Algorithmica
https://doi.org/10.1007/s00453-018-0460-3

The Minimum Feasible Tileset Problem

Yann Disser[1] · **Stefan Kratsch**[2] · **Manuel Sorge**[3,4]

Received: 3 August 2017 / Accepted: 17 May 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

An extended abstract of this article appeared at the *12th Workshop on Approximation and Online Algorithms*, Wrocław, September 2014 [10]. 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).

✉ Manuel Sorge
sorge@post.bgu.ac.il

Yann Disser
disser@mathematik.tu-darmstadt.de

Stefan Kratsch
kratsch@informatik.hu-berlin.de

1 Institut für Mathematik, Graduate School CE, TU Darmstadt, Darmstadt, Germany

2 Institut für Informatik, Humboldt-Universität zu Berlin, Berlin, Germany

3 Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Berlin, Germany

4 Department of Industrial Engineering and Management, Ben Gurion University of the Negev, Beer Sheva, Israel

Published online: 29 May 2018

the MINIMUM FEASIBLE TILESET problem is fixed-parameter tractable both when parameterized with the number of scenarios and with the number of symbols.

Keywords Set packing · Approximation algorithm · APX hardness · Parameterized complexity · Minimum feasible tileset

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, …) [2,24,26]. 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 (motivated by board games where two-sided map tiles are often used to allow for a variety of scenarios, 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 $\mathcal{S}$.

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 $\phi\colon \mathcal{T} \to F$, 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

$$\boxed{\begin{matrix}A\\B\end{matrix}}, \boxed{\begin{matrix}B\\C\end{matrix}}, \boxed{\begin{matrix}1\\2\end{matrix}}, \boxed{\begin{matrix}2\\3\end{matrix}}.$$

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 $n \in \mathbb{N}$, $F = \{1, \ldots, 3n\}$, 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

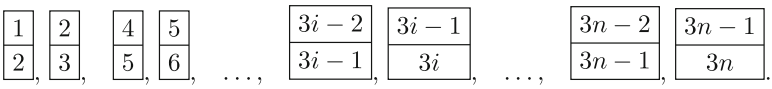

Clearly, this tileset is feasible for all scenarios $\{a, b\} \in \mathcal{S}$ where a and b do not occur together on a tile (because we can select different tiles for them). In the remaining scenarios, the symbols a and b occur together on a tile, that is, for some $i \in \{1, \ldots, n\}$ we have $a = 3i - 2$ and $b = 3i - 1$ or $a = 3i - 1$ and $b = 3i$ (up to symmetry). These scenarios are fine too, because all numbers $3i - 1$ occur twice on the tiles. It will become clear later (see Example 2), 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 Sect. 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 tileset 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 tileset (Sect. 2). This lets us view MINIMUM FEASIBLE TILESET 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 (Sect. 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 (Sect. 4). We show that the problem is fixed-parameter tractable with respect to the number of scenarios (Sect. 5) and the number of symbols (Sect. 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 loosing the information about the size of the optimum solution and such a compression to $O(|F|^{d-\epsilon})$ bits is unlikely (Sect. 6). Finally, we provide a preliminary result on the relevant variant of MINIMUM FEASIBLE TILESET 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 (Sect. 6).

Related Work The problem most closely related to MINIMUM FEASIBLE TILESET 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., [1,7,30] and [8,22] for some recent results). The main difference between the two problems is that SET PACKING is a maximization problem whereas MINIMUM FEASIBLE TILESET seeks to minimize the size of a feasible tileset—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 TILESET we get a nontrivial polynomial-time algorithm via integer linear programming (see Sect. 5).

As alluded to above, the MINIMUM FEASIBLE TILESET 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. [3] 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 [11,14,18]. Abstracting from orientations, MINIMUM FEASIBLE TILESET belongs to a class of problems in which we are given a set of symbols and a family of scenarios and we are looking for a graph with symbols as vertices that fulfills a certain constraint for each scenario. The constraint where the subgraph induced by each scenario has to be connected is well-studied [4,5,12,17,19,31]. In particular, it is NP-hard to find the minimum number of edges needed [12] and to decide whether a planar solution [4,19] or a solution of treewidth at most three [17] exists.

Preliminaries For some positive integer $\ell \in \mathbb{N}$, we denote $[\ell] := \{1, 2, \ldots, \ell\}$. 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 [27], for example. For the relevant notions from parameterized complexity and approximation complexity we refer to textbooks [13,25,28] and Refs. [6,29], 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, $\mathcal{S}$ a family of scenarios over F, and $\mathcal{T}$ a tileset feasible for $\mathcal{S}$. There is a tileset $\mathcal{T}' \subseteq \binom{F}{2}$ feasible for $\mathcal{S}$ 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 $S \subset F$ 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). In the following, we use the notation $\phi[\mathcal{T}] = \bigcup_{T \in \mathcal{T}} \phi(T)$.

(i) $\Rightarrow$ (ii): Assume that $\mathcal{T}$ is feasible for S and let $\phi\colon \mathcal{T} \to F$ 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 $C_1, \ldots, C_k$ denote the connected components of $G(\mathcal{T})$ and assume that there are feasible orientations $\vec{G}_1, \ldots, \vec{G}_k$ for $S \cap C_1, \ldots, S \cap C_k$, respectively. Since $S \subseteq \bigcup_i C_i$, we obtain a feasible orientation $\vec{G}$ of $G(\mathcal{T})$ for S as $\vec{G}_1[C_1] \uplus \cdots \uplus \vec{G}_k[C_k]$. We define the mapping $\phi : \mathcal{T} \to F$ 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. □

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 $G(\mathcal{T})$. □

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

Let $C_1, \ldots, C_k$ 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 tileset $\mathcal{T}'$ 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.) Each connected component C_i, $i \in [k]$, contains at least $|C_i|$ edges since it is connected and contains a cycle. Since we introduced $|C| \leq \sum_{i=1}^{k} |C_i|$ 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 $S \backslash C$. This implies that there is still a feasible orientation of $G(\mathcal{T}')$ for $S \backslash C$. 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 tileset $\mathcal{T}''$ 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 tileset $\mathcal{T}''$. 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 $S \backslash C$ 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 $S \backslash C$, applying Lemma 2 we obtain that $\mathcal{T}''$ is feasible for S. Clearly, $G(\mathcal{T}'')$ is a forest, as required. □

Intuitively, Lemmas 2 and 3 imply that only the partition of the symbols induced by the component structure of the graph of a tileset 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 tileset 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 "⇒": 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.

"⇐": 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 $\mathcal{S}$. □

Example 2 We can now observe that for the instance $(F = \{1, \ldots, 3n\}, \mathcal{S} = \binom{F}{2})$ 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 most n connected components. Each component contains at least $\ell - 1$ tiles if it is of size ℓ. In other words, for each connected component, we save at most one tile compared with simply using one tile for each symbol. Thus, a feasible tileset for $\mathcal{S}$ contains at least $3n - n = 2n$ 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 "⇒": 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 $C_1, \ldots, C_k$ of $G(\mathcal{T}')$ induce a partition $\mathcal{P}'$ which we claim to be admissible: Indeed, by Theorem 1 we have $C_i \nsubseteq S$ 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.

"⇐": Let $\mathcal{P} = \{P_1, \ldots, P_p\}$ 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 $P_i \nsubseteq S$ 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. □

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 [20]:

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 (Theorem 2) We give a PTAS-reduction from MAXIMUM BOUNDED 3-DIMENSIONAL MATCHING [6]. 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 $f\colon (0, 1) \to \mathbb{Q}$ such that for every $r \in (0, 1)$ we have $f(r) > 1$ and for an arbitrary given $f(r)$-approximate feasible tileset $\mathcal{T}$ 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 $g\colon (0, 1) \to \mathbb{N}$ 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, $\mathcal{S} := \{S \subseteq F \mid |S| \leq g(r) \wedge \forall E \in D\colon E \backslash S \neq \emptyset\}$. 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 tileset, 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 $\mathrm{opt}_{3\mathrm{DM}}$ denote the size of an optimal solution to the MAXIMUM BOUNDED 3- DIMENSIONAL MATCHING instance, we thus have $|\mathcal{P}^*| \geq \mathrm{opt}_{3\mathrm{DM}}$ for an admissible partition $\mathcal{P}^*$ containing the maximum number of parts. Lemma 4 implies that $|\mathcal{P}^*| = n - \mathrm{opt}_{\mathrm{FT}}$, where $\mathrm{opt}_{\mathrm{FT}}$ denotes the size of an optimal solution for the MINIMUM FEASIBLE TILESET instance. Rearranging terms hence yields $\mathrm{opt}_{\mathrm{FT}} = n - |\mathcal{P}^*|$. Because of $|\mathcal{P}^*| \geq \mathrm{opt}_{3\mathrm{DM}}$, we have

$$\mathrm{opt}_{\mathrm{FT}} \;\leq\; n - \mathrm{opt}_{3\mathrm{DM}}\,. \tag{1}$$

Now let $\mathcal{T}$ be an arbitrary $f(r)$-approximate feasible tileset $\mathcal{T}$ 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, $|\mathcal{T}| \leq f(r)\,\mathrm{opt}_{\mathrm{FT}}$ and hence, $n - |\mathcal{P}| \leq f(r)\,\mathrm{opt}_{\mathrm{FT}}$. Applying Inequality (1) we thus obtain

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

We create a partition $\mathcal{P}^1 = \mathcal{P}^1_3 \cup \mathcal{P}^1_{g(r)+1}$ from $\mathcal{P}$ as follows. Obtain $\mathcal{P}^1_3$ 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 $\mathcal{P}^1_3$. To create $\mathcal{P}^1_{g(r)+1}$, if $|F \setminus \bigcup_{P \in \mathcal{P}^1_3} P| < g(r) + 1$, then put $F \setminus \bigcup_{P \in \mathcal{P}^1_3} P$ into $\mathcal{P}^1_{g(r)+1}$ as a single set of size at most $g(r)$. In this case we call $\mathcal{P}^1_{g(r)+1}$ *degenerate*. Otherwise, put $\mathcal{P}^1_{g(r)+1}$ to be an arbitrary partition of $F \setminus \bigcup_{P \in \mathcal{P}^1_3} P$ 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 $\mathcal{P}^1_3$ is admissible for $\mathcal{S}$ because each triple in $\mathcal{P}^1_3$ 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

$$n - |\mathcal{P}^1| \;\leq\; f(r)\left(n - \mathrm{opt}_{3\mathrm{DM}}\right). \tag{3}$$

Note that the three-element sets in $\mathcal{P}^1$, i.e., $\mathcal{P}^1_3$, form a three-dimensional matching for the instance (X, Y, Z, D). The sets in $\mathcal{P}^1_3$ will form our r-approximate three-dimensional matching M after one further augmentation step. The aim of this augmentation is to bound $|\mathcal{P}^1_3|$ by a function of $|\mathcal{P}^1_{g(r)+1}|$. 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 $\bigcup_{P\in\mathcal{P}^1_3} P$, then add this triple to $\mathcal{P}^1_3$. If we now have $|F\setminus\bigcup_{P\in\mathcal{P}^1_3} P| < g(r)+1$, then replace the sets in $\mathcal{P}^1_{g(r)+1}$ by the single set $F\setminus\bigcup_{P\in\mathcal{P}^1_3} P$. Otherwise, replace $\mathcal{P}^1_{g(r)+1}$ by an arbitrary partition of $F\setminus\bigcup_{P\in\mathcal{P}^1_3} P$ 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 $\mathcal{P}^1_{g(r)+1}$ was degenerate before applying them. Otherwise, we have $|F\setminus\bigcup_{P\in\mathcal{P}^1_3} P| > g(r)+3$ before applying the modification. Hence, each set in $\mathcal{P}^1_{g(r)+1}$ had size at least $g(r)+1\geq 4$. Since we moved only three elements from these sets to $\mathcal{P}^1_3$ and afterwards repartition the remaining elements with sets of size at least 4, the total number of sets cannot decrease.

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

$$n-|\mathcal{P}^2| \;\leq\; f(r)\left(n-\mathrm{opt}_{3\mathrm{DM}}\right). \tag{4}$$

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

We now claim that $12|\mathcal{P}^2_3|/(g(r)+1)\geq|\mathcal{P}^2_{g(r)+1}|$. If $\mathcal{P}^1_{g(r)+1}$ is degenerate and if this relation does not hold, then $|F|$ is upper bounded by $3|\mathcal{P}^2_3|+g(r)-1 < 3(g(r)+1)/12+g(r)-1$, 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 $\mathcal{P}^1_{g(r)+1}$ is degenerate.

If $\mathcal{P}^1_{g(r)+1}$ is not degenerate, we claim that the above modification of $\mathcal{P}^1$ is applicable as long as $12|\mathcal{P}^1_3| < (g(r)+1)|\mathcal{P}^1_{g(r)+1}|$, 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 $\mathcal{P}^1_3$, there are at most twelve elements of F whose incident triples in D cannot be added to $\mathcal{P}^1_3$, because they overlap with E. Hence, if $12|\mathcal{P}^1_3| < (g(r)+1)|\mathcal{P}^1_{g(r)+1}|$, then there exists at least one element of F whose incident triples do not overlap with any triple in $\mathcal{P}^1_3$. This means that at least one triple will be added to $\mathcal{P}^1_3$ 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 $12|\mathcal{P}^2_3|/(g(r)+1)\geq|\mathcal{P}^2_{g(r)+1}|$. We thus have

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

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

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

Thus,

$$\frac{g(r)+1}{g(r)+13} \cdot \left((1-f(r))n + f(r)\,\mathrm{opt}_{3\mathrm{DM}}\right) \leq |M|. \tag{5}$$

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 $3\,\mathrm{opt}_{3\mathrm{DM}}$ elements—or it is in a triple together with a matched element—of this type there are at most $12\,\mathrm{opt}_{3\mathrm{DM}}$ elements because each element is in at most three triples in D. Hence, $n \leq 3\,\mathrm{opt}_{3\mathrm{DM}} + 12\,\mathrm{opt}_{3\mathrm{DM}} = 15\,\mathrm{opt}_{3\mathrm{DM}}$. Note that $1 - f(r) < 0$. Inequality (5) thus implies

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

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,

$$\begin{aligned} x &:= -rg(r) - 13r + 15g(r) + 15 - 14g(r) - 14 \\ &= -rg(r) - 13r + g(r) + 1 = (1-r)g(r) - 13r + 1. \end{aligned}$$

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,

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

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 [16]. 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. □

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 Sect. 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 $\bar{F}_i(F') = \{C \in \binom{F'}{i} \mid \forall S \in \mathcal{S}\colon C \not\subseteq S\}$ 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_{root} to be an arbitrary element of $\bigcup_{t \in \mathcal{T}_2 \cup \mathcal{T}_3} t$ **else** take f_{root} to be an arbitrary element of 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 The same modification that we 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 the 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 tileset $\mathcal{T}^\star$. 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 $\mathcal{T}^\star$. 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 $G^\star := G(\mathcal{T}^\star)$ be the graph induced by $\mathcal{T}^\star$ and $\bar{F}_i^\star$ be the set of connected components of size $i \in \{2, \ldots, |F|\}$ in $G^\star$. By Lemma 1, we may assume that $G^\star$ is a forest. Furthermore, because each symbol appears in at least one scenario, graph $G^\star$ 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 $\mathcal{T}^\star$, we may distribute all the available gold by giving $4/3 \cdot \frac{i-1}{i}$ gold to each symbol in a component of $\bar{F}_i^\star$, for all $i \in \{2, \ldots, |F|\}$. 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 $G^\star$ *sufficiently paid*, if all its symbols are sufficiently paid. Let $F^\star_{\geq 4} := F \setminus \bigcup_{C \in \bar{F}_2^\star \cup \bar{F}_3^\star} C$ be the set of all symbols not in components of size two or three in $G^\star$. In paying the symbols we will maintain the invariant that each element of $\bar{F}_2^\star \cup \bar{F}_3^\star \cup F^\star_{\geq 4}$ is

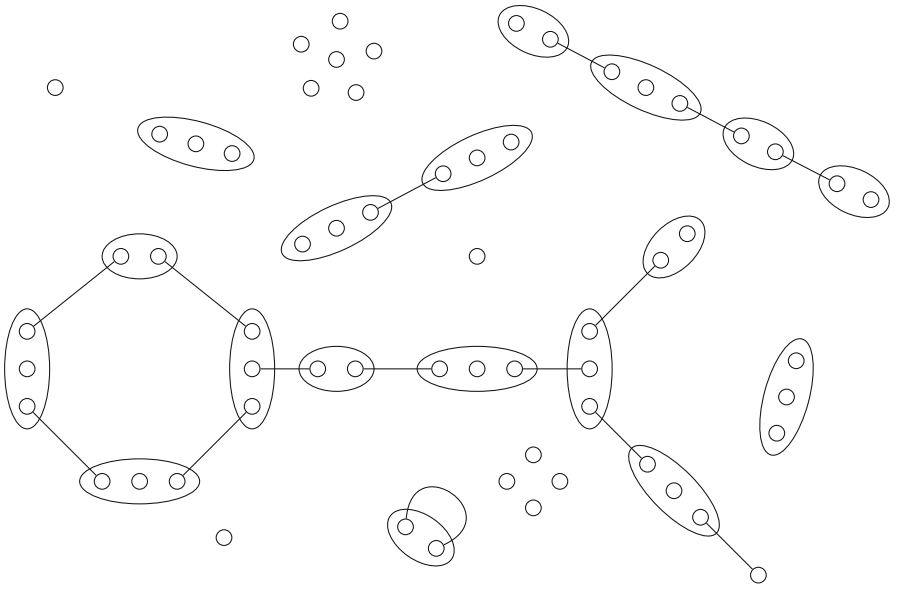

Fig. 1 Illustration of the graph H that has as its nodes the components of sizes 2 and 3 in $G^\star$ and all symbols that appear in components of other sizes. An edge between symbols corresponds to a tile in $\mathcal{T}_2$

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 $\bar{F}_2^\star \cup \bar{F}_3^\star$ as its vertices, as well as the symbols that are not part of these components, i.e., $V = \bar{F}_2^\star \cup \bar{F}_3^\star \cup F_{\geq 4}^\star$ (see Fig. 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 $G^\star$ 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 the path $P \in \mathcal{P}$ of length k as follows (cf. Fig. 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

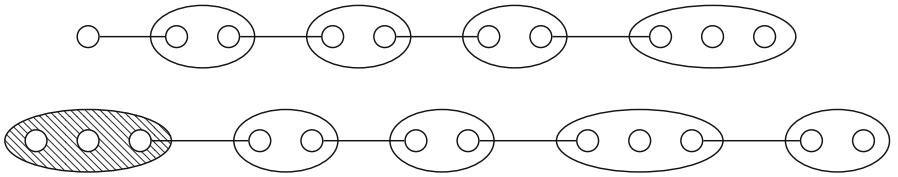

Fig. 2 Illustration of our procedure for paying all symbols represented by the nodes along a path in graph H. Shaded components have been sufficiently paid for previously. For the top path we have $N_1^{\text{end}} = N_3^{\text{end}} = 1$, $N_2 = 3$, $N_0^{\text{end}} = N_2^{\text{end}} = N_3 = 0$. For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

before. We 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

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

Using the fact that each endpoint of P that contributes to N_1^{end} represents a symbol that is part of a component in $G^\star$ 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

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

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 $p_1, p_2 \in \bar{F}_2^\star$ by an edge which corresponds to a 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 $\mathcal{T}^\star$ is feasible, no component of $G^\star$ is contained in a single scenario (Theorem 1), and thus $p_1, p_2 \in \bar{F}_2^\star \subseteq \bar{F}_2(F)$. 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

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

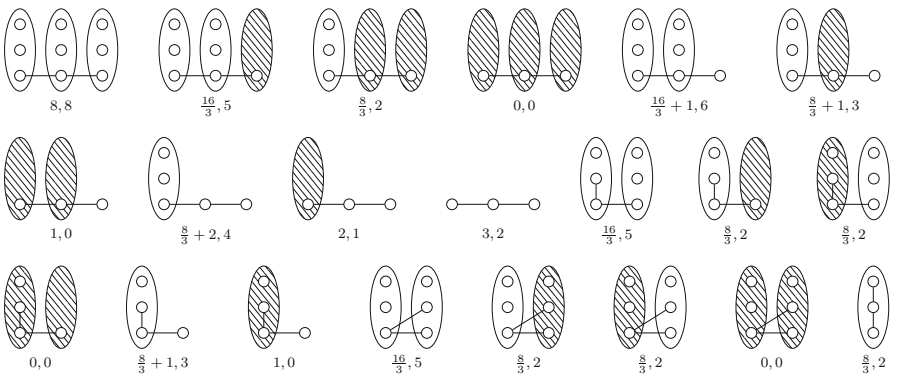

Fig. 3 Possible intersections of components of $G(\mathcal{T}_3)$ (arcs) and $G^\star$ (ellipses). Shaded components have been sufficiently paid previously. Configurations are labeled by the available gold C^+ and the required gold C^-. Symmetrical configurations are omitted

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 $G^\star$. If $|C| = 2$, then the gold available among the two symbols is $C^+ = \frac{4}{3}$, while we require only $C^- = 1$ unit of gold. If $|C| = 3$, then 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 $\bar{F}_2^\star$ are sufficiently paid. In the next step we ensure that all components of $\bar{F}_3^\star$ are sufficiently paid. By construction, every element of $\bar{F}_3^\star$, 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 $\bar{F}_3^\star$ 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. Fig. 3 in the following). Let $\mathcal{C}_3 \subseteq \bar{F}_3^\star$ be the set of components of size 3 in $G^\star$ 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 $\{f_1, f_2, f_3\} \cap F^\star_{\geq 4}$ that are not yet sufficiently paid. Since all components in $\bar{F}_2^\star$ 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 $\bar{F}_2^\star \cup \bar{F}_3^\star \cup F^\star_{\geq 4}$ 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 $\bar{F}_2^\star \cup \bar{F}_3^\star$ 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 $\bar{F}_2^\star \cup \bar{F}_3^\star \cup F_{\geq 4}^\star$ have been sufficiently paid and the proof is complete. □

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 with the 3 tiles required in the optimum solution, i.e., Eqs. (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 $G^\star$ illustrated in Fig. 3 may lead to 8 tiles in our solution compared with the 6 tiles required in the optimum solution. (iii) Each symbol of a component of size 4 in $G^\star$ 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 [7]). 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 $(F, \mathcal{S}, \ell)$ 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 $(F, \mathcal{S}, \ell)$ 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 $(F, \mathcal{S}, \ell)$ admits a feasible tileset with at most ℓ tiles. The ILP uses $\mathcal{O}(|\mathcal{S}|^{|\mathcal{S}|})$ variables. Lenstra [23] 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 [15], modifying an algorithm by Kannan [21].

Theorem 4 (Frank and Tardos [15]) *In $\mathcal{O}^*(p^{\mathcal{O}(p)})$ time we can decide whether a given ILP with p variables is feasible.*[1]

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 (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 $(F, \mathcal{S}, \ell)$ with k scenarios $\mathcal{S} = \{S_1, \ldots, S_k\}$ 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 $\mathcal{I} = \{I_1, \ldots, I_s\} \in \Pi(I)$ 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 $I := I_1 \cup \ldots \cup I_s$ and have the pattern $\mathcal{I}$ associated

[1] Recall that $\mathcal{O}^*$-notation ignores factors that are polynomial in the input size.

with them, in the following way: exactly s tiles, say, $T_1, \ldots, T_s$, 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:

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

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 $1, \ldots, k-1$ 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 $k^k + 1$.

2. 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:

$$x_{\emptyset,\{1\}},\ x_{\emptyset,\{2\}},\ x_{\emptyset,\{3\}},\ x_{\emptyset,\{1,2\}},\ x_{\emptyset,\{1,3\}},\ x_{\emptyset,\{2,3\}},\ x_{\emptyset,\{1,2,3\}},$$
$$x_{\{1\},\{2\}},\ x_{\{1\},\{3\}},\ x_{\{1\},\{2,3\}},\ x_{\{2\},\{3\}},\ x_{\{2\},\{1,3\}},\ x_{\{3\},\{1,2\}}.$$

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 $[k]\backslash(I \cup J)$), 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

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

We compare the number of tiles that provide one of their symbols for the 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.

3. 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 $(F, \mathcal{S}, \ell)$ of MINIMUM FEASIBLE TILESET admits a feasible tileset $\mathcal{T}$ of minimum cardinality $|\mathcal{T}| \leq \ell$. Since $\mathcal{T}$ is feasible for each scenario $S_i \in \mathcal{S}$, we may let $\varphi_i \colon S_i \to \mathcal{T}$ 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.

1. *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

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

   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 the *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 the 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 T we determine a pattern for each symbol $s \in F$. We let $\mathcal{T}(s) := \{T \in \mathcal{T} \mid \exists i \in [k] : \varphi_i(s) = T\} = \{T_1, \ldots, T_r\}$, 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 $\{I_1, \ldots, I_r\}$ 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 the *pattern* $\{I_1, \ldots, I_r\} \in \Pi(I)$.
   For each $I \in [k]$ and each partition $\mathcal{I} \in \Pi(I)$ we set $y_{\mathcal{I}}$ to the number of symbols in F with the 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 $(F, \mathcal{S}, \ell)$ is feasible and fix a feasible assignment to the x- and y-variables. We derive a feasible tileset 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 the symbols in I is matched by the sum of the 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 the pattern $\mathcal{I}$ matches the corresponding variable $y_{\mathcal{I}}$. For a symbol with the assigned pattern $\mathcal{I} = \{I_1, \ldots, I_r\}$ the intention is to use r tiles $T_1, \ldots, T_r$ 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 the x-variables, and thereby ensure that the final tileset 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 the 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 the 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 the pattern $\mathcal{I} = \{I_1, \ldots, I_r\}$ contributes to r constraints (10), one for each $I_1, \ldots, I_r$. Accordingly, these constraints enforce the correct sum of the corresponding variables $x_{I_1,\cdot}, \ldots, x_{I_r,\cdot}$. (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 the 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. □

6 Bounded Number of Symbols

We now analyze the influence of the number of symbols $|F|$ on the complexity of solving an instance $(F, \mathcal{S}, \ell)$ of the decision variant of Minimum Feasible Tileset. (That is, as in Sect. 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 Sect. 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 the structural insights of Sect. 2 we are able to improve on this naive running time.

Theorem 6 *Any instance $(F, \mathcal{S}, \ell)$ 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 (Theorem 6) We describe a dynamic programming algorithm for solving an arbitrary instance $(F, \mathcal{S}, \ell)$. 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 $i \in [|F|]$ for which there is a partition of D into i sets $D_1, \ldots, D_i$ such that no scenario contains any set in $\{D_1, \ldots, D_i\}$ 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 $(F, \mathcal{S}, \ell)$ 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.)

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

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. □

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 NP ⊆ 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 $\mathcal{O}(|F|^{d-\varepsilon})$. Note that this restricted case has an essentially matching upper bound of $|\mathcal{S}| < (|F|+1)^d = \mathcal{O}(|F|^d)$.[2] As a consequence, there is no reduction to size polynomial in $|F|$ for the general MINIMUM FEASIBLE TILESET problem: any size $\mathcal{O}(|F|^c)$ 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* NP ⊆ coNP/poly.

To prove Theorem 7 we employ a similar result by Dell and Marx [8] for EXACT COVER BY d- SETS, which is defined as follows.[3]

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 [8] is given in terms of the size k of an exact d-set cover. Clearly, $k = \frac{|X|}{d}$ and, thus, we have $\mathcal{O}(k^{d-\varepsilon}) = \mathcal{O}(|X|^{d-\varepsilon})$ and may instead phrase the result in terms of $|X|$. Furthermore, their result builds on work by Dell and van Melkebeek [9] 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 [8]) *Let $d \geq 3$ and ε be a positive real. There is no polynomial-time algorithm that reduces every instance $(X, \mathcal{C})$ of* EXACT COVER BY d- SETS *to an equivalent instance of size $\mathcal{O}(|X|^{d-\varepsilon})$ (possibly with respect to a different problem), unless* NP ⊆ coNP/poly.

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

[2] A 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.

[3] Dell 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* $(F, \mathcal{S}, \ell)$ *with* $F = X$ *and scenario size at most* d.

Proof (Lemma 5, Sketch) 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 $\mathcal{S} = \binom{X}{d-1} \cup (\binom{X}{d} \backslash \mathcal{C})$. 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 $\binom{X}{d} \backslash \mathcal{C}$ 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. □

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 $S\colon F \to \mathbb{N}$ 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 $(F, \mathcal{S}, \ell)$ 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 $(F, \mathcal{S}, \ell)$ 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 $I \subset F$ 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:

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

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 $\mathcal{O}^*(k^{\mathcal{O}(k^2)})$. □

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 Sect. 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. Bansal, N., Caprara, A., Sviridenko, M.: A new approximation method for set covering problems, with applications to multidimensional bin packing. SIAM J. Comput. **39**(4), 1256–1278 (2009)
2. Bezzo, N., Cortez, R.A., Fierro, R.: Exploiting heterogeneity in robotic networks. In: Redundancy in Robot Manipulators and Multi-Robot Systems. Lecture Notes in Electrical Engineering, vol. 57, pp. 53–75. Springer, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-33971-4_4
3. Biedl, T., Chan, T., Ganjali, Y., Hajiaghayi, M., Wood, D.: Balanced vertex-orderings of graphs. Discrete Appl. Math. **148**(1), 27–48 (2005)
4. Buchin, K., van Kreveld, M.J., Meijer, H., Speckmann, B., Verbeek, K.: On planar supports for hypergraphs. J. Graph Algorithms Appl. **15**(4), 533–549 (2011)
5. Chen, J., Komusiewicz, C., Niedermeier, R., Sorge, M., Suchý, O., Weller, M.: Polynomial-time data reduction for the subset interconnection design problem. SIAM J. Discrete Math. **29**(1), 1–25 (2015)
6. Crescenzi, P.: A short guide to approximation preserving reductions. In: Proceedings of the Twelfth Annual IEEE Conference on Computational Complexity (CCC)
7. Cygan., M.: 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), pp. 509–518 (2013)
8. Dell, H., Marx, D.: Kernelization of packing problems. In: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 68–81 (2012)

9. Dell, H., van Melkebeek, D.: Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. J. ACM **61**(4), 23:1–23:27 (2014)
10. Disser, Y., Kratsch, S., Sorge, M.: The minimum feasible tileset problem. In: Proceedings of the 12th Workshop on Approximation and Online Algorithms (WAOA '14), LNCS, vol. 8952, pp. 144–155. Springer, Heidelberg (2014)
11. Disser, Y., Matuschke, J.: Degree-constrained orientations of embedded graphs. J. Comb. Optim. **31**(2), 758–773 (2016)
12. Du, D.-Z., Miller, Z.: Matroids and subset interconnection design. SIAM J. Discrete Math. **1**(4), 416–424 (1988)
13. Flum, J., Grohe, M.: Parameterized Complexity Theory. Springer, Berlin (2006)
14. Frank, A., Gyárfás, A.: How to orient the edges of a graph. Coll. Math. Soc. Janos Bolyai **18**, 353–364 (1976)
15. Frank, A., Tardos, É.: An application of simultaneous diophantine approximation in combinatorial optimization. Combinatorica **7**(1), 49–65 (1987)
16. Garey, M.R., Johnson, D.S.: Computers and Intractability. A Guide to the Theory of NP-Completeness. W.H Freeman and Company, San Francisco (1979)
17. Gottlob, G., Greco, G.: Decomposing combinatorial auctions and set packing problems. J. ACM **60**(4), 24 (2013)
18. Hakimi, S.: On the degrees of the vertices of a directed graph. J. Frankl. Inst. **279**(4), 290–308 (1965)
19. Johnson, D.S., Pollak, H.O.: Hypergraph planarity and the complexity of drawing Venn diagrams. J. Graph Theory **11**(3), 309–325 (1987)
20. Kann, V.: Maximum bounded 3-dimensional matching is MAX SNP-complete. Inf. Process. Lett. **37**(1), 27–35 (1991)
21. Kannan, R.: Minkowski's convex body theorem and integer programming. Math. Oper. Res. **12**, 415–440 (1987)
22. Koutis, I.: Faster algebraic algorithms for path and packing problems. In: Proceedings of the 35th International Colloquium on Automata (ICALP), pp. 575–586 (2008)
23. Lenstra, H.W.: Integer programming with a fixed number of variables. Math. Oper. Res. **8**, 538–548 (1983)
24. Lundh, R., Karlsson, L., Saffiotti, A.: Autonomous functional configuration of a network robot system. Robot. Auton. Syst. **56**(10), 819–830 (2008)
25. Marek, C., Fomin, F.V., Kowalik, $\mathcal{L}$., Lokshtanov, D., Marx, D., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Berlin (2015)
26. Mittal, S.: A survey of techniques for architecting and managing asymmetric multicore processors. ACM Comput. Surv. **48**(3), 1–38 (2016)
27. Reinhard, D.: Graph Theory. Graduate Texts in Mathematics, vol. 173, 5th edn. Springer, Berlin (2016)
28. Rodney, G., Downey, R.G., Fellows, M.R.: Fundamentals of Parameterized Complexity. Texts in Computer Science. Springer, Berlin (2013)
29. Schuurman, P., Woeginger, G.J.: Approximation schemes—a tutorial. http://www.win.tue.nl/~gwoegi/papers/ptas.pdf
30. Sviridenko, M., Ward, J.: Large neighborhood local search for the maximum set packing problem. In: 40th International Colloquium on Automata, Languages, and Programming (ICALP), pp. 792–803 (2013)
31. van Bevern, R., Kanj, I., Komusiewicz, C., Niedermeier, R., Sorge,M.: Twins in subdivision drawings of hypergraphs. In: Proceedings of the 24th International Symposium on Graph Drawing& Network Visualization, LNCS, vol. 9801, pp. 67–80. Springer, Heidelberg (2016)