficiently paid all components in $\bar{F}_2^* \cup \bar{F}_3^*$ using gold only from these components. This means that all remaining symbols that are not sufficiently paid yet have at least $\frac{4}{3} \cdot \frac{4-1}{4} = 1$ gold available, which we can use to pay these symbols themselves. Now all elements of $\bar{F}_2^* \cup \bar{F}_3^* \cup F_{\geq 4}^*$ have been sufficiently paid and the proof is complete. \square

Our analysis of Algorithm *A* is tight in three different spots: (i) A path of length 1 in the graph H defined above that visits a component of size 2 and a component of size 3 of the optimum solution \mathcal{T} may lead to 4 tiles in our solution compared to the 3 tiles required in the optimum solution, i.e., Equations (8) and (9) coincide if $N_2^{\text{end}} = N_3^{\text{end}} = 1$ and all other terms vanish. (ii) The first intersection of a component of $G(\mathcal{T}_3)$ with components of G^* illustrated in Figure 3 may lead to 8 tiles in our solution compared to the 6 tiles required in the optimum solution. (iii) Each symbol of a component of size 4 in G^* might result in a single tile for this symbol only, in which case the optimum solution requires 3 tiles for the symbols of the component, while our solution requires 4 tiles. To improve Algorithm *A* we have to address each of these three bottlenecks. For (i), we either would have to alter the matching \mathcal{T}_2 to prevent the described situation, or combine the analysis to account for the loss in other places. The aspect (ii) can easily be prevented by employing a more sophisticated set packing algorithm (e.g., the $(4/3 + \varepsilon)$ -approximation of Cygan [6]). Finally, to avoid (iii), we would need to pack sets of size 4 similarly to our packing of sets of size 3. In addition to requiring one more level of analysis, this would also complicate the other levels, as we would have to include sets of size 4 in our reasoning there.

5 Bounded number of scenarios

In this section, we prove that MINIMUM FEASIBLE TILESET can be solved in polynomial time when the number $|\mathcal{S}|$ of scenarios is some constant. For convenience, for the course of this section, we switch to the decision variant of MINIMUM FEASIBLE TILESET. That is, we equip each instance (F, \mathcal{S}) of MINIMUM FEASIBLE TILESET with an additional integer ℓ , and we ask, whether there is a feasible tileset for \mathcal{S} with at most ℓ tiles. Clearly, solving the decision variant in polynomial time implies that also the optimization variant is solvable in polynomial time. We provide an algorithm that solves any instance (F, \mathcal{S}, ℓ) in time $f(|\mathcal{S}|)|(F, \mathcal{S}, \ell)|^c$, i.e., in time $\mathcal{O}(|(F, \mathcal{S}, k)|^c)$ for bounded values of $|\mathcal{S}|$. In other words, MINIMUM FEASIBLE TILESET is fixed-parameter tractable with respect to the number of scenarios.

Our algorithm works by first translating the input instance (F, \mathcal{S}, ℓ) into an integer linear program (ILP) in such a way that the ILP is feasible (i.e., contains at least one integer point) if and only if (F, \mathcal{S}, ℓ) admits a feasible tileset with at most ℓ tiles. The ILP uses $\mathcal{O}(|\mathcal{S}|^{|\mathcal{S}|})$ variables. Lenstra [22] proved that deciding feasibility of any ILP is fixed-parameter tractable with respect to the number of variables; the currently fastest algorithm was obtained by Frank and Tardos [14], modifying an algorithm by Kannan [20].

Theorem 4 (Frank and Tardos [14]). *In $\mathcal{O}^*(p^{\mathcal{O}(p)})$ time we can decide whether a given ILP with p variables is feasible.*

Using this, we can prove the following result.

Theorem 5. MINIMUM FEASIBLE TILESET *on instances with at most k scenarios can be solved in time $\mathcal{O}^*(k^{\mathcal{O}(k^{k+1})})$.*

Intuitively, a bounded number of scenarios also implies a bound on the number of different subsets of scenarios in which a tile can appear. Thus, one would like to forget the actual identities of the symbols and only remember *how many* symbols appear, say, exactly in scenarios S_1 , S_5 , and S_6 . It appears, however, that grouping symbols in this way is insufficient since symbols from the same group can nevertheless have different patterns for how they are provided by tiles: E.g., one tile could provide such a symbol in all three scenarios S_1 , S_5 , and S_6 , whereas other symbols of the same group might need three separate tiles for S_1 , S_5 , and S_6 . To cope with this, the constructed ILP has separate variables for all partitions of scenario subsets as well as variables for all ways of using a tile (recall that a tile has two symbols, meaning that it has two disjoint subsets of the scenario that express when either symbol is provided by the tile).

Proof of Theorem 5. We formulate MINIMUM FEASIBLE TILESET as an ILP and employ Kannan’s algorithm. Intuitively, each tile contributes both of its symbols to different (disjoint) subsets of the scenarios. For example, if we have 5 scenarios, a tile might contribute one of its symbols to scenarios 1 and 4, the other to scenarios 3 and 5, and neither to scenario 2. Each tile is associated with such a pattern of how it contributes to scenarios, and one part of the variables of our ILP track the number of tiles having each of the possible patterns. On the other hand, each symbol has a pattern associated with it, depending on which occurrences of the symbol are provided by the same tile. In our example, a symbol appearing in scenarios 1, 2, and 4 might be provided by the same tile in scenarios 1 and 4, and by a different tile in scenario 2. The remaining variables of the ILP track the number of symbols having each of the possible patterns. We provide exchange arguments to show that enforcing correct totals for these variables by linear constraints is sufficient to ensure that a feasible assignment of tiles to symbols exists for each scenario.

ILP formulation. To make our description precise, let an instance (F, \mathcal{S}, ℓ) with k scenarios $\mathcal{S} = \{S_1, \dots, S_k\}$ be given. For brevity, we refer to a subset of \mathcal{S} by the corresponding index set. For every subset $I \subseteq [k]$ of scenarios we count the number of symbols that occur exactly in these scenarios and denote this number by $c_I = |\bigcap_{i \in I} S_i \setminus \bigcup_{i \notin I} S_i|$. The family of all partitions of I is denoted by $\Pi(I)$. The ILP is constructed as follows.

1. For each set $I \subseteq [k]$ and each partition $\mathcal{I} = \{I_1, \dots, I_s\} \in \Pi(I)$ we introduce a variable $y_{\mathcal{I}}$. The intention is that variable $y_{\mathcal{I}}$ counts the number of symbols that occur (exactly) in scenarios $I := I_1 \cup \dots \cup I_s$ and have pattern \mathcal{I} associated with them, in the following way: Exactly s tiles, say, T_1, \dots, T_s , are used for such a symbol and the symbol is provided by tile T_i in the scenarios I_i .

For each I we add a constraint that enforces the total number of patterns to equal the number c_I of symbols that occur in the scenarios I :

$$c_I = \sum_{\mathcal{I} \in \Pi(I)} y_{\mathcal{I}} \quad \forall I \subseteq [k].$$

For example, if $I = \{1, 2, 3\}$, the following variables are created:

$$y_{\{\{1,2,3\}\}}, y_{\{\{1,2\},\{3\}\}}, y_{\{\{1,3\},\{2\}\}}, y_{\{\{2,3\},\{1\}\}}, y_{\{\{1\},\{2\},\{3\}\}}.$$

The number of y -variables equals the number of subpartitions of the set $[k]$. This is upper bounded by $k^k + 1$: We can k -color all subpartitions other than the partition into

singletons by using color k for all unused elements and colors $1, \dots, k-1$ for the elements of each set in the partition (only the partition into singletons has k sets). Thus, we get an injective mapping of all but one subpartition into the k colorings of $[k]$; this gives a total of $k^k + 1$.

- For the tiles, we introduce variables $x_{I,J}$ for all $I, J \subseteq [k]$ with $I \cap J = \emptyset$ and $I \cup J \neq \emptyset$; for convenience we identify $x_{I,J} = x_{J,I}$. Intuitively, the variable $x_{I,J}$ stands for the number of tiles that provide one of their symbols for scenarios I and the other symbol for scenarios J . For example, for $k = 3$ we create the following variables:

$$\begin{aligned} &x_{\emptyset,\{1\}}, x_{\emptyset,\{2\}}, x_{\emptyset,\{3\}}, x_{\emptyset,\{1,2\}}, x_{\emptyset,\{1,3\}}, x_{\emptyset,\{2,3\}}, x_{\emptyset,\{1,2,3\}}, \\ &x_{\{1\},\{2\}}, x_{\{1\},\{3\}}, x_{\{1\},\{2,3\}}, x_{\{2\},\{3\}}, x_{\{2\},\{1,3\}}, x_{\{3\},\{1,2\}}. \end{aligned}$$

The number of x -variables is $\frac{3^k-1}{2}$ corresponding to all partitions of $[k]$ into three sets (i.e., I , J , and $[k] \setminus (I \cup J)$), without $I = J = \emptyset$, and identifying $x_{I,J}$ with $x_{J,I}$.

We add constraints that enforce that the number of tiles of each pattern match the sum of the corresponding y -variables. Concretely, we add

$$\sum_{\substack{I \subseteq J \subseteq [k] \\ \mathcal{J} \in \Pi(J) \\ I \in \mathcal{J}}} y_{\mathcal{J}} = \sum_{J \subseteq [k] \setminus I} x_{I,J} \quad \forall I \subseteq [k], I \neq \emptyset. \quad (10)$$

We compare the number of tiles that provide one of their symbols for scenarios in I with the number of symbols that have I in their pattern. For the set of scenarios J such symbols appear in we must have $I \subseteq J \subseteq [k]$, and we need partitions $\mathcal{J} \in \Pi(J)$ that contain I .

- As a final constraint we enforce that the total number of used tiles is no more than ℓ . To this end, we simply sum over all x -variables and add

$$\frac{1}{2} \sum_{\substack{I, J \subseteq [k] \\ I \cap J = \emptyset \\ I \cup J \neq \emptyset}} x_{I,J} \leq \ell.$$

This completes our construction. We use $p \leq k^k + 1 + \frac{3^k-1}{2} = \mathcal{O}(k^k)$ variables and, thus, Kannan's algorithm decides feasibility of our ILP in time $\mathcal{O}^*(p^{\mathcal{O}(p)}) = \mathcal{O}^*((k^k)^{\mathcal{O}(k^k)}) = \mathcal{O}^*(k^{\mathcal{O}(k^{k+1})})$.

Correctness. First assume that the given instance (F, \mathcal{S}, ℓ) of MINIMUM FEASIBLE TILESET admits a feasible tiling \mathcal{T} of minimum cardinality $|\mathcal{T}| \leq \ell$. Since \mathcal{T} is feasible for each scenario $S_i \in \mathcal{S}$, we may let $\varphi_i: S_i \rightarrow \mathcal{T}$ be an injective function that assigns each symbol in S_i a unique tile in \mathcal{T} that can provide it. We specify feasible values for the x - and y -variables.

- x -variables.* Each tile $T \in \mathcal{T}$ has two symbols, say, $T = \{s, s'\}$, and, hence, for each $i \in [k]$ it is the image of at most one of s and s' . Formally, let

$$\begin{aligned} I &:= \{i \in [k] \mid \varphi_i(s) = T\}, \\ J &:= \{j \in [k] \mid \varphi_j(s') = T\}. \end{aligned}$$

That is, the set I contains all scenarios for which tile T provides symbol s , and J is the analogue for symbol s' . Since the functions φ_i are injective, we must have that $I \cap J = \emptyset$.

We have $I \cup J \neq \emptyset$ as otherwise T would not be used for any scenario, contradicting the minimality of \mathcal{T} . We say that tile T has *pattern* $\{I, J\}$.

For each $I, J \subseteq [k]$ with $I \cap J = \emptyset$ and $I \cup J \neq \emptyset$, we set $x_{I,J}$ to the number of tiles with pattern $\{I, J\}$. Clearly, the constraint forcing the total value of the x -variables to be at most ℓ is fulfilled since $|\mathcal{T}| \leq \ell$.

2. *y-variables*. Similarly to the tiles in \mathcal{T} we determine a pattern for each symbol $s \in F$. We let $\mathcal{T}(s) := \{T \in \mathcal{T} \mid \exists i \in [k] : \varphi_i(s) = T\} = \{T_1, \dots, T_r\}$, i.e., the set of tiles that provide s in at least one scenario. Let $I \subseteq [k]$ be the set of scenarios containing s . We define a partition $\{I_1, \dots, I_r\}$ of I by

$$I_p := \{i \in [k] \mid \varphi_i(s) = T_p\},$$

for all $p \in [r]$. We say that symbol s has *pattern* $\{I_1, \dots, I_r\} \in \Pi(I)$.

For each $I \subseteq [k]$ and each partition $\mathcal{I} \in \Pi(I)$ we set $y_{\mathcal{I}}$ to the number of symbols in F with pattern \mathcal{I} . Clearly, this fulfills the constraint that all y -variables whose pattern is a partition of some set $I \subseteq [k]$ equals the total number c_I of symbols that occur exactly among the scenarios in I .

It remains to verify that the constraint relating x - and y -variables is satisfied. To this end, let us fix some $I \subseteq [k]$, $I \neq \emptyset$, and consider the constraint

$$\sum_{\substack{I \subseteq J \subseteq [k] \\ \mathcal{J} \in \Pi(J) \\ I \in \mathcal{J}}} y_{\mathcal{J}} = \sum_{J \subseteq [k] \setminus I} x_{I,J}.$$

For each tile $T \in \mathcal{T}$ that contributes to the right-hand-side, there must be a unique symbol s in F , such that $\varphi_i(s) = T$ if and only if $i \in I$. For this symbol, we have $T \in \mathcal{T}(s)$, the set of scenarios J containing s satisfies $I \subseteq J \subseteq [k]$, and I is part of the pattern of s . Hence, s contributes to the left-hand-side. Conversely, if s is a symbol contributing to the left-hand-side, then I must be part of the pattern of s . This means that there is a unique tile $T \in \mathcal{T}$, such that $\varphi_i(s) = T$ if and only if $i \in I$. This tile has I in its pattern and thus contributes to the right-hand-side. Overall, the contribution to both sides is equal, and our assignment to x - and y -variables is feasible, as claimed.

Now, assume that the ILP constructed from (F, \mathcal{S}, ℓ) is feasible and fix a feasible assignment to the x - and y -variables. We derive a feasible tiling for all scenarios in \mathcal{S} . The set of all symbols can be partitioned according to the scenarios $I \subseteq [k]$ that each symbol appears in. The total count c_I of symbols in I is matched by the sum of y -variables that are indexed by the partitions $\mathcal{I} \in \Pi(I)$. We arbitrarily assign to each symbol with scenario set I a pattern $\mathcal{I} \in \Pi(I)$ under the sole constraint that the total number of symbols with pattern \mathcal{I} matches the corresponding variable $y_{\mathcal{I}}$. For a symbol with assigned pattern $\mathcal{I} = \{I_1, \dots, I_r\}$ the intention is to use r tiles T_1, \dots, T_r that are each responsible for one set $I_p \in \mathcal{I}$.

We will use a number of tiles that exactly matches the sum of x -variables, and thereby ensure that the final tiling has cardinality at most ℓ . We do not pick symbols for each tile but, according to the x -variables, we pick for each tile two disjoint sets of scenarios in which its two symbols will be used. Concretely, exactly $x_{I,J}$ tiles will be used in I -scenarios for one symbol and in J scenarios for their other symbol, i.e., we use $x_{I,J}$ tiles of pattern $\{I, J\}$. Recall that $I \cap J = \emptyset$ and that the sum of these variables does not exceed the maximum number of allowed tiles ℓ .

Finally, we assign symbols to tiles according to symbol and tile patterns in a canonical way. Specifically, symbols whose pattern contains some fixed $I \subseteq [k]$ are assigned to tiles that contain I in their pattern. By constraint (10) the number of symbols and the number of tiles are equal. Note that each tile is used for two disjoint sets $I, J \subseteq [k]$ and each variable $x_{I,J}$ appears in two (10)-constraints (for I and for J). Thus, each tile with pattern $\{I, J\}$ is assigned two symbols, one requiring the tile for the scenarios in I and the other requiring it the ones in J . Similarly, a symbol with pattern $\mathcal{I} = \{I_1, \dots, I_r\}$ contributes to r constraints (10), one for each I_1, \dots, I_r . Accordingly, these constraints enforce the correct sum of the corresponding variables $x_{I_1, \dots, I_r, \cdot}$. (Recall that we identified $x_{I,J}$ with $x_{J,I}$.)

We argue that the constructed tileset is indeed feasible for all scenarios $S_i \in \mathcal{S}$. Consider any symbol $s \in S_i$ with pattern \mathcal{J} . Since s appears in S_i , we have $i \in I \in \mathcal{J}$ for some set I . By the above, we know that there is a tile T containing s that has I as a part of its pattern $\{I, J\}$. Since, by definition, $I \cap J = \emptyset$, we have $i \notin J$ and may safely use T for symbol s in scenario S_i . \square

6 Bounded number of symbols

We now analyze the influence of the number of symbols $|F|$ on the complexity of solving an instance (F, \mathcal{S}, ℓ) of the decision variant of MINIMUM FEASIBLE TILESET. (That is, as in section Section 5, we want to decide whether there is a feasible tileset for \mathcal{S} with at most ℓ tiles.) It is easy to see that the problem becomes solvable in polynomial time when F is bounded: The instance is trivial if $\ell \geq |F|$ since, in that case, we can afford to dedicate a separate tile for each symbol. Otherwise, there are only $\mathcal{O}(|F|^{2\ell}) \subseteq \mathcal{O}(|F|^{2|F|})$ ways to fix ℓ tiles. As mentioned in Section 3, each candidate tileset can be verified by solving a bipartite matching problem for each scenario, on a graph that has an edge between each symbol in the scenario and every tile containing that symbol. This yields an overall runtime of $\mathcal{O}^*(|F|^{2|F|})$, and, hence, fixed-parameter tractability in $|F|$. Using structural insights of Section 2 we are able to improve on this naive running time.

Theorem 6. *Instances (F, \mathcal{S}, ℓ) of the decision variant of MINIMUM FEASIBLE TILESET can be solved in time $\mathcal{O}^*(3^{|F|})$.*

Note that, as every symbol occurs in a scenario, $\ell \geq |F|/2$. Hence, Theorem 6 gives a fixed-parameter algorithm also for parameter ℓ .

Proof of Theorem 6. We describe a dynamic programming algorithm for solving an instance (F, \mathcal{S}, ℓ) . Recall that we may assume $\ell < |F|$; otherwise the instance is trivial. Our algorithm uses a table M of size $2^{|F|}$ that is indexed by subsets $D \subseteq F$, with each entry taking integer values from $[|F|] \cup \{-\infty\}$. At the end of the computation, each entry $M(D)$ will be set to $-\infty$ if $D \subseteq S$ for some scenario $S \in \mathcal{S}$, and otherwise to the maximum integer $i \in [0, |F|]$ for which there is a partition of D into i sets D_1, \dots, D_i such that no scenario contains any set in $\{D_1, \dots, D_i\}$ as a subset.

In the end, by Theorem 1, the entry $M(F)$ contains the maximum number of components in the graph corresponding to a feasible tileset. Accordingly, every corresponding tileset \mathcal{T} has minimum cardinality. Hence, and since each connected component C in the graph (F, \mathcal{T}) is composed of $|C| - 1$ tiles, the instance (F, \mathcal{S}, ℓ) admits a tileset of size ℓ if and only if $M(F) \geq |F| - \ell$.

We fill out the entries of the table in order of increasing subset sizes. Each entry is computed via the following recurrence. (Note that the 1 in the maximum taken over subsets D' of D stands

for the trivial partition of D into just one set. This is the best value in case that no split into at least two sets can be found such that both sets are not subsets of scenarios.)

$$M(D) = \begin{cases} -\infty, & \text{if } D \subseteq S \text{ for some } S \in \mathcal{S}, \\ \max_{\substack{D' \subset D \\ 2 \leq |D'| \leq |D|/2}} \{1, M(D') + M(D \setminus D')\}, & \text{otherwise.} \end{cases}$$

Thus, for each $D \subseteq F$ that is not a subset of a scenario we need to compute the maximum of $M(D') + M(D \setminus D')$ over less than $2^{|D|}$ subsets D' of D . By the well-known binomial theorem the total number of evaluations taken over all $D \subseteq F$ can be upper bounded by $3^{|F|}$ giving us the claimed runtime. \square

After this fixed-parameter tractability result, and taking into account the trivial bound of $2^{|F|}$ for the number of scenarios (giving a worst-case size of instances of $\mathcal{O}(2^{|F|}|F|)$), it is natural to ask whether polynomial-time preprocessing can simplify input instances to size polynomial in $|F|$. We show that this is impossible unless $\text{NP} \subseteq \text{coNP/poly}$ (and the polynomial hierarchy collapses). More generally, we prove that for the restricted case d -MINIMUM FEASIBLE TILESET, where scenarios have size at most d , no polynomial-time algorithm can achieve a size of $\mathcal{O}(k^{d-\varepsilon})$. Note that this restricted case has an essentially matching upper bound of $|\mathcal{S}| < (|F| + 1)^d = \mathcal{O}(|F|^d)$.¹ As a consequence there is no reduction to size polynomial in $|F|$ for the general MINIMUM FEASIBLE TILESET problem: Any size $\mathcal{O}(k^c)$ preprocessing for MINIMUM FEASIBLE TILESET could be used for d -MINIMUM FEASIBLE TILESET, for any $d > c$, and violate the lower bound.

Theorem 7. *Let $d \geq 3$ and ε be a positive real. There is no polynomial-time algorithm that reduces every instance of d -MINIMUM FEASIBLE TILESET to an equivalent instance (possibly of a different problem) of size $\mathcal{O}(|F|^{d-\varepsilon})$, unless $\text{NP} \subseteq \text{coNP/poly}$.*

To prove Theorem 7 we employ a similar result by Dell and Marx [7] for EXACT COVER BY d -SETS, which is defined as follows.²

EXACT COVER BY d -SETS

Input: A universe X and a family \mathcal{C} of d -element sets $C \in \binom{X}{d}$.

Problem: Is there an *exact d -set cover* for X , i.e., a partition of X into a family $\mathcal{C}' \subseteq \mathcal{C}$ of disjoint sets?

Note that the original result by Dell and Marx [7] is given in terms of the size k of an exact d -set cover. Clearly, $k = \frac{|U|}{d}$ and, thus, we have $\mathcal{O}(k^{d-\varepsilon}) = \mathcal{O}(|U|^{d-\varepsilon})$ and may instead phrase the result in terms of $|U|$. Furthermore, their result builds on work by Dell and van Melkebeek [8] and, thus, extends to any polynomial time algorithms (rather than just problem kernels as mentioned there) whose output instances can be with respect to a different problem. We give the following paraphrased version of the result.

Theorem 8 (Dell and Marx [7]). *Let $d \geq 3$ and ε be a positive real. There is no polynomial-time algorithm that reduces every instance (U, \mathcal{H}) of EXACT COVER BY d -SETS to an equivalent instance of size $\mathcal{O}(|U|^{d-\varepsilon})$ (possibly with respect to a different problem), unless $\text{NP} \subseteq \text{coNP/poly}$.*

The following lemma, together with Theorem 8, directly implies Theorem 7.

¹A compression to $\mathcal{O}(|F|^d)$ size can be achieved by specifying one bit for each possible scenario in \mathcal{S} and setting it to one if the scenario is present and zero otherwise.

²Dell and Marx called this problem PERFECT d -SET MATCHING.

Lemma 5. *There is a polynomial-time reduction from EXACT COVER BY d -SETS to MINIMUM FEASIBLE TILESET such that instances (X, \mathcal{C}) are mapped to instances (F, \mathcal{S}, ℓ) with $F = X$ and scenario size at most d .*

Proof sketch of Lemma 5. The proof is similar to the proof of NP-hardness in Theorem 2. Given an instance of EXACT COVER BY d -SETS with universe X and a family \mathcal{C} we construct an instance $(F, \mathcal{S}, (d-1)n/d)$ of MINIMUM FEASIBLE TILESET with $F = X$ and $n = |X|$. Applying the equivalence of finding a feasible tileset of size $(d-1)n/d$ and finding an admissible partition for \mathcal{S} of size n/d then gives Lemma 5.

To construct the instance of MINIMUM FEASIBLE TILESET, we simply set $\mathcal{S} = \binom{X}{d-1} \cup (\binom{X}{d} \setminus \mathcal{C})$. Similarly to the reduction used for Theorem 2, the scenarios $\binom{X}{d-1}$ enforce that every admissible partition contains only parts of size exactly d . The constraints $\binom{X}{d} \setminus \mathcal{C}$ enforce that only sets of \mathcal{C} occur in an admissible partition. Hence, each admissible partition is also an exact d -set cover and vice-versa. \square

We now consider a more general setting: In the GENERALIZED MINIMUM FEASIBLE TILESET problem we are also given a set of symbols and a set of scenarios, but here each scenario may be a *multi-set* of symbols (or, equivalently, each scenario is a function $S: F \rightarrow \mathbb{N}$ indicating the number of copies of each symbol f needed for S). We prove that GENERALIZED MINIMUM FEASIBLE TILESET can be solved in time $\mathcal{O}^*(|F|^{\mathcal{O}(|F|^2)})$. Note that for this problem the solution size ℓ may be much larger than $|F|$ and similarly the number of scenarios cannot in general be bounded in $|F|$.

Theorem 9. GENERALIZED MINIMUM FEASIBLE TILESET can be solved in time $\mathcal{O}^*(|F|^{\mathcal{O}(|F|^2)})$, i.e., it is fixed-parameter tractable with respect to $|F|$.

Proof. Let (F, \mathcal{S}, ℓ) be an instance of GENERALIZED MINIMUM FEASIBLE TILESET and let $k := |F|$. We will construct an integer linear program (ILP) with $\binom{k}{2} = \mathcal{O}(k^2)$ variables and $\mathcal{O}^*(2^k)$ constraints that is feasible if and only if (F, \mathcal{S}, ℓ) admits a feasible tileset with at most ℓ tiles. Using Kannan's algorithm (Theorem 4) then completes the proof.

We introduce one variable $x_{s,s'} \geq 0$ for each possible tile type, i.e., for each pair of symbols $s, s' \in \binom{F}{2}$. We interpret $x_{s,s'}$ as the number of tiles of type s, s' that the solution will contain. We begin with the constraint ensuring that we do not use more than ℓ tiles overall:

$$\sum_{\{s,s'\} \in \binom{F}{2}} x_{s,s'} \leq \ell$$

We need to add constraints to the ILP to ensure that the resulting assignment to the $x_{s,s'}$ -variables corresponds to a feasible tileset, i.e., that each scenario S can be implemented using the corresponding numbers of tiles of each type. This is the case if and only if there is a matching from the symbols in S to the tiles that cover all symbols in S . Clearly, in order not to use too many variables, we do not want to compute a (one-sided perfect) matching for each scenario S . By Hall's Theorem, it is instead sufficient to ensure that for each subset $I \subseteq F$ of symbols appearing at least once in scenario S there are at least that many tiles involving these symbols. If $c_{s,S}$ denotes the number of occurrences of symbol s in scenario S , we obtain the following constraints:

$$\sum_{\{s,s'\} \cap I \neq \emptyset} x_{s,s'} \geq \sum_{s \in I} c_{s,S} \quad \forall S \in \mathcal{S}, \forall I \subseteq F$$

In total we use $\binom{k}{2} = \mathcal{O}(k^2)$ variables and $1 + m \cdot 2^k + \binom{k}{2}$ constraints. Using Kannan’s algorithm for testing feasibility of an ILP with p variables in time $\mathcal{O}^*(p^{\mathcal{O}(p)})$ (Theorem 4) we get a total running time of $\mathcal{O}^*(k^{\mathcal{O}(k^2)})$. \square

7 Conclusion

We initiated the study of the MINIMUM FEASIBLE TILESET problem and exposed an interesting combinatorial structure. We proved the problem to be NP-complete even in the restricted case with scenarios of size at most three and APX-hard in general. On the positive side, we showed that the MINIMUM FEASIBLE TILESET problem admits a $4/3$ -approximation algorithm and that it is fixed-parameter tractable with respect to the number of scenarios and number of symbols. The latter algorithm works also for the GENERALIZED MINIMUM FEASIBLE TILESET problem where each scenario can contain multiple copies of a symbol and we believe that it can be further generalized to work also for the original assignment problem where also tiles of larger (but constant) size are allowed. It would be interesting to see whether our other positive results transfer to this more general setting. We note that our approximation algorithm relies heavily on the structural observations from Section 2 which do not seem to generalize well. Our integer linear program for a fixed number of scenarios does not seem easily adaptable either.

References

- [1] N. Bansal, A. Caprara, and M. Sviridenko. A new approximation method for set covering problems, with applications to multidimensional bin packing. *SIAM Journal on Computing*, 39(4):1256–1278, 2009.
- [2] T. Biedl, T. Chan, Y. Ganjali, M. Hajiaghayi, and D. Wood. Balanced vertex-orderings of graphs. *Discrete Applied Mathematics*, 148(1):27–48, 2005.
- [3] K. Buchin, M. J. van Kreveld, H. Meijer, B. Speckmann, and K. Verbeek. On planar supports for hypergraphs. *Journal of Graph Algorithms and Applications*, 15(4):533–549, 2011.
- [4] J. Chen, C. Komusiewicz, R. Niedermeier, M. Sorge, O. Suchý, and M. Weller. Polynomial-time data reduction for the subset interconnection design problem. *SIAM Journal on Discrete Mathematics*, 29(1):1–25, 2015.
- [5] P. Crescenzi. A short guide to approximation preserving reductions. In *Proceedings of the Twelfth Annual IEEE Conference on Computational Complexity (CCC)*, pages 262–273. IEEE Computer Society, 1997.
- [6] M. Cygan. Improved approximation for 3-dimensional matching via bounded pathwidth local search. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 509–518, 2013.
- [7] H. Dell and D. Marx. Kernelization of packing problems. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 68–81, 2012.
- [8] H. Dell and D. van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *Journal of the ACM*, 61(4):23:1–23:27, 2014.

- [9] Y. Disser, S. Kratsch, and M. Sorge. The minimum feasible tileset problem. In *Proceedings of the 12th Workshop on Approximation and Online Algorithms (WAOA '14)*, volume 8952 of *LNCS*, pages 144–155. Springer, 2014.
- [10] Y. Disser and J. Matuschke. Degree-constrained orientations of embedded graphs. *Journal of Combinatorial Optimization*, 31(2):758–773, 2016.
- [11] D.-Z. Du and Z. Miller. Matroids and subset interconnection design. *SIAM Journal on Discrete Mathematics*, 1(4):416–424, 1988.
- [12] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [13] A. Frank and A. Gyárfás. How to orient the edges of a graph. *Colloquia mathematica societatis Janos Bolyai*, 18:353–364, 1976.
- [14] A. Frank and É. Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987.
- [15] M. R. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [16] G. Gottlob and G. Greco. Decomposing combinatorial auctions and set packing problems. *Journal of the ACM*, 60(4):24, 2013.
- [17] S. Hakimi. On the degrees of the vertices of a directed graph. *Journal of the Franklin Institute*, 279(4):290–308, 1965.
- [18] D. S. Johnson and H. O. Pollak. Hypergraph planarity and the complexity of drawing Venn diagrams. *Journal of Graph Theory*, 11(3):309–325, 1987.
- [19] V. Kann. Maximum bounded 3-dimensional matching is MAX SNP-complete. *Information Processing Letters*, 37(1):27–35, 1991.
- [20] R. Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of Operations Research*, 12:415–440, 1987.
- [21] I. Koutis. Faster algebraic algorithms for path and packing problems. In *Proceedings of the 35th International Colloquium on Automata (ICALP)*, pages 575–586, 2008.
- [22] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.
- [23] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [24] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 5th edition, 2016.
- [25] Rodney G. Downey and Micheal R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- [26] P. Schuurman and G. J. Woeginger. Approximation schemes—a tutorial.

- [27] M. Sviridenko and J. Ward. Large neighborhood local search for the maximum set packing problem. In *40th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 792–803, 2013.
- [28] R. van Bevern, I. Kanj, C. Komusiewicz, R. Niedermeier, and M. Sorge. Twins in subdivision drawings of hypergraphs. In *Proceedings of the 24th International Symposium on Graph Drawing & Network Visualization*, volume 9801 of *LNCS*, pages 67–80. Springer, 2016.