Discrete Optimization 31 (2019) 56–78

Contents lists available at ScienceDirect

Discrete Optimization

www.elsevier.com/locate/disopt

Additive stabilizers for unstable graphs

Karthekeyan Chandrasekaran[a], Corinna Gottschalk[b], Jochen Könemann[c,1], Britta Peis[b], Daniel Schmand[b,*], Andreas Wierz[b]

[a] University of Illinois Urbana-Champaign, Department of Industrial and Enterprise Systems Engineering, United States
[b] RWTH Aachen University, School of Business and Economics, Germany
[c] University of Waterloo, Department of Combinatorics & Optimization, Canada

A R T I C L E I N F O

Article history:
Received 9 August 2017
Received in revised form 1 August 2018
Accepted 15 August 2018
Available online 1 September 2018

Keywords:
Graph stabilization
Gallai–Edmonds decomposition
Matching
Vertex cover
Approximation algorithm
Hardness

A B S T R A C T

A weighted graph is called stable if the maximum weight of an integral matching equals the cost of a minimum-weight fractional vertex cover. We address the following question: how can we modify a given unstable graph in the least intrusive manner in order to achieve stability? Previous works have addressed stabilization through addition or deletion of the smallest possible number of edges/vertices. In this work we investigate the following more fine-grained *additive* stabilization strategy: given a graph $G = (V, E)$ with unit edge weights; find non-negative $c \in \mathbb{R}^E$ with minimum $\sum_e c_e$ such that adding c_e to the unit edge weight of each $e \in E$ yields a stable graph.

We provide the first super-constant hardness of approximation results for any graph stabilization problem: (i) unless the current best-known algorithm for the *densest-k-subgraph* problem can be improved, there is no $o(|V|^{1/24})$-approximation for additive stabilizers; (ii) the additive stabilizer problem has no $o(\log |V|)$ approximation unless $P = NP$.

On the algorithmic side, we present (iii) a polynomial time algorithm with approximation factor at most $\sqrt{|V|}$ for a super-class of the instances generated in our hardness proofs, (iv) an algorithm to solve min additive stabilizer in factor-critical graphs exactly in polynomial time, and (v) an algorithm to solve min additive stabilizer in arbitrary graphs exactly in time exponential in the size of the Tutte set. Our main tools are the Gallai–Edmonds decomposition and structural results for the problem that reduce the continuous decision domain to a discrete decision domain.

1. Introduction

Over the last two decades, algorithmic game theory has established itself as a vibrant and rich subarea of theoretical computer science as is evidenced by several recent books (e.g., see [1–3]). A crucial driver in

* Corresponding author.
E-mail addresses: karthe@illinois.edu (K. Chandrasekaran), gottschalk@oms.rwth-aachen.de (C. Gottschalk), jochen@uwaterloo.ca (J. Könemann), peis@oms.rwth-aachen.de (B. Peis), schmand@oms.rwth-aachen.de (D. Schmand), wierz@oms.rwth-aachen.de (A. Wierz).
[1] Supported by an NSERC Discovery Grant.

https://doi.org/10.1016/j.disopt.2018.08.003

this development is the increasingly networked structure of today's society, and the impact this development has on the day-to-day interactions that humans engage in. Finding and analysing graph-theoretic models for such networks is at the heart of the field of network exchange theory [4,5].

Social networks, and the interaction of individuals in those also motivated our work. Specifically, our interest started with [6], where Kleinberg & Tardos introduced *network bargaining* as a natural extension of Nash's classical two-player bargaining game [7] to the network setting. An instance of a network bargaining game is an edge weighted graph $(G = (V, E), w : E \to \mathbb{R}_+)$. The players in the game correspond to the vertices in the graph. Each $\{u, v\} \in E$ corresponds to a potential deal of given value $w_{uv} \geq 0$. Each player is allowed to interact with the neighbours to agree upon a sharing of the value on the edge between them and eventually arrive at a deal with at most one of her neighbours. Therefore, outcomes in network bargaining correspond to *matchings* $M \subseteq E$, and an *allocation* $y \in \mathbb{R}^V_+$ of $w(M) = \sum_{e \in M} w_e$ to the players. In particular, an outcome satisfies $y_u + y_v = w_{uv}$ for all $\{u, v\} \in M$, and $y_u = 0$ if u is not incident to an edge of M (u is *exposed*).

Kleinberg and Tardos introduced the concept of *stability*, and call an allocation y *stable* if $y_u + y_v \geq w_{uv}$ for *all* edges $\{u, v\} \in E$. Naturally extending Nash's bargaining solution, the authors define the *outside option* α_u of a player u given an allocation y as the largest value that u can *extract* from one of its neighbours. An allocation y is then deemed to be *balanced* if the value of each matching edge $\{u, v\} \in M$ is split according to Nash's bargaining condition: each player $a \in \{u, v\}$ receives its outside option α_a, and the remaining value of w_{uv} is divided equally among the players u and v. One of Kleinberg and Tardos' main results is that balanced outcomes exist in a given network bargaining instance if and only if stable ones exist, and these can be computed efficiently.

Network bargaining is closely related to the cooperative *matching game* introduced by Shapley and Shubik [8], where the player set once more corresponds to the vertices of an underlying graph $G = (V, E)$, and the characteristic function assigns the maximum weight of a matching in $G[S]$ to each set $S \subseteq V$ of vertices. The *core* of an instance of this game consists of allocations $y \in \mathbb{R}^V_+$ of the weight $\nu(G, w)$ of a maximum-weight matching to the players such that $y_u + y_v \geq w_{uv}$ for all $\{u, v\} \in E$. Hence, core allocations exactly correspond to stable allocations in network bargaining [9].

We recall the classical maximum weight matching LP that has a variable x_e for each edge $e \in E$. In the LP we use $x(\delta(v))$ as a convenient short-hand for $\sum_{e \in \delta(v)} x_e$, where $\delta(v)$ denotes the set of edges incident to v. We also use $\mathbb{0}$ and $\mathbb{1}$ for vectors of 0s and 1s, respectively, of appropriate dimensions.

$$\nu_f(G, w) := \max \left\{ \sum_{e \in E} w_e x_e \;:\; x(\delta(v)) \leq 1 \;\text{ for all } v \in V, x \geq \mathbb{0} \right\}. \tag{P}$$

The linear programming dual of (P) has a variable y_v for each vertex $v \in V$, and a covering constraint for each edge $e \in E$:

$$\tau_f(G, w) := \min \left\{ \sum_{v \in V} y_v \;:\; y_u + y_v \geq w_{uv} \text{ for all } \{u, v\} \in E, y \geq \mathbb{0} \right\}. \tag{D}$$

Feasible solutions of (P) and (D) will henceforth be referred to as *fractional matchings* and *fractional w-vertex covers*, respectively. We will denote the integer optimum solution values to the above linear programs by dropping the subscript f. In the unit-weight special case, where $w = \mathbb{1}$, we will omit the argument w from the ν and τ notation for brevity. An immediate observation is that a given instance of network bargaining has a stable outcome if and only if the core of the corresponding matching game is non-empty, that is, $\nu(G, w) = \nu_f(G, w) = \tau_f(G, w)$, where the second equality follows from linear programming duality. In other words, stable outcomes exist if and only if the linear program (P) admits an integral optimum solution. We will call a (possibly weighted) graph *stable* if the induced network bargaining instance admits a stable outcome.

In [10], Bock et al. proposed the following problem: given an *unstable* graph G, modify G in the least *intrusive* way in order to attain stability. The authors focused on the concrete question of removing the smallest number of edges from G so that the resulting graph is stable. Bock et al. showed that this problem is as hard to approximate as the vertex cover problem, even if the underlying graph is factor critical (i.e., even if deleting any vertex from G yields a graph with a perfect matching). The authors complemented this negative result by presenting an algorithm whose approximation-factor is proportional to the *sparsity* of the graph.

Concurrent to our work, Ahmadian et al. [11] and Ito et al. [12] proposed a *vertex-stabilizer* problem: given a graph $G = (V, E)$, find a minimum-cardinality set of vertices $S \subseteq V$ such that $G[V \setminus S]$ is stable. Both papers presented a combinatorial polynomial-time exact algorithm for this problem, and showed that the min cost variants of vertex stabilization are NP-hard. Ito et al. [12] proposed stabilizing a graph by adding a minimum number of vertices or edges. They showed that both of these problems are polynomial-time solvable. However, the minimum cost variant of stabilization by edge addition is NP-hard.

In this work, we consider a more *nimble* and in a sense *continuous* way of stabilizing a given unstable graph $G = (V, E)$. Instead of deleting/adding vertices/edges, we consider adding a small *subsidy* to a carefully chosen subset of the edges in order to create a stable weighted graph. The subsidy should be thought of as an additional incentive deployed by a central authority in order to achieve stability. A natural goal for the central authority would then be to minimize the total subsidy doled out in the stabilization process.

Definition 1.1. Given an undirected graph $G = (V, E)$ with unit edge weights, a *fractional additive stabilizer* is a vector $c \in \mathbb{R}^E_+$ such that $(G, \mathbb{1} + c)$ is stable. In the *minimum fractional additive stabilizer* problem (MFASP), the goal is to find a fractional stabilizer of smallest weight $\mathbb{1}^T c$.

We emphasize that we do not allow the addition of edges in MFASP, but are restricted to add weight to existing edges. We further note that the weight increases in MFASP need not be integral, and can take on arbitrary non-negative real values.

1.1. Our contributions

Several variants of graph stabilization are known to be NP-hard. Stronger hardness of approximation results have been elusive so far, and the gaps between them and the known approximation factors are large. In this work, we show strong approximation-hardness results, and nearly matching positive results.

In the following theorem, we relate MFASP to the *densest k-subgraph* problem (DkS) in which one is given a graph $G = (V, E)$, and the goal is to find a subset S of at most k vertices in V such that the graph $G[S]$ induced by S has the maximum number of edges. DkS is known not to admit a polynomial-time approximation scheme, assuming $NP \not\subseteq \cap_{\epsilon>0} BPTIME(2^{n^\epsilon})$ [13]; it is also known not to admit an efficient $n^{1/(\log\log n)^c}$-approximation for a constant $c > 0$ assuming the exponential time hypothesis [14]. On the other hand, the best known performance guarantee of any approximation algorithm is only $\approx O(|V|^{1/4})$ [15]. The true approximability of DkS is widely believed to lie closer to the known upper bound than to the known hardness lower-bound.

Theorem 1.2. *A polynomial time approximation algorithm with approximation factor $o(|V|^{1/24})$ for MFASP would lead to a polynomial time $o(|V|^{1/4})$-approximation for Densest k-Subgraph (DkS). Furthermore, there is no $o(\log(|V|))$-approximation algorithm for MFASP unless $P = NP$.*

For unit weights, it is well-known that (P) has an integral solution if and only if the set of *inessential* vertices X (those vertices that are exposed by some maximum matching) forms an independent set (e.g., see

[16,17]). Let Y be the set of neighbours of X in G, and $Z = V \setminus (X \cup Y)$. The triple (X, Y, Z) is called the *Gallai–Edmonds decomposition* of G [18–20] and the set Y is known as the *Tutte set*. As we will see later, the optimization problem given by an instance of MFASP naturally decomposes into two subproblems: that of picking a maximum matching between the vertices in Y and the factor critical components in $G[X]$, and that of picking a maximum matching in each of the components of $G[X]$. Our two hardness results in Theorem 1.2 demonstrate the hardness of each of these subproblems.

The instances generated in the hardness proofs of Theorem 1.2 have the property that $G[X]$ has no singleton factor-critical components. In the following positive result, we show an approximation algorithm for such hard instances. Let OPT denote the optimum stabilization cost of the given instance.

Theorem 1.3. *Let $G = (V, E)$ be a graph with Gallai–Edmonds decomposition (X, Y, Z). If all factor critical components of $G[X]$ have size greater than one then there is a $\min\{OPT, \sqrt{|V|}\}$-approximation algorithm for MFASP in G.*

The algorithmic result in Theorem 1.3 in conjunction with the hardness result in Theorem 1.2 assuming the hardness of approximation of DkS shows nearly matching approximability and inapproximability results.

While stabilization by min-edge deletion is already NP-hard and APX-hard in factor-critical graphs [10], we give a polynomial time algorithm to solve MFASP in factor-critical graphs.

Theorem 1.4. *There exists a polynomial-time algorithm to solve MFASP in factor-critical graphs.*

We further exploit the efficient solvability of MFASP in factor-critical graphs to present an exact algorithm for MFASP in general graphs whose running time is exponential only in the size of the Tutte set. Thus, our algorithm can be viewed as a *fixed parameter algorithm* (e.g., see [21]) where the parameter is the size of the Tutte set.

Theorem 1.5. *There exists an algorithm to solve MFASP for a graph $G = (V, E)$ with Gallai–Edmonds decomposition $V = X \cup Y \cup Z$ in time $O(2^{|Y|}poly(|V|))$.*

We conclude by giving a conditional approximation algorithm that achieves a $(k + 1)/2$-approximation when the number of non-trivial factor-critical components in the Gallai–Edmonds-decomposition exceeds the size of the Tutte set by a multiplicative factor of at least $1 + (1/k)$ (see Theorem 7.1).

1.2. Further related work

Various ways of modifying a given graph to achieve a property have been studied in the literature, but most previous works seem to consider *monotone* properties; i.e., properties that are closed under vertex- and edge-removal (e.g., see [22,23]). An exception is the work by Mishra et al. [24] who recently studied several ways to modify a given input graph into a so called *König–Egerváry* graph (KEG); i.e., a graph in which the size of a maximum matching equals that of a minimum vertex cover. Mishra et al. studied the problems of removing the smallest number of vertices (or, edges, respectively) to produce a KEG, and showed that these problems are NP-hard. Furthermore, both vertex and edge-deletion variants do not admit $O(1)$-approximation algorithms assuming the Unique Games Conjecture. The edge-deletion problem is also NP-hard to approximate within a factor of 2.88. For the vertex deletion problem, the authors gave an $O(\log n \log\log n)$-approximation algorithm. Finally, the authors showed that one can obtain KEG in a given graph $G = (V, E)$ with at least $3|E|/5$ edges.

In [25], Könemann et al. addressed a closely related problem of finding a minimum-cardinality set of edges to remove from a graph G such that the resulting graph has a fractional vertex cover of value at most $\nu(G)$.

We note that the resulting graph here may not be stable. While this problem is known to be NP-hard [26], Könemann et al. gave an efficient algorithm to find approximate solutions in sparse graphs.

2. Preliminaries

In the rest of the paper, we will only work with unit-weight graphs as input instances for MFASP. However, the results hold for uniform weights since scaling preserves stability as well as our results.

We emphasize the following fact that is implicit from our earlier discussion. A graph G is stable iff there is a maximum matching M and $y \in \mathbb{R}^V_+$ such that the characteristic vector χ_M of M and y form an optimal pair of solutions for (P) and (D). Then, a direct consequence of complementary slackness is that $y_v = 0$ if v is M-exposed, as well as $y_v + y_u = w_{uv}$ for all $\{u, v\} \in M$. A feasible solution to a MFASP instance $G = (V, E)$ is determined by a vector $c \in \mathbb{R}^E_+$ such that there is a matching M and a fractional $(\mathbb{1} + c)$-vertex cover y satisfying $\sum_{e \in M}(1+c_e) = \sum_{v \in V} y_v$. Sometimes, if we also need M and y in the context, we also call the triple (M, y, c) a solution to the MFASP instance. Moreover, such a matching M will be a maximum $(\mathbb{1}+c)$-weight matching. Note that calculating c and (M, y, c) are computationally equivalent, since calculating M and y can be done in polynomial time, once c is known. We emphasize that we use w_e to refer to the total edge weight of an edge e, while c_e to refer to the weight added for stabilizing.

We recall the following properties of the Gallai–Edmonds decomposition (as defined in Section 1.1) (e.g., see [27,28]): Let $G = (V, E)$ and $V = X \cup Y \cup Z$ be the Gallai–Edmonds decomposition of G. Then

(i) every maximum matching in G contains a perfect matching in $G[Z]$,
(ii) every connected component in $G[X]$ is factor-critical,
(iii) every maximum matching exposes at most one vertex in every connected component of $G[X]$, and
(iv) every maximum matching matches the vertices in Y to distinct components of $G[X]$.

We say that a component in $G[X]$ is non-trivial if it contains more than one vertex.

3. Structural results

In this section, we exhibit important structural properties of optimal solutions to MFASP. We will exploit these properties throughout in the design of algorithms and in our hardness results.

Theorem 3.1. *Let $G = (V, E)$ be an instance of MFASP and (M^*, y^*, c^*) be any optimal solution. Then the following hold:*

(i) $c^*_e = 0$ *for all* $e \in E \setminus M^*$, *and* $1 + c^*_{uv} = y^*_u + y^*_v$ *for* $\{u, v\} \in M^*$,
(ii) $0 \le c^*_e \le 1$ *for all edges* $e \in M^*$, *and*
(iii) $|M^*| = \nu(G)$, *i.e.,* M^* *is a maximum cardinality matching in* G.

In addition to the above properties, there always exists an optimal solution (M^*, y^*, c^*) of MFASP where

(iv) *both c^* and y^* are half-integral, and*
(v) *the support of y^* contains the Tutte set.*

Subsequently, Theorem 3.1 will allow us to focus on MFASP solutions in which $y_v \ge 1/2$ for all vertices v in the Tutte set. Note also that we may assume w.l.o.g. that the set Z in the Gallai–Edmonds decomposition of our graph is empty. If that is not the case, one could first consider the graph without Z and then extend the stabilizer using a perfect matching on Z without additional cost. This is done by setting $c_e := 0$ for every edge $e = \{u, v\} \in E[Z] \cup \delta(Z)$ and $y_v := 1/2$ for all $v \in Z$. For proving Theorem 3.1, we will often use the complementary slackness conditions for the linear programs (P) and (D). For a feasible primal–dual solution pair (M^*, y^*), we say that y^* satisfies complementary slackness with M^*, if the conditions $(u,v) \in M^* \Rightarrow y^*_u + y^*_v = 1 + c^*_{uv}$ and $y^*_u + y^*_v > 1 + c^*_{uv} \Rightarrow (u,v) \notin M^*$ are fulfilled.

Proof of Theorem 3.1. Since c^* is a fractional stabilizer for G, by the discussion in Section 2, it follows that M^* and y^* form a primal–dual solution pair. By complementary slackness, we immediately have that $1 + c^*_{uv} = y^*_u + y^*_v$ for all edges $\{u,v\}$ in M^*.

Now consider an edge $\{u,v\} \in E \setminus M^*$, and suppose that $c^*_{uv} > 0$. Then we may decrease c^*_{uv} to zero: since y^* is still a feasible fractional $(\mathbb{1}+c^*)$-vertex cover, and y^* satisfies complementary slackness with M^*, we obtain a better fractional additive stabilizer, thus contradicting the optimality of c^*. Thus, for every edge $e \in E \setminus M^*$, we have $c^*_e = 0$.

Now, consider an edge $\{u,v\} \in M^*$. By complementary slackness, we have $y^*_u + y^*_v = 1 + c^*_{uv}$. If $c^*_{uv} > 1$, then we obtain (c', y') where $c'_{uv} := 1$, $c'_e := c^*_e$ for every edge $e \in E \setminus \{uv\}$ and $y'_u := 1$, $y'_v := 1$, $y'_i := y^*_i$ for every vertex $i \in V \setminus \{u,v\}$. The resulting solution y' is a feasible fractional $(\mathbb{1}+c')$-vertex cover and y' satisfies complementary slackness with M^*. Thus, c' is a fractional additive stabilizer. We note that $\sum_{e\in E} c'_e < \sum_{e\in E} c^*_e$, a contradiction to the optimality of c^*. It follows that $0 \le c^*_e \le 1$ for all edges $e \in M^*$, as desired.

To show that M^* is a maximum cardinality matching in G, suppose for the sake of contradiction, the cardinality of M^* is strictly less than the cardinality of a maximum matching in G. We know that M^* and y^* satisfy complementary slackness. Since, by our assumption, M^* is not a maximum cardinality matching in G, there exists an M^*-augmenting path P. Let u_s and u_e denote the first and last vertices in the path P, respectively. Since y^* is a minimal fractional $(\mathbb{1}+c^*)$-vertex cover, and u_s and u_e are exposed in M^*, we have

$$y^*_u + y^*_v \ge 1 \qquad \forall\ \{u,v\} \in P \setminus M^*, \tag{1}$$

$$y^*_u + y^*_v = 1 + c^*_{uv} \qquad \forall\ \{u,v\} \in M^* \cap P, \tag{2}$$

$$y^*_{u_s} = 0 = y^*_{u_e}. \tag{3}$$

Let N be the matching obtained by taking the symmetric difference of M^* and P. Let us obtain new weights as follows:

$$c'_{uv} := \begin{cases} c^*_{uv} \ \text{ if } \{u,v\} \in E \setminus P \\ y^*_u + y^*_v - 1 \ \text{ if } \{u,v\} \in N \cap P \\ 0 \ \text{ if } \{u,v\} \in P \setminus N \end{cases}$$

We now show that the weight of the matching N w.r.t. $(\mathbb{1}+c')$ is identical to that of matching M^* w.r.t. $(\mathbb{1}+c^*)$:

$$\begin{aligned}
\sum_{e\in N}(1+c'_e) - \sum_{e\in M^*}(1+c^*_e) &= \sum_{\{u,v\}\in N\cap P}(1+(y^*_u+y^*_v-1)) - \sum_{\{u,v\}\in M^*\cap P}(1+c^*_e) \\
&= \sum_{\{u,v\}\in M^*\cap P}(y^*_u+y^*_v) - \sum_{\{u,v\}\in M^*\cap P}(1+c^*_e) \\
&= \sum_{\{u,v\}\in M^*\cap P}(1+c^*_e) - \sum_{\{u,v\}\in M^*\cap P}(1+c^*_e) \\
&= 0.
\end{aligned}$$

The second and third inequality are due to Eqs. (3) and (2). By Definition of c', we have that y^* is a feasible fractional $(\mathbb{1}+c')$-vertex cover in G. Moreover, by the construction of N and c', the $(\mathbb{1}+c')$-weight of matching N is equal to the sum $\sum_{v\in V} y^*_v$. Because of the LP duality relation between the two values, N is a matching of maximum $(\mathbb{1}+c')$-weight, and y^* is a minimum fractional $(\mathbb{1}+c')$-vertex cover. Hence, c' is a fractional additive stabilizer. Next we note that

$$\begin{aligned}\sum_{e\in E} c'_e - \sum_{e\in E} c^*_e &= \sum_{\{u,v\}\in N\cap P} (y^*_u + y^*_v - 1) - \sum_{\{u,v\}\in M^*\cap P} c^*_{uv} \\ &= \sum_{\{u,v\}\in N\cap P} (y^*_u + y^*_v) - |N\cap P| - \sum_{\{u,v\}\in M^*\cap P} c^*_{uv} \\ &= \sum_{\{u,v\}\in M^*\cap P} (y^*_u + y^*_v) - |N\cap P| - \sum_{\{u,v\}\in M^*\cap P} c^*_{uv} \\ &= |M^*\cap P| - |N\cap P| \\ &= -1.\end{aligned}$$

The third equality is due to (3), the fourth follows from (2). Hence, c' is a fractional additive stabilizer whose weight is smaller than that of c^*, a contradiction to the optimality of c^*.

We next show Property (iv) of the theorem. Let $\bar{c}$ be a minimum fractional additive stabilizer. By Property (iii), we know that there exists a maximum matching in G that is also a maximum $(\mathbb{1}+\bar{c})$-weight matching. Let M^* be such a matching. We consider the following linear program, called $LP(G, M^*)$:

$$\begin{aligned} \min \sum_{e\in M^*} c_e & & (4)\\ y_u + y_v = c_{uv} + 1 & \quad \forall\ \{u,v\}\in M^* \\ y_u + y_v \geq 1 & \quad \forall\ \{u,v\}\in E\setminus M^* \\ y_u = 0 & \quad \forall\ u\in V,\ u \text{ is exposed by } M^* \\ c, y \geq \mathbb{0} & \end{aligned}$$

If (c, y) is an optimal solution of $LP(G, M^*)$, then c gives a minimum fractional additive stabilizer for G. In order to show that c is a fractional additive stabilizer, it is sufficient to find a fractional $(\mathbb{1}+c)$-vertex cover y that satisfies complementary slackness conditions with M^*. But, by the constraints in $LP(G, M^*)$, it is clear that y satisfies complementary slackness conditions with M^*. Furthermore, c is a minimum fractional additive stabilizer, since otherwise, we could derive a contradiction to the optimality of $\bar{c}$. Thus, it is sufficient to show that there exists a half-integral optimal solution (c^*, y^*) of $LP(G, M^*)$.

We observe that if G is bipartite, then for every matching $\bar{M}$ in G, the extreme point solutions to $LP(G, \bar{M})$ are integral since the constraint matrix is totally unimodular and the right-hand side is integral.

Now, suppose $G = (V, E)$ is non-bipartite. For this case we will use a standard trick already used in a paper of Nemhauser and Trotter [29]. We construct a bipartite graph $G' = (V_1 \cup V_2, E')$ as follows: for each vertex $u \in V$, we introduce vertices $u_1 \in V_1$, $u_2 \in V_2$ and for each edge $\{u, v\} \in E$, we introduce edges $\{u_1, v_2\}, \{u_2, v_1\}$ in E'. For each matching edge $\{u, v\} \in M^*$, we include edges $\{u_1, v_2\}, \{u_2, v_1\}$ in M'. Thus M' is a matching in G' that exposes u_1 and u_2 for every vertex $u \in V$ that is exposed by M^*. Let (c', y') be an integral optimal solution of $LP(G', M')$. Let (c^*, y^*) be obtained by setting $c^*_{uv} := 1/2(c'_{u_1v_2} + c'_{u_2v_1})\ \forall\ \{u, v\} \in M^*$ and $y^*_u := 1/2(y'_{u_1} + y'_{u_2})\ \forall\ u \in V$. Clearly, (c^*, y^*) is half-integral. It remains to show that (c^*, y^*) is an optimum to $LP(G, M^*)$:

The feasibility of the solution (c^*, y^*) for $LP(G, M^*)$ is easy to verify. We note that $\sum_{e\in M^*} c^*_e = 1/2\sum_{e\in M'} c'_e$. We will prove optimality. Suppose (c^*, y^*) is not optimal for $LP(G, M^*)$. Then there exist $(\tilde{c}, \tilde{y})$ feasible for $LP(G, M^*)$ such that $\sum_{e\in M^*} \tilde{c}_e < \sum_{e\in M^*} c^*_e$. Consider the solution $(\tilde{c}', \tilde{y}')$ obtained by setting $\tilde{c}'_{u_1v_2} = \tilde{c}'_{u_2v_1} = \tilde{c}_{uv}$ for every $\{u, v\} \in M^*$ and $\tilde{y}'_{u_1} = \tilde{y}'_{u_2} = \tilde{y}_u$ for every $u \in V$. The resulting solution $(\tilde{c}', \tilde{y}')$ is feasible to $LP(G', M')$. Moreover $\sum_{e\in M'} \tilde{c}'_e = 2\sum_{e\in M^*} \tilde{c}_e < 2\sum_{e\in M^*} c^*_e = \sum_{e\in M'} c'_e$, a contradiction to the optimality of (c', y'). Hence, property (iv) is shown.

In order to prove (v), consider the Gallai–Edmonds decomposition of G given by $V = X \cup Y \cup Z$. Let c^* be a half-integral minimum fractional additive stabilizer for G. Let M^* be a maximum $(\mathbb{1} + c^*)$-weight matching and y^* be a half-integral minimum fractional $(\mathbb{1} + c^*)$-vertex cover (such y^* and c^* exist by (iv)). Suppose that $y^*_v = 0$ for some $v \in Y$. We will construct a half-integral fractional additive stabilizer c' without increasing the cost and a fractional $(\mathbb{1} + c')$-vertex cover y' that satisfies complementary slackness with M^* and has $y'_w > 0$ for each node w of the Tutte set.

Since M^* is maximum, every node of Y is matched. For $v \in Y$, we denote by S_v the factor-critical component in $G[X]$ which is matched to v and by s_v the vertex matched to v. Let $Y' := \{v \in Y : y^*_v = 0\}$. We set $c'_e := 0\ \forall e \in \bigcup_{v \in Y'} (E(S_v) \cup \{v, s_v\})$ and $c'_e := c^*_e$ otherwise. It is clear that c' is half-integral and the cost of c' cannot be more than that of c^*. We define

$$y'_w := \begin{cases} 1/2, & \text{if } w \in Y' \text{ or } w \in \displaystyle\bigcup_{v \in Y'} V(S_v), \\ y^*_w & \text{else.} \end{cases}$$

We note that y' is a fractional $(\mathbb{1} + c')$-vertex cover for the following reason: since all vertices in $Y' \cup (\bigcup_{v \in Y'} V(S_v))$ have y'-value $1/2$, and all edges in $E[(\bigcup_{v \in Y'} V(S_v)) \cup Y']$ have c'-value 0, it follows that the covering constraints on edges in $E[(\bigcup_{v \in Y'} V(S_v)) \cup Y']$ are satisfied. Since $c'_e = c^*_e$ on all edges outside $E[(\bigcup_{v \in Y'} V(S_v)) \cup Y']$, and since $y'_u \geq y^*_u$ for all vertices $u \in V \setminus (\bigcup_{v \in Y'} V(S_v))$, the feasibility of y' w.r.t. c' is implied by the feasibility of y^* w.r.t. c^* for all edges outside $E[(\bigcup_{v \in Y'} V(S_v)) \cup Y']$ that are not incident to vertices in $\bigcup_{v \in Y'} V(S_v)$. Taking into account that no edges exist between vertices in $\bigcup_{v \in Y'} V(S_v)$ and vertices in $\bigcup_{v \in Y \setminus Y'} V(S_v)$, it remains to check the covering constraints for edges $\{u, t\}$ between some vertex u in $\bigcup_{v \in Y'} V(S_v)$ and some vertex t in $Y \setminus Y'$. However, such edges cannot be matching edges, so $c'_{ut} = 0$. Since $y'_u = 1/2$ and $y'_t \geq 1/2$, the feasibility of y' follows as a consequence. Finally it follows by construction that y' satisfies complementary slackness with M^*: for every matching edge $e = \{u, v\} \in M^* \cap E[(\bigcup_{v \in Y'} V(S_v)) \cup Y']$, we have $c'_e = 0$ and $y'_u = y'_w = 1/2$. We know by the Gallai–Edmonds-Decomposition that no matching edges exist between a vertex in $(\bigcup_{v \in Y'} V(S_v)) \cup Y'$ and $V \setminus ((\bigcup_{v \in Y'} V(S_v)) \cup Y')$. Moreover, for every matching edge e in $E[(\bigcup_{v \in Y \setminus Y'} V(S_v)) \cup (Y \setminus Y')]$, we know that $c'_e = c^*_e$ and that y' takes the same values on the end vertices as y^*, implying that y' satisfies complementary slackness with M^*. $\square$

Theorem 3.1 has some important structural consequences that we will turn to now. In the following two lemmas, we consider a graph $G = (V, E)$ with Gallai–Edmonds decomposition $X \cup Y \cup Z$. We let (M, y, c) be a feasible MFASP solution satisfying the properties of Theorem 3.1.

We first provide a lower-bound on the cost incurred by any feasible MFASP solution satisfying the properties of Theorem 3.1 on a factor critical component in $G[X]$.

Lemma 3.2. *Let K be a non-trivial component in $G[X]$ with a vertex u such that $y_u = 0$. If u is exposed by M, then $\sum_{e \in E(K)} c_e \geq 1$. On the other hand, if K is matched to Y by an edge $e' = \{v, w\}$ with $w \in Y$, then $\sum_{e \in E(K)} c_e + c_{e'} \geq y_w$.*

Proof. In this proof, we use an equivalent definition of factor-critical graphs: A graph is factor-critical if and only if it has an odd ear-decomposition. Furthermore, the initial vertex of the ear-decomposition can be chosen arbitrarily [30]. An ear-decomposition of a graph G is a sequence $r, P_1, \ldots, P_k$ with $G = (\{r\}, \emptyset) + P_1 + \cdots + P_k$ such that P_i is either a path where exactly the endpoints belong to $\{r\} \cup V(P_1) \cup \cdots \cup V(P_{i-1})$ or a circuit where exactly one of its vertices belongs to $\{r\} \cup V(P_1) \cup \ldots V(P_{i-1})$. An ear-decomposition is called odd if all P_i have odd length. If M' is a maximum matching in K that exposes u, then we can choose an odd ear decomposition with initial vertex u such each of the ears P_i are odd alternating paths or cycles whose first and last edges are not part of M'.

Assume first that u is M-exposed, and consider an odd ear decomposition of K with u being used as the initial vertex. Using the existence of an odd ear decomposition we deduce that there is an odd-length alternating circuit C with vertices

$$u_0 = u, u_1, \ldots, u_{2t}, u_{2t+1} = u.$$

Since y is a feasible (fractional) vertex cover, and $y_u = 0$, we immediately have $y_{u_1}, y_{u_{2t}} \geq 1$. By Theorem 3.1(i) we know have that $1 + c_{u_{2t'-1}u_{2t'}} = y_{u_{2t'-1}} + y_{u_{2t'}}$ for all $1 \leq t' \leq t$, as well as $y_{u_{2t'}} + y_{u_{2t'+1}} \geq 1$ for all $1 \leq t' < t$. Hence,

$$t + \sum_{e \in E(K)} c_e \geq \sum_{t'=1}^{t} (1 + c_{\{u_{2t'-1}, u_{2t'}\}}) = \sum_{i=1}^{2t} y_{u_i} \geq$$
$$2 + \left(\sum_{i=2}^{2t-1} y_{u_i} \right) = 2 + \sum_{t'=1}^{t-1} (y_{u_{2t'}} + y_{u_{2t'+1}}) \geq 2 + (t-1) = t+1,$$

which proves the first statement.

Now, let us consider the case where K is matched by an edge $\{v, w\}$ where w is in Y. We recall that we have $y_w \in \{1/2, 1\}$. If $y_v = 0$, then the proof is identical to that of the first statement. If $y_v = 1$, clearly, $c_{e'} = y_w$. Therefore, we may assume that $y_v = 1/2$ and thus $c_{e'} = y_w - 1/2$. It remains to show that $\sum_{e \in E(K)} c_e \geq 1/2$.

Using the existence of an odd ear decomposition once more, we deduce the existence of an odd-length, M-alternating u, v-path in K (in which the first and last edges are not in M). Let $u = v_0, v_1, \ldots, v_{2t+1} = v$ be the vertex set of such a path P.

Suppose for the sake of contradiction $c_e = 0$ for all $e \in E(K) \supseteq P$. Since y is a vertex cover and the matching edges are tight w.r.t. y, it follows that $y_{v_i} = 1$ for odd $i \leq 2t$ and $y_{v_i} = 0$ for even i. But that implies $y_{v_{2t}} + y_v = y_v = 1/2$, a contradiction. Hence, $\sum_{e \in E(K)} c_e > 0$. Since c is half-integral, it follows that $\sum_{e \in E(P)} c_e \geq 1/2$.

Lemma 3.2 has the following nice consequence. □

Lemma 3.3. *Let (M, y, c) be a (not necessarily optimum) solution for MFASP in G fulfilling the properties of Theorem 3.1 and $(y, c) = \text{argmin}\{\mathbb{1}^T c : (M, y, c)$ is feasible for MFASP$\}$. Consider a non-trivial factor-critical component K in $G[X]$. If there exists an edge between K and Y in M, then $y_v = 1/2$ for every $v \in V(K)$ and $c(e) = 0$ for every $e \in E(K)$.*

Proof. Suppose we have a vertex $u \in V(K)$ with $y_u = 0$. Let $f = \{v, w\}$ with $w \in Y$ be the edge between K and Y in M. Now, we can obtain another solution (M, y', c') by setting $y'_a := 1/2$ for every $a \in V(K)$, $y'_a := y_a$ for every $a \in V(G) \setminus V(K)$ and $c'_e := 0$ for every $e \in E(K)$, $c'_f = y_w - 1/2$ and $c'_e := c_e$ for every $e \in E(G) \setminus E(K)$. The feasibility of (M, y', c') follows from the assumptions about (M, y, c). Moreover by Lemma 3.2, we have

$$\begin{aligned}
\mathbb{1}^T c' &= \sum_{e \in E(K)} c'(e) + c'(f) + \sum_{e \in E(G) \setminus (E(K) \cup \{f\})} c'(e) \\
&= 0 + y_w - \frac{1}{2} + \sum_{e \in E(G) \setminus (E(K) \cup \{f\})} c(e) \\
&\leq \sum_{e \in E(K)} c(e) + c(f) + \sum_{e \in E(G) \setminus (E(K) \cup \{f\})} c(e) - \frac{1}{2} \\
&< \mathbb{1}^T c,
\end{aligned}$$

thus contradicting the optimality of c.

Thus, we may assume that every vertex $u \in V(K)$ has $y_u \in \{1/2, 1\}$. In this case, again by optimality of c, we obtain that $y_u = 1/2$ for every $u \in V(K)$ and $c_e = 0$ for every $e \in E(K)$. □

4. Inapproximability

In this section, we will show that MFASP is hard to approximate in general graphs and thus prove Theorem 1.2. The first part and the second part of Theorem 1.2 will be derived as consequences of Propositions 4.1 and 4.5, respectively.

Reduction from densest k-Subgraph. We start by showing that MFASP is at least as hard as the Densest k-Subgraph Problem in a certain approximation preserving sense. Recall that, in an instance of DkS, we are given a graph $G = (V, E)$, and we are asked to find a set $S \subseteq V$ of at most k vertices such that the induced subgraph $G[S]$ has a maximum number of edges. In [31] Hajiaghayi and Jain related the hardness of DkS to that of the *minimum k-edge coverage* problem (MkEC). In an instance of MkEC, one is once more given a graph $G = (V, E)$, and an integer parameter k. The goal is now to find a set of k edges that spans the smallest number of vertices. Alternatively, we are looking for the smallest set $S \subseteq V$ that induces at least k edges. Hajiaghayi and Jain showed that an f-approximation for MkEC yields a $2f^2$-approximation for DkS. We now relate MkEC and MFASP.

Proposition 4.1. *There exists a polynomial-time algorithm that takes as input an instance (G, k) of MkEC and constructs another instance H of MFASP, where $|V(H)| \leq 7|V(G)|^3$, such that an f-approximate solution for MFASP in H can be used to obtain a $2f$-approximate solution for MkEC in (G, k) in polynomial time.*

Suppose we are given an instance of MkEC with n vertices. Using the reduction stated in the above theorem, this translates to an instance of MFASP with $\hat{n} \leq 7n^3$ vertices. Suppose that we had an approximation algorithm for MFASP that produced a solution with approximation factor $o(\hat{n}^{\alpha}) = o(n^{3\alpha})$. Via the result of Hajiaghayi and Jain, we would then immediately obtain an algorithm for DkS with performance ratio $o(n^{6\alpha})$. This would lead to an improvement over the best known upper bound on the approximability of DkS for $\alpha = 1/24$, proving the first part of Theorem 1.2.

Proof of Proposition 4.1. Let $G = (V, E)$ and $k < |E|$ be an instance of MkEC. We construct an instance $\hat{G}$ of MFASP whose Gallai–Edmonds decomposition (GED) has a specific form. Recall from Theorem 3.1 that (M^*, y^*, c^*) is an optimum MFASP solution only if M^* is a maximum cardinality matching in the underlying graph. The instance $\hat{G}$ will encode the problem of picking k edges for MkEC as the problem of identifying k factor-critical components in the GED that are to be exposed by the matching in a solution for MFASP. An illustration can be found in Fig. 1.

Let Y_1 be a copy of vertex set V; we will later show that Y_1 is a part of the Tutte set Y of the GED of the constructed graph. Furthermore, for each edge $e = \{v, w\} \in E$, we add a triangle, and we let Δ denote the collection of these triangles. We will later show that each triangle will form a component of X in the GED. We connect each node of a triangle corresponding to an edge $\{v, w\} \in E$ to the vertices v and w in Y_1. As any maximum matching matches each vertex of the Tutte set to a distinct factor-critical component, we modify the instance such that the number of triangles is exactly $|V| + k$. To achieve this, we either add vertices to Y_1 which are connected to all vertices of all triangles or we add triangles all of whose vertices are connected to all vertices in Y_1.

While MkEC allows choosing any collection of k edges, there may exist a collection of k triangles in our current graph such that the remaining triangles cannot be matched perfectly to Y_1. To remedy the situation, we add $q - 1$ copies $Y_2, \ldots, Y_q$ of Y_1, where q will be chosen later, and connect each vertex of Y_i

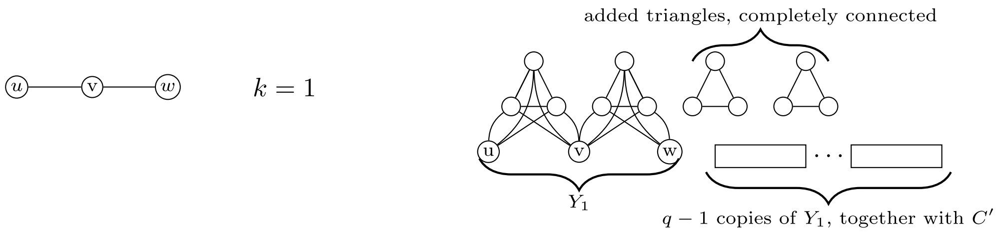

Fig. 1. Instance of MkEC and schema of a corresponding MFASP instance $\hat{G}$.

($i \in \{2, \dots, q\}$) with the same nodes as the corresponding node in the "original" set Y_1. We will later show that all these copies belong to the Tutte set Y of the Gallai–Edmonds decomposition. Moreover, we add $|Y_1| * (q-1)$ triangles and connect all their vertices to all vertices in $Y' := Y_1 \cup \dots \cup Y_q$. Call this set of newly added triangles C'.

The following two lemmas describe the relevant structure of the construction.

Lemma 4.2. *Let $q \geq \max_{v \in V(G)} |\delta(v)|$. Then for any choice of k triangles, there is a perfect matching between Y' and the triangles that were not chosen.*

Proof. Let $\mathcal{E}$ be the set of k triangles that we wish to expose. We construct a matching $M_{\mathcal{E}}$ that exposes precisely $\mathcal{E}$. For each triangle in $\Delta \setminus \mathcal{E}$ corresponding to some edge $\{u, v\} \in E$ match the triangle to a currently exposed copy of u. Note that q is at least the maximum degree in G, and hence this process matches all triangles in $\Delta \setminus \mathcal{E}$.

Let $\widetilde{\Delta}$ be the collection of triangles not in $\Delta \cup \mathcal{E}$, and let $\widetilde{Y'}$ be the collection of $M_{\mathcal{E}}$ exposed vertices in Y'. Clearly, $|\widetilde{\Delta}| = |\widetilde{Y'}|$, and the graph induced by the edges between $\widetilde{Y'}$ and the vertices of triangles in $\widetilde{\Delta}$ is complete bipartite. Thus, picking any $\widetilde{Y'}$-perfect matching in this graph and adding its edges to $M_{\mathcal{E}}$ yields the desired matching exposing $\mathcal{E}$. $\square$

Lemma 4.3. *The Gallai–Edmonds decomposition of $\hat{G}$ is given by $Y = Y'$, $X = V(\hat{G}) \setminus Y'$, $Z = \emptyset$.*

Proof. In any graph $G = (V, E)$, the size of a maximum matching can be characterized by the *Tutte–Berge formula* [32]:

$$2\nu(G) = |V| - \max_{W \subseteq V} (q_G(W) - |W|)$$

where $q_G(W)$ denotes the number of components with an odd number of nodes in $G[V \setminus W]$.

By plugging $W = Y'$ into the Tutte–Berge formula, we see that a maximum matching has size at most $\frac{V(\hat{G})-k}{2} = 2|Y'| + k$. A matching of size $2|Y'| + k$ exists by Lemma 4.2.

By Lemma 4.2, every vertex $v \in V(\hat{G}) \setminus Y'$ is inessential, that is, there exists a maximum matching in $\hat{G}$ exposing v. Now, suppose there was a maximum matching M exposing a vertex $v \in Y'$. We know $|M| = 2|Y'| + k$, but any matching can contain at most $|Y'| + k$ edges of $E[V(\hat{G}) \setminus (Y')]$, as that is the number of triangles. All other edges have one endpoint in Y'. Thus, if M exposes v, then $|M| < |Y'| + k + |Y'|$, which is a contradiction. $\square$

Together, Lemmas 4.2 and 4.3 imply that for any choice of k unmatched factor-critical components, there is a maximum matching exposing exactly one vertex in these k components and conversely, every maximum matching is of this form. We have shown in Theorem 3.1 that it suffices to consider stabilizers (M, y, c) where M is a maximum matching, y and c are half-integral and y is positive on the Tutte set. Once the set of unmatched components is fixed, we can see how to obtain an optimal stabilizer for this situation: Start with a matching between the matched components and Y' and extend it arbitrarily to a maximum matching M. Let $K \subset V(\hat{G})$ be the set of vertices in triangles not matched to Y'. We set $c_e = 1$ for the matching edge in each of these triangles, $y_v = 1$ for both matched vertices within the triangle and $y_v = 0$ for the remaining vertices in K. For $v \in Y'$, we set $y_v = 1$ if $v \in N(K)$ and $y_v = 1/2$ otherwise. By Lemma 3.3, $y_v = 1/2$ is then optimal for all vertices v in matched triangles. Consequently, $c(e) = 1/2$ if $e \in M \cap \delta(N(K))$ and $c(e) = 0$ for the remaining edges. In total, $\mathbb{1}^T c = k + 1/2|N(K)|$.

If the unmatched components correspond to a set E' of edges in G, then $N(K) \cap Y_1$ corresponds exactly to the vertices in G spanned by E'. Consequently, the cost of the stabilizer consists of k (for the unmatched triangles) and q times the number of spanned vertices in G as the neighbourhoods of the copies of Y_1 are identical. Suppose a component that does not correspond to an edge in G is unmatched. Then, $y_v = 1$ for all $v \in Y'$ and therefore $c(E) = k + 1/2|Y'|$. Thus, we can modify the solution by choosing to expose components corresponding to edges instead without increasing the cost. W.l.o.g. we modify any solution to MFASP this way. Then, we have the following Lemma:

Lemma 4.4. *G has a solution of MkEC of size at most x if and only if $\hat{G}$ has a MFASP of cost at most $k + qx/2$.*

We next show that Lemma 4.4 yields a factor-preserving hardness. If $k > 0$, then we have $x \geq 2$. Moreover, set $q = \max\{k, \max_{v \in G}|\delta(v)|\}$. Let x^* be the value of an optimal solution for MkEC, then the optimal value of MFASP is $k + \frac{qx^*}{2}$. Suppose there was an f-approximation for MFASP. This would yield a stabilizer solution of cost $k + \frac{qx}{2} \leq f(k + \frac{qx^*}{2})$ for some x. We observe that $\frac{qx}{2} \leq (f-1)k + f\frac{qx^*}{2} \leq qf(1 + \frac{x^*}{2}) \leq qfx^*$. Therefore, we have a $2f$-approximate solution of MkEC which proves Proposition 4.1. □

While DkS is believed to be difficult, there are no strong inapproximability results known. Next, we show set-cover-hardness for MFASP, which leads to a stronger inapproximability result.

Reduction from set cover. We exploit a different aspect of MFASP for this reduction: We could look at MFASP as a problem consisting of two subproblems: How to choose the matched factor-critical components and, having fixed those, how to choose the matching within the unmatched components and thus decide the y-values. In the previous reduction, the difficulty is completely encoded in the first subproblem—once we chose the matched components, the second subproblem was easy. In the following reduction, we consider a construction where the matched components are the same for any reasonably good solution and the difficulty is encoded in the second subproblem. We recall that an instance of set cover is specified by $(\mathcal{S}, \mathcal{X})$ where $\mathcal{X}$ is the ground set, and $\mathcal{S}$ is a collection of subsets of $\mathcal{X}$; the goal is to find a smallest number of sets in $\mathcal{S}$ that jointly cover $\mathcal{X}$.

Proposition 4.5. *There exists a polynomial time algorithm that takes as input an instance $(\mathcal{S}, \mathcal{X})$ with $|\mathcal{X}| = n$ of set cover and constructs another instance H of MFASP, where $|V(H)| \leq 6(n|\mathcal{S}|)^3$, such that an f-approximate solution for MFASP in H can be used to obtain a $2f$-approximate solution for the set cover instance in polynomial time.*

We now show that Proposition 4.5 implies the second part of Theorem 1.2. We recall that there is no $(\log(n) - \epsilon)$-approximation algorithm for Set Cover, even if the number of sets $|\mathcal{S}|$ is at most n^2, unless $P = NP$ [33]. Proposition 4.5 implies that there is no $(1/2)(\log(n) - \epsilon)$-approximation for MFASP, where n is the number of elements of the corresponding Set Cover instance. Now let $\hat{n}$ denote the number of vertices in the MFASP instance constructed in the proof of Proposition 4.5. We have $\hat{n} \leq 6(n|S|)^3 \leq 6n^9$. Hence, for $\hat{n}$ being the number of vertices in an MFASP instance, we conclude that unless $P = NP$, there is no approximation algorithm for MFASP with approximation factor better than $(1/18 - \epsilon)\log(\hat{n})$ for sufficiently large $\hat{n}$.

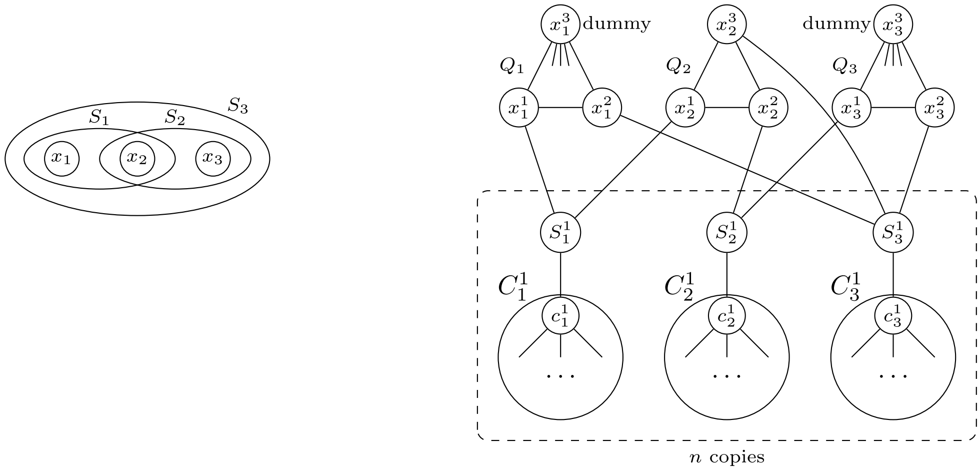

Fig. 2. Set cover instance and constructed MFASP instance $\hat{G}$.

Proof of Proposition 4.5. Let $(\mathcal{S}, \mathcal{X})$ be an instance of Set Cover, and define the *frequency* μ_i of element x_i as the number of sets in $\mathcal{S}$ that contain x_i. Without loss of generality, $\mu_i > 1$. Otherwise, the only set containing an item x_i has to be part of any solution, so it suffices to consider instances with $\mu_i > 1$ for all $i \in [n]$.

We construct a graph $\hat{G}$ with a specific Gallai–Edmonds decomposition. Our goal will be to decide whether a set is included in a set cover based on the y-values of the Tutte set. For each set S_j, create n vertices $S_j^1, \dots, S_j^n$. Let $Y' = \{S_j^i : i \in [n], j \in [m]\}$. This will be our Tutte set. For each S_j^i create a clique C_j^i of size $2N+1$ with $N > (nm)^2$ with a designated vertex c_j^i and add an edge $\{S_j^i, c_j^i\}$. Later we will show that each of these constructed cliques is a factor-critical component of the Gallai–Edmonds decomposition. It is not hard to see that the maximum-cardinality matching in any optimum MFASP solution must match all of the constructed large cliques, or incur a large stabilization cost inside the exposed clique.

Hence, the purpose of the large cliques described above is to ensure that the following factor-critical components are exposed in the matching of an optimum MFASP solution: For each element x_i with μ_i odd, construct an odd cycle Q_i consisting of μ_i vertices $x_i^1, \dots, x_i^{\mu_i}$. For each element with μ_i even, construct an odd cycle Q_i consisting of $\mu_i + 1$ vertices $x_i^1, \dots, x_i^{\mu_i+1}$, where the vertex $x_i^{\mu_i+1}$ is a dummy vertex. Let $\hat{S}_{(1,i)}, \dots, \hat{S}_{(\mu_i,i)}$ denote the sets in $\mathcal{S}$ containing x_i (choose the order arbitrarily). Consider the $n\mu_i$ copies of the corresponding vertices in Y' and add edges $\{x_i^k, \hat{S}_{(k,i)}^\ell\}\ \forall\ \ell \in [n],\ \forall\ k \in [\mu_i]\ \forall\ i \in [n]$. That is, add an edge between the kth vertex for x_i and all copies of the kth set in the list. For every $i \in [n]$ with μ_i even, add edges between $x_i^{\mu_i+1}$ and all vertices in Y'. Let the resulting graph be $\hat{G} = (\hat{V}, \hat{E})$. (See Fig. 2).

We now analyse the structure of the instance we built:

Claim 4.6. *The Gallai–Edmonds decomposition of G is $X \cup Y \cup Z$ where $Z = \emptyset$, $Y = Y'$ and $X = \hat{V} \setminus Y'$.*

Proof. Using the Tutte–Berge formula for the set $W = \{S_j^i : 1 \le j \le m, 1 \le i \le n\}$, we see that a maximum matching has size at most $\mathrm{nm}(N+1) + \sum_{i=1}^n \lfloor \mu_i/2 \rfloor$. Clearly, a matching of this size exposing a vertex v can be constructed for any $v \in C_j^i$ or $v \in Q_i$. Moreover, a matching of this size cannot be constructed by exposing a vertex $v = S_j^i$: If so, such a vertex S_j^i would belong to a factor-critical component K in the Gallai–Edmonds decomposition. Moreover, such a component K will also contain C_j^i since the vertices in C_j^i are in X as argued earlier and we have an edge between S_j^i and the vertex $c_j^i \in C_j^i$. But factor-critical graphs are 2-edge connected, and removing the edge (S_j^i, c_j^i) would separate the graph K into two components, a contradiction. $\square$

Claim 4.7. *Let $T \subseteq [m]$ denote the indices corresponding to a set cover. Then there exists a feasible solution to the MFASP instance $\hat{G}$ whose stabilizer cost is at most $n(1+|T|/2)$.*

Proof. For each $i \in [n]$, let k_i denote an arbitrarily chosen index in T such that the set S_{k_i} contains the element x_i. Consider a matching $\bar{M}$ obtained by matching S_j^i with c_j^i $\forall j \in [m], i \in [n]$, picking a perfect matching of the rest of the clique vertices $V(C_j^i) \setminus \{c_j^i\}$, exposing $x_i^{k_i}$ and picking a perfect matching of the rest of the vertices in each odd cycle Q_i. Obtain a fractional vertex cover solution $\bar{y}$ as follows: For every $i \in [n]$, let $\bar{y}_{S_j^i} = 1$ if $j \in T$ and $\bar{y}_{S_j^i} = 1/2$ if $j \in [m] \setminus T$. For each $i \in [n]$, set $\bar{y}_{x_{k_i}} = 0$, $\bar{y}_{x_k} = 1$ for the two copies of x_k in Q_i that are adjacent to x_{k_i} and $\bar{y}_{x_k} = 1/2$ for the other vertices in Q_i.

Obtain the solution $\bar{c}$ as $\bar{c}_{uv} = \bar{y}_u + \bar{y}_v - 1$ for every $uv \in \bar{M}$. Then the solution $(\bar{M}, \bar{y}, \bar{c})$ is a feasible solution to the MFASP instance $\hat{G}$. Moreover, the cost of the stabilizer $\mathbb{1}^T \bar{c}$ is $n(1+|T|/2)$. □

Let (M, y, c) be an f-approximate feasible solution to the MFASP instance $\hat{G}$. As we shall now see, such a solution may be assumed to be structurally nice.

Claim 4.8. *We can assume the following properties.*

1. *M matches S_j^i to C_j^i for every $j \in [m]$, $i \in [n]$ and $y_v = 1/2$ for every $v \in V(C_j^i)$.*
2. *$y_{S_j^i} = y_{S_j^1}$ for every $i \in [n]$, $j \in [m]$.*

Proof. We split the proof into two parts and show both properties separately.

1. If this was not the case, then there is at least one clique C_j^i with a vertex v with $y_v = 0$. Thus, $y_w = 1$ for all $w \in V(C_j^i) \setminus v$. By the complementary slackness condition given in Section 2, we have $y_j + y_k = 1 + c_{jk}$ for each matching edge $\{j,k\} \in M$. Thus, $\sum_{e \in M \cap E(C_j^i)} c_e = |M \cap E(C_j^i)| = N$. However, we note that $T = [m]$ is a feasible set cover and by Claim 4.7, this gives a feasible stabilizer of cost at most $nm/2 + n$.
2. If $y_{S_j^i} \neq y_{S_j^1}$, then consider $i_0 = \operatorname{argmin}_{i \in [n]} \sum_{j \in [m]} y_{S_j^i}$. Since the neighbourhood of $\{S_1^{i_0}, \ldots, S_m^{i_0}\}$ is identical to that of $\{S_1^i, \ldots, S_m^i\}$ for any $i \in [n]$, letting $y_{S_j^i} = y_{S_j^{i_0}}$ for all $j \in [m]$ does not affect feasibility, and does not increase the stabilization cost. □

Recall that we say that a vertex is M-exposed if the vertex is not incident to an edge in M. We can conclude that the set of M-exposed vertices contains exactly one vertex in each odd cycle Q_i. Moreover, we can assume an exposed vertex is not a dummy vertex $x_i^{\mu_i+1}$. Otherwise, we could change that without increasing the cost of the stabilizer.

Lemma 4.9. *Let X be the M-exposed vertices. Let $P := \{j \in [m] : S_j^1 \in N(X)\}$. Then $\{S_j : j \in P\}$ is a set cover of cardinality at most $2f|P^*|$, where P^* is the set of indices corresponding to the optimal set cover.*

Proof. We first show that P is indeed a set cover. Consider an element x_i. By Claim 4.8, there exists $k \in [\mu_i]$ such that x_i^k is M-exposed. There exists a set S_r^1 adjacent to x_i^k. Hence $r \in P$ and thus S_r covers x_i.

It remains to bound the cardinality of P. Since M exposes exactly one vertex in each odd cycle Q_i, by complementary slackness conditions, $\sum_{e \in Q_i} c_e \geq 1$. Thus, $\sum_{i \in [n]} \sum_{e \in Q_i} c_e \geq n$. Let $r \in P$. So S_r^1 is adjacent to an M-exposed vertex in Q_i and hence $y_{S_r^1} \geq 1$. By Claim 4.8, we have that $y_{S_r^i} \geq 1$ for every $i \in [n]$. Since $y_v = 1/2$ for every $v \in V(C_j^i)$ (using Lemma 3.3), we have that $c_{S_r^i, c_r^i} \geq 1/2$. Thus, $\sum_{r \in P} \sum_{i \in [n]} c_{S_r^i, c_r^i} \geq |P|n/2$. Thus, the cost of the stabilizer $\mathbb{1}^T c \geq n(1 + |P|/2)$. Hence, $|P| \leq 2(\mathbb{1}^T c/n - 1) \leq 2(f \mathbb{1}^T c^*/n - 1)$.

By Claim 4.7, we have that the cost of the optimal stabilizer $\mathbb{1}^T c^*$ is at most $n(1 + |P^*|/2)$. Thus, $|P| \leq f|P^*| + 2(f-1) \leq 2f|P^*|$. □

Proposition 4.5 follows by Lemma 4.9. □

5. An OPT-approximation in graphs with no singletons

In this section, we present an algorithm that achieves a $\min\{\sqrt{n}, OPT\}$-approximation factor in graphs whose Gallai–Edmonds decomposition has no trivial factor critical components. As a subroutine, we use an algorithm to solve MFASP in factor-critical graphs. We will mention how to solve MFASP in factor-critical graphs by solving an appropriately chosen LP. We restate the main theorem of the section for convenience:

Theorem 1.3. *Let $G = (V, E)$ be a graph with Gallai–Edmonds decomposition (X, Y, Z). If all factor critical components of $G[X]$ have size greater than one then there is a $\min\{OPT, \sqrt{|V|}\}$-approximation algorithm for MFASP in G.*

In the remainder of this section, we prove Theorem 1.3. We describe the algorithm as part of the proof. We refer the reader to Section 5.1 for a pseudocode. Fix an optimum solution (M^*, y^*, c^*) satisfying the properties given in Theorem 3.1. Then, by Lemma 3.3, we have $c^*_e = 0$ for e in a component that is M^*-covered. Moreover, as mentioned before, w.l.o.g. $Z = \emptyset$. As usual, $OPT := \sum_{e\in E} c^*(e)$. Let r denote the difference between the number of components in $G[X]$ and the number of vertices in Y. As M^* is a maximum matching, the properties of the Gallai–Edmonds decomposition imply that M^* exposes exactly r vertices, at most one in each component of $G[X]$. Further, M^* matches at most one vertex of a component to a vertex in Y, while the rest are matched within the component.

For each factor-critical component K, we compute a lower bound on the cost of an optimum stabilizer where the matching exposes K.

Lemma 5.1. *Let K be a (non-trivial) factor-critical component in $G[X]$. For each vertex w in K, let $\ell_{K,w}$ denote the optimum value of the following LP (in which we use $N_G(S)$ for the set of vertices in $V \setminus S$ that have a neighbour in S):*

$$\begin{aligned}
\ell_{K,w} := \min \quad & \sum_{v\in V(K)\cup N_G(V(K))} y_v - \left(\frac{|V(K)|-1}{2}\right) - \frac{|N_G(V(K))|}{2} \\
s.t.\ & y_i + y_j \geq 1\ \forall\ \{i,j\} \in E[V(K) \cup N_G(V(K))] \\
& y_i \geq 1/2\ \forall\ i \in N_G(V(K)) \\
& y_i \geq 0\ \forall\ i \in V(K) \\
& y_w = 0
\end{aligned}$$

Let $f(K) := \min_{w\in V(K)} \ell_{K,w}$. If M^* exposes a vertex in K then $\mathbb{1}^T c^* \geq r - 1 + f(K)$.

Proof. We recall that (M^*, y^*, c^*) is an optimum MFASP solution that satisfies the properties of Theorem 3.1. We then have

$$\mathbb{1}^T c^* = \sum_{e\in M^*} c^*_e = \sum_{K'\neq K: K' \text{ is } M^*\text{-exposed}} \sum_{e\in M^*\cap E(K')} c^*_e + \sum_{e\in M^*\cap E(K)} c^*_e + \sum_{e\in M^*\cap \delta_G(Y)} c^*_e.$$

The first double-sum on the right-hand side is at least $r - 1$ by Lemma 3.2 (this lemma can be applied because the factor-critical components are non-trivial). By complementary slackness conditions as mentioned in Section 2, we know that for every edge $\{i, j\} \in M^*$, we have $1 + c^*_{ij} = y^*_i + y^*_j$. Moreover, for an M^*-exposed vertex u, we have $y^*_u = 0$. As M^* exposes one vertex in K,

$$\sum_{e\in M^*\cap E(K)} c^*_e = \sum_{\{i,j\}\in M^*\cap E(K)} (y^*_i + y^*_j - 1) = \sum_{v\in V(K)} y^*_v - \left(\frac{|V(K)|-1}{2}\right).$$

If $\{i,j\} \in M^*$ with $i \in Y$, then $j \in X$ is a vertex in a factor-critical component that is matched by M^*. By Lemma 3.3, we have that $y_j^* = 1/2$ and by Theorem 3.1, we have $y_i^* \geq 1/2$. Hence,

$$\sum_{e \in M^* \cap \delta_G(Y)} c_e^* = \sum_{\{i,j\} \in M^* : i \in Y, j \in X} (y_i^* + y_j^* - 1) = \sum_{i \in Y} \left(y_i^* - \frac{1}{2} \right) \geq \sum_{i \in N_G(V(K))} y_i^* - \frac{|N_G(V(K))|}{2}.$$

Let w be a M^*-exposed vertex in K. Then, y^* restricted to the vertices $V(K) \cup N_G(V(K))$ is a feasible solution to the LP corresponding to $\ell_{K,w}$. Combining the three relations, we get that

$$\begin{aligned}\sum_{e \in M^*} c_e^* &\geq r - 1 + \sum_{v \in V(K) \cup N_G(V(K))} y_v^* - \left(\frac{|V(K)| - 1}{2} \right) - \frac{|N_G(V(K))|}{2} \\ &\geq r - 1 + \ell_{K,w} \geq r - 1 + f(K). \quad \square\end{aligned}$$

Remark 5.2 (*Proof of Theorem 1.4*). Lemma 5.1 can be used to obtain an optimal solution for MFASP in factor-critical graphs and thereby prove Theorem 1.4. We note that factor-critical graphs are the special case where G consists of one component K (and thus $N(V(K)) = \emptyset$). In that case an optimum stabilizer can be obtained by computing $f(K)$, choosing any matching M^* exposing $w^* := \mathrm{argmin}_{w \in V(K)} \ell_{K,w}$ and setting c^* to fulfil complementary slackness (i.e., if y^* is a solution for ℓ_{K,w^*}, then set $c^*(uv) := y_u^* + y_v^* - 1$ for every $uv \in M^*$ and $c^*(uv) := 0$ for every $uv \in E \setminus M$). The resulting solution is feasible for MFASP and its cost is equal to the lower bound given by Lemma 5.1 and hence optimal.

We continue with the proof of Theorem 1.3. In order to identify a suitable matching to stabilize, we build an auxiliary graph G' as follows: Contract each component K in $G[X]$ to a pseudo-vertex v_K and assign edge weight $w_e := f(K)$ for all edges e incident to the contracted vertex v_K. Compute a matching M in G' of maximum weight covering Y.

Lemma 5.3. *The cost $\mathbb{1}^T c^*$ of an optimum stabilizer (M^*, c^*, y^*) is at least*

$$r - 1 + \max_{K : v_K \text{ is } M\text{-exposed}} f(K).$$

Proof. Let $K_1 = \mathrm{argmax}_{K : v_K \text{ is } M\text{-exposed}} f(K)$. If M^* exposes K_1, then Lemma 5.1 proves the claim. So, we may assume that M^* matches K_1. Consider M^* restricted to the edges in the bipartite graph G'. Both M and M^* are maximum cardinality matchings in G' and v_{K_1} is M-exposed. So, we have an M-alternating path P starting from v_{K_1} and ending at another vertex corresponding to a contracted factor-critical component. Let $P = v_{K_1}, b_1, v_{K_2}, b_2, \ldots, v_{K_{t-1}}, b_{t-1}, v_{K_t}$ for some $t \geq 1$ and v_{K_t} is M^*-exposed. Since M is a maximum weight matching, we have $\sum_{e \in M \setminus P} w_e \leq \sum_{e \in M \cap P} w_e$. Thus, $\sum_{i=1}^{t-1} f(K_i) \leq \sum_{i=2}^{t} f(K_i)$ and we have that $f(K) = f(K_1) \leq f(K_t)$. Thus, the cost of the stabilizer c^* is at least $r - 1 + f(K_t) \geq r - 1 + f(K_1)$. $\square$

We now stabilize M. For each M-exposed vertex v_K, let

$$w_K := \arg\min_{w \in V(K)} \ell_{K,w},$$

and let $\overline{y}^{w_K}$ denote the solution y achieving the optimum for ℓ_{K,w_K}. Extend M inside each factor-critical component K: if v_K is matched by M using edge $\{u,b\}$ where $u \in V(K), b \in Y$, then extend M using a matching in K that exposes u. If v_K is exposed by M, extend M using a matching in K that exposes w_K. Let $\overline{M}$ denote the resulting matching.

For each vertex v_K matched by M, set $\overline{y}_u := 1/2$ for all vertices $u \in V(K)$. For each vertex v_K that is exposed by M, set $\overline{y}_u := \overline{y}_u^{w_K}$ for all vertices $u \in V(K)$. For each vertex $b \in Y$ that is adjacent to a

M-exposed v_K, set $\overline{y}_b := \max_{K: v_K \text{ is } M\text{-exposed}} \overline{y}_b^{w_K}$. For each vertex $b \in Y$ with no adjacent M-exposed v_K, set $\overline{y}_b := 1/2$. We note that these choices produce a feasible fractional vertex cover because no trivial factor-critical components exist.

Finally, set $\overline{c}(uv) := \overline{y}_u + \overline{y}_v - 1$ for edges $\{u, v\} \in \overline{M}$ and $\overline{c}(uv) := 0$ for edges $\{u, v\} \in E \setminus \overline{M}$. We next show that the solution $(\overline{M}, \overline{y}, \overline{c})$ is a feasible solution.

Lemma 5.4. *$(\overline{M}, \overline{y}, \overline{c})$ is a feasible solution to MFASP.*

Proof. By construction, $\overline{M}$ is a matching and $\sum_{e\in\overline{M}}(1 + \overline{c}_e) = \sum_{e\in\overline{M}}(\overline{y}_u + \overline{y}_v)$. It remains to show that $\overline{y}$ is a feasible fractional w-vertex cover for $w_e = 1 + \overline{c}_e$ for every $e \in E$.

Consider an edge $e = \{u, v\} \in E$. If $e \in \overline{M}$, then $\overline{y}_u + \overline{y}_v = 1 + \overline{c}_{uv}$. Let $e \in E \setminus \overline{M}$. For such edges, we have $\overline{c}_e = 0$ and hence $1 + \overline{c}_e = 1$.

We distinguish several cases. If $e \in K$ where v_K is matched by M, then $\overline{y}_u = \overline{y}_v = 1/2$ and hence $\overline{y}_u + \overline{y}_v = 1$. If $e \in K$ where v_K is exposed by M, then $\overline{y}_u + \overline{y}_v = \overline{y}_u^{w_K} + \overline{y}_v^{w_K} \geq 1$ by the feasibility of the solution $\overline{y}^{w_K}$ to the LP corresponding to $\ell_{K,w}$. If $e \in \delta_G(Y)$, then let $u \in Y, v \in V(K)$. If $v \in V(K)$ where v_K is matched by M, then $\overline{y}_v = 1/2$ and moreover $\overline{y}_u \geq 1/2$ and hence $\overline{y}_u + \overline{y}_v \geq 1$. If $v \in V(K)$ where v_K is exposed by M, then $\overline{y}_v = \overline{y}_v^{w_K}$ and $\overline{y}_u = \max_{w_K: v_K \text{ is } M\text{-exposed}} \overline{y}_u^{w_K} \geq \overline{y}_u^{w_K}$. By the feasibility of the solution $\overline{y}^{w_K}$ to the LP corresponding to $\ell_{K,w}$, we have that $\overline{y}_u + \overline{y}_v \geq \overline{y}_u^{w_K} + \overline{y}_v^{w_K} \geq 1$. $\square$

We now bound the cost of the constructed solution $(\overline{M}, \overline{y}, \overline{c})$.

Lemma 5.5. *The cost $\mathbb{1}^T\overline{c}$ of the stabilizer $(\overline{M}, \overline{y}, \overline{c})$ is at most $(\sum_{e\in E} c_e^*)^2$.*

Proof. Let $\mathcal{K}$ be the set of components such that v_K is M-exposed. The cost of $(\overline{M}, \overline{y}, \overline{c})$ is

$$\sum_{e\in\overline{M}} \overline{c}_e = \sum_{K\in\mathcal{K}} \sum_{\{u,v\}\in\overline{M}\cap K} (\overline{y}_u + \overline{y}_v - 1) + \sum_{u\in Y, v\in X: \{u,v\}\in\overline{M}} (\overline{y}_u + \overline{y}_v - 1) \tag{5}$$

We begin with a bound on the first term in Eq. (5). Let $K \in \mathcal{K}$. Since v_K is exposed by M, we have

$$\sum_{uv\in\overline{M}\cap K} (\overline{y}_u + \overline{y}_v - 1) = \sum_{u\in V(K)} \overline{y}_u^{w_K} - \left(\frac{|V(K)| - 1}{2}\right).$$

We next bound the second term in Eq. (5). For this, we note that $\overline{y}_v = 1/2$ for $v \in X$ with $\{u, v\} \in \overline{M}$. Let $Y' := \{u \in Y : u \text{ is adjacent to an M-exposed vertex}\}$. For $u \in Y \setminus Y'$, we have $\overline{y}_u = 1/2$ and such vertices do not contribute to the sum. Thus, the second term is

$$\begin{aligned}
\sum_{u\in Y}\left(\overline{y}_u - \frac{1}{2}\right) &= \sum_{u\in Y'}\left(\overline{y}_u - \frac{1}{2}\right) = \sum_{u\in Y'}\left(\max_{K\in\mathcal{K}: u\in N_G(V(K))}\left(\overline{y}_u^{w_K} - \frac{1}{2}\right)\right) \\
&\leq \sum_{u\in Y'}\left(\sum_{K\in\mathcal{K}: u\in N_G(V(K))}\left(\overline{y}_u^{w_K} - \frac{1}{2}\right)\right) = \sum_{K\in\mathcal{K}} \sum_{u\in Y'\cap N_G(V(K))}\left(\overline{y}_u^{w_K} - \frac{1}{2}\right) \\
&= \sum_{K\in\mathcal{K}} \sum_{u\in Y\cap N_G(V(K))}\left(\overline{y}_u^{w_K} - \frac{1}{2}\right) = \sum_{K\in\mathcal{K}}\left(\sum_{u\in N_G(V(K))} \overline{y}_u^{w_K} - \frac{|N_G(V(K))|}{2}\right).
\end{aligned}$$

Therefore,

$$\sum_{e\in\overline{M}} \overline{c}_e \leq \sum_{K\in\mathcal{K}}\left(\sum_{v\in V(K)\cup N_G(V(K))} \overline{y}_v^{w_K} - \frac{|V(K)| - 1}{2} - \frac{|N_G(V(K))|}{2}\right)$$

$$\begin{aligned}
&\leq r \max_{K\in\mathcal{K}} f(K) \qquad \text{(since } |\mathcal{K}| = r \text{ and each term in the sum is at most } \max_{K\in\mathcal{K}} f(K)\text{)}\\
&\leq \left(\frac{r + \max_{K\in\mathcal{K}} f(K)}{2}\right)^2 \qquad \text{(since arithmetic mean is at least geometric mean)}\\
&\leq \left(\frac{1 + \sum_{e\in E} c_e^*}{2}\right)^2 \qquad \text{(by Lemma 5.3)}\\
&\leq \left(\sum_{e\in E} c_e^*\right)^2. \quad \square
\end{aligned}$$

The final inequality above uses the fact that $\sum_{e\in E} c_e^* \geq 1$ by Lemma 3.2.

If $OPT > \sqrt{n}$, any solution fulfilling the properties of Theorem 3.1 is a $\sqrt{n}$-approximation as the cost of any such solution is bounded by $\nu(G) \leq n/2$ (setting $c_e := 1$ for every edge e in a maximum matching gives a feasible solution to MFASP). Therefore, Lemmas 5.5 and 5.4 and the construction of $(\overline{M}, \overline{y}, \overline{c})$ imply Theorem 1.3.

5.1. Approximation algorithm of Section 5

Algorithm in graphs with no singletons

1. For each factor-critical component K in $G[X]$:

   (a) For each vertex w in K, solve following LP:

   $$\begin{aligned}
   \ell_{K,w} := \min \sum_{v\in V(K)\cup N_G(V(K))} & y_v - \left(\frac{|V(K)|-1}{2}\right) - \frac{|N_G(V(K))|}{2}\\
   y_i + y_j &\geq 1 \ \forall\ \{i,j\} \in E[V(K)\cup N_G(V(K))]\\
   y_i &\geq 1/2 \ \forall\ i \in N_G(V(K))\\
   y_i &\geq 0 \ \forall\ i \in V(K)\\
   y_w &= 0
   \end{aligned}$$

   (b) Let $f(K) := \min_{w\in V(K)} \ell_{K,w}$.

2. Construct an auxiliary bipartite graph G' from G as follows: Contract each component K in $G[X]$ to a pseudo-vertex v_K and assign edge weight $w_e := f(K)$ for all edges e incident to the contracted vertex v_K. Delete edges in $E[Y]$.

3. Compute a matching M in G' of maximum weight covering Y.

4. For each M-exposed vertex v_K, let $w_K := \operatorname{argmin}_{w\in V(K)} \ell_{K,w}$, let $\overline{y}^{w_K}$ denote the solution y achieving the optimum for ℓ_{K,w_K}.

5. Identify a matching $\overline{M}$: Extend M inside each factor-critical component K: if v_K is matched by M using edge u, b where $u \in V(K), b \in Y$, then extend M using a matching in K that exposes u. If v_K is exposed by M, extend M using a matching in K that exposes w_K. Let $\overline{M}$ denote the resulting matching.

6. Identify a fractional vertex cover $\overline{y}$: For each vertex v_K matched by M, set $\overline{y}_u = 1/2$ for all vertices $u \in V(K)$. For each vertex v_K that is exposed by M, set $\overline{y}_u = \overline{y}_u^{w_K}$ for all vertices $u \in V(K)$. For each vertex $b \in Y$ that is adjacent to a M-exposed v_K, set $\overline{y}_b = \max_{K:v_K \text{ is } M\text{-exposed}} \overline{y}_b^{w_K}$. For each vertex $b \in Y$ with no adjacent M-exposed v_K, set $\overline{y}_b = 1/2$.

7. Identify a feasible MFASP solution $\overline{c}$: Set $\overline{c}(uv) = \overline{y}_u + \overline{y}_v - 1$ for edges $\{u, v\} \in \overline{M}$ and $\overline{c}(uv) = 0$ for edges $\{u, v\} \in E \setminus \overline{M}$.
8. Return $(\overline{M}, \overline{y}, \overline{c})$.

---

6. An exact algorithm for MFASP

In this section, we describe an exact algorithm to solve MFASP in arbitrary graphs (Theorem 1.5). The algorithm is based on the Gallai–Edmonds decomposition $V(G) = X \cup Y \cup Z$ of a graph G and uses a polynomial time exact algorithm to solve MFASP in the factor-critical components in $G[X]$. The runtime of our algorithm grows exponentially only in the size of the Tutte set. Thus, our algorithm is fixed parameter tractable when the parameter is the size of the Tutte set. In particular, the resulting algorithm runs in polynomial-time if the size of the Tutte set is bounded by $O(\log n)$.

Outline of the algorithm. By Theorem 3.1, we know that there exists a subset $S^* \subseteq Y$ and a half-integral minimum fractional stabilizer c^* with a half-integral minimum fractional $(1 + c^*)$-vertex cover solution y^* such that $y^*_v = 1$ for all $v \in S^*$, and $y^*_v = 1/2$ for all $v \in Y \setminus S^*$. This motivates the following subproblem: given a set $\hat{S} \subseteq Y$, find a minimum fractional stabilizer c which admits a minimum fractional $(1 + c)$-vertex cover y satisfying $y_v = 1$ if $v \in \hat{S}$ and $y_v = 1/2$ if $v \in Y \setminus \hat{S}$. Or, decide that no such solution exists. In Section 6.1, we present a polynomial-time algorithm for this subproblem. Repeatedly applying this algorithm to all subsets of the Tutte set and searching for the optimal one gives the optimal solution to MFASP and implies Theorem 1.5.

6.1. Algorithm to find the optimal stabilizer knowing the subset of Tutte vertices with y-value one

Let G be a graph with Gallai–Edmonds decomposition X, Y, Z. In this section, we focus on the following problem: Given a set $\hat{S} \subseteq Y$, find a minimum cost fractional additive stabilizer (M, y, c) among those which have $y_v = 1$ if $v \in \hat{S}$, and $y_v = \frac{1}{2}$ if $v \in Y \setminus \hat{S}$. Let $f(\hat{S}) = \sum_{e \in E} c_e$ denote the cost of such a solution.

We give an overview of the algorithm to compute $f(\hat{S})$. (For a formal description see Algorithm $MFASP(\hat{S})$.) Let (M, y, c) denote the triple of an optimal solution corresponding to $f(\hat{S})$. Let us examine the structure of the optimal solution (M, y, c). Recall that we are restricting c and y to be half-integral.

Finding an optimum with knowledge of matching edges between Y and X. Let us focus on the matching edges in M that link Y to X and argue that it is sufficient to know these links to find an optimal solution c. We consider a component $K \in G[X]$ matched to some vertex in $v \in Y$ by M and distinguish two cases.

(i) K is non-trivial. By Lemma 3.3, we have $c(e) = 0$ for $e \in E[K]$ and $c(\delta(V(K))) = y_v - 1/2$.

(ii) Suppose $K = \{u\}$. If u is not incident to any vertex $w \in Y \setminus \hat{S}$, then we may assign $y_u = 0$ thereby incurring a cost of $c_e = 0$ and this is optimal. Otherwise, the feasible y assigns $y_u = 1/2$ and as before incurs an optimal cost of $y_v - 1/2$ over $\delta(V(K))$.

Next let us consider $K \in G[X]$ that is not matched to any vertex in Y.

(i) Suppose there are no edges $\{v, u\} \in \delta(V(K))$ that are incident to a vertex $u \in Y \setminus \hat{S}$. Then the optimal fractional additive stabilizer c restricted to the set of edges in $E(K) \cup \delta(V(K))$ should also be an optimal fractional additive stabilizer for K itself and vice-versa. Therefore, the stabilizer values on these edges can be computed using the exact algorithm for factor-critical graphs.

(ii) Suppose there are edges $\{v, u\} \in \delta(V(K))$ that are incident to a vertex $u \in Y \setminus \hat{S}$. In this case, the vertex cover values y should satisfy the covering constraints for the edges in $\delta(V(K))$. In particular, $y_v \geq 1/2$ for vertices $v \in V(K)$ which have neighbours in $Y \setminus \hat{S}$. As a consequence, the optimal stabilizer restricted to the set of edges in $E(K) \cup \delta(V(K))$ may not be the optimal stabilizer for K itself. However, the optimal fractional additive stabilizer restricted to the set of edges in $E(K) \cup \delta(V(K))$ should also be an optimal fractional additive stabilizer for a modified graph $\tilde{K}$ obtained from K by adding an extra loop $\{v, v\}$ to each vertex $v \in V(K)$ linked to a vertex $u \in Y \setminus \hat{S}$. Conversely, we can modify an optimal fractional additive stabilizer c over the set of edges in $E(K) \cup \delta(V(K))$ to take the same values as an optimal fractional additive stabilizer for $\tilde{K}$ without losing optimality. Further, we note that we can compute a minimum fractional additive stabilizer in $\tilde{K}$ by calculating the values $\ell_{\tilde{K},w}$ as defined in Lemma 5.1 for every node w that is not incident to a loop in $\tilde{K}$ and output the best.

We observe that if every vertex $v \in V(K)$ has an edge adjacent to a vertex $u \in Y \setminus \hat{S}$, then this necessitates $y_v \geq 1/2$ for every vertex in K and therefore K must necessarily be matched to a vertex in Y.

Computing matching edges between Y and X. From the above discussion, it is clear that the cost of the solution $f(\hat{S})$ does not depend on the precise choice of the edges used to match the components of $G[X]$ to Y but only depends on which components of $G[X]$ are matched by M. Therefore, we can also identify the edges between Y and X in an optimal matching M as follows: Let us denote by $\kappa(K, \hat{S})$ the cost of the stabilizer over the edges in $E[K] \cup \delta(V[K])$ if K is not matched to Y (as observed before, we can compute $\kappa(K, \hat{S})$ by applying an exact algorithm to the factor-critical graph $\tilde{K}$, for example the one presented in Section 5). If K must necessarily be matched, then we set $\kappa(K, \hat{S})$ to infinity. Let T denote the trivial components in $G[X]$ all of whose neighbours are in $\hat{S}$. Let us construct a weighted bipartite graph H from G as follows: Delete Z, delete the edges between vertices in Y, and contract each component K of $G[X]$ to a vertex v_K; replace the multi-edges by a single edge to make it a simple graph and for a vertex $u \in Y$ that is adjacent to some node in K, we introduce weight $\kappa(K, \hat{S})$ on the edge $\{u, v_K\}$.

By the above discussion, a maximum weight matching N in H covering all vertices in Y gives the edges of an optimal matching M between Y and X. Therefore,

$$f(\hat{S}) = \frac{1}{2}\Big(|\hat{S}| - |\{K \in T : N \text{ covers } v_K\}|\Big) + \sum_{K \in G[X]:\ v_K \text{ is exposed by } N} \kappa(K, \hat{S}).$$

Hence, we find a maximum weight matching in H to compute $f(\hat{S})$.

Remark 6.1. Between matching a component in T or a non-trivial factor-critical component, M prefers the latter choice by Lemma 3.2. Thus, assigning $\kappa(K, \hat{S}) = 0$ for $K \in T$ implicitly assumes that components in T are only matched if there is no other choice.

Algorithm $MFASP(\hat{S})$**.**

---

1. For each factor-critical component K in $G[X]$ compute the cost $\kappa(K, \hat{S})$ needed to stabilize K and the edges linking K to vertices in $Y \setminus \hat{S}$ in case K would not be matched to Y. (We discussed above that this can be done in polynomial time.) Let (M^K, c^K, y^K) be an optimal stabilizer for $K \cup \delta(V(K))$ among those with $M^K \cap \delta(K) = \emptyset$.

2. Shrink the components K in $G[X]$ to pseudo-vertices v_K, assign the weight $\kappa(K, \hat{S})$ to all edges linking a Tutte vertex to pseudo-vertex v_K, and compute a bipartite matching $\hat{M}$ of maximum weight covering Y (this is possible in polynomial time).

   If no feasible solution exists, that is, if there exists an unmatched component K with $\kappa(K, \hat{S}) = \infty$, STOP and RETURN INFEASIBLE.

3. Obtain a maximum matching in G by extending $\hat{M}$ as follows:

- for each component K not matched to Y add the matching edges in M^K to $\hat{M}$;
- for each component K having v_K matched to Y, pick a vertex in K that has the corresponding matching edge adjacent to it, say v, and add a maximum matching in K that exposes v to $\hat{M}$;
- for each component C in $G[Z]$ (we note that all these components are even), add an arbitrary perfect matching to $\hat{M}$;

4. Obtain a fractional additive stabilizer as follows:
   - $\hat{c}_e = c_e^K$ for all components K in $G[X]$ that are not matched to Y,
   - $\hat{c}_e = \frac{1}{2}$ for each matching edge $e \in \hat{M}$ linked to a Tutte vertex $v \in \hat{S}$, except if $e = \{v, w\}$ for some vertex w that is a trivial component in $G[X]$ with $N_G(w) \subseteq \hat{S}$, and
   - $\hat{c}_e = 0$ else.
5. Obtain a fractional $(1 + \hat{c})$-vertex cover $\hat{y}$ that satisfies complementary slackness with $\hat{M}$ as follows:
   - $\hat{y}_v = \frac{1}{2}$ for all vertices in Z, all vertices in $Y \setminus \hat{S}$ and all vertices in components K in $G[X]$ that are matched to Y except if v is a trivial component in $G[X]$ with $N_G(v) \subseteq \hat{S}$,
   - $\hat{y}_v = 1$ for all $v \in \hat{S}$,
   - $\hat{y}_v = y_v^K$ for all vertices in components K in $G[X]$ that are not matched to Y, and
   - $\hat{y}_v = 0$ for all vertices v that are trivial components in $G[X]$ with $N_G(v) \subseteq \hat{S}$ and matched to Y.
6. Return $(\hat{M}, \hat{y}, \hat{c})$ and $f(\hat{S}) := \sum_{e \in E} \hat{c}_e$;

Remark 6.2. As mentioned earlier, not every possible choice of $\hat{S}$ has a fractional additive stabilizer c which has $y_v = 1$ if $v \in \hat{S}$, and $y_v = 1/2$ if $v \in Y \setminus \hat{S}$ for a half-integral minimum fractional $(1 + c)$-vertex cover y. For example, consider a graph where $Z = \emptyset$, $Y = \{v\}$ and $G[X]$ consists of two triangles whose nodes are all connected to v. Then y_v must be 1 and $\hat{S} = \emptyset$ is not feasible. The algorithm detects these cases in Step 2.

7. Approximation for graphs with many nontrivial components

We can also use the algorithm that computes $f(\hat{S})$ from the previous section to obtain an approximation algorithm for graphs that have a large number of non-trivial factor-critical components in the Gallai–Edmonds decomposition.

Theorem 7.1. *For a graph G with Gallai–Edmonds decomposition $V(G) = X \cup Y \cup Z$, let $\mathcal{C}^+$ denote the set of nontrivial components in X. If $|\mathcal{C}^+| \geq (1 + 1/k)|Y|$ for $k > 0$, then there is a $(k/2 + 1)$-approximation algorithm for MFASP.*

Proof. Let (M^*, y^*, c^*) be an optimal solution for MFASP with cost $c^*(E)$ and $X_1, \ldots, X_r$ be the non-trivial components of $G[X]$. We first note that the number of unmatched non-trivial components of X is at least $\lceil (1/k)\,|Y| \rceil$. We know by Lemma 3.2 that the optimal stabilizer pays at least 1 over the edges in each of these components. This yields

$$c^*(E) = \sum_{e \in E} c_e^* \geq \sum_{i=1}^{r} \sum_{e \in X_i} c_e^* \geq \frac{|Y|}{k}.$$

Let $\hat{S} = Y$, i.e., fix $y_t = 1$ for all vertices t in the Tutte set, and calculate an optimal solution $(\hat{M}, \hat{y}, \hat{c})$ corresponding to $f(\hat{S})$. We observe that the optimal solution corresponding to $f(\hat{S})$ can be computed efficiently using the algorithm from Section 6.1. Recall that $y^*_v \geq 1/2$ for each vertex v in the Tutte set. Therefore,

$$\begin{aligned}\hat{c}(E) = \hat{y}(V) - |\hat{M}| &= \sum_{v\in V}\hat{y}_v - |M^*| = \sum_{v\in X}\hat{y}_v + \sum_{v\in Y}\hat{y}_v - |M^* \setminus E[Z]| \\ &\leq \sum_{v\in X} y^*_v + \frac{|Y|}{2} + \sum_{v\in Y} y^*_v - |M^*\setminus E[Z]| \\ &\leq \sum_{v\in X} y^*_v + \left(\frac{k}{2}\right)c^*(E) + \sum_{v\in Y} y^*_v - |M^*\setminus E[Z]| = \left(\frac{k}{2}+1\right)c^*(E),\end{aligned}$$

which finishes the proof. In the first inequality above, we have used $\sum_{v\in X}\hat{y}_v \leq \sum_{v\in X} y^*_v$ since y^* restricted to X is also feasible for the auxiliary problem with y_v fixed to one on all Tutte vertices. $\square$

References

[1] G. Chalkiadakis, E. Elkind, M. Wooldridge, Computational aspects of cooperative game theory, Synth. Lect. Artif. Intell. Mach. Learn. 5 (6) (2011) 1–168.

[2] N. Nisan, T. Roughgarden, É. Tardos, V. Vazirani, Algorithmic Game Theory, Cambridge University Press, Cambridge, UK, 2007.

[3] Y. Shoham, K. Leyton-Brown, Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations, Cambridge University Press, 2008.

[4] D. Easley, J. Kleinberg, Networks, Crowds, and Markets: Reasoning about a Highly Connected World, Cambridge University Press, 2010.

[5] M. Jackson, Social and Economic Networks, Princeton University Press, 2008.

[6] J. Kleinberg, É. Tardos, Balanced outcomes in social exchange networks, in: Proc. 40th Annual ACM Sympos. Theory Comput., 2008, pp. 295–304.

[7] J. Nash Jr., The bargaining problem, Econometrica (1950) 155–162.

[8] L. Shapley, M. Shubik, The assignment game I: The core, Internat. J. Game Theory 1 (1) (1971) 111–130.

[9] M. Bateni, M. Hajiaghayi, N. Immorlica, H. Mahini, The cooperative game theory foundations of network bargaining games, in: Automata, Languages and Programming, Springer, 2010, pp. 67–78.

[10] A. Bock, K. Chandrasekaran, J. Könemann, B. Peis, L. Sanità, Finding small stabilizers for unstable graphs, Math. Program. 154 (1–2) (2015) 173–196.

[11] S. Ahmadian, H. Hosseinzadeh, L. Sanità, Stabilizing network bargaining games by blocking players, in: Integer Programming and Combinatorial Optimization - 18th International Conference, IPCO 2016, Liège, Belgium, June 1-3, 2016, Proceedings, 2016, pp. 164–177.

[12] T. Ito, N. Kakimura, N. Kamiyama, Y. Kobayashi, Y. Okamoto, Efficient stabilization of cooperative matching games, Theoret. Comput. Sci. 677 (2017) 69–82.

[13] S. Khot, Ruling out ptas for graph min-bisection, dense k-subgraph, and bipartite clique, SIAM J. Comput. 36(4) (2005) 1025–1071.

[14] P. Manurangsi, Almost-polynomial ratio ETH-hardness of approximating densest K-subgraph, in: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC '17), 2017, pp. 954–961.

[15] A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, A. Vijayaraghavan, Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k-subgraph, in: STOC 2010, 2010, pp. 201–210.

[16] E. Balas, Integer and fractional matchings, North-Holland Math. Stud. 59 (1981) 1–13.

[17] J. Uhry, Sur le problème du couplage maximal [French], Rev. Française Autom., Inf. Rech. Opér. Rech. Opér. 9 (3) (1975) 13–20.

[18] J. Edmonds, Maximum matching and a polyhedron with 0, l-vertices, J. Res. Nat. Bur. Stand. B 69 (1965) 125–130.

[19] J. Edmonds, Matroid partition, in: G. Dantzig, A. Veinott (Eds.), Mathematics of the Decision Sciences, AMS, Providence, RI, 1968, pp. 335–345.

[20] T. Gallai, Maximale Systeme unabhängiger Kanten [German], Magyar Tud. Akad. Mat. Kut. Int. Közl 9 (1964) 401–413.

[21] R. Downey, M. Fellows, Parametrized Complexity, Springer, 1999.

[22] N. Alon, A. Shapira, Every monotone graph property is testable, SIAM J. Comput. 38 (2) (2008) 505–522.

[23] N. Alon, A. Shapira, B. Sudakov, Additive approximation for edge-deletion problems, in: Proc. 46th Annual IEEE Sympos. Foundations Comput. Sci., 2005, pp. 419–428.

[24] S. Mishra, V. Raman, S. Saurabh, S. Sikdar, C. Subramanian, The Complexity of König subgraph problems and above-guarantee vertex cover, Algorithmica 61 (4) (2011) 857–881.

[25] J. Könemann, K. Larson, D. Steiner, Network bargaining: Using approximate blocking sets to stabilize unstable instances, Theory Comput. Syst. 57 (3) (2015) 655–672.

[26] P. Biró, M. Bomhoff, P.A. Golovach, W. Kern, D. Paulusma, Solutions for the stable roommates problem with payments, Theoret. Comput. Sci. 540 (2014) 53–61.

[27] B. Korte, J. Vygen, Combinatorial Optimization: Theory and Algorithms, Springer, Berlin, Germany, 2002.
[28] A. Schrijver, Combinatorial Optimization: Polyhedra and Efficiency, Vol. 24, Springer, Berlin, Germany, 2003.
[29] G.L. Nemhauser, L.E. Trotter Jr., Vertex packings: Structural properties and algorithms, Math. Program. 8 (1) (1975) 232–248.
[30] L. Lovász, A note on factor-critical graphs, Studia Sci. Math. Hungar. 7 (1972) 279–280.
[31] M. Hajiaghayi, K. Jain, The prize-collecting generalized steiner tree problem via a new approach of primal-dual schema, in: SODA 2006, 2006, pp. 631–640.
[32] C. Berge, Sur le couplage maximum d'un graph [French], C.R. Hebd. Seances Acad. Sci. (Paris) (1958) 258–259.
[33] I. Dinur, D. Steurer, Analytical approach to parallel repetition, in: Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014, 2014, pp. 624–633.