

Finding small stabilizers for unstable graphs

Adrian Bock¹, Karthekeyan Chandrasekaran², Jochen Könemann³, Britta Peis⁴, Laura Sanità³

¹ EPFL, Switzerland. adrianaloysius.bock@epfl.ch

² Harvard University, USA. karthe@seas.harvard.edu

³ University of Waterloo, Canada. {jochen,laura.sanita}@uwaterloo.ca

⁴ RWTH Aachen University, Germany. britta.peis@oms.rwth-aachen.de

Abstract. An undirected graph $G = (V, E)$ is *stable* if its inessential vertices (those that are exposed by at least one maximum matching) form a stable set. We call a set of edges $F \subseteq E$ a *stabilizer* if its removal from G yields a stable graph. In this paper we study the following natural edge-deletion question: given a graph $G = (V, E)$, can we find a minimum-cardinality stabilizer?

Stable graphs play an important role in cooperative game theory. In the classic *matching game* introduced by Shapley and Shubik [19] we are given an undirected graph $G = (V, E)$ where vertices represent players, and we define the *value* of each subset $S \subseteq V$ as the cardinality of a maximum matching in the subgraph induced by S . The *core* of such a game contains all *fair* allocations of the *value* of V among the players, and is well-known to be non-empty iff graph G is *stable*. The stabilizer problem addresses the question of how to modify the graph to ensure that the core is non-empty.

We show that this problem is vertex-cover hard. We then prove that there is a minimum-cardinality stabilizer that avoids some maximum matching of G . We use this insight to give efficient approximation algorithms for sparse graphs and for regular graphs.

1 Introduction

Given an undirected graph $G = (V, E)$, a subset of edges $M \subseteq E$ is a *matching* if every vertex $v \in V$ is incident to at most one edge in M . Dually, a subset of vertices $U \subseteq V$ is called *vertex cover* if every edge has at least one endpoint in U . The corresponding optimization problems of finding a matching and vertex cover of largest and smallest size, respectively, have a rich history in the field of Combinatorial Optimization. Relaxing canonical integer programming formulations for these problems yields the following primal-dual pair of linear programs:

$$\nu_f(G) := \max\{\mathbf{1}^T x : x(\delta(v)) \leq 1 \forall v \in V, x \geq 0\} \quad (\text{P})$$

where $\delta(v)$ denotes the set of edges incident to v , and

$$\tau_f(G) := \min\{\mathbf{1}^T y : y_u + y_v \geq 1 \forall uv \in E, y \geq 0\}. \quad (\text{D})$$

We will henceforth refer to feasible solutions of (P) and (D) as *fractional matchings* and *vertex covers*, respectively. An application of duality theory easily yields

$$\nu(G) \leq \nu_f(G) = \tau_f(G) \leq \tau(G)$$

where $\nu(G)$ and $\tau(G)$ denote the size of a maximum matching and a minimum vertex cover, respectively.

In this paper, we study graphs G with the property $\nu(G) = \tau_f(G)$. The class of such graphs, known as *stable* graphs, subsumes the well-studied class of *König-Egerváry* (KEG) graphs (e.g., see [20, 13, 14, 15]) for which $\nu(G) = \tau(G)$. Stable graphs arise quite naturally in the study of cooperative *matching games* introduced by Shapley and Shubik in their seminal paper [19]. An instance of this game is associated with an undirected graph $G = (V, E)$ where vertices represent players. We define the *value* of each subset $S \subseteq V$ as the cardinality of a maximum matching in the subgraph $G[S]$ induced by S , and the *core* of the game consists of all *stable* allocations of total value $\nu(G)$ among the vertices in V in which no coalition of vertices has an incentive to deviate. This is formally defined as

$$\text{core}(G) := \left\{ y \in \mathbb{R}_+^V : \sum_{v \in S} y_v \geq \nu(G[S]) \ \forall S \subseteq V, \sum_{v \in V} y_v = \nu(G) \right\}.$$

It is well-known (e.g., see [8]) that $\text{core}(G)$ is non-empty iff G is stable.

Matching games in turn are closely related to *network bargaining*, a natural, recent generalization of Nash's famous bargaining solution [16] to networks due to Kleinberg and Tardos [11]. Here, we are given an undirected graph $G = (V, E)$ whose vertices correspond to players, and whose edges correspond to potential unit-value deals between the incident players. Each player is allowed to engage in at most one deal with one of its neighbors. Hence, a permissible outcome is naturally associated with a matching M among the vertices of G , as well as an allocation $y \in \mathbb{R}_+^V$ of $|M|$ among M 's endpoints. Kleinberg and Tardos define an allocation to be *stable* if $y_u + y_v \geq 1$ for all $uv \in E$. The authors further define an *outside option* α_u for each vertex $u \in V$ as

$$\alpha_u := \max\{1 - y_v : uv \in \delta(u) \setminus M\},$$

and say that an outcome (M, y) is *balanced* if for every edge $uv \in M$, the surplus $1 - \alpha_u - \alpha_v$ is split evenly among u and v . The main result in [11] is that an instance of network bargaining has a stable outcome iff it has a balanced one. One now realizes (see also [5]) that a stable outcome exists iff the core of the underlying matching game instance is non-empty and hence iff G is stable.

In this paper, we focus on *unstable* instances of the matching game, where the core is empty. Our motivating goal is to establish strategies for *stabilizing* such instances in the *least intrusive way*; i.e., we would like to alter the input graph in few places and ideally maintain the value of the grand coalition formed by the set of vertices V in the process. The following natural edge-deletion *stabilizer*

problem formalizes this: given a graph $G = (V, E)$, find the smallest edge set $F \subseteq E$ such that the subgraph $G \setminus F := (V, E \setminus F)$ is stable.

Stable graphs form a proper superclass of KEGs which in turn form a superclass of bipartite graphs. Readers familiar with the literature of bipartite graphs would immediately recognize that the stabilizer problem closely resembles the optimization problems of deleting the minimum number of edges to convert a given graph into a KEG or a bipartite graph, both of which have been well-studied in the literature (e.g., see [1, 15]). The investigation of structural properties of unstable graphs has a long history (e.g., see [21, 3, 17]), but there are few algorithmic results on how to convert an unstable graph to a stable graph. Biró et al. [6] recently studied the minimum stabilizer problem in the weighted setting, where maximum-weight matchings are considered instead of maximum-cardinality matchings. The authors showed that the problem is NP-hard in this case, and leave the complexity of the question in the unweighted setting open.

1.1 Our results

We first show that removing a minimum stabilizer from a given graph G does not reduce the cardinality of the maximum matching. Hence the value of the grand coalition of the associated matching game remains the same.

Theorem 1. *For every minimum stabilizer F , we have $\nu(G \setminus F) = \nu(G)$.*

The proof of this theorem is algorithmic: given any stabilizer F , we can efficiently find a maximum matching M in G and a stabilizer F' such that $F' \subseteq F$ and $M \cap F' = \emptyset$. The last equality implies that M is still a maximum matching in $G \setminus F'$. The result motivates the following intermediate *M-stabilizer* problem: given a maximum matching M , find a minimum-cardinality stabilizer F_M that is disjoint from M . In the network bargaining setting, this question asks how to convert a specific maximum matching into one with a stable allocation through minimal edge deletions in the underlying network. Biró et al. [6] previously showed that this problem is NP-hard. We strengthen the hardness result and complement it with a tight algorithmic counterpart.

Theorem 2. *The M-stabilizer problem is NP-hard, and no efficient $(2 - \varepsilon)$ -approximation algorithm exists for any $\varepsilon > 0$ assuming the Unique Games Conjecture [10]. Furthermore, the M-stabilizer problem admits an efficient 2-approximation algorithm.*

The hardness proof employs an approximation preserving reduction from vertex cover. The approximation algorithm uses linear programming, and one shows that a suitable linear programming relaxation for the problem has a half-integral optimal solution. Turning to the stabilizer problem, we first extend the hardness result obtained for *M-stabilizers* answering the open question in [6]. Interestingly, our hardness result holds for *factor-critical* graphs (see next subsection for the definition).

Theorem 3. *The stabilizer problem is NP-hard. Furthermore, no efficient $(2 - \varepsilon)$ -approximation algorithm exists for any $\varepsilon > 0$ assuming the Unique Games Conjecture [10].*

Theorems 1 and 2 suggest that the crux of the hardness of the stabilizer problem lies in finding the *right* maximum matching that survives the removal of a minimum stabilizer. Once such a matching is found one could indeed *simply* apply our 2-approximation for the M -stabilizer problem. However, not every maximum matching survives the removal of a minimum stabilizer. In fact, for two different maximum matchings M and M' , the cardinality of F_M and $F_{M'}$ can differ by a factor of $\Omega(|V|)$ even on a planar factor-critical graph. In Section 3.1, we present an approximation algorithm whose approximation factor depends on the *sparsity* of the graph. We say that a graph $G = (V, E)$ is ω -sparse if $|E(S)| \leq \omega |S|$ for all vertex subsets $S \subseteq V$.

Theorem 4. *There exists an efficient $O(\omega)$ -approximation algorithm for the stabilizer problem, where ω is the sparsity of the input graph.*

We note that the above result implies a constant factor approximation algorithm for graphs with constant sparsity, e.g., *planar graphs*. We do not know whether a constant factor approximation algorithm can be developed for arbitrary graphs. However, we give a 2-approximation algorithm for regular graphs (graphs where all vertex degrees are equal). In the network bargaining setting, this gives a 2-approximation algorithm to stabilize networks in which every player has the same number of potential deals to make.

Theorem 5. *There exists an efficient 2-approximation algorithm for the stabilizer problem in regular graphs.*

The analysis of our algorithm nicely combines some classic results about matchings and vertex covers such as the structure of basic solutions of (P) and (D) and the Edmonds-Gallai decomposition.

The proof of Theorem 1 is presented in Section 2, and that of Theorems 4 and 5 are presented in Section 3. The proofs of Theorems 2 and 3 are deferred to the full version of the paper [7].

1.2 Related work

The problem of removing vertices or edges from a graph in order to attain a certain graph property is natural, and thus not surprisingly, its variants have been studied extensively. Much of the work on deletion problems addresses *monotone* graph properties (e.g., see [22, 2]) that are invariant under edge-removal or vertex-removal. Crucially, graph stability is not a monotone property as one easily verifies: the triangle is not stable, and adding a single pendant edge to one of its vertices yields a stable graph.

Our work is closely related to that of Mishra et al. [15] on vertex-removal and edge-removal problems to attain the König-Egerváry graph property. Similar

to stability, KEG is not a monotone property. Mishra et al. showed that it is NP-hard to approximate the corresponding edge-deletion problem to within a factor of 2.88. Assuming the Unique Games Conjecture, no constant-factor approximation may exist for the problem. We note that the reductions used in [15] will likely not be helpful for proving hardness for the stabilizer problem as the graphs constructed in the reduction are stable. On the positive side, the authors show that, for a given graph $G = (V, E)$ one can efficiently find a KEG (and hence stable) subgraph with at least $3|E|/5$ edges.

The recent paper by Könemann et al. [12] addressed the related, NP-hard problem of finding a minimum-cardinality *blocking set* in an input graph $G = (V, E)$. Here one wants to find a set of edges $F \subseteq E$ such that $G \setminus F$ has a fractional vertex cover of size at most $\nu(G)$. Importantly, the resulting graph $G \setminus F$ is not required to be stable; indeed, the cardinality of a minimum blocking set can differ from the cardinality of a minimum stabilizer by a factor of $\Omega(|V|)$.

1.3 Preliminaries

Given an undirected graph G and a matching M in G , a path is called *M-alternating* if it alternates edges from M and those from $E \setminus M$. An odd cycle of length $2k + 1$ in which exactly k edges are in M is called an *M-blossom*. An *M-flower* is an even *M-alternating* path from an exposed vertex to a vertex u such that there exists a blossom through u . For a subset of vertices $S \subseteq V$, we use $E(S)$ to denote the set of edges in the graph induced by S and $G[S]$ to denote the subgraph induced by S . A graph $G = (V, E)$ is called *factor-critical* if for all $v \in V$, $G[V \setminus \{v\}]$ has a perfect matching; i.e., a matching that does not expose any vertex. A vertex v is called *inessential* for G if there exists a maximum matching M that exposes v , and *essential* otherwise. In this paper, we will also use the following characterization of stable graphs.

Theorem 6 ([11]). *The following are equivalent: (i) G is stable, (ii) The set of inessential vertices of G forms a stable set, (iii) G contains no M -flower for any maximum matching M . Moreover, if G is not stable, then G contains an M -flower for every maximum matching M .*

Given a graph G , the Edmonds-Gallai decomposition is a partition of its vertex set into three parts $B(G)$, $C(G)$, $D(G)$, where $B(G)$ is the set of inessential vertices, the set $C(G)$ consists of the neighbors of $B(G)$ and $D(G) = V \setminus (B(G) \cup C(G))$. We list several standard but useful properties. For a proof see, e.g., [18].

Theorem 7. *Given a graph G , the Edmonds-Gallai decomposition of the graph $B(G), C(G), D(G)$ can be computed in polynomial time. Further, we have the following properties.*

1. *Each component of $G[B(G)]$ is factor-critical.*
2. *Every maximum matching M in G exposes at most one vertex in each component K of $G[B(G)]$.*

3. If U is a non-trivial factor-critical component in $G[B(G)]$ (i.e., a factor-critical component with more than one vertex), then $\nu(G \setminus E(U)) < \nu(G)$.

The following proposition is a consequence of the Edmonds-Gallai decomposition theorem, which follows from classic results by Balas [3] and Pulleyblank [17]. We include its proof in the full version of this paper [7].

Proposition 1. *Let M be a maximum matching in G that also matches the maximum possible number of isolated vertices in $G[B(G)]$. Let k be the number of non-trivial factor-critical components with at least one vertex exposed by M . Then $k = 2(\nu_f(G) - \nu(G))$.*

2 Maximum matchings and minimum stabilizers

We first show that the deletion of any minimum stabilizer does not alter the cardinality of the maximum matching.

Proof (of Theorem 1). Let F be a minimum stabilizer. Find a maximum matching M in G such that $|M \cap F|$ is minimum. Suppose $|M \cap F| \neq 0$.

Consider $G' := G \setminus (F \setminus M)$, the graph obtained by removing all the edges of $F \setminus M$ from G . Clearly M is still a maximum matching in G' . However, since F is minimum, G' is not stable. By Theorem 6, this implies that there exists an M -flower in G' starting at an M -exposed vertex w .

Suppose the M -flower contains an edge $uv \in F$. Then, $uv \in M$, since all other edges from F have been removed in G' . Therefore, we can find an even M -alternating path P from w to either u or v . Switching along the edges of this path, we obtain another maximum matching $M' = M \Delta P$ in G with $|F \cap M'| < |F \cap M|$, a contradiction.

It follows that the M -flower does not contain any edge from F , and therefore the M -flower also exists in $G \setminus F$. However, since $G \setminus F$ is stable, this implies that $M \setminus F$ is not a maximum matching in $G \setminus F$. Apply Edmonds' maximum matching algorithm on the graph $G \setminus F$ initialized with the matching $M \setminus F$, and construct an $M \setminus F$ -alternating tree starting with the exposed vertex w . There are two possibilities: either we find an augmenting path P or a frustrated tree rooted at w . In the first case, the path P starts with w and ends with a $M \setminus F$ -exposed vertex, say w' . However, such a path cannot exist in G because M is a maximum matching, and therefore w' must have been incident to an edge $f \in M \cap F$. Also, note that the path P is in $G \setminus F$. Hence, $P + f$ is an even M -alternating path in G containing exactly one edge in $M \cap F$. Switching along the edges of this path, we obtain another maximum matching $M' = M \Delta P$ in G with $|F \cap M'| < |F \cap M|$, a contradiction.

The only remaining possibility is that we find a frustrated tree T rooted at w . Let $G[T] = (V_T, E_T)$ be the graph induced by all vertices in the frustrated tree T (after expanding pseudonodes). In this case, $M \cap E_T$ is a maximum matching in $G[T]$, and the M -flower is contained in E_T . However, if we continue Edmonds' algorithm, it would remove the vertices of the frustrated tree, and continue

running in the resulting subgraph to find a maximum matching. Therefore it ends by computing a maximum matching M^* in $G \setminus F$ with $M^* \cap E_T = M \cap E_T$. Therefore, we have a M^* -flower in $G \setminus F$, again a contradiction. \square

We remark here that the above proof is algorithmic, therefore given a stabilizer F , we can find in polynomial time a maximum matching M in G and another stabilizer $F' \subseteq F$ such that $M \cap F' = \emptyset$. The first step of computing a maximum matching M in G with minimum intersection with F can be done by assigning a cost of one to the edges in F , zero to the rest of the edges, and computing a min-cost matching in G of cardinality $\nu(G)$.

We next prove a lower bound on the cardinality of a stabilizer.

Theorem 8. *For every minimum stabilizer F , we have $|F| \geq 2(\nu_f(G) - \nu(G))$.*

Proof. Let $B(G), C(G), D(G)$ denote the Edmonds-Gallai decomposition and let M be a maximum matching in G that also matches the maximum possible number of isolated vertices in $G[B(G)]$. Let U_1, \dots, U_k denote the non-trivial components in $G[B(G)]$ with at least one vertex exposed by M . Let F be a minimum stabilizer and $H = G \setminus F$. For each component U_1, \dots, U_k , at least one vertex $v_i \in U_i$ becomes essential in H . Suppose not, then all vertices of some U_i are inessential in H . This implies that F contains all edges in $G[U_i]$. Thus, by Theorem 7, we have that $\nu(H) < \nu(G)$, a contradiction to Theorem 1.

Pick a maximum matching N in H . Then, N will cover all these vertices v_1, \dots, v_k that are essential in H . Since $G[U_i]$ is factor-critical and M matches all but one vertex in U_i using edges in $G[U_i]$, we may assume without loss of generality, that M misses all these vertices. The graph $M \Delta N$ is a disjoint union of even cycles and even paths since $|M| = |N| = \nu(G)$. Consider the k disjoint paths starting at the vertices v_1, \dots, v_k in $M \Delta N$. We observe that at least one of the M edges in each of these paths should belong to F , otherwise we can obtain a maximum matching in H that exposes the starting vertex v_i , thus contradicting $v_i \notin B(H)$. Hence $|F| \geq k$. The result follows by Proposition 1. \square

3 Finding small stabilizers

In this section, we return to the problem of finding small stabilizers. The following two sections present algorithms for the problem in sparse, and regular graphs, respectively.

3.1 An $O(\omega)$ -approximation algorithm for sparse graphs

Before proving Theorem 4, we state and prove the following lemma that is the main ingredient of our algorithm.

Lemma 1. *Let G be a graph with $\nu_f(G) > \nu(G)$. There exists an efficient algorithm to find a set of edges L with $|L| = O(\omega)$, such that*

$$(i) \quad \nu(G \setminus L) = \nu(G),$$

$$(ii) \nu_f(G \setminus L) \leq \nu_f(G) - \frac{1}{2}.$$

In other words, Lemma 1 shows that we can find a small subset of edges to remove from G without decreasing the size of the maximum matching but reducing the size of the minimum fractional vertex cover. The proof of Lemma 1 will use two classic results on the structure of fractional and integral matchings.

Theorem 9. [4] *Every basic feasible solution to (P) has components equal to 0, 1 or $\frac{1}{2}$, and the edges with half integral components induce vertex disjoint cycles.*

Theorem 10. [3, 21] *Let \hat{x} be a maximum fractional matching in a graph G having half integral fractional components for a minimum number of odd cycles C_1, \dots, C_q . Let $\hat{M} := \{e \in E : \hat{x}_e = 1\}$ and M_i be a maximum matching in C_i . Then $M = \hat{M} \cup M_1 \cup \dots \cup M_q$ is a maximum matching in G . Moreover, such \hat{x} and M can be found in time polynomial in the number of vertices.*

We are now ready to prove Lemma 1.

Proof (Proof of Lemma 1). Consider \hat{x} and M as in Theorem 10 for the graph G . By duality theory, there exists a fractional vertex cover y with $\mathbf{1}^T y = \mathbf{1}^T \hat{x}$ satisfying complementary slackness conditions with \hat{x} . Moreover, we can always find such a vector y with half integral components (e.g., see [9]). We will give an efficient algorithm to find a vertex u with the following properties:

- (a) $y_u = \frac{1}{2}$,
- (b) $L_u := \{uv : y_w = \frac{1}{2}\}$ satisfies $\nu(G \setminus L_u) = \nu(G)$ and $|L_u| \leq 4\omega$.

First, let us argue that (a) + (b) implies the result. Assume we can find such a vertex u . The only non-trivial conclusion that needs to be verified is that $\nu_f(G \setminus L_u) \leq \nu_f(G) - 1/2$. Consider the vector y' defined as $y'_v = y_v$ for all $v \neq u$ and $y'_u = 0$ otherwise. Then y' is a fractional vertex cover for $G \setminus L_u$ (vertex u cannot be adjacent to vertices with y -value zero because y is a fractional vertex cover for G).

Now let us prove that a vertex u satisfying (a) + (b) can be found efficiently. Consider an arbitrary cycle in \hat{x} , e.g., C_1 . Since $\hat{x}_e > 0$ for every edge $e = uv$ in C_1 , it follows that the vertex cover constraint is tight (i.e., $y_u + y_v = 1$ holds) for all edges in C_1 , and therefore $y_v = \frac{1}{2}$ for all vertices in C_1 .

Set $H := C_1$, and *mark* all vertices in C_1 . Note that C_1 is an odd cycle, therefore if we remove any subset of edges incident to one marked vertex in H , then we do not decrease the size of a maximum integral matching in the resulting graph. Repeat the following process, which will maintain a collection of four invariants for the graph H : (i) Every vertex in H has y -value $\frac{1}{2}$, (ii) removing any subset of edges incident to one marked vertex of H does not decrease the size of a maximum matching, (iii) from any marked vertex, there is an even-length M -alternating path to C_1 , (iv) at least half of the vertices of H are marked. All properties clearly hold initially when H consists of C_1 only.

1. If there is a marked vertex in H with $|L_u| \leq 4\omega$, then u satisfies properties (a) and (b). STOP.

2. Otherwise, consider an arbitrary marked vertex u in H that is adjacent to a vertex $w \notin H$ with $y_w = \frac{1}{2}$. Such a w must be matched in M as otherwise, we could obtain an M -augmenting path in G by concatenating wu , the even length M -alternating u, C_1 -path guaranteed by property (iii) and an appropriate even-length alternating path along C_1 to the M -exposed vertex on C_1 .
3. Let z be the vertex matched to w by M . By complementary slackness, $y_z = \frac{1}{2}$. Add w and z to H and mark z . Go to 1.

It is straight-forward to verify that properties (i)–(iv) continue to hold throughout the execution of the above process. Thus, it only remains to show that we can always find a vertex w as specified in Step 2 above; i.e., if all marked vertices u have $|L_u| > 4\omega$, then there exists a marked vertex in H that is adjacent to a vertex $w \notin H$ with $y_w = 1/2$. Suppose not. Consider the subgraph $G[H]$ induced by the vertices in H . This subgraph has the property that the degree of every marked vertex u in $G[H]$ is at least $|L_u| > 4\omega$. However, by (iv), the number of marked vertices is more than half the total number of vertices in $G[H]$. This contradicts the ω -sparsity of $V(H)$ in G . Finally, it is easy to see that the above process runs in polynomial time.

With this Lemma at hand, we are now ready to prove our main theorem. We will use the following algorithm:

Algorithm 1.

INITIALIZE $G' = G$.

FOR $i = 1, \dots, 2(\nu_f(G) - \nu(G))$:

1. Let L be the set of edges returned by the algorithm in Lemma 1 when its input is the current graph G' .
 2. Set $G' \leftarrow G' \setminus L$.
 3. If G' is stable, STOP.
-

Proof (Proof of Theorem 4).

Let G be an unstable graph. We use Algorithm 1. We will now prove that (a) whenever the above algorithm stops, the current graph G' is stable, and (b) the total number of edges removed during the complete execution of the algorithm is $O(\omega) \cdot |F^*|$, where F^* is a minimum stabilizer. Clearly (a) + (b) implies the result.

First, let us argue about stability. If the algorithm stops in step (iii) for some iteration $i < 2(\nu_f(G) - \nu(G))$, this is clear. So we may assume that the algorithm stops after performing all $2(\nu_f(G) - \nu(G))$ iterations. The graph G' output at this point has $\nu_f(G') \leq \nu_f(G) - \frac{1}{2}(2(\nu_f(G) - \nu(G))) = \nu(G) = \nu(G')$. This is because, by Lemma 1, in each iteration the size of a minimum fractional vertex cover decreases by at least $\frac{1}{2}$ while the size of the maximum matching is maintained. Hence, by definition of stability, G' is stable.

By Lemma 1, in each iteration we remove $O(\omega)$ edges and the total number of iterations is at most $2(\nu_f(G) - \nu(G))$. The bound on the approximation factor

follows from Theorem 8. The running time bound also follows since the number of applications of the algorithm in Lemma 1 is at most $2(\nu_f(G) - \nu(G)) \leq |F^*| \leq |E|$ times. \square

We end the section with an observation about our algorithm that will be useful for our approximation results on regular graphs.

Proposition 2. *The stabilizer output by Algorithm 1 has size at most $2(\nu_f(G) - \nu(G)) \cdot \Delta(G)$, where $\Delta(G)$ is the maximum degree of a vertex in G .*

Proof. In each iteration of the algorithm, we remove a subset of edges incident to some vertex. Therefore we remove at most $\Delta(G)$ edges in each iteration. Further, the number of iterations is at most $2(\nu_f(G) - \nu(G))$. \square

3.2 A 2-approximation algorithm for regular graphs

In this section, we give a 2-approximation algorithm for solving the stabilizer problem in regular graphs.

Proof. (Proof of Theorem 5) We use Algorithm 1. Consider a d -regular graph G , i.e., a graph where every vertex has degree d . Let $k := 2(\nu_f(G) - \nu(G))$. By Proposition 2, the size of F output by the algorithm is at most kd . We complete the proof by showing that every stabilizer in G is of size at least $kd/2$.

Consider the Edmonds-Gallai decomposition of G , namely $B(G), C(G), D(G)$. Let S denote the isolated vertices in $G[B]$. Consider a maximum matching M in G that also matches the maximum possible number of vertices in S . By Proposition 1, the number of non-trivial factor-critical components in $G[B(G)]$ with at least one vertex exposed by M is equal to k .

Let S_u denote the vertices in S that are exposed by M . We first observe that the size $\nu(G)$ of the maximum matching in G is $(|V| - k - |S_u|)/2$. Consider the following primal and dual linear programs.

$$\begin{array}{ll}
 \min \sum_{e \in E} z_e & (\mathcal{P}) \\
 y_u + y_v + z_{uv} \geq 1 \quad \forall uv \in E \\
 \sum_{u \in V} y_u = \nu(G) \\
 y, z \geq 0
 \end{array}
 \qquad
 \begin{array}{ll}
 \max \sum_{e \in E} \alpha_e - \gamma \nu(G) & (\mathcal{D}) \\
 \alpha(\delta(u)) \leq \gamma \quad \forall u \in V \\
 0 \leq \alpha \leq 1
 \end{array}$$

By setting z to be the indicator vector of the minimum stabilizer, we can obtain y such that (y, z) is a feasible solution to the primal program. This is because, if z is the indicator vector of a stabilizer in G , then by definition there exists a fractional vertex cover y in $G \setminus \text{Support}(z)$ with size equal to $\nu(G \setminus \text{Support}(z))$. We also know by Theorem 1 that for every minimum stabilizer F , $\nu(G \setminus F) = \nu(G)$.

Thus, the primal program is a relaxation of the stabilizer problem. Consequently, the objective value of any feasible solution to the dual program is a lower bound on the size of a minimum stabilizer. We will provide a dual feasible solution with objective value at least $kd/2$.

Consider the dual solution ($\gamma = d$, $\alpha_e = 1 \forall e \in E$). Since the graph is d -regular we have that $\alpha(\delta(u)) = d$. Thus, all dual constraints are satisfied and hence, it is a dual feasible solution. The objective value is

$$\sum_{e \in E} \alpha_e - \gamma \nu(G) = \frac{d|V|}{2} - d \left(\frac{|V| - k - |S_u|}{2} \right) = d \left(\frac{k + |S_u|}{2} \right) \geq \frac{kd}{2}.$$

□

Concluding Remarks. We conclude the paper with a remark about the linear program (\mathcal{P}). If we add the integrality constraints on the z variables, we obtain an integer program (IP) and it follows by our result that the integrality gap of the resulting IP is at most 2 for d -regular graphs. Könemann et al. [12] proved a $\Theta(n)$ -bound on the integrality gap of the IP for general graphs. However, the resulting IP is *not* a formulation for our minimum stabilizer problem, since the integral optimum solution of the IP could be $\Omega(n)$ away from the size of a minimum stabilizer for arbitrary graphs (not necessarily regular). In order to obtain a formulation for our stabilizer problem, we could introduce additional variables x and impose the existence of a matching in $G \setminus \text{Support}(z)$ of size $\nu(G)$:

$$\begin{aligned} & \min \sum_{e \in E} z_e \\ & y_u + y_v + z_{uv} \geq 1 \quad \forall uv \in E, \\ & \sum_{u \in V} y_u = \nu(G), \\ & x(\delta(v)) \leq 1 \quad \forall v \in V, \quad \sum_{e \in E} x_e = \nu(G), \quad x(E[S]) \leq \frac{|S| - 1}{2} \quad \forall S \subseteq V, |S| \text{ odd}, \\ & x_e + z_e \leq 1 \quad \forall e \in E, \\ & x, y, z \geq 0, \quad x, z \text{ integral}. \end{aligned}$$

However, we can show a lower bound of $\Omega(n)$ on the integrality gap of the above formulation. We refer the reader to the full version of the paper [7] for the example that exhibits the integrality gap.

References

- [1] A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev, *$O(\sqrt{\log n})$ approximation algorithms for min UnCut, min 2CNF deletion, and directed cut problems*, Proceedings, ACM Symposium on Theory of Computing, 2005, pp. 573–581.

- [2] Noga Alon, Asaf Shapira, and Benny Sudakov, *Additive approximation for edge-deletion problems*, Annals of mathematics **170** (2009), 371–411.
- [3] E. Balas, *Integer and fractional matchings*, Annals of Discrete Mathematics **11** (1981), 1–13.
- [4] M. L. Balinski, *On maximum matching, minimum covering and their connections*, Proceedings of the Princeton symposium on mathematical programming (1970).
- [5] M. Bateni, M. Hajiaghayi, N. Immorlica, and H. Mahini, *The cooperative game theory foundations of network bargaining games*, Proc. International Colloquium on Automata, Languages and Processing, 2010, pp. 67–78.
- [6] P. Biró, M. Bomhoff, P. A. Golovach, W. Kern, and D. Paulusma, *Solutions for the stable roommates problem with payments*, WG, 2012, pp. 69–80.
- [7] A. Bock, K. Chandrasekaran, J. Könemann, B. Peis, and L. Sanità, *Finding small stabilizers for unstable graphs*, Manuscript in preparation, 2014.
- [8] G. Chalkiadakis, E. Elkind, and M. Wooldridge, *Computational aspects of cooperative game theory*, Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers, 2011.
- [9] D. Hochbaum, *Approximation algorithms for the set covering and vertex cover problems*, SIAM Journal on Computing **11** (1982), no. 3, 555–556.
- [10] S. Khot, *On the power of unique 2-Prover 1-Round games*, Proceedings, ACM Symposium on Theory of Computing, 2002, pp. 767–775.
- [11] J. M. Kleinberg and É. Tardos, *Balanced outcomes in social exchange networks*, Proceedings, ACM Symposium on Theory of Computing, 2008, pp. 295–304.
- [12] J. Könemann, K. Larson, and D. Steiner, *Network bargaining: Using approximate blocking sets to stabilize unstable instances*, Proceedings, Symposium on Algorithmic Game Theory, 2012, pp. 216–226.
- [13] E. Korach, *Flowers and Trees in a Ballet of K_4 , or a Finite Basis Characterization of Non-König-Egerváry Graphs*, Tech. Report 115, IBM Israel Scientific Center, 1982.
- [14] E. Korach, T. Nguyen, and B. Peis, *Subgraph characterization of Red/Blue-Split Graph and König Egerváry Graphs*, Proceedings, ACM-SIAM Symposium on Discrete Algorithms, 2006, pp. 842–850.
- [15] S. Mishra, V. Raman, S. Saurabh, S. Sikdar, and C. R. Subramanian, *The complexity of König subgraph problems and above-guarantee vertex cover*, Algorithmica **61** (2011), no. 4, 857–881.
- [16] J. Nash, *The bargaining problem*, Econometrica **18** (1950), 155–162.
- [17] W. R. Pulleyblank, *Fractional matchings and the edmonds-gallai theorem*, Discrete Appl. Math. **16** (1987), 51–58.
- [18] A. Schrijver, *Combinatorial optimization*, Springer, New York, 2003.
- [19] L. S. Shapley and M. Shubik, *The assignment game: the core*, International Journal of Game Theory **1** (1971), no. 1, 111–130.
- [20] F. Sterboul, *A characterization of the graphs in which the transversal number equals the matching number*, J. Combin. Theory Ser. B (1979), 228–229.
- [21] J. P. Uhry, *Sur le problème du couplage maximal*, RAIRO **3** (1975), 13–20.
- [22] M. Yannakakis, *Node- and edge-deletion NP-complete problems*, Proceedings, ACM Symposium on Theory of Computing, 1978, pp. 253–264.