A polynomial kernel for FEEDBACK ARC SET on Bipartite Tournaments

Pranabendu Misra[1], Venkatesh Raman[2], M. S. Ramanujan[2] and Saket Saurabh[2]

[1] Chennai Mathematical Institute, Chennai, India.
pranabendu@cmi.ac.in
[2] The Institute of Mathematical Sciences, Chennai, India.
{vraman|msramanujan|saket}@imsc.res.in

Abstract. In the k-FEEDBACK ARC/VERTEX SET problem we are given a directed graph D and a positive integer k and the objective is to check whether it is possible to delete at most k arcs/vertices from D to make it acyclic. Dom et al.[*CIAC, 2006*] initiated a study of the FEEDBACK ARC SET problem on bipartite tournaments (k-FASBT) in the realm of parameterized complexity. They showed that k-FASBT can be solved in time $O(3.373^k n^6)$ on bipartite tournaments having n vertices. However, until now there was no known polynomial sized problem kernel for k-FASBT. In this paper we obtain a cubic vertex kernel for k-FASBT. This completes the kernelization picture for the FEEDBACK ARC/VERTEX SET problem on tournaments and bipartite tournaments, as for all other problems polynomial kernels were known before. We obtain our kernel using a non-trivial application of "independent modules" which could be of independent interest.

1 Introduction

In the k-FEEDBACK ARC/VERTEX SET problem we are given a directed graph D and a positive integer k and the objective is to check whether it is possible to delete at most k arcs/vertices from D to make it acyclic. These are classical NP-complete problems and have been extensively studied in every paradigm that deals with coping with NP-hardness including approximation and parameterized algorithms. These problems have several applications and in many of these applications we can restrict our inputs to more structured directed graphs like *tournaments* (orientation of a complete graph) and *bipartite tournaments* (orientation of a complete bipartite graph).

FEEDBACK ARC SET on tournaments is useful in *rank aggregation*. In *rank aggregation* we are given several rankings of a set of objects, and we wish to produce a single ranking that on average is as consistent as possible with the given ones, according to some chosen measure of consistency. This problem has been studied in the context of voting [9, 13], machine learning [12], and search engine ranking [17]. A natural consistency measure for rank aggregation is the number of pairs that occur in a different order in the two rankings. This leads to *Kemeny rank aggregation* [24, 25], a special case of a weighted version of k-FEEDBACK ARC SET on tournaments (k-FAST). Similarly, FEEDBACK ARC SET on bipartite tournaments finds its usefulness in applications that require establishment of mappings between "ontologies", we refer to [28] for further details.

In this paper, we consider k-FEEDBACK ARC SET problem on bipartite tournaments. More precisely the problem we study is the following:

k-FEEDBACK ARC SET IN BIPARTITE TOURNAMENTS (k-FASBT)
Input: A bipartite tournament $H = (X \cup Y, E)$ and a positive integer k.
Parameter: k
Question: Does there exist a subset $F \subseteq E$ of at most k arcs whose removal makes H acyclic?

In the last few years several algorithmic results have appeared around the k-FEEDBACK ARC SET problem on tournaments and bipartite tournaments. Speckenmeyer [29] showed that FEEDBACK VERTEX SET is NP-complete on tournaments. k-FAST was conjectured to be NP-complete for a long time [5], and only recently was this proven. Ailon et al. [2] gave a randomized reduction from FEEDBACK ARC SET on general directed graphs, which was independently derandomized by Alon [3] and Charbit et al. [11]. From approximation perspective, initially a constant factor approximation was obtained for k-FAST in [32] and later it was shown in [26] that it admits a polynomial time approximation scheme. Now we turn to problems on bipartite tournaments. Cai et al. [10] showed that FEEDBACK VERTEX SET on bipartite tournaments is NP-complete. They have also established a min-max theorem for feedback vertex set on bipartite tournaments. However, only recently Guo et al. [20] showed that k-FASBTis NP-complete. k-FASBT is also known to admit constant factor approximation algorithms [21, 33].

These problems are also well studied in parameterized complexity. In this area, a problem with input size n and a parameter k is said to be fixed parameter tractable (FPT) if there exists an algorithm to solve this problem in time $f(k) \cdot n^{O(1)}$, where f is an arbitrary function of k. Raman and Saurabh [27] showed that k-FAST is FPT by obtaining an algorithm running in time $O(2.415^k \cdot k^{4.752} + n^{O(1)})$. Recently, Alon et al. [4] have improved this result by giving an algorithm for k-FAST running in time $O(2^{O(\sqrt{k}\log^2 k)} + n^{O(1)})$. Moreover, a new algorithm due to Karpinsky and Schudy [23] with running time $O(2^{O(\sqrt{k})} + n^{O(1)})$ improves again the complexity of k-FAST. Dom et al. [16] obtained an algorithm with running time $O(3.373^k n^6)$ for k-FASBT based on a new forbidden subgraph characterization. In this paper we investigate k-FASBT from the view point of kernelization, currently one of the most active subfields of parameterized algorithms.

A parameterized problem is said to admit a *polynomial kernel* if there is a polynomial time algorithm, called a *kernelization* algorithm, that reduces the input instance to an instance whose size is bounded by a polynomial $p(k)$ in k, while preserving the answer. This reduced instance is called a $p(k)$ *kernel* for the problem. Kernelization has been at the forefront of research in parameterized complexity in the last couple of years, leading to various new polynomial kernels as well as tools to show that several problems do not have a polynomial kernel under some complexity-theoretic assumptions [6–8, 14, 15, 18, 31]. In this paper we continue the current theme of research on kernelization and obtain a *cubic vertex* kernel for k-FASBT. That is, we give a polynomial time algorithm which given an input instance (H, k) to k-FASBT obtains an equivalent instance (H', k') on $O(k^3)$ vertices. This completes the kernelization picture for the FEEDBACK ARC/VERTEX SET problem on tournaments and bipartite tournaments, as for all other problems polynomial kernels were known before. Feedback Vertex Set admits a kernel with $O(k^2)$ and $O(k^3)$ vertices on tournaments and bipartite tournaments, respectively [1, 4, 16]. However, only recently an $O(k)$ vertex kernel for k-FAST was obtained [6]. We obtain our kernel for k-FASBT using a data reduction rule that applies independent modules in a non trivial way. Previously, clique and transitive modules were used in obtaining kernels for Cluster Editing and k-FAST [6, 19].

2 Preliminaries

A bipartite tournament $H = (X \cup Y, E)$ is an orientation of a complete bipartite graph, meaning its vertex set is the union of two independent disjoint sets X and Y and for every pair of vertices $u \in X$ and $v \in Y$ there is exactly one arc between them. Some times we will use $V(H)$ to denote $X \cup Y$. For a vertex $v \in H$, we define $N^+(v) = \{w \in H | (v, w) \in E\}$ and $N^-(v) = \{u \in H | (u, v) \in E\}$. Essentially, the sets $N^+(v)$ and $N^-(v)$ are the set of out-neighbors and in-neighbors of v respectively. By C_4 we denote a directed cycle of length 4. Given a digraph $D = (V, E)$ and a subset W of either V or E, by $D \setminus W$ we denote the digraph obtained by deleting W from either V or E.

Given a directed graph $D = (V, E)$, a subset $F \subseteq E$ of arcs is called a feedback arc set *fas* if $D \setminus F$ is a directed acyclic graph. A feedback arc set F is called *minimal fas* if none of the proper subsets of F is a fas. Given a directed graph $D = (V, E)$ and a set F of arcs in E define $D\{F\}$ to be the directed graph obtained from D by reversing all arcs of F. In our arguments we will need the following folklore characterization of minimal feedback arc sets in directed graphs.

Proposition 1. *Let $D = (V, A)$ be a directed graph and F be a subset of A. Then F is a minimal feedback arc set of D if and only if $D\{F\}$ is a directed acyclic graph.*

3 Modular Partitions

A *Module* of a directed graph $D = (V, E)$ is a set $S \subseteq V$ such that $\forall u, v \in S$, $N^+(u) \setminus S = N^+(v) \setminus S$, and $N^-(u) \setminus S = N^-(v) \setminus S$. Essentially, every vertex in S has the same set of in-neighbors and out-neighbors outside S. The empty set and the whole of the vertex set are called *trivial modules.* We always mean non-trivial modules unless otherwise stated. Every vertex forms a singleton module. A *maximal module* is a module such that we cannot extend it by adding any vertex. The *modular partition*, $\mathcal{P}$, of a directed graph D, is a partition of the vertex set V into $(V_1, V_2, \ldots, V_\ell)$, such that every V_i is a maximal module. Now we look at some simple properties of the modules of a bipartite tournament.

Lemma 1. *Let $H = (X \cup Y, E)$ be a bipartite tournament, and S be any non-trivial module of H. Then S is an independent set. Thus $S \subset X$ or $S \subset Y$.*

Proof. If $|S| = 1$ then it is obvious. So we assume that $|S| \geq 2$. But S cannot contain two vertices x and y such that $x \in X$ and $y \in Y$ because their neighborhoods excluding

S are different. Hence, S contains vertices form only one of the partitions. Hence, S is an independent set. Now because H is a bipartite tournament we have that $S \subseteq X$ or $S \subseteq Y$. □

Consider the modular partition $\mathcal{P} = \mathbb{A} \cup \mathbb{B}$ of $X \cup Y$, where $\mathbb{A} = \{A_1, A_2, ...\}$ is a partition of X and $\mathbb{B} = \{B_1, B_2, ...\}$ is a partition of Y. Let $A \in \mathbb{A}$ and $B \in \mathbb{B}$ be two modules. Since A and B are modules, all the arcs between them are either directed from A to B (denoted by $A \to B$) or B to A (denoted by $B \to A$).

Lemma 2 ([30]). *For a bipartite tournament $H = (X \cup Y, E)$, a modular partition is unique and it can be computed in $O(|X \cup Y| + |E|)$.*

Now we define the well known notion of quotient graph of a modular partition.

Definition 1. *To a modular partition $\mathcal{P} = \{V_1, \dots, V_\ell\}$ of a directed graph D, we associate a quotient directed graph $\mathcal{D}_\mathcal{P}$, whose vertices $\{v_1, \dots, v_\ell\}$ are in one to one correspondence with the parts of $\mathcal{P}$. There is an arc (v_i, v_j) from vertex v_i to vertex v_j of $\mathcal{D}_\mathcal{P}$ if and only if $V_i \to V_j$. We denote its vertex set by $V(\mathcal{D}_\mathcal{P})$ and the edge set by $E(\mathcal{D}_\mathcal{P})$.*

We conclude this section with the observation that for a bipartite tournament H, the quotient graph $\mathcal{H}_\mathcal{P}$ corresponding to the modular partition $\mathcal{P}$, is a bipartite tournament. We refer to the recent survey of Habib and Paul [22] for further details and other algorithmic applications of modular decomposition.

4 Reductions and Kernel

In this section we show that k-FASBT admits a polynomial kernel with $O(k^3)$ vertices. We provide a set of reduction rules and assume that at each step we use the first possible applicable rule. After each reduction rule we discuss its soundness, that is, we prove that the input and output instances are equivalent. If no reduction rule can be used on an instance (H, k), we claim that $|V(H)|$ is bounded by $O(k^3)$. Through out this paper, whenever we say cycle, we mean directed cycle; whenever we speak of a fas, we mean a *minimal fas*, unless otherwise stated.

4.1 Data Reduction Rules

We start with some simple observations regarding the structure of directed cycles in bipartite tournaments.

Lemma 3 ([16]). *A bipartite tournament H has a cycle C if and only if it has a C_4.*

Lemma 4. [⋆][3] *Let $H = (V, E)$ be a bipartite tournament. If $v \in V$ is part of some cycle C, then it is also part of some C_4.*

We now have the following simple rule.

[3] Proofs of results labelled with ⋆ have been omitted due to lack of space.

Reduction Rule 1 *If $v \in V$ is not part of any cycle in H then remove v, that is, return an instance $(H \setminus v, k)$.*

Lemma 5. [$\star$] *Reduction Rule 1 is sound and can be applied in polynomial time.*

Reduction Rule 2 *If there is an arc $e \in E$, such that there are at least $k+1$ C_4's, which pairwise have only e in common, then reverse e and reduce k by 1. That is, return the instance $(H \setminus \{e\}, k-1)$.*

Lemma 6. *Reduction Rule 2 is sound and can be applied in polynomial time.*

Proof. Such an arc e must be in any fas of size $\leq k$. Otherwise, we have to pick at least one distinct arc for each of the cycles, and there are $\geq k+1$ of them.

To find such arcs, we use the idea of an *opposite arc*. Two vertex disjoint arcs $e_1 = (a, b)$ and $e_2 = (c, d)$ are called opposite arcs if (a, b, c, d) forms a C_4. Consider the set of opposite arcs of an arc e; this can be easily computed in polynomial time. It is easy to see that an arc e has $\geq k+1$ vertex disjoint opposite arcs if and only it there are $\geq k+1$ C_4 which pairwise have only e in common. Hence, we only need to check if the size of the opposite set, for any arc e, is larger than k and reverse e if that is the case. $\square$

Let $H = (X \cup Y, E)$ be a bipartite tournament. From now onwards we fix the unique modular partition $\mathcal{P} = \mathbb{A} \cup \mathbb{B}$ of $X \cup Y$, where $\mathbb{A} = \{A_1, A_2, ...\}$ is a partition of X and $\mathbb{B} = \{B_1, B_2, ...\}$ is a partition of Y. Next we show how a C_4 interacts with the modular partition.

Lemma 7. *Let H be a bipartite tournament, then any C_4 in H has each of its vertices in different modules of $\mathcal{P}$.*

Proof. Let u and v be any two vertices of a C_4 in H. If u and v are from different partitions of H then by Lemma 1, they cannot be in the same module. And if they are from the same partition, then there is some vertex w from the other partition that comes between them in the cycle as $u \to w \to v$, and hence in the outgoing neighbourhood of one but the incoming neighbourhood of the other. So they cannot be in the same module. $\square$

The main reduction rule that enables us to obtain an $O(k^3)$ kernel is based on the following crucial lemma. It states that there exists an optimum solution where all arcs between two modules are either a part of the solution or none of them are.

Lemma 8. *Let X_1 and Y_1 be two modules of a bipartite tournament $H(X \cup Y, E)$ such that $X_1 \subseteq X$ and $Y_1 \subseteq Y$ and let $E(X_1, Y_1)$ be the set of arcs between these two modules. Let F be any minimal fas of H. Then there exists an fas F^* such that $|F^*| \leq |F|$ and $E(X_1, Y_1) \subseteq F^*$ or $E(X_1, Y_1) \cap F^* = \phi$.*

Proof. Let $X_1 = \{x_1, x_2, ...x_r\}$ and $Y_1 = \{y_1, y_2, ...y_s\}$. Let $e = (x_1, y_1)$ be an arc of $E(X_1, Y_1)$. We define, what we call a mirroring operation with respect to the arc e, that produces a solution F' that contains all the arcs of the module if e was in F and does not contain any arc of the module if e was not in F; i.e. the mirroring operation mirrors the intersection of the solution F with the arc e, and the arcs incident on the end points of e, to all arcs of the module. To obtain F^* we will mirror that arc of the module that has the smallest intersection with F among the arcs incident on its end points. We define this operation formally below.

The operation $mirror(F, e, E(X_1, Y_1))$ returns a subset $F'(e)$ of arcs obtained as follows. Let F_{XY} is the set of all arcs in F which have at least one end point in $X_1 \cup Y_1$. Let $e \in E(X_1, Y_1)$ and define $Ext(e) = \{f \in F_{XY} | f \cap e \neq \phi, f \notin E(X_1, Y_1)\}$, i.e. $Ext(e)$ is the set of all external edges in F_{XY} which are incident on endpoints of e. Let $Ext = \cup_{i=1}^{rs} Ext(e_i)$.

$$
\begin{aligned}
F_1(e) &= \{(x_i, y) | (x_1, y) \in Ext(e)\} \cup \{(y_i, x) | (y_1, x) \in Ext(e)\} \\
&\quad \cup \{(y, x_i) | (y, x_1) \in Ext(e)\} \cup \{(x, y_i) | (x, y_1) \in Ext(e)\}, \\
F_2(e) &= \begin{cases} E(X_1, Y_1) & e \in F \\ \phi & e \notin F \end{cases}
\end{aligned}
$$

For an edge e define, $F'(e) = F - F_{XY} \cup F_1(e) \cup F_2(e)$. We claim that $F'(e)$ for any arc $e \in E(X_1, Y_1)$ is a feedback arc set, and that there exists an $e \in E(X_1, Y_1)$ such that $|F'(e)| \leq |F|$, and we will output F^* as $F'(e)$ for that e.

Claim 1: $F'(e)$ is a feedback arc set for any arc $e \in E(X_1, Y_1)$.

Proof. Suppose not, and let C be a cycle in the graph $H\{F'(e)\}$, and let $e = (x_s, y_t)$ and replace any vertex $x_j, j \neq s$ by x_s and any vertex $y_j, j \neq y_t$ by y_t in the cycle C to get C' (See Fig. 1). We claim that C' is a closed walk (from which a cycle can be obtained), in $H\{F\}$ contradicting the fact that F was a fas for H.

That C' is a closed walk is clear because for any arc (u, x_j) or (x_j, u) incident on $x_j, j \neq s$, the corresponding arcs (u, x_s) or (x_s, u) exists as X_1 is a module. Similarly for any arc (u, y_j) or (y_j, u) incident on $y_j, j \neq t$, the corresponding arcs (u, y_t) or (y_t, u) exists as Y_1 is a module.

Note that C' does not contain arcs incident on $x_j, j \neq s$ or $y_j, j \neq t$ and the only arcs affected by the mirroring operation to obtain F' are the arcs incident on $x_j, j \neq s$ or $y_j, j \neq t$. Hence C' is a walk in $H\{F\}$. This proves Claim 1. □

Claim 2: There exists an arc $e \in E(X_1, Y_1)$ such that $|F'(e)| \leq |F|$.

Proof. Let $e_1, e_2, ...e_{rs}$ be the arcs in $E(X_1, Y_1)$. By the definition of F', we have

$$
\sum_{i=1}^{rs} |F'(e_i)| = (rs)|F| - (rs)|F_{XY}| + \sum_{i=1}^{rs} |F_1(e_i)| + \sum_{i=1}^{rs} |F_2(e_i)|.
$$

We argue that the last three terms cancel out. For $x \in X_1$, define $Ext(x) = \{f \in Ext | f \cap x \neq \phi\}$. Whenever an edge incident on x is mirrored by other arcs, $Ext(x)$ is repeated on each of the r vertices in X_1 and there are s edges in $E(X_1, Y_1)$ which are incident on x. Therefore $Ext(x)$ is repeated (rs) times. Similarly we can define $Ext(y)$ and show that it is repeated (rs) times. Therefore the third term $\sum_{i=1}^{rs} |F_1(e_i)| =$

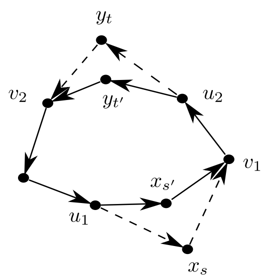

Fig. 1. Constructing a closed walk from the cycle C.

$rs|Ext|$. And whenever e is in F, it is repeated (rs) times, therefore the last term $\sum_{i=1}^{rs} |F_2(e_i)| = rs|F_{XY} - Ext|$. Therefore the sum of the last two terms is nothing but (rs) times $|F_{XY}|$.

So we have $\sum_{i=1}^{rs} |F'(e_i)| = (rs)|F|$. Thus there is an arc $e \in E(X_1, Y_1)$ such that $|F'(e)| \leq |F|$. Note that to find such an i, we can simply find that i that has the smallest intersection between F and the arcs incident on its end points. □

This completes the proof of Lemma 8. □

From this point, we will only consider fas which either contains all the arcs between two modules or contains none of them. An easy consequence of this lemma is that if S is a module of size at least $k+1$, then a fas of size at most k contains no arc incident on S. This is because between any other module R and S there are at least $k+1$ arcs and they all cannot be in the fas.

Reduction Rule 3 *In* $\mathbb{A} \cup \mathbb{B}$ *truncate all modules of size greater than* $k+1$ *to* $k+1$. *In other words if* S *is a module of size more than* $k+1$ *then delete all but* $k+1$ *of its vertices arbitrarily.*

Lemma 9. *Reduction Rule 3 is sound.*

Proof. Once we have a modular partition, we can see which modules have size $\geq k+1$ and arbitrarily delete vertices from them to reduce their size to $k+1$.

Next we show that the rule is correct. Suppose H is a bipartite tournament and H' is obtained from H by deleting a single vertex v from some large module S which has $\geq k+2$ vertices. Consider a fas F' of H' of size $\leq k$ and let $S' = S \setminus \{v\}$, which is a module in H'. Since $|S'| \geq k+1$, no arc incident on S' is in F'. Now reverse the arcs of F' in H. If F' is not a fas of H then there is a cycle C of length 4 in H. If C does not contain v then it will be present in H' with F' reversed, which is a contradiction. Otherwise the C contains v. Let $u \in S'$ and consider the cycle C' of length 4 obtained from C by replacing v with u. C' is present in H' with F' reversed, which is again a contradiction.

Now we can use induction on the number of vertices deleted from a module, to show that truncating a single module does not change the solution. We can then use induction on the number of truncated modules to show that the rule is correct. □

4.2 Analysis of Kernel Size

We apply the above Reduction Rules 1, 2 and 3 exhaustively (until no longer possible) and obtain a reduced instance (H', k') of k-FASBT from the initial instance (H, k). Observe that we will not keep applying these rules indefinitely to an instance because either the size of the graph or the parameter k drops. For *brevity we abuse notation* and denote the reduced instance also by (H, k). As before we fix the unique modular partition $\mathcal{P} = \mathbb{A} \cup \mathbb{B}$ of $X \cup Y$, where $\mathbb{A} = \{A_1, A_2, ...\}$ is a partition of X and $\mathbb{B} = \{B_1, B_2, ...\}$ is a partition of Y. Let $\mathcal{H}_\mathcal{P}$ be the corresponding quotient graph.

The following lemma lists some properties of the quotient graph.

Lemma 10. [$\star$] *Let $\mathcal{H}_\mathcal{P}$ be the quotient graph of H. Then the following hold:*

(a) Every vertex in $\mathcal{H}_\mathcal{P}$ is part of some C_4.
(b) For each arc $e \in E(\mathcal{H}_\mathcal{P})$ there are $\leq k$ C_4 in $\mathcal{H}_\mathcal{P}$ which pairwise have e as the only is the common arc.
(c) If H has a fas of size $\leq k$ then $\mathcal{H}_\mathcal{P}$ has a fas of size $\leq k$.
(d) For all $u, v \in V(\mathcal{H}_\mathcal{P})$, $N^+(u) \neq N^+(v)$ or equivalently $N^-(u) \neq N^-(v)$.

Let $\mathcal{H}_\mathcal{P}$ be the quotient graph of H. Let F be an fas of $\mathcal{H}_\mathcal{P}$ of size at most k. Let T be the topological sort of $\mathcal{H}_\mathcal{P}$ obtained after reversing the arcs of F. If we arrange the vertices of $\mathcal{H}_\mathcal{P}$ from left to right in order of T, then only the arcs of F go from right to left. For an arc $e \in F$, the span of e is the set of vertices which lie between the endpoints of e in T. We call a vertex v an *affected vertex* if there is some arc in F which is incident on v. Otherwise we call v an *unaffected vertex.*

Lemma 11. *If v is an unaffected vertex, then there exists an arc $e \in F$ such that v is in the span of e.*

Proof. Let v be any unaffected vertex in $\mathcal{H}_\mathcal{P}$. By lemma 10(a) there is a C_4 u, v, w, t which contains v. Since v is unaffected the order $u \to v \to w$ is fixed in T. Therefore the arc $t \to u$ goes from right to left in T. Hence, v is contained in the span of (t, u). □

Lemma 12. *Let u and v be two unaffected vertices from the same vertex partition of the bipartite tournament $\mathcal{H}_\mathcal{P}$ such that u occurs before v in T. Then there exists w from the other partition which lies between u and v in T.*

Proof. Since $\mathcal{H}_\mathcal{P}$ is the quotient graph of H and u and v are different modules, therefore there is some vertex w in the other partition such that $(u, w), (w, v) \in E(\mathcal{H}_\mathcal{P})$. And these two arcs are not in the fas F as u and v are unaffected. Hence, w comes between u and v in T. □

Lemma 13. *There are at most $2k+2$ unaffected vertices in the span of any arc $e \in F$.*

Proof. Let the span of e contain $2k + 3$ unaffected vertices. Then without loss of generality assume that $k + 2$ of these vertices come from the partition A. Then there are at least $k + 1$ vertices of partition B which lie between each pair of consecutive vertices of A by Lemma 12. This gives us $k + 1$ C_4's which pairwise have only e in common. But then the graph H contains an arc e' for which there are $k+1$ C_4's which pairwise have only e' in common. This contradicts the fact that H was reduced with respect to Rule 2. □

Reduction Rule 4 *If $\mathcal{H}_{\mathcal{P}}$ contains more than $2k^2 + 2k$ vertices then return NO.*

Lemma 14. *Reduction rule 4 is sound.*

Proof. By Lemma 10(c), if H has a fas of size $\leq k$, then $\mathcal{H}_{\mathcal{P}}$ has a fas of size $\leq k$. If there are more than $2k^2 + 2k$ vertices in $\mathcal{H}_{\mathcal{P}}$, then there is some arc e whose span contains more than $2k + 2$ unaffected vertices, this contradicts Lemma 13. □

Hence, we may assume that $\mathcal{H}_{\mathcal{P}}$ contains $O(k^2)$ vertices.

Lemma 15. *H contains $O(k^3)$ vertices.*

Proof. Each vertex in $\mathcal{H}_{\mathcal{P}}$ corresponds to a module in H and any module in H has size at most $k + 1$ (due to Reduction Rule 3). Hence, there are $O(k^3)$ vertices in H. □

Lemma 15 implies the following theorem.

Theorem 1. k-FASBT *admits a polynomial kernel with $O(k^3)$ vertices.*

5 Conclusion

In this paper we give a polynomial kernel for k-FASBT with $O(k^3)$ vertices. It is an interesting open problem to obtain a kernel with $O(k)$ or even $O(k^2)$ vertices. Our result adds to a small list of problems for which modular based approach has turned out to be useful. It will be interesting to find more applications of graph modules for kernelization. We also remark that the kernel obtained in the paper generalizes to multi-partite tournaments.

References

1. Abu-Khzam, F.N.: A kernelization algorithm for d-hitting set. Journal of Computer and System Sciences 76(7), 524 – 531 (2010).
2. Ailon, N., Charikar, M., Newman, A.: Aggregating inconsistent information: ranking and clustering. In: ACM Symposium on Theory of Computing (STOC). pp. 684–693 (2005).
3. Alon, N.: Ranking tournaments. SIAM J. Discrete Math. 20(1), 137–142 (2006).
4. Alon, N., Lokshtanov, D., Saurabh, S.: Fast FAST. In: ICALP. LNCS, vol. 5555, pp. 49–58. Springer (2009).
5. Bang-Jensen, J., Thomassen, C.: A polynomial algorithm for the 2-path problem for semicomplete digraphs. SIAM J. Discrete Math. 5(3), 366–376 (1992).
6. Bessy, S., Fomin, F.V., Gaspers, S., Paul, C., Perez, A., Saurabh, S., Thomassé, S.: Kernels for feedback arc set in tournaments. In: FSTTCS. pp. 37–47 (2009).
7. Bodlaender, H.L., Downey, R.G., Fellows, M.R., Hermelin, D.: On problems without polynomial kernels. J. Comput. Syst. Sci 75(8), 423–434 (2009).
8. Bodlaender, H.L., Fomin, F.V., Lokshtanov, D., Penninkx, E., Saurabh, S., Thilikos, D.M.: (Meta) Kernelization. In: FOCS. pp. 629–638 (2009).
9. Borda, J.: Mémoire sur les élections au scrutin. Histoire de l'Académie Royale des Sciences (1781).
10. Cheng Cai, M., Deng, X., Zang, W.: A min-max theorem on feedback vertex sets. Math. Oper. Res. 27(2), 361–371 (2002).

11. Charbit, P., Thomassé, S., Yeo, A.: The minimum feedback arc set problem is NP-hard for tournaments. Combin. Probab. Comput. 16(1), 1–4 (2007).
12. Cohen, W.W., Schapire, R.E., Singer, Y.: Learning to order things. In: Advances in neural information processing systems (NIPS). pp. 451–457 (1997).
13. Condorcet, M.: Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix (1785).
14. Dell, H., van Melkebeek, D.: Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. In: STOC. pp. 251–260. ACM (2010).
15. Dom, M., Lokshtanov, D., Saurabh, S.: Incompressibility through colors and IDs. In: ICALP. LNCS, vol. 5555, pp. 378–389. Springer (2009).
16. Dom, M., Guo, J., Hüffner, F., Niedermeier, R., Truß, A.: Fixed-parameter tractability results for feedback set problems in tournaments. J. Discrete Algorithms 8(1), 76–86 (2010).
17. Dwork, C., Kumar, R., Naor, M., Sivakumar, D.: Rank aggregation methods for the web. In: World Wide Web Conference (WWW) (2001).
18. Fomin, F.V., Lokshtanov, D., Saurabh, S., Thilikos, D.M.: Bidimensionality and kernels. In: SODA. pp. 503–510 (2010).
19. Guo, J.: A more effective linear kernelization for cluster editing. Theor. Comput. Sci 410(8-10), 718–726 (2009).
20. Guo, J., Hüffner, F., Moser, H.: Feedback arc set in bipartite tournaments is NP-Complete. Inf. Process. Lett. 102(2-3), 62–65 (2007).
21. Gupta, S.: Feedback arc set problem in bipartite tournaments. Inf. Process. Lett. 105(4), 150–154 (2008).
22. Habib, M., Paul, C.: A survey of the algorithmic aspects of modular decomposition. Computer Science Review 4(1), 41–59 (2010).
23. Karpinski, M., Schudy, W.: Faster algorithms for feedback arc set tournament, kemeny rank aggregation and betweenness tournament. CoRR abs/1006.4396 (2010).
24. Kemeny, J.: Mathematics without numbers. Daedalus 88, 571–591 (1959).
25. Kemeny, J., Snell, J.: Mathematical models in the social sciences. Blaisdell (1962).
26. Kenyon-Mathieu, C., Schudy, W.: How to rank with few errors. In: ACM Symposium on Theory of Computing (STOC). pp. 95–103 (2007).
27. Raman, V., Saurabh, S.: Parameterized algorithms for feedback set problems and their duals in tournaments. Theor. Comput. Sci 351(3), 446–458 (2006).
28. Sanghvi, B., Koul, N., Honavar, V.: Identifying and eliminating inconsistencies in mappings across hierarchical ontologies. In: OTM Conferences (2). Lecture Notes in Computer Science, vol. 6427, pp. 999–1008 (2010).
29. Speckenmeyer, E.: On feedback problems in digraphs. In: Workshop on Graph-Theoretic Concepts in Computer Science (WG). Lecture Notes in Computer Science, vol. 411, pp. 218–231. Springer (1989).
30. Tedder, M., Corneil, D.G., Habib, M., Paul, C.: Simpler linear-time modular decomposition via recursive factorizing permutations. In: ICALP (1). pp. 634–645. Lecture Notes in Computer Science, Springer (2008).
31. Thomassé, S.: A $4k^2$ kernel for feedback vertex set. ACM Transactions on Algorithms 6(2) (2010).
32. van Zuylen, A., Hegde, R., Jain, K., Williamson, D.P.: Deterministic pivoting algorithms for constrained ranking and clustering problems. In: ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 405–414 (2007).
33. van Zuylen, A.: Linear programming based approximation algorithms for feedback set problems in bipartite tournaments. Theor. Comput. Sci. 412(23), 2556–2561 (2011).

6 Appendix: Omitted Proofs

Proof of Lemma 4

Proof. We prove the statement of the lemma by induction on the length of C. If C is a C_4, then there is nothing to prove. Suppose the length of C is longer than 4 and assume that the statement of the lemma holds for all cycles shorter than C. Let $\dots v, a, b, c \dots$ be a subpath of C. If the arc between c and v is (c, v) then, v, a, b, c is a C_4 containing v and if it is (v, c), then we have a shorter cycle which contains v and excludes a and b. By the induction hypothesis, there is a C_4 which contains v. □

Proof of Lemma 5

Proof. Let $H' = H \setminus v$ and F' be a fas of H'. Observe that, deleting v doesn't affect any cycle in H and, every cycle of H is present in H'.

Suppose F' is not an fas of H. Hence, after reversing the arcs of F' in H, there is still a cycle left. But this cycle does not involve v, so it is also present in H' with F' reversed. Therefore there is a cycle in H', which is not removed by F'. Hence, F' is not a fas of H', which is a contradiction.

And by lemma 4, to determine if a vertex is part of any cycle, we only need to check if v is part of some C_4. This can easily be done in polynomial time. □

Proof of Lemma 10

Proof. (a) Consider a vertex $u \in \mathcal{H}_{\mathcal{P}}$. Let S be the corresponding module in H. Consider a vertex $x \in S$. If x is not part of any cycle then rule 1 applies, which is a contradiction. Otherwise x is part of some cycle in H and by lemma 4 it is part of some C_4 x, y, z, p in H. By lemma 7, each of x, y, z, p lies in a different module. Let v, w, t be vertices in $\mathcal{H}_{\mathcal{P}}$ corresponding to the modules containing y, z, p respectively. Then by definition of $E(\mathcal{H}_{\mathcal{P}})$, (u, v, w, t) form a C_4 in $\mathcal{H}_{\mathcal{P}}$.

(b) Suppose there were an arc $e \in E(\mathcal{H}_{\mathcal{P}})$ which is the only pairwise common arc of $\geq k+1$ C_4 in $\mathcal{H}_{\mathcal{P}}$. Consider such a collection of C_4 in $\mathcal{H}_{\mathcal{P}}$. For each vertex in this collection, there is a module in H. Pick one vertex from each module and the arcs between these vertices in H. Let $e' \in H$ be the arc corresponding to e. This gives us a collection of $\geq k+1$ C_4 in H, which have e' in common pairwise. But then Rule 2 applies.

(c) Let F' be any fas of H. Therefore $|F'| \leq k$. Let F be the corresponding set of arcs in $\mathcal{H}_{\mathcal{P}}$, so $|F| \leq k$. Every arc $(u, v) \in F$ corresponds to the set of all arcs between the modules S_u and S_v in F'. We reverse F in $\mathcal{H}_{\mathcal{P}}$ and F' in H respectively. If F is not a fas of $\mathcal{H}_{\mathcal{P}}$, then there is a $C_4 = (u, v, w, t) \in \mathcal{H}_{\mathcal{P}}$ with F reversed. Then consider the vertices $x, y, z, p \in H$ where $x \in S_u, y \in S_v, z \in S_w, p \in S_t$. By Lemma 8 the arcs $(u, v), (v, w), (w, t), (t, u)$ in $\mathcal{H}_{\mathcal{P}}$ with F reversed, imply the arcs $(x, y), (y, z), (z, p), (p, x)$ in H with F' reversed. This implies that F' is not a fas of H which is a contradiction.

(d) $\mathcal{H}_{\mathcal{P}}$ is a Bipartite Tournament. Hence, if for some $u, v \in \mathcal{H}_{\mathcal{P}}$, $N^+(u) = N^+(v) \implies N^-(u) = N^-(v)$. But then S_u is not a maximal module, because we can add any vertex from S_v to it. This is a contradiction to the fact that $\mathbb{A} \cup \mathbb{B}$ was a maximal modular partition. □